



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

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: Ethernet Switching

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 xi

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

Part 1

Layer 2 Bridging and VLANs

Chapter 1

Bridging and VLANs—Overview 3

Understanding Bridging and VLANs on EX Series Switches	3
Ethernet LANs, Transparent Bridging, and VLANs	3
How Bridging Works	4
Types of Switch Ports	6
IEEE 802.1Q Encapsulation and Tags	6
Assignment of Traffic to VLANs	6
Ethernet Switching Tables	7
Layer 2 and Layer 3 Forwarding of VLAN Traffic	7
GVRP and MVRP	7
Routed VLAN Interface	8
Understanding Private VLANs on EX Series Switches	9
Understanding Virtual Routing Instances on EX Series Switches	10
Understanding Redundant Trunk Links on EX Series Switches	11
Understanding Q-in-Q Tunneling on EX Series Switches	13
How Q-in-Q Tunneling Works	14
Disabling MAC Address Learning	14
Mapping C-VLANs to S-VLANs	15
All-in-One Bundling	15
Many-to-One Bundling	15
Mapping a Specific Interface	15
Routed VLAN Interfaces on Q-in-Q VLANs	16
Limitations for Q-in-Q Tunneling	16

Chapter 4	Verifying Bridging and VLAN Configuration	115
	Verifying That a Series of Tagged VLANs Has Been Created	115
	Verifying That Virtual Routing Instances Are Working	117
	Verifying That Q-in-Q Tunneling Is Working	118
	Verifying That a Private VLAN Is Working	118
	Monitoring Ethernet Switching	120
	Monitoring GVRP	121
	Verifying That MVRP Is Working Correctly	122
Chapter 5	Troubleshooting Bridging and VLAN Configuration	125
	Troubleshooting Ethernet Switching	125
	MAC Address in the Switch's Ethernet Switching Table Is Not Updated After a MAC Address Move	125
Chapter 6	Configuration Statements for Bridging and VLANs	127
	[edit ethernet-switching-options] Configuration Statement Hierarchy	127
	[edit interfaces] Configuration Statement Hierarchy	129
	[edit protocols] Configuration Statement Hierarchy	130
	[edit routing-instances] Configuration Statement Hierarchy	136
	[edit vlans] Configuration Statement Hierarchy	136
	arp	138
	bridge-priority	139
	customer-vlans	140
	description	141
	disable	141
	disable (MVRP)	142
	dot1q-tunneling (Ethernet Switching)	142
	dot1q-tunneling (VLANs)	143
	drop-threshold	144
	ether-type	145
	ethernet-switching-options	146
	filter	149
	group-name	150
	gvrp	151
	instance-type	152
	interface	152
	interface (MVRP)	153
	interface	154
	interface	154
	interface	155
	interfaces	156
	join-timer	156
	join-timer (MVRP)	157
	l3-interface	158
	layer2-protocol-tunneling	159

leave-timer	160
leave-timer (MVRP)	161
leaveall-timer	162
leaveall-timer (MVRP)	163
mac-limit	164
mac-table-aging-time	165
members	166
mvrp	167
native-vlan-id	168
no-dynamic-vlan	169
no-local-switching	169
no-mac-learning	170
no-mac-learning	171
port-mode	171
primary-vlan	172
redundant-trunk-group	173
registration	174
routing-instances	174
shutdown-threshold	175
vlan	176
vlan-id	177
vlan-range	177
vlangs	178

Chapter 7

Operational Mode Commands for Bridging and VLANs 181

clear ethernet-switching layer2-protocol-tunneling error	182
clear ethernet-switching layer2-protocol-tunneling statistics	183
clear gvrp statistics	184
clear mvrp statistics	185
show ethernet-switching interfaces	186
show ethernet-switching layer2-protocol-tunneling interface	189
show ethernet-switching layer2-protocol-tunneling statistics	191
show ethernet-switching layer2-protocol-tunneling vlan	194
show ethernet-switching mac-learning-log	196
show ethernet-switching statistics aging	198
show ethernet-switching statistics mac-learning	200
show ethernet-switching table	203
show gvrp	209
show gvrp statistics	211
show mvrp	213
show mvrp dynamic-vlan-memberships	215
show mvrp statistics	216
show redundant-trunk-group	218
show vlangs	219

About This Topic Collection

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

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](http://www.juniper.net/support/requesting_support.html) .

Part 1

Layer 2 Bridging and VLANs

- Bridging and VLANs—Overview on page 3
- Examples: Bridging and VLAN Configuration on page 21
- Configuring Bridging and VLANs on page 95
- Verifying Bridging and VLAN Configuration on page 115
- Troubleshooting Bridging and VLAN Configuration on page 125
- Configuration Statements for Bridging and VLANs on page 127
- Operational Mode Commands for Bridging and VLANs on page 181

Chapter 1

Bridging and VLANs—Overview

- Understanding Bridging and VLANs on EX Series Switches on page 3
- Understanding Private VLANs on EX Series Switches on page 9
- Understanding Virtual Routing Instances on EX Series Switches on page 10
- Understanding Redundant Trunk Links on EX Series Switches on page 11
- Understanding Q-in-Q Tunneling on EX Series Switches on page 13
- Understanding Multiple VLAN Registration Protocol (MVRP) on EX Series Switches on page 17
- Understanding Layer 2 Protocol Tunneling on EX Series Switches on page 19

Understanding Bridging and VLANs on EX Series Switches

Network switches use Layer 2 bridging protocols to discover the topology of their LAN and to forward traffic toward destinations on the LAN.

This topic explains the following concepts regarding bridging and VLANs on Juniper Networks EX Series Ethernet Switches:

- Ethernet LANs, Transparent Bridging, and VLANs on page 3
- How Bridging Works on page 4
- Types of Switch Ports on page 6
- IEEE 802.1Q Encapsulation and Tags on page 6
- Assignment of Traffic to VLANs on page 6
- Ethernet Switching Tables on page 7
- Layer 2 and Layer 3 Forwarding of VLAN Traffic on page 7
- GVRP and MVRP on page 7
- Routed VLAN Interface on page 8

Ethernet LANs, Transparent Bridging, and VLANs

Ethernet is a data link layer technology, as defined by Layer 2 of the Open Systems Interconnection (OSI) model of communications protocols. Ethernet was first standardized by the IEEE in 1982, in IEEE 802.3. Ethernet is used to create LANs. The network devices, called *nodes*, on the LAN transmit data in bundles that are generally called frames or packets.

Each node on a LAN has a unique identifier so that it can be unambiguously located on the network. Ethernet uses the Layer 2 media access control (MAC) address for this purpose. MAC addresses are hardware addresses that are programmed (“burned”) into the Ethernet processor in the node.

A characteristic of Ethernet is that nodes on a LAN can transmit data frames at any time. However, the physical connecting cable between the nodes—either coaxial, copper-based (Category 5), or optical cable—can carry only a single stream of data at a time. One result of this design is that when two nodes transmit at the same time, their frames can collide on the cable and generate an error. Ethernet uses a protocol called carrier-sense multiple access with collision detection (CSMA/CD) to detect frame collisions. If a node receives a collision error message, it stops transmitting immediately and waits for a period of time before trying to send the frame again. If the node continues to detect collisions, it progressively increases the time between retransmissions in an attempt to find a time when no other data is being transmitted on the LAN. The node uses a backoff algorithm to calculate the increasing retransmission time intervals.

Ethernet LANs were originally implemented for small, simple networks that carried primarily text. Over time, LANs have become larger and more complex; the type of data they carry has grown to include voice, graphics, and video; and the increased speed of Ethernet interfaces on LANs has resulted in exponential increases in traffic on the network.

The IEEE 802.1D-2004 standard addresses some of the problems caused by the increase in LAN and complexity. This standard defines *transparent bridging* (generally called simply bridging). Bridging divides a single physical LAN (a single *broadcast domain*) into two or more virtual LANs, or VLANs. Each *VLAN* is a collection of network nodes that are grouped together to form separate broadcast domains. On an Ethernet network that is a single LAN, all traffic is forwarded to all nodes on the LAN. On VLANs, frames whose origin and destination are in the same VLAN are forwarded only within the local VLAN. Frames that are not destined for the local VLAN are the only ones forwarded to other broadcast domains. VLANs thus limit the amount of traffic flowing across the entire LAN, reducing the possible number of collisions and packet retransmissions within a VLAN and on the LAN as a whole.

On an Ethernet LAN, all network nodes must be physically connected to the same network. On VLANs, the physical location of the nodes is not important, so you can group network devices in any way that makes sense for your organization, such as by department or business function, types of network nodes, or even physical location. Each VLAN is identified by a single IP subnetwork and by standardized IEEE 802.1Q encapsulation (discussed below).

How Bridging Works

The transparent bridging protocol allows a switch to learn information about all the nodes on the LAN, including nodes on all the different VLANs. The switch uses this information to create address-lookup tables, called *Ethernet switching tables* that it consults when forwarding traffic to or toward a destination on the LAN.

Transparent bridging uses five mechanisms to create and maintain Ethernet switching tables on the switch:

- Learning
- Forwarding
- Flooding
- Filtering
- Aging

The first bridging mechanism is *learning*. When a switch is first connected to an Ethernet LAN or VLAN, it has no information about other nodes on the network. The switch goes through a learning process to obtain the MAC addresses of all the nodes on the network. It stores these in the Ethernet switching table. To learn MAC addresses, the switch reads all packets that it detects on the LAN or on the local VLAN, looking for MAC addresses of sending nodes. It places these addresses into its Ethernet switching table, along with two other pieces of information—the interface (or port) on which the traffic was received and the time when the address was learned.

The second bridging mechanism is *forwarding*. Switches forward traffic, passing it from an incoming interface to an outgoing interface that leads to or toward the destination. To forward frames, the switch consults the Ethernet switching table to see whether the table contains the MAC address corresponding to the frames' destination. If the Ethernet switching table contains an entry for the desired destination address, the switch sends the traffic out the interface associated with the MAC address. The switch also consults the Ethernet switching table in the same way when transmitting frames that originate on devices connected directly to the switch. If the Ethernet switching table does not contain an entry for the desired destination address, the switch uses flooding, which is the third bridging mechanism.

Flooding is how the switch learns about destinations not in its Ethernet switching table. If this table has no entry for a particular destination MAC address, the switch floods the traffic out all interfaces except the interface on which it was received. (If traffic originates on the switch, the switch floods it out all interfaces.) When the destination node receives the flooded traffic, it sends an acknowledgment packet back to the switch, allowing it to learn the MAC address of the node and to add the address to its Ethernet switching table.

Filtering, the fourth bridging mechanism, is how broadcast traffic is limited to the local VLAN whenever possible. As the number of entries in the Ethernet switching table grows, the switch pieces together an increasingly complete picture of the VLAN and the larger LAN—of which nodes are in the local VLAN and which are on other network segments. The switch uses this information to filter traffic. Specifically, for traffic whose source and destination MAC addresses are in the local VLAN, filtering prevents the switch from forwarding this traffic to other network segments.

Finally, the switch uses *aging*, the fifth bridging mechanism, to keep the entries in the Ethernet switching table current. For each MAC address in the Ethernet switching table, the switch records a timestamp of when the information about the network node was learned. Each time the switch detects traffic from a MAC address, it updates the timestamp. A timer on the switch periodically checks the timestamp, and if it is older than a user-configured value, the switch removes the node's MAC address from the Ethernet switching table. This aging process ensures that the switch tracks only active nodes on the network and that it is able to flush out network nodes that are no longer available.

Types of Switch Ports

The ports, or interfaces, on a switch operate in either access mode or trunk mode.

An interface in access mode connects to a network device, such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. The interface itself belongs to a single VLAN. The frames transmitted over an access interface are normal Ethernet frames. By default, when you boot a switch and use the factory-default configuration, or when you boot the switch and do not explicitly configure a port mode, all interfaces on the switch are in access mode.

Trunk interfaces handle traffic for multiple VLANs, multiplexing the traffic for all those VLANs over the same physical connection. Trunk interfaces are generally used to interconnect switches to one another.

IEEE 802.1Q Encapsulation and Tags

To identify which VLAN traffic belongs to, all frames on an Ethernet VLAN are identified by a tag, as defined in the IEEE 802.1Q standard. These frames are *tagged* and are encapsulated with 802.1Q tags.

For a simple network that has only a single VLAN, all traffic has the same 802.1Q tag.

When an Ethernet LAN is divided into VLANs, each VLAN is identified by a unique 802.1Q tag. The tag is applied to all frames so that the network nodes receiving the frames know which VLAN the frames belong to. Trunk ports, which multiplex traffic among a number of VLANs, use the tag to determine the origin of frames and where to forward them.

VLANs 0 and 4095 are reserved by the Juniper Networks JUNOS Software, so you cannot use them in your network.

Assignment of Traffic to VLANs

You assign traffic to a particular VLAN in one of the following ways:

- By interface (port) on the switch. You specify that all traffic received on a particular interface on the switch is assigned to a specific VLAN. If you use the default factory switch settings, all traffic received on an access interface is untagged. This traffic is part of a default VLAN, but it is not tagged with an 802.1Q tag. When configuring the switch, you specify which VLAN to assign the traffic to. You configure the VLAN either by using a VLAN number (called a VLAN ID) or by using a name, which the switch translates into a numeric VLAN ID.
- By MAC address. You can specify that all traffic received from a specific MAC address be forwarded to a specific egress interface (next hop) on the switch. This method is administratively cumbersome to configure manually, but it can be useful when you are using automated databases to manage the switches on your network.



NOTE: If a Juniper Networks EX4200 Ethernet Switch is interconnected with other switches in a Virtual Chassis configuration, each individual switch that is included as a member of the configuration is identified with a member ID. The member ID functions as an FPC slot number. When you are configuring interfaces for a Virtual Chassis configuration, you specify the appropriate member ID (0 through 9) as the *slot* element of the interface name.

The default factory settings for a Virtual Chassis configuration include FPC 0 as a member of the default VLAN because FPC 0 is configured as part of the `ethernet-switching` family. In order to include FPC 1 through FPC 9 in the default VLAN, add the `ethernet-switching` family to the configurations for those interfaces.

Ethernet Switching Tables

As EX Series switches learn the MAC addresses of the devices on local VLANs, they store them in the bridge on the switch. With each MAC address, the Ethernet switching table stores and associates the name of the interface (or port) on which the switch learned that address. The switch uses the information in this table when forwarding packets toward their destination.

Layer 2 and Layer 3 Forwarding of VLAN Traffic

To pass traffic within a VLAN, the switch uses Layer 2 forwarding protocols, including IEEE 802.1Q, Spanning Tree Protocol (STP), and GARP VLAN Registration Protocol (GVRP).

To pass traffic between two VLANs, the switch uses standard Layer 3 routing protocols, such as static routing, OSPF, and RIP. On EX Series switches, the same interfaces that support Layer 2 bridging protocols also support Layer 3 routing protocols, providing multilayer switching.

GVRP and MVRP

The GARP VLAN Registration Protocol (GVRP) and Multiple VLAN Registration Protocol (MVRP) are used to manage dynamic VLAN registration in a LAN.

GVRP is an application protocol of the Generic Attribute Registration Protocol (GARP) and is defined in the IEEE 802.1Q standard. GVRP learns VLANs on a particular 802.1Q trunk interface and adds the corresponding trunk interface to the VLAN if the advertised VLAN is preconfigured on the switch.

MVRP is an application protocol of the Multiple Registration Protocol (MRP) and is defined in the IEEE 802.1ak standard. MRP and MVRP were designed by IEEE to perform the same functions as GARP and GVRP while overcoming some GARP and GVRP limitations, in particular limitations involving bandwidth usage and convergence times in large networks with large numbers of VLANs. MVRP was created by IEEE as a replacement application for GVRP.

The VLAN registration information sent by MVRP and GVRP includes the current VLANs membership—that is, which switches are members of which VLANs—and

which switch interfaces are in which VLAN. GVRP and MVRP share all VLAN information configured on a local switch.

MVRP can also be used to dynamically create VLANs, which are VLANs created on one switch and propagated to other switches as part of the MVRP message exchange process.

As part of ensuring that VLAN membership information is current, GVRP and MVRP remove switches and interfaces from the VLAN information when those switches and interfaces become unavailable. Pruning VLAN information has these benefits:

- Limits the network VLAN configuration to active participants only, reducing network overhead.
- Targets the scope of broadcast, unknown unicast, and multicast (BUM) traffic to interested devices only.

Routed VLAN Interface

In a traditional network, broadcast domains consist of either physical interfaces connected to a single switch or logical interfaces connected to one or more switches through VLAN configurations. Switches send traffic to hosts that are part of the same broadcast domain, but routers are needed to route traffic from one broadcast domain to another and to perform other Layer 3 functions such as traffic engineering. EX Series switches use a Layer 3 routed VLAN interface (RVI) named `vlan` to perform these routing functions, using it to route data to other Layer 3 interfaces. The RVI functions as a logical router, eliminating the need for having both a switch and a router.

The RVI (the `vlan` interface) must be configured as part of a broadcast domain or virtual private LAN service (VPLS) routing instance for Layer 3 traffic to be routed out of it. The RVI supports IPv4, IPv6, MPLS, and IS-IS traffic. At least one Layer 2 logical interface must be operational for the RVI to be operational. You must configure a broadcast domain or VPLS routing instance for the RVI just as you would configure a VLAN on the switch. Multicast data, broadcast data, or unicast data is switched between ports within the same RVI broadcast domain or VPLS routing instance. The RVI routes data that is destined for the switch's media access control (MAC) address.

Jumbo frames of up to 9216 bytes are supported on an RVI. 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.

See “Configuring Routed VLAN Interfaces (CLI Procedure)” on page 99.

To learn more about configuring routing protocols and policies, see the *JUNOS Software Routing Protocols Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos100/index.html>.

- Related Topics**
- Understanding Layer 2 Protocol Tunneling on EX Series Switches on page 19
 - Understanding Multiple VLAN Registration Protocol (MVRP) on EX Series Switches on page 17
 - Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21
 - Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 28
 - Example: Configure Automatic VLAN Administration Using GVRP on page 46
 - Example: Connecting an Access Switch to a Distribution Switch on page 36

Understanding Private VLANs on EX Series Switches

The private VLAN (PVLAN) feature on Juniper Networks EX Series Ethernet Switches allows an administrator to split a broadcast domain into multiple isolated broadcast subdomains, essentially putting a VLAN inside a VLAN. Just like regular VLANs, PVLANS are isolated on Layer 2 and require that a Layer 3 device be used to route traffic among them. Private VLANs are useful for restricting the flow of broadcast and unknown unicast traffic and for limiting the communication between known hosts.



NOTE: Configuring a voice over IP (VoIP) VLAN on PVLAN interfaces is not supported.

In a private VLAN, one VLAN is designated the primary VLAN, and other VLANs are nested inside that VLAN as secondary VLANs.

- **Primary**—A VLAN used to forward frames downstream to isolated and community VLANs.
- **Isolated**—A secondary VLAN that receives packets only from the primary VLAN and forwards frames upstream to the primary VLAN.
- **Community**—A secondary VLAN that transports frames among community interfaces within the same community and forwards frames upstream to the primary VLAN.

Private VLANs provide IP address conservation and efficient allocation of those IP addresses. In a typical network, VLANs usually correspond to a single IP subnet. In private VLANs, the hosts in all the secondary VLANs still belong to the same IP subnet as the subnet allocated to the primary VLAN. Hosts within the secondary VLAN are numbered out of IP subnets associated with the primary VLAN, and their IP subnet masking information reflects that of the primary VLAN subnet. Any primary routed VLAN interfaces (RVIs) perform functions similar to proxy ARP to enable communication between hosts that are members of a different secondary VLAN.



NOTE: If you enable `no-mac-learning` on a primary VLAN, all isolated VLANs in that private VLAN inherit that setting. If you want to disable MAC address learning on any community VLANs, you must configure `no-mac-learning` on each of those VLANs.

- Related Topics**
- Understanding Bridging and VLANs on EX Series Switches on page 3
 - Example: Configuring a Private VLAN on an EX Series Switch on page 68
 - Creating a Private VLAN (CLI Procedure) on page 105

Understanding Virtual Routing Instances on EX Series Switches

Virtual routing instances allow administrators to divide a Juniper Networks EX Series Ethernet Switch into multiple independent virtual routers, each with its own routing table. Splitting a device into many virtual routing instances isolates traffic traveling across the network without requiring multiple devices to segment the network.

You can use virtual routing instances to isolate customer traffic on your network and to bind customer-specific instances to customer-owned interfaces.

EX Series switches support up to 256 virtual routing instances. Virtual routing and forwarding (VRF) is often used in conjunction with Layer 3 subinterfaces, allowing traffic on a single physical interface to be differentiated and associated with multiple virtual routers. Each logical Layer 3 subinterface can belong to only one routing instance.

- Related Topics**
- Understanding Layer 3 Subinterfaces
 - Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 73
 - Configuring Virtual Routing Instances (CLI Procedure) on page 105

Understanding Redundant Trunk Links on EX Series Switches

In a typical enterprise network comprised of distribution and access layers, a redundant trunk link provides a simple solution for network recovery when a trunk port goes down. Traffic is routed to another trunk port, keeping network convergence time to a minimum. You can configure a maximum of 16 redundant trunk groups on a standalone switch or on a Virtual Chassis.

To configure a redundant trunk link, create a redundant trunk group. The redundant trunk group is configured on the access switch, and contains two links: a primary or active link, and a secondary link. If the active link fails, the secondary link automatically starts forwarding data traffic without waiting for normal STP convergence.

Data traffic is forwarded only on the active link. Data Traffic on the secondary link is dropped and shown as dropped packets when you issue the operational mode command `show interfaces xe-interface-name extensive`.

While data traffic is blocked on the secondary link, Layer 2 control traffic is still permitted. For example, an LLDP session can be run between two Juniper Networks EX Series Ethernet Switches on the secondary link.

STP is enabled by default on EX Series switches to create a loop-free topology. When trunk links are placed in a redundant group, they cannot be part of an STP topology. The Juniper Networks JUNOS Software for EX Series switches does not allow an interface to be in a redundant trunk group and in an STP topology at the same time. However, STP can continue operating in other parts of the network. For example, STP may continue operating between the distribution switches and linking them to the enterprise core.

Figure 1 on page 12 shows three switches in a basic topology for redundant trunk links. Switch 1 and Switch 2 make up the distribution layer, and Switch 3 makes up the access layer. Switch 3 is connected to the distribution layer through trunk ports `ge-0/0/9.0` (Link 1) and `ge-0/0/10.0` (Link 2). Link 1 and Link 2 are in a redundant trunk group called `group1`. Link 1 is designated as the primary link. Traffic flows between Switch 3 in the access layer and Switch 1 in the distribution layer through Link 1. While Link 1 is active, Link 2 blocks traffic.

Figure 1: Redundant Trunk Group, Link 1 Active





Figure 2 on page 13 illustrates how the redundant trunk link topology works when the primary link goes down.

Figure 2: Redundant Trunk Group, Link 2 Active

Link 1 is down between Switch 3 and Switch 1. Link 2 takes over as the active link. Traffic between the access layer and the distribution layer is automatically switched to Link 2 between Switch 1 and Switch 2.

- Related Topics**
- Example: Configuring Redundant Trunk Links for Faster Recovery on page 61
 - `redundant-trunk-group`

Understanding Q-in-Q Tunneling on EX Series Switches

Q-in-Q tunneling allows service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites. Using Q-in-Q tunneling, providers can also segregate or bundle customer traffic into fewer VLANs or different VLANs by adding another layer of 802.1Q tags. Q-in-Q tunneling is useful when customers have overlapping VLAN IDs, because the customer's 802.1Q (dot1Q) VLAN tags are prepended by the service VLAN (S-VLAN) tag. The Juniper Networks JUNOS Software implementation of Q-in-Q tunneling supports the IEEE 802.1ad standard.

This topic describes:

- How Q-in-Q Tunneling Works on page 14
- Disabling MAC Address Learning on page 14
- Mapping C-VLANs to S-VLANs on page 15
- Routed VLAN Interfaces on Q-in-Q VLANs on page 16
- Limitations for Q-in-Q Tunneling on page 16

How Q-in-Q Tunneling Works

In Q-in-Q tunneling, as a packet travels from a customer VLAN (C-VLAN) to a service provider's VLAN, a customer-specific 802.1Q tag is added to the packet. This additional tag is used to segregate traffic into service-provider-defined service VLANs (S-VLANs). The original customer 802.1Q tag of the packet remains and is transmitted transparently, passing through the service provider's network. As the packet leaves the S-VLAN in the downstream direction, the extra 802.1Q tag is removed.

When Q-in-Q tunneling is enabled on Juniper Networks EX Series Ethernet Switches, trunk interfaces are assumed to be part of the service provider network and access interfaces are assumed to be customer facing. An access interface can receive both tagged and untagged frames in this case.

An interface can be a member of multiple S-VLANs. You can map one C-VLAN to one S-VLAN (1:1) or multiple C-VLANs to one S-VLAN (N:1). Packets are double-tagged for an additional layer of segregating or bundling of C-VLANs. C-VLAN and S-VLAN tags are unique; so you can have both a C-VLAN 101 and an S-VLAN 101, for example. You can limit the set of accepted customer tags to a range of tags or to discrete values. Class-of-service (CoS) values of C-VLANs are unchanged in the downstream direction. You may, optionally, copy ingress priority and CoS settings to the S-VLAN. Using private VLANs, you can isolate users to prevent the forwarding of traffic between user interfaces even if the interfaces are on the same VLAN.

You can use the **native** option to specify an S-VLAN for untagged and priority tagged packets when using many-to-one bundling and mapping a specific interface approaches to map C-VLANs to S-VLANs. Otherwise the packets are discarded. The **native** option is not available for all-in-one bundling because there is no need to specify untagged and priority tagged packets when all packets are mapped to the C-VLAN. See the Mapping C-VLANs to S-VLANs section of this document for information on the methods of mapping C-VLANs to S-VLANs.

Firewall filters allow you to map an interface to a VLAN based on a policy. Using firewall filters to map an interface to a VLAN is useful when you want a subset of traffic from a port to be mapped to a selected VLAN instead of the designated VLAN. To configure a firewall filter to map an interface to a VLAN, the **vlan** option has to be configured as part of the firewall filter and the **mapping policy** option must be specified in the interface configuration for each logical interface using the filter.

Disabling MAC Address Learning

In a Q-in-Q deployment, customer packets from downstream interfaces are transported without any changes to source and destination MAC addresses. You can disable MAC address learning at both the interface level and the VLAN level. Disabling MAC address learning on an interface disables learning for all the VLANs of which that interface is a member. When you disable MAC address learning on a VLAN, MAC addresses that have already been learned are flushed.

If you disable MAC address learning on an interface or a VLAN, you cannot include MAC move limiting or 802.1X authentication in that same VLAN configuration.

When a routed VLAN interface (RVI) is associated with either an interface or a VLAN on which MAC address learning is disabled, the Layer 3 routes resolved on that VLAN or that interface are not resolved with the Layer 2 component. This results in routed packets flooding all the interfaces associated with the VLAN.

Mapping C-VLANs to S-VLANs

There are three ways to map C-VLANs to an S-VLAN:

- All-in-one bundling—Use the `dot1q-tunneling` option to map without specifying customer VLANs. All packets from all access interfaces are mapped to the S-VLAN.
- Many-to-one bundling—Use the `customer-vlans` option to specify which C-VLANs are mapped to the S-VLAN.
- Mapping a specific interface—Use the `mapping` option to indicate a specific S-VLAN for a given C-VLAN. The specified C-VLAN applies to only one VLAN and not all access interfaces as in the cases of all-in-one and many-to-one bundling.

If you configure multiple methods, the switch gives priority to mapping a specific interface, then to many-to-one bundling, and last to all-in-one bundling. However, you cannot have overlapping rules for the same C-VLAN under a given approach.

- All-in-One Bundling on page 15
- Many-to-One Bundling on page 15
- Mapping a Specific Interface on page 15

All-in-One Bundling

All-in-one bundling maps all packets from all access interfaces to the S-VLAN. All-in-one bundling is configured using the `dot1q-tunneling` option without specifying customer VLANs.

When all-in-one bundling is used, all packets leaving the C-VLAN, including untagged and priority tagged packets, enter the S-VLAN.

Many-to-One Bundling

Many-to-one bundling is used to specify which C-VLANs are mapped to an S-VLAN. Many-to-one bundling is configured using the `customer-vlans` option.

Many-to-one bundling is used when you want a subset of the C-VLANs on the access switch to be part of the S-VLAN. When using many-to-one bundling, untagged and priority tagged packets can be mapped to the S-VLAN when the `native` option is specified along with the `customer-vlans` option.

Mapping a Specific Interface

Use the mapping a specific interface approach when you want to assign an S-VLAN to a specific C-VLAN on an interface. The mapping a specific interface configuration only applies to the configured interface, not to all access interfaces as in the cases of the all-in-one bundling and many-to-one bundling approaches. The mapping a

specific interface approach is configured using the **mapping** option to indicate a specific S-VLAN for a given C-VLAN.

The mapping a specific interface approach has two suboptions for treatment of traffic: **swap** and **push**. When traffic that is mapped to a specific interface is pushed, the packet retains its tag as it moves between the S-VLAN and C-VLAN and an additional VLAN tag is added to the packet. When traffic that is mapped to a specific interface is swapped, the incoming tag is replaced with a new VLAN tag. Using the **swap** option is also referred to as VLAN ID translation.

It might be useful to have S-VLANs that provide service to multiple customers. Each customer will typically have its own S-VLAN plus access to one or more S-VLANs that are used by multiple customers. A specific tag on the customer side is mapped to an S-VLAN. Typically, this functionality is used to keep data from different customers separate or to provide individualized treatment of the packets on a certain interface.

Routed VLAN Interfaces on Q-in-Q VLANs

Routed VLAN interfaces (RVIs) are supported on Q-in-Q VLANs.

Packets arriving on an RVI that is using Q-in-Q VLANs will get routed regardless of whether the packet is single or double tagged. The outgoing routed packets contain an S-VLAN tag only when exiting a trunk interface; the packets exit the interface untagged when exiting an access interface.

Limitations for Q-in-Q Tunneling

Q-in-Q tunneling does not support most access port security features. There is no per-VLAN (customer) policing or per-VLAN (outgoing) shaping and limiting with Q-in-Q tunneling unless you configure these security features using firewall filters.

- Related Topics**
- Understanding Bridging and VLANs on EX Series Switches on page 3
 - Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 65
 - Configuring Q-in-Q Tunneling (CLI Procedure) on page 106

Understanding Multiple VLAN Registration Protocol (MVRP) on EX Series Switches

You can configure Multiple VLAN Registration Protocol (MVRP) on Juniper Networks EX Series Ethernet Switches. The primary purpose of MVRP is to manage dynamic VLAN registration in a LAN. In managing dynamic VLAN registration, MVRP also prunes VLAN information. MVRP can also be used to dynamically create VLANs in switching networks.

MVRP is an application protocol of the Multiple Registration Protocol (MRP) and is defined in the IEEE 802.1ak standard. MRP and MVRP were designed by IEEE to perform the same functions as Generic Attribute Registration Protocol (GARP) and GARP VLAN Registration Protocol (GVRP) while overcoming some GARP and GVRP limitations, in particular limitations involving bandwidth usage and convergence time in large networks with large numbers of VLANs.

MVRP was created by IEEE as a replacement application for GVRP. MVRP and GVRP cannot be run concurrently to share VLAN information in a switching network.

This topic describes:

- How MVRP Works on EX Series Switches on page 17
- Basics of MVRP on EX Series Switches on page 18
- MVRP Registration Modes on page 18
- MRP Timers on page 18
- MRP VLAN Messages on page 18
- MVRP Limitations on page 19

How MVRP Works on EX Series Switches

The VLAN registration information sent by MVRP protocol data units (PDUs) includes the current VLANs membership—that is, which switches are members of which VLANs—and which switch interfaces are in which VLAN. MVRP shares all information in the PDU with all switches participating in MVRP in the switching network.

MVRP stays synchronized using these PDUs. The MVRP PDUs are sent to other switches on the network only when an MVRP state change occurs. The switches in the network participating in MVRP receive these PDUs during state changes and update their MVRP states accordingly. MVRP timers dictate when PDUs can be sent and when switches receiving MVRP PDUs can update their MVRP information.

VLAN information is distributed as part of the MVRP message exchange process and can be used to dynamically create VLANs, which are VLANs created on one switch and propagated to other switches as part of the MVRP message exchange process. Dynamic VLAN creation using MVRP is enabled by default but can be disabled.

As part of ensuring that VLAN membership information is current, MVRP removes switches and interfaces from the VLAN information when they become unavailable. Pruning VLAN information has these benefits:

- Limits the network VLAN configuration to active participants only, reducing network overhead.
- Targets the scope of broadcast, unknown unicast, and multicast (BUM) traffic to interested devices only.

Basics of MVRP on EX Series Switches

MVRP is disabled by default on all EX Series switches. You can configure MVRP on EX Series switch interfaces to participate in MVRP for the switching network. MVRP can only be enabled on trunk interfaces, and dynamic VLAN configuration through MVRP is enabled by default when MVRP is enabled.

MVRP Registration Modes

The MVRP registration mode defines whether an interface does or does not participate in MVRP.

The following MVRP registration modes are configurable:

- forbidden—The interface does not register and does not participate in MVRP.
- normal—The interface accepts MVRP messages and participates in MVRP. This is the default registration mode setting.

MRP Timers

MVRP registration and updates are controlled by timers that are part of the MRP protocol. These timers are set on a per-interface basis and define when MVRP PDUs can be sent and when MVRP information can be updated on a switch.

The following timers are used to control the operation of MVRP:

- Join timer—Controls the interval for the next MVRP PDU transmit opportunity.
- Leave timer—Controls the period of time that an interface on the switch waits in the Leave state before changing to the unregistered state.
- LeaveAll timer—Controls the frequency with which the interface generates LeaveAll messages.



BEST PRACTICE: Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

MRP VLAN Messages

MVRP uses MRP messages to register and declare MVRP states for a switch and to inform the switching network that a switch is leaving MVRP. These messages are communicated as part of the PDU to communicate the state of a particular switch interface on the switching network to the other switches in the network.

The following messages are communicated for MVRP:

- Empty—VLAN information is not being declared and is not registered.
- In—VLAN information is not being declared but is registered.
- JoinEmpty—VLAN information is being declared but not registered.
- JoinIn—VLAN information is being declared and is registered.
- Leave—VLAN information that was previously registered is being withdrawn.
- LeaveAll—All registrations will be de-registered. Participants that want to participate in MVRP will need to re-register.
- New—VLAN information is new and possibly not previously registered.

MVRP Limitations

MVRP does not work with any versions of Spanning Tree Protocol (STP) except Rapid Spanning Tree Protocol (RSTP).

- Related Topics**
- Understanding Bridging and VLANs on EX Series Switches on page 3
 - Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 76
 - Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 109

Understanding Layer 2 Protocol Tunneling on EX Series Switches

Layer 2 protocol tunneling (L2PT) allows you to send Layer 2 protocol data units (PDUs) across a service provider network and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain. This feature is useful when you want to run Layer 2 protocols on a network that includes switches located at remote sites that are connected across a service provider network.

This topic includes:

- Layer 2 Protocols Supported by L2PT on EX Series Switches on page 19
- How L2PT Works on page 20
- L2PT Basics on EX Series Switches on page 20

Layer 2 Protocols Supported by L2PT on EX Series Switches

L2PT on EX Series switches supports the following Layer 2 protocols:

- Cisco Discovery Protocol (CDP)
- GARP VLAN Registration Protocol (GVRP)
- Link Layer Discovery Protocol (LLDP)
- Multiple VLAN Registration Protocol (MVRP)

- Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP)
- VLAN Spanning Tree Protocol (VSTP)
- VLAN Trunking Protocol (VTP)



NOTE: CDP and VTP cannot be configured on EX Series switches. L2PT does, however, tunnel CDP and VTP PDUs.

How L2PT Works

L2PT works by encapsulating Layer 2 PDUs, tunneling them across a service provider network, and decapsulating them for delivery to their destination switches. L2PT encapsulates Layer 2 PDUs by enabling the ingress provider edge (PE) device to rewrite the PDUs' destination media access control (MAC) addresses before forwarding them onto the service provider network. The devices in the service provider network treat these encapsulated PDUs as multicast Ethernet packets. Upon receipt of these PDUs, the egress PE devices decapsulate them by replacing the destination MAC addresses with the address of the Layer 2 protocol that is being tunneled before forwarding the PDUs to their destination switches.

L2PT Basics on EX Series Switches

L2PT is enabled on a per-VLAN basis. When you enable L2PT on a VLAN, all access interfaces are considered to be customer-facing interfaces, all trunk interfaces are considered to be service provider network-facing interfaces, and the specified Layer 2 protocol is disabled on the access interfaces. L2PT only acts on logical interfaces of the family `ethernet-switching`.



NOTE: Access interfaces in an L2PT-enabled VLAN should not receive L2PT-tunneled PDUs. If an access interface does receive L2PT-tunneled PDUs, it might mean that there is a loop in the network. As a result, the interface will be shut down.

L2PT is configured under the `[edit vlans vlan-name dot1q-tunneling]` hierarchy level, meaning Q-in-Q tunneling is (and must be) enabled. If L2PT is not enabled, Layer 2 PDUs are handled in the same way they were handled before L2PT was enabled.



NOTE: If the switch receives untagged Layer 2 control PDUs to be tunneled, then you must configure the switch to map untagged (native) packets to an L2PT-enabled VLAN. For more information on assigning untagged packets to VLANs, see “Understanding Q-in-Q Tunneling on EX Series Switches” on page 13 and “Configuring Q-in-Q Tunneling (CLI Procedure)” on page 106.

- Related Topics**
- Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 88
 - Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 65

Chapter 2

Examples: Bridging and VLAN Configuration

- Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21
- Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 28
- Example: Connecting an Access Switch to a Distribution Switch on page 36
- Example: Configure Automatic VLAN Administration Using GVRP on page 46
- Example: Configuring Redundant Trunk Links for Faster Recovery on page 61
- Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 65
- Example: Configuring a Private VLAN on an EX Series Switch on page 68
- Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 73
- Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 76
- Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 88

Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch

EX Series switches use bridging and virtual LANs (VLANs) to connect network devices in a LAN—desktop computers, IP telephones, printers, file servers, wireless access points, and others—and to segment the LAN into smaller bridging domains. The switch's default configuration provides a quick setup of bridging and a single VLAN.

This example describes how to configure basic bridging and VLANs for an EX Series switch:

- Requirements on page 22
- Overview and Topology on page 22
- Configuration on page 23
- Verification on page 27

Requirements

This example uses the following software and hardware components:

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

Before you set up bridging and a VLAN, be sure you have:

- Installed your EX Series switch. See *Installing and Connecting an EX3200 or EX4200 Switch*.
- Performed the initial switch configuration. See *Connecting and Configuring an EX Series Switch (J-Web Procedure)*.

Overview and Topology

EX Series switches connect network devices in an office LAN or a data center LAN to provide sharing of common resources such as printers and file servers and to enable wireless devices to connect to the LAN through wireless access points. Without bridging and VLANs, all devices on the Ethernet LAN are in a single broadcast domain, and all the devices detect all the packets on the LAN. Bridging creates separate broadcast domains on the LAN, creating VLANs, which are independent logical networks that group together related devices into separate network segments. The grouping of devices on a VLAN is independent of where the devices are physically located in the LAN.

To use an EX Series switch to connect network devices on a LAN, you must, at a minimum, configure bridging and VLANs. If you simply power on the switch and perform the initial switch configuration using the factory-default settings, bridging is enabled on all the switch's interfaces, all interfaces are in access mode, and all interfaces belong to a VLAN called **default**, which is automatically configured. When you plug access devices—such as desktop computers, Avaya IP telephones, file servers, printers, and wireless access points—into the switch, they are joined immediately into the **default** VLAN and the LAN is up and running.

The topology used in this example consists of one EX4200-24T switch, which has a total of 24 ports. Eight of the ports support Power over Ethernet (PoE), which means they provide both network connectivity and electric power for the device connecting to the port. To these ports, you can plug in devices requiring PoE, such as Avaya VoIP telephones, wireless access points, and some IP cameras. (Avaya phones have a built-in hub that allows you to connect a desktop PC to the phone, so the desktop and phone in a single office require only one port on the switch.) The remaining 16 ports provide only network connectivity. You use them to connect devices that have their own power sources, such as desktop and laptop computers, printers, and servers. Table 1 details the topology used in this configuration example.

Table 1: Components of the Basic Bridging Configuration Topology

Property	Settings
----------	----------

Table 1: Components of the Basic Bridging Configuration Topology (continued)

Switch hardware	EX4200-24T 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 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 (no PoE required)	ge-0/0/8 through ge-0/0/12
Connections to file servers (no PoE required)	ge-0/0/17 and ge-0/0/18
Connections to integrated printer/fax/copier machines (no PoE required)	ge-0/0/19 through ge-0/0/20
Unused ports (for future expansion)	ge-0/0/13 through ge-0/0/16, and ge-0/0/21 through ge-0/0/23

Configuration

CLI Quick Configuration By default, after you perform the initial configuration on the EX4200 switch, switching is enabled on all interfaces, a VLAN named **default** is created, and all interfaces are placed into this VLAN. You do not need to perform any other configuration on the switch to set up bridging and VLANs. To use the switch, simply plug the Avaya IP phones into the PoE-enabled ports **ge-0/0/1** through **ge-0/0/7**, and plug in the PCs, file servers, and printers to the non-PoE ports, **ge-0/0/8** through **ge-0/0/12** and **ge-0/0/17** through **ge-0/0/20**.

Step-by-Step Procedure To configure bridging and VLANs:

1. Make sure the switch is powered on.
2. Connect the wireless access point to switch port **ge-0/0/0**.
3. Connect the seven Avaya phones to switch ports **ge-0/0/1** through **ge-0/0/7**.
4. Connect the five PCs to ports **ge-0/0/8** through **ge-0/0/12**.
5. Connect the two file servers to ports **ge-0/0/17** and **ge-0/0/18**.
6. Connect the two printers to ports **ge-0/0/19** and **ge-0/0/20**.

Results Check the results of the configuration:

```
[edit]
user@switch> show configuration
## Last commit: 2008-03-06 00:11:22 UTC by triumph
version 9.0;
system {
  root-authentication {
```

```
        encrypted-password "$1$urmA7AFM$x5SaGEUOdSI3u1K/iITGh1"; ##
        SECRET-DATA
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
    commit {
        factory-settings {
            reset-chassis-lcd-menu;
            reset-virtual-chassis-configuration;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/1 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/2 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/3 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/4 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/5 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/6 {
        unit 0 {
            family ethernet-switching;
        }
    }
}
```

```
    }  
  }  
  ge-0/0/7 {  
    unit 0 {  
      family ethernet-switching;  
    }  
  }  
  ge-0/0/8 {  
    unit 0 {  
      family ethernet-switching;  
    }  
  }  
  ge-0/0/9 {  
    unit 0 {  
      family ethernet-switching;  
    }  
  }  
  ge-0/0/10 {  
    unit 0 {  
      family ethernet-switching;  
    }  
  }  
  ge-0/0/11 {  
    unit 0 {  
      family ethernet-switching;  
    }  
  }  
  ge-0/0/12 {  
    unit 0 {  
      family ethernet-switching;  
    }  
  }  
  ge-0/0/13 {  
    unit 0 {  
      family ethernet-switching;  
    }  
  }  
  ge-0/0/14 {  
    unit 0 {  
      family ethernet-switching;  
    }  
  }  
  ge-0/0/15 {  
    unit 0 {  
      family ethernet-switching;  
    }  
  }  
  ge-0/0/16 {  
    unit 0 {  
      family ethernet-switching;  
    }  
  }  
  ge-0/0/17 {  
    unit 0 {  
      family ethernet-switching;  
    }  
  }
```

```
}
ge-0/0/18 {
  unit 0 {
    family ethernet-switching;
  }
}
ge-0/0/19 {
  unit 0 {
    family ethernet-switching;
  }
}
ge-0/0/20 {
  unit 0 {
    family ethernet-switching;
  }
}
ge-0/0/21 {
  unit 0 {
    family ethernet-switching;
  }
}
ge-0/0/22 {
  unit 0 {
    family ethernet-switching;
  }
}
ge-0/0/23 {
  unit 0 {
    family ethernet-switching;
  }
}
ge-0/1/0 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/1/0 {
  unit 0 {
    family ethernet-switching;
  }
}
ge-0/1/1 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/1/1 {
  unit 0 {
    family ethernet-switching;
  }
}
ge-0/1/2 {
  unit 0 {
    family ethernet-switching;
  }
}
}
```

```

ge-0/1/3 {
  unit 0 {
    family ethernet-switching;
  }
}
protocols {
  lldp {
    interface all;
  }
  rstp;
}
poe {
  interface all;
}

```

Verification

To verify that switching is operational and that a VLAN has been created, perform these tasks:

- Verifying That the VLAN Has Been Created on page 27
- Verifying That Interfaces Are Associated with the Proper VLANs on page 27

Verifying That the VLAN Has Been Created

Purpose Verify that the VLAN named `default` has been created on the switch.

Action List all VLANs configured on the switch:

```
user@switch> show vlans
```

Name	Tag	Interfaces
default		ge-0/0/0.0*, ge-0/0/1.0, ge-0/0/2.0, ge-0/0/3.0, ge-0/0/4.0, ge-0/0/5.0, ge-0/0/6.0, ge-0/0/7.0, ge-0/0/8.0*, ge-0/0/9.0, ge-0/0/10.0, ge-0/0/11.0*, ge-0/0/12.0, ge-0/0/13.0, ge-0/0/14.0, ge-0/0/15.0, ge-0/0/16.0, ge-0/0/17.0, ge-0/0/18.0, ge-0/0/19.0*, ge-0/0/20.0, ge-0/0/21.0, ge-0/0/22.0, ge-0/0/23.0, ge-0/1/0.0*, ge-0/1/1.0*, ge-0/1/2.0*, ge-0/1/3.0*
mgmt		me0.0*

Meaning The `show vlans` command lists the VLANs configured on the switch. This output shows that the VLAN `default` has been created.

Verifying That Interfaces Are Associated with the Proper VLANs

Purpose Verify that Ethernet switching is enabled on switch interfaces and that all interfaces are included in the VLAN.

Action List all interfaces on which switching is enabled:

```
user@switch> show ethernet-switching interfaces
```

Interface	State	VLAN members	Blocking
ge-0/0/0.0	up	default	unblocked
ge-0/0/1.0	down	default	blocked - blocked by STP/RTG
ge-0/0/2.0	down	default	blocked - blocked by STP/RTG
ge-0/0/3.0	down	default	blocked - blocked by STP/RTG
ge-0/0/4.0	down	default	blocked - blocked by STP/RTG
ge-0/0/5.0	down	default	blocked - blocked by STP/RTG
ge-0/0/6.0	down	default	blocked - blocked by STP/RTG
ge-0/0/7.0	down	default	blocked - blocked by STP/RTG
ge-0/0/8.0	up	default	unblocked
ge-0/0/9.0	down	default	blocked - blocked by STP/RTG
ge-0/0/10.0	down	default	blocked - blocked by STP/RTG
ge-0/0/11.0	up	default	unblocked
ge-0/0/12.0	down	default	blocked - blocked by STP/RTG
ge-0/0/13.0	down	default	blocked - blocked by STP/RTG
ge-0/0/14.0	down	default	blocked - blocked by STP/RTG
ge-0/0/15.0	down	default	blocked - blocked by STP/RTG
ge-0/0/16.0	down	default	blocked - blocked by STP/RTG
ge-0/0/17.0	down	default	blocked - blocked by STP/RTG
ge-0/0/18.0	down	default	blocked - blocked by STP/RTG
ge-0/0/19.0	up	default	unblocked
ge-0/0/20.0	down	default	blocked - blocked by STP/RTG
ge-0/0/21.0	down	default	blocked - blocked by STP/RTG
ge-0/0/22.0	down	default	blocked - blocked by STP/RTG
ge-0/0/23.0	down	default	blocked - blocked by STP/RTG
ge-0/1/0.0	up	default	unblocked
ge-0/1/1.0	up	default	unblocked
ge-0/1/2.0	up	default	unblocked
ge-0/1/3.0	up	default	unblocked
me0.0	up	mgmt	unblocked

Meaning The `show ethernet-switching interfaces` command lists all interfaces on which switching is enabled (in the **Interfaces** column), along with the VLANs that are active on the interfaces (in the **VLAN members** column). The output in this example shows all the connected interfaces, `ge-0/0/0` through `ge-0/0/12` and `ge-0/0/17` through `ge-0/0/20` and that they are all part of VLAN `default`. Notice that the interfaces listed are the logical interfaces, not the physical interfaces. For example, the output shows `ge-0/0/0.0` instead of `ge-0/0/0`. This is because JUNOS Software creates VLANs on logical interfaces, not directly on physical interfaces.

- Related Topics**
- Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 28
 - Example: Connecting an Access Switch to a Distribution Switch on page 36
 - Example: Configure Automatic VLAN Administration Using GVRP on page 46
 - Understanding Bridging and VLANs on EX Series Switches on page 3

Example: Setting Up Bridging with Multiple VLANs for EX Series Switches

To segment traffic on a LAN into separate broadcast domains, you create separate virtual LANs (VLANs) on an EX Series switch. Each VLAN is a collection of network nodes. When you use VLANs, frames whose origin and destination are in the same VLAN are forwarded only within the local VLAN, and only frames not destined for

the local VLAN are forwarded to other broadcast domains. VLANs thus limit the amount of traffic flowing across the entire LAN, reducing the possible number of collisions and packet retransmissions within the LAN.

This example describes how to configure bridging for an EX Series switch and how to create two VLANs to segment the LAN:

- Requirements on page 29
- Overview and Topology on page 29
- Configuration on page 30
- Verification on page 34

Requirements

This example uses the following hardware and software components:

- One EX4200-48P Virtual Chassis switch
- JUNOS Release 9.0 or later for EX Series switches

Before you set up bridging and VLANs, be sure you have:

- Installed the EX Series switch. See *Installing and Connecting an EX3200 or EX4200 Switch*.
- Performed the initial switch configuration. See *Connecting and Configuring an EX Series Switch (J-Web Procedure)*.

Overview and Topology

EX Series switches connect all devices in an office or data center into a single LAN to provide sharing of common resources such as printers and file servers and to enable wireless devices to connect to the LAN through wireless access points. The default configuration creates a single VLAN, and all traffic on the switch is part of that broadcast domain. Creating separate network segments reduces the span of the broadcast domain and allows you to group related users and network resources without being limited by physical cabling or by the location of a network device in the building or on the LAN.

This example shows a simple configuration to illustrate the basic steps for creating two VLANs on a single switch. One VLAN, called **sales**, is for the sales and marketing group, and a second, called **support**, is for the customer support team. The sales and support groups each have their own dedicated file servers, printers, and wireless access points. For the switch ports to be segmented across the two VLANs, each VLAN must have its own broadcast domain, identified by a unique name and tag (VLAN ID). In addition, each VLAN must be on its own distinct IP subnet.

The topology for this example consists of one EX4200-48P switch, which has a total of 48 Gigabit Ethernet ports, all of which support Power over Ethernet (PoE). Most of the switch ports connect to Avaya IP telephones. The remainder of the ports connect to wireless access points, file servers, and printers.

Table 2: Components of the Multiple VLAN Topology

Property	Settings
Switch hardware	EX4200-48P, 48 Gigabit Ethernet ports, all PoE-enabled (ge-0/0/0 through ge-0/0/47)
VLAN names and tag IDs	sales, tag 100 support, tag 200
VLAN subnets	sales: 192.0.2.0/25 (addresses 192.0.2.1 through 192.0.2.126) support: 192.0.2.128/25 (addresses 192.0.2.129 through 192.0.2.254)
Interfaces in VLAN sales	Avaya IP telephones: ge-0/0/3 through ge-0/0/19 Wireless access points: ge-0/0/0 and ge-0/0/1 Printers: ge-0/0/22 and ge-0/0/23 File servers: ge-0/0/20 and ge-0/0/21
Interfaces in VLAN support	Avaya IP telephones: ge-0/0/25 through ge-0/0/43 Wireless access points: ge-0/0/24 Printers: ge-0/0/44 and ge-0/0/45 File servers: ge-0/0/46 and ge-0/0/47
Unused interfaces	ge-0/0/2 and ge-0/0/25

This configuration example creates two IP subnets, one for the sales VLAN and the second for the support VLAN. The switch bridges traffic within a VLAN. For traffic passing between two VLANs, the switch routes the traffic using a Layer 3 routing interface on which you have configured the address of the IP subnet.

To keep the example simple, the configuration steps show only a few devices in each of the VLANs. Use the same configuration procedure to add more LAN devices.

Configuration

Configure Layer 2 switching for two VLANs:

CLI Quick Configuration To quickly configure Layer 2 switching for the two VLANs (sales and support) and to quickly configure Layer 3 routing of traffic between the two VLANs, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/0/0 unit 0 description "Sales wireless access point port"
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/3 unit 0 description "Sales phone port"
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/22 unit 0 description "Sales printer port"
set interfaces ge-0/0/22 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/20 unit 0 description "Sales file server port"
set interfaces ge-0/0/20 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/24 unit 0 description "Support wireless access point port"
set interfaces ge-0/0/24 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/26 unit 0 description "Support phone port"
set interfaces ge-0/0/26 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/44 unit 0 description "Support printer port"
```

```

set interfaces ge-0/0/44 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/46 unit 0 description "Support file server port"
set interfaces ge-0/0/46 unit 0 family ethernet-switching vlan members support
set interfaces vlan unit 0 family inet address 192.0.2.0/25
set interfaces vlan unit 1 family inet address 192.0.2.128/25
set vlans sales l3-interface vlan.0
set vlans sales vlan-id 100
set vlans support vlan-id 200
set vlans support l3-interface vlan.1

```

Step-by-Step Procedure

Configure the switch interfaces and the VLANs to which they belong. By default, all interfaces are in access mode, so you do not have to configure the port mode.

1. Configure the interface for the wireless access point in the sales VLAN:

```

[edit interfaces ge-0/0/0 unit 0]
user@switch# set description "Sales wireless access point port"
user@switch# set family ethernet-switching vlan members sales

```

2. Configure the interface for the Avaya IP phone in the sales VLAN:

```

[edit interfaces ge-0/0/3 unit 0]
user@switch# set description "Sales phone port"
user@switch# set family ethernet-switching vlan members sales

```

3. Configure the interface for the printer in the sales VLAN:

```

[edit interfaces ge-0/0/22 unit 0]
user@switch# set description "Sales printer port"
user@switch# set family ethernet-switching vlan members sales

```

4. Configure the interface for the file server in the sales VLAN:

```

[edit interfaces ge-0/0/20 unit 0]
user@switch# set description "Sales file server port"
user@switch# set family ethernet-switching vlan members sales

```

5. Configure the interface for the wireless access point in the support VLAN:

```

[edit interfaces ge-0/0/24 unit 0]
user@switch# set description "Support wireless access point port"
user@switch# set family ethernet-switching vlan members support

```

6. Configure the interface for the Avaya IP phone in the support VLAN:

```

[edit interfaces ge-0/0/26 unit 0]
user@switch# set description "Support phone port"
user@switch# set family ethernet-switching vlan members support

```

7. Configure the interface for the printer in the support VLAN:

```

[edit interfaces ge-0/0/44 unit 0]
user@switch# set description "Support printer port"
user@switch# set family ethernet-switching vlan members support

```

8. Configure the interface for the file server in the support VLAN:

```
[edit interfaces ge-0/0/46 unit 0]
user@switch# set description "Support file server port"
user@switch# set family ethernet-switching vlan members support
```

9. Create the subnet for the sales broadcast domain:

```
[edit interfaces]
user@switch# set vlan unit 0 family inet address 192.0.2.1/25
```

10. Create the subnet for the support broadcast domain:

```
[edit interfaces]
user@switch# set vlan unit 1 family inet address 192.0.2.129/25
```

11. Configure the VLAN tag IDs for the sales and support VLANs:

```
[edit vlans]
user@switch# set sales vlan-id 100
user@switch# set support vlan-id 200
```

12. To route traffic between the sales and support VLANs, define the interfaces that are members of each VLAN and associate a Layer 3 interface:

```
[edit vlans]
user@switch# set sales l3-interface vlan.0
user@switch# set support l3-interface vlan.1
```

Display the results of the configuration:

```
user@switch> show configuration
interfaces {
  ge-0/0/0 {
    unit 0 {
      description "Sales wireless access point port";
      family ethernet-switching {
        vlan members sales;
      }
    }
  }
  ge-0/0/3 {
    unit 0 {
      description "Sales phone port";
      family ethernet-switching {
        vlan members sales;
      }
    }
  }
  ge-0/0/22 {
    unit 0 {
      description "Sales printer port";
```

```

        family ethernet-switching {
            vlan members sales;
        }
    }
}
ge-0/0/20 {
    unit 0 {
        description "Sales file server port";
        family ethernet-switching {
            vlan members sales;
        }
    }
}
ge-0/0/24 {
    unit 0 {
        description "Support wireless access point port";
        family ethernet-switching {
            vlan members support;
        }
    }
}
ge-0/0/26 {
    unit 0 {
        description "Support phone port";
        family ethernet-switching {
            vlan members support;
        }
    }
}
ge-0/0/44 {
    unit 0 {
        description "Support printer port";
        family ethernet-switching {
            vlan members support;
        }
    }
}
ge-0/0/46 {
    unit 0 {
        description "Support file server port";
        family ethernet-switching {
            vlan members support;
        }
    }
}
vpls {
    unit 0 {
        family inet address 192.0.2.0/25;
    }
    unit 1 {
        family inet address 192.0.2.128/25;
    }
}
}
vpls {
    sales {

```

```

        vlan-id 100;
        interface ge-0/0/0.0;
        interface ge-0/0/3.0;
        interface ge-0/0/20.0;
        interface ge-0/0/22.0;
        I3-interface vlan 0;
    }
    support {
        vlan-id 200;
        interface ge-0/0/24.0;
        interface ge-0/0/26.0;
        interface ge-0/0/44.0;
        interface ge-0/0/46.0;
        I3-interface vlan 1;
    }
}

```



TIP: To quickly configure the sales and support VLAN interfaces, issue the `load merge` terminal command, then copy the hierarchy and paste it into the switch terminal window.

Verification

Verify that the “sales” and “support” VLANs have been created and are operating properly, perform these tasks:

- Verifying That the VLANs Have Been Created and Associated to the Correct Interfaces on page 34
- Verifying That Traffic Is Being Routed Between the Two VLANs on page 35
- Verifying That Traffic Is Being Switched Between the Two VLANs on page 35

Verifying That the VLANs Have Been Created and Associated to the Correct Interfaces

Purpose Verify that the VLANs sales and support have been created on the switch and that all connected interfaces on the switch are members of the correct VLAN.

Action List all VLANs configured on the switch:

Use the operational mode commands:

```

user@switch> show vlans
Name      Tag      Interfaces
default
          ge-0/0/1.0, ge-0/0/2.0, ge-0/0/4.0, ge-0/0/5.0,
          ge-0/0/6.0, ge-0/0/7.0, ge-0/0/8.0, ge-0/0/9.0,
          ge-0/0/10.0*, ge-0/0/11.0, ge-0/0/12.0, ge-0/0/13.0*,
          ge-0/0/14.0, ge-0/0/15.0, ge-0/0/16.0, ge-0/0/17.0,
          ge-0/0/18.0, ge-0/0/19.0, ge-0/0/21.0, ge-0/0/23.0*,
          ge-0/0/25.0, ge-0/0/27.0, ge-0/0/28.0, ge-0/0/29.0,
          ge-0/0/30.0, ge-0/0/31.0, ge-0/0/32.0, ge-0/0/33.0,

```

```

                                ge-0/0/34.0, ge-0/0/35.0, ge-0/0/36.0, ge-0/0/37.0,
                                ge-0/0/38.0, ge-0/0/39.0, ge-0/0/40.0, ge-0/0/41.0,
                                ge-0/0/42.0, ge-0/0/43.0, ge-0/0/45.0, ge-0/0/47.0,
                                ge-0/1/0.0*, ge-0/1/1.0*, ge-0/1/2.0*, ge-0/1/3.0*

sales      100
                                ge-0/0/0.0*, ge-0/0/3.0, ge-0/0/20.0, ge-0/0/22.0

support    200
                                ge-0/0/0.24, ge-0/0/26.0, ge-0/0/44.0, ge-0/0/46.0*

mgmt
                                me0.0*

```

Meaning The `show vlans` command lists all VLANs configured on the switch and which interfaces are members of each VLAN. This command output shows that the `sales` and `support` VLANs have been created. The `sales` VLAN has a tag ID of 100 and is associated with interfaces `ge-0/0/0.0`, `ge-0/0/3.0`, `ge-0/0/20.0`, and `ge-0/0/22.0`. VLAN `support` has a tag ID of 200 and is associated with interfaces `ge-0/0/24.0`, `ge-0/0/26.0`, `ge-0/0/44.0`, and `ge-0/0/46.0`.

Verifying That Traffic Is Being Routed Between the Two VLANs

Purpose Verify routing between the two VLANs.

Action List the Layer 3 routes in the switch's Address Resolution Protocol (ARP) table:

```

user@switch> show arp
MAC Address      Address      Name      Flags
00:00:0c:06:2c:0d  192.0.2.3   vlan.0    None
00:13:e2:50:62:e0  192.0.2.11  vlan.1    None

```

Meaning Sending IP packets on a multiaccess network requires mapping from an IP address to a MAC address (the physical or hardware address). The ARP table displays the mapping between the IP address and MAC address for both `vlan.0` (associated with `sales`) and `vlan.1` (associated with `support`). These VLANs can route traffic to each other.

Verifying That Traffic Is Being Switched Between the Two VLANs

Purpose Verify that learned entries are being added to the Ethernet switching table.

Action List the contents of the Ethernet switching table:

```

user@switch> show ethernet-switching table

Ethernet-switching table: 8 entries, 5 learned
VLAN      MAC address      Type      Age Interfaces
default   *                Flood     - All-members
default   00:00:05:00:00:01 Learn     - ge-0/0/10.0
default   00:00:5e:00:01:09 Learn     - ge-0/0/13.0
default   00:19:e2:50:63:e0 Learn     - ge-0/0/23.0
sales     *                Flood     - All-members
sales     00:00:5e:00:07:09 Learn     - ge-0/0/0.0

```

```

support          *          Flood          - All-members
support          00:00:5e:00:01:01 Learn          - ge-0/0/46.0

```

Meaning The output shows that learned entries for the `sales` and `support` VLANs have been added to the Ethernet switching table, and are associated with interfaces `ge-0/0/0.0` and `ge-0/0/46.0`. Even though the VLANs were associated with more than one interface in the configuration, these interfaces are the only ones that are currently operating.

- Related Topics**
- Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21
 - Example: Connecting an Access Switch to a Distribution Switch on page 36
 - Example: Configure Automatic VLAN Administration Using GVRP on page 46
 - Understanding Bridging and VLANs on EX Series Switches on page 3

Example: Connecting an Access Switch to a Distribution Switch

In large local area networks (LANs), you commonly need to aggregate traffic from a number of access switches into a distribution switch.

This example describes how to connect an access switch to a distribution switch:

- Requirements on page 36
- Overview and Topology on page 37
- Configuring the Access Switch on page 38
- Configuring the Distribution Switch on page 43
- Verification on page 45

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, one EX3200-24P, which has twenty-four 1-Gigabit Ethernet ports, all of which support Power over Ethernet (PoE), and an uplink module with four 1-Gigabit Ethernet ports.
- JUNOS Release 9.0 or later for EX Series switches

Before you connect an access switch to a distribution switch, be sure you have:

- Installed the two switches. See *Installing and Connecting an EX3200 or EX4200 Switch*.

- Performed the initial software configuration on both switches. See Connecting and Configuring an EX Series Switch (J-Web Procedure).

Overview and Topology

In a large office that is spread across several floors or buildings, or in a data center, 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 access switch to a distribution switch.

In the topology, the LAN is segmented into two VLANs, one for the sales department and the second for the support team. One 1-Gigabit Ethernet port on the access switch's uplink module connects to the distribution switch, to one 1-Gigabit Ethernet port on the distribution switch.

Figure 3 on page 37 shows one EX4200 switch that is connected to the three access switches.

Figure 3: Topology for Configuration




Table 3 on page 37 explains the components of the example topology. The example shows how to configure one of the three access switches. The other access switches could be configured in the same manner.

Table 3: Components of the Topology for Connecting an Access Switch to a Distribution Switch

Property	Settings
Access switch hardware	EX3200-24P, 24 1-Gigabit Ethernet ports, all PoE-enabled (ge-0/0/0 through ge-0/0/23); one 4-port 1-Gigabit Ethernet uplink module (EX-UM-4SFP)

Table 3: Components of the Topology for Connecting an Access Switch to a Distribution Switch (continued)

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	sales, tag 100 support, tag 200
VLAN subnets	sales: 192.0.2.0/25 (addresses 192.0.2.1 through 192.0.2.126) support: 192.0.2.128/25 (addresses 192.0.2.129 through 192.0.2.254)
Trunk port interfaces	On the access switch: ge-0/1/0 On the distribution switch: ge-0/0/0
Access port interfaces in VLAN sales (on access switch)	Avaya IP telephones: ge-0/0/3 through ge-0/0/19 Wireless access points: ge-0/0/0 and ge-0/0/1 Printers: ge-0/0/22 and ge-0/0/23 File servers: ge-0/0/20 and ge-0/0/21
Access port interfaces in VLAN support (on access switch)	Avaya IP telephones: ge-0/0/25 through ge-0/0/43 Wireless access points: ge-0/0/24 Printers: ge-0/0/44 and ge-0/0/45 File servers: ge-0/0/46 and ge-0/0/47
Unused interfaces on access switch	ge-0/0/2 and ge-0/0/25

Configuring the Access Switch

To configure the access switch:

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

```
[edit]
set interfaces ge-0/0/0 unit 0 description "Sales Wireless access point port"
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/3 unit 0 description "Sales phone port"
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/22 unit 0 description "Sales printer port"
set interfaces ge-0/0/22 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/20 unit 0 description "Sales file server port"
set interfaces ge-0/0/20 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/24 unit 0 description "Support wireless access point port"
set interfaces ge-0/0/24 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/26 unit 0 description "Support phone port"
set interfaces ge-0/0/26 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/44 unit 0 description "Support printer port"
set interfaces ge-0/0/44 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/46 unit 0 description "Support file server port"
set interfaces ge-0/0/46 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/1/0 unit 0 description "Uplink module port connection to
distribution switch"
set interfaces ge-0/1/0 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/1/0 unit 0 family ethernet-switching native-vlan-id 1
set interfaces ge-0/1/0 unit 0 family ethernet switching vlan members [sales
support]
set interfaces vlan unit 0 family inet address 192.0.2.1/25
```

```

set interfaces vlan unit 1 family inet address 192.0.2.129/25
set vlans sales interface ge-0/0/0.0
set vlans sales interface ge-0/0/3.0
set vlans sales interface ge-0/0/22.0
set vlans sales interface ge-0/0/20.0
set vlans sales l3-interface vlan.0
set vlans sales vlan-id 100
set vlans sales vlan-description "Sales VLAN"
set vlans support interface ge-0/0/24.0
set vlans support interface ge-0/0/26.0
set vlans support interface ge-0/0/44.0
set vlans support interface ge-0/0/46.0
set vlans support vlan-id 200
set vlans support l3-interface vlan.1
set vlans support vlan-description "Support VLAN"

```

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

1. Configure the 1-Gigabit Ethernet interface on the uplink module to be the trunk port that connects to the distribution switch:

```

[edit interfaces ge-0/1/0 unit 0]
user@access-switch# set description "Uplink module port connection to
distribution switch"
user@access-switch# set ethernet-switching port-mode trunk

```

2. Specify the VLANs to be aggregated on the trunk port:

```

[edit interfaces ge-0/1/0 unit 0]
user@access-switch# set ethernet-switching vlan members [ sales support
]

```

3. Configure the VLAN ID to use for packets that are received with no dot1q tag (untagged packets):

```

[edit interfaces ge-0/1/0 unit 0]
user@access-switch# set ethernet-switching native-vlan-id 1

```

4. Configure the sales VLAN:

```

[edit vlans sales]
user@access-switch# set vlan-description "Sales VLAN"
user@access-switch# set vlan-id 100
user@access-switch# set l3-interface vlan.0

```

5. Configure the support VLAN:

```

[edit vlans support]
user@access-switch# set vlan-description "Support VLAN"
user@access-switch# set vlan-id 200
user@access-switch# set l3-interface vlan.1

```

6. Create the subnet for the sales broadcast domain:

```
[edit interfaces]
user@access-switch# set vlan unit 0 family inet address 192.0.2.1/25
```

7. Create the subnet for the support broadcast domain:

```
[edit interfaces]
user@access-switch# set vlan unit 1 family inet address 192.0.2.129/25
```

8. Configure the interfaces in the sales VLAN:

```
[edit interfaces]
user@access-switch# set ge-0/0/0 unit 0 description "Sales wireless access
point port"
user@access-switch# set ge-0/0/0 unit 0 family ethernet-switching vlan
members sales
user@access-switch# set ge-0/0/3 unit 0 description "Sales phone port"
user@access-switch# set ge-0/0/3 unit 0 family ethernet-switching vlan
members sales
user@access-switch# set ge-0/0/20 unit 0 description "Sales file server
port"
user@access-switch# set ge-0/0/20 unit 0 family ethernet-switching vlan
members sales
user@access-switch# set ge-0/0/22 unit 0 description "Sales printer port"
user@access-switch# set ge-0/0/22 unit 0 family ethernet-switching vlan
members sales
```

9. Configure the interfaces in the support VLAN:

```
[edit interfaces]
user@access-switch# set ge-0/0/24 unit 0 description "Support wireless
access point port"
user@access-switch# set ge-0/0/24 unit 0 family ethernet-switching vlan
members support
user@access-switch# set ge-0/0/26 unit 0 description "Support phone port"
user@access-switch# set ge-0/0/26 unit 0 family ethernet-switching vlan
members support
user@access-switch# set ge-0/0/44 unit 0 description "Support printer
port"
user@access-switch# set ge-0/0/44 unit 0 family ethernet-switching vlan
members support
user@access-switch# set ge-0/0/46 unit 0 description "Support file server
port"
user@access-switch# set ge-0/0/46 unit 0 family ethernet-switching vlan
members support
```

10. Configure descriptions and VLAN tag IDs for the sales and support VLANs:

```
[edit vlans]
user@access-switch# set sales vlan-description "Sales VLAN"
user@access-switch# set sales vlan-id 100
user@access-switch# set support vlan-description "Support VLAN"
user@access-switch# set support vlan-id 200
```

11. To route traffic between the sales and support VLANs and associate a Layer 3 interface with each VLAN:

```
[edit vlans]
user@access-switch# set sales 13-interface vlan.0
user@access-switch# set support 13-interface vlan.1
```

Results Display the results of the configuration:

```
user@access-switch> show
interfaces {
  ge-0/0/0 {
    unit 0 {
      description "Sales wireless access point port";
      family ethernet-switching {
        vlan members sales;
      }
    }
  }
  ge-0/0/3 {
    unit 0 {
      description "Sales phone port";
      family ethernet-switching {
        vlan members sales;
      }
    }
  }
  ge-0/0/20 {
    unit 0 {
      description "Sales file server port";
      family ethernet-switching {
        vlan members sales;
      }
    }
  }
  ge-0/0/22 {
    unit 0 {
      description "Sales printer port";
      family ethernet-switching {
        vlan members sales;
      }
    }
  }
  ge-0/0/24 {
    unit 0 {
      description "Support wireless access point port";
      family ethernet-switching {
        vlan members support;
      }
    }
  }
  ge-0/0/26 {
    unit 0 {
      description "Support phone port";
      family ethernet-switching {
        vlan members support;
      }
    }
  }
}
```

```

    }
  }
}
ge-0/0/44 {
  unit 0 {
    description "Support printer port";
    family ethernet-switching {
      vlan members sales;
    }
  }
}
ge-0/0/46 {
  unit 0 {
    description "Support file server port";
    family ethernet-switching {
      vlan members support;
    }
  }
}
ge-0/1/0 {
  unit 0 {
    description "Uplink module port connection to distribution switch";
    family ethernet-switching {
      port-mode trunk;
      vlan members [ sales support ];
      native-vlan-id 1;
    }
  }
}
vlan {
  unit 0 {
    family inet address 192.0.2.1/25;
  }
  unit 1 {
    family inet address 192.0.2.129/25;
  }
}
vllans {
  sales {
    vlan-id 100;
    vlan-description "Sales VLAN";
    I3-interface vlan.0;
  }
  support {
    vlan-id 200;
    vlan-description "Support VLAN";
    I3-interface vlan.1;
  }
}
}

```



TIP: To quickly configure the distribution switch, issue the `load merge terminal` command, then copy the hierarchy and paste it into the switch terminal window.

Configuring the Distribution Switch

To configure the distribution switch:

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

```
set interfaces ge-0/0/0 description "Connection to access switch"
set interfaces ge-0/0/0 ethernet-switching port-mode trunk
set interfaces ge-0/0/0 ethernet-switching vlan members [ sales support ]
set interfaces ge-0/0/0 ethernet-switching native-vlan-id 1
set interfaces vlan unit 0 family inet address 192.0.2.2/25
set interfaces vlan unit 1 family inet address 192.0.2.130/25
set vlans sales vlan-description "Sales VLAN"
set vlans sales vlan-id 100
set vlans sales l3-interface vlan.0
set vlans support vlan-description "Support VLAN"
set vlans support vlan-id 200
set vlans support l3-interface vlan.1
```

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

1. Configure the interface on the switch to be the trunk port that connects to the access switch:

```
[edit interfaces ge-0/0/0 unit 0]
user@distribution-switch# set description "Connection to access switch"
user@distribution-switch# set ethernet-switching port-mode trunk
```

2. Specify the VLANs to be aggregated on the trunk port:

```
[edit interfaces ge-0/0/0 unit 0]
user@distribution-switch# set ethernet-switching vlan members [ sales
support ]
```

3. Configure the VLAN ID to use for packets that are received with no dot1q tag (untagged packets):

```
[edit interfaces]
user@distribution-switch# set ge-0/0/0 ethernet-switching native-vlan-id 1
```

4. Configure the sales VLAN:

```
[edit vlans sales]
user@distribution-switch# set vlan-description "Sales VLAN"
user@distribution-switch# set vlan-id 100
user@distribution-switch# set l3-interface vlan.0
```

5. Configure the support VLAN:

```
[edit vlans support]
user@distribution-switch# set vlan-description "Support VLAN"
user@distribution-switch# set vlan-id 200
```

```
user@distribution-switch# set l3-interface vlan.1
```

6. Create the subnet for the sales broadcast domain:

```
[edit interfaces]
user@distribution-switch# set vlan unit 0 family inet address 192.0.2.2/25
```

7. Create the subnet for the support broadcast domain:

```
[edit interfaces]
user@distribution-switch# set vlan unit 1 family inet address
192.0.2.130/25
```

Results Display the results of the configuration:

```
user@distribution-switch> show
interfaces {
  ge-0/0/0 {
    description "Connection to access switch";
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan members [ sales support ];
        native-vlan-id 1;
      }
    }
  }
  vlan {
    unit 0 {
      family inet address 192.0.2.2/25;
    }
    unit 1 {
      family inet address 192.0.2.130/25;
    }
  }
}
vlans {
  sales {
    vlan-id 100;
    vlan-description "Sales VLAN";
    l3-interface vlan.0;
  }
  support {
    vlan-id 200;
    vlan-description "Support VLAN";
    l3-interface vlan.1;
  }
}
```



TIP: To quickly configure the distribution switch, issue the `load merge terminal` command, then copy the hierarchy and paste it into the switch terminal window.

Verification

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

- Verifying the VLAN Members and Interfaces on the Access Switch on page 45
- Verifying the VLAN Members and Interfaces on the Distribution Switch on page 45

Verifying the VLAN Members and Interfaces on the Access Switch

Purpose Verify that the `sales` and `support` have been created on the switch.

Action List all VLANs configured on the switch:

```
user@switch> show vlans
```

Name	Tag	Interfaces
default		ge-0/0/1.0, ge-0/0/2.0, ge-0/0/4.0, ge-0/0/5.0, ge-0/0/6.0, ge-0/0/7.0, ge-0/0/8.0*, ge-0/0/9.0, ge-0/0/10.0, ge-0/0/11.0*, ge-0/0/12.0, ge-0/0/13.0, ge-0/0/14.0, ge-0/0/15.0, ge-0/0/16.0, ge-0/0/17.0, ge-0/0/18.0, ge-0/0/19.0*, ge-0/0/21.0, ge-0/0/23.0, ge-0/0/25.0, ge-0/0/27.0*, ge-0/0/28.0, ge-0/0/29.0, ge-0/0/30.0, ge-0/0/31.0*, ge-0/0/32.0, ge-0/0/33.0, ge-0/0/34.0, ge-0/0/35.0*, ge-0/0/36.0, ge-0/0/37.0, ge-0/0/38.0, ge-0/0/39.0*, ge-0/0/40.0, ge-0/0/41.0, ge-0/0/42.0, ge-0/0/43.0*, ge-0/0/45.0, ge-0/0/47.0, ge-0/1/1.0*, ge-0/1/2.0*, ge-0/1/3.0*
sales	100	ge-0/0/0.0*, ge-0/0/3.0, ge-0/0/20.0, ge-0/0/22.0, ge-0/1/0.0*,
support	200	ge-0/0/24.0*, ge-0/0/26.0, ge-0/0/44.0, ge-0/0/46.0,
mgmt		me0.0*

Meaning The output shows the `sales` and `support` VLANs and the interfaces associated with them.

Verifying the VLAN Members and Interfaces on the Distribution Switch

Purpose Verify that the `sales` and `support` have been created on the switch.

Action List all VLANs configured on the switch:

```
user@switch> show vlans
```

Name	Tag	Interfaces
default		ge-0/0/1.0, ge-0/0/2.0, ge-0/0/3.0, ge-0/0/4.0, ge-0/0/5.0, ge-0/0/6.0, ge-0/0/7.0*, ge-0/0/8.0, ge-0/0/9.0, ge-0/0/10.0*, ge-0/0/11.0, ge-0/0/12.0, ge-0/0/13.0, ge-0/0/14.0, ge-0/0/15.0, ge-0/0/16.0, ge-0/0/17.0, ge-0/0/18.0*, ge-0/0/19.0, ge-0/0/20.0, ge-0/0/21.0, ge-0/0/22.0*, ge-0/0/23.0, ge-0/1/1.0*, ge-0/1/2.0*, ge-0/1/3.0*
sales	100	ge-0/0/0.0*
support	200	ge-0/0/0.0*
mgmt		me0.0*

Meaning The output shows the **sales** and **support** VLANs associated to interface **ge-0/0/0.0**. Interface **ge-0/0/0.0** is the trunk interface connected to the access switch.

- Related Topics**
- Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21
 - Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 28
 - Example: Configure Automatic VLAN Administration Using GVRP on page 46
 - Understanding Bridging and VLANs on EX Series Switches on page 3

Example: Configure Automatic VLAN Administration Using GVRP

As a network expands and the number of clients and VLANs increases, VLAN administration becomes complex, and the task of efficiently configuring VLANs on multiple EX Series switches becomes increasingly difficult. To automate VLAN administration, you can enable GARP VLAN Registration Protocol (GVRP) on the network.



NOTE: Only trunk interfaces can be enabled for GVRP.

This example describes how to use GVRP to automate administration of VLAN membership changes within your network:

- Requirements on page 47
- Overview and Topology on page 47
- Configuring VLANs and GVRP on Access Switch A on page 49
- Configuring VLANs and GVRP on Access Switch B on page 52

- Configuring VLANS and GVRP on the Distribution Switch on page 56
- Verification on page 58

Requirements

This example uses the following hardware and software components:

- Two EX3200 access switches
- One EX4200 distribution switch
- JUNOS Release 9.0 or later for EX Series switches

Before you configure GVRP on the access switches and on the distribution switch, be sure you have:

- Performed the initial software configuration on the switches. See *Connecting and Configuring an EX Series Switch (J-Web Procedure)*.
- Configured the VLANs on both the access switches and on the distribution switch. (Dynamic VLAN configuration is not supported.)
- Configured a trunk interface on all the switches.

Overview and Topology

When you are setting up your network, you should configure all VLANs on all switches, even though some switches are not actively participating in a VLAN. Then enable GVRP on the trunk interface of each switch. GVRP ensures that the VLAN membership information on the trunk interface is updated as the switch's access interfaces become active or inactive in the configured VLANs.

You do not need to take an extra step of explicitly binding a VLAN to the trunk interface. When GVRP is enabled, the trunk interface advertises all the VLANs that are active (bound to access interfaces) on that switch. A GVRP-enabled trunk interface does not advertise VLANs that have been configured on the switch but that are not currently bound to an access interface. Thus, GVRP provides the benefit of reducing network overhead—by limiting the scope of broadcast, unknown unicast, and multicast (BUM) traffic to interested devices only.

This example shows a network with three VLANs: finance, sales, and lab.

Access Switch A has been configured to support all three VLANs and all three VLANs are active, bound to interfaces that are connected to personal computers:

- `ge-0/0/1`— Connects PC1 as member of finance vlan, VLAN ID 100
- `ge-0/0/2`— Connects PC2 as member of lab vlan, VLAN ID 200
- `ge-0/0/3`— Connects PC3 as member of sales vlan, VLAN ID 300

Access Switch B has also been configured to support three VLANs. However, currently only two VLANs are active, bound to interfaces that are connected to personal computers:

- ge-0/0/0— Connects PC4 as member of finance vlan, VLAN ID 100
- ge-0/0/1— Connects PC5 as member of lab vlan, VLAN ID 200

The Distribution Switch is also configured to support the three VLANs (finance, lab, sales). However, the Distribution Switch does not have any access interfaces that are connecting devices as members of these VLANs. The Distribution Switch has two trunk interfaces:

- xe-0/1/1— Connects Distribution Switch to Access Switch A.
- xe-0/1/0— Connects Distribution Switch to Access Switch B.

Figure 4 on page 48 shows GVRP configured on two access switches and one distribution switch.

Figure 4: GVRP Configured on Two Access Switches and One Distribution Switch for Automatic VLAN Administration




Table 4: Components of the Network Topology

Property	Settings
Switch hardware	<ul style="list-style-type: none"> ■ Access Switch A—EX3200 switch ■ Access Switch B—EX3200 access switch ■ Distribution Switch—EX4200 switch

Table 4: Components of the Network Topology (continued)

VLAN names and tag IDs	finance, tag 100 lab, tag 200 sales, tag 300
Interfaces	<p>Access Switch A Interfaces</p> <ul style="list-style-type: none"> ■ ge-0/0/1— Connects PC1 to Access Switch A. ■ ge-0/0/2— Connects PC2 to Access Switch A. ■ ge-0/0/3— Connects PC3 to Access Switch A. ■ xe-0/1/1— Connects Access Switch A to Distribution Switch. (trunk) <p>Access Switch B Interfaces</p> <ul style="list-style-type: none"> ■ ge-0/0/0— Connects PC4 to Access Switch B. ■ ge-0/0/1— Connects PC5 to Access Switch B. ■ xe-0/1/0— Connects Access Switch B to Distribution Switch. (trunk) <p>Distribution Switch Interfaces</p> <ul style="list-style-type: none"> ■ xe-0/1/1— Connects Distribution Switch to Access Switch A. (trunk) ■ xe-0/1/0— Connects Distribution Switch to Access Switch B. (trunk)

When VLAN access interfaces become active or inactive, GVRP ensures that the updated information is advertised on the trunk interface. Thus, the Distribution Switch does not forward traffic to inactive VLANs.

Configuring VLANs and GVRP on Access Switch A

To configure three VLANs on the switch, bind access interfaces for PC1, PC2, and PC3 to the VLANs (finance, lab, sales), and enable GVRP on the trunk interface of Access Switch A, perform these tasks:

CLI Quick Configuration To quickly configure Access Switch A to support the three VLANs, bind interfaces for the three PCs to the appropriate VLANs, and enable GVRP on the trunk interface, copy the following commands and paste them into the switch terminal window of Switch A:

```
[edit]
set vlans finance vlan-id 100
set vlans lab vlan-id 200
set vlans sales vlan-id 300
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members finance
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members lab
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members sales
set interfaces xe-0/1/1 unit 0 family ethernet-switching port-mode trunk
set protocols gvrp interface xe-0/1/1.0
```



NOTE: As we recommend, default GVRP timers are used in this example. The default values associated with each GVRP timer are: 200 ms for the `join-timer`, 600 ms for the `leave-timer`, and 1000 cs (10000 ms) for the `leaveall-timer`. Modifying timers to inappropriate values may cause an imbalance in the operation of GVRP. Refer to IEEE 802.1D [2004] Clause 12 for more information. The timer values are displayed when you use the `show gvrp` command to verify that GVRP is enabled. For more information on the timers, see `gvrp` and its associated configuration statements.

Step-by-Step Procedure To configure Access Switch A to support the three VLANs, bind interfaces for the three PCs to the appropriate VLANs, and enable GVRP on the trunk interface, copy the following commands and paste them into the switch terminal window of Switch A:

1. Configure the finance VLAN:

```
[edit]
user@Access-Switch-A# set vlans finance vlan-id 100
```

2. Configure the lab VLAN:

```
[edit]
user@Access-Switch-A# set vlans lab vlan-id 200
```

3. Configure the sales VLAN:

```
[edit]
user@Access-Switch-A# set vlans sales vlan-id 300
```

4. Configure an Ethernet interface as a member of the finance VLAN:

```
[edit]
user@Access-Switch-A# set interfaces ge-0/0/1 unit 0 family
  ethernet-switching vlan members finance
```

5. Configure an Ethernet interface as a member of the lab VLAN:

```
[edit]
user@Access-Switch-A# set interfaces ge-0/0/2 unit 0 family
  ethernet-switching vlan members lab
```

6. Configure an Ethernet interface as a member of the sales VLAN:

```
[edit]
user@Access-Switch-A# set interfaces ge-0/0/3 unit 0 family
  ethernet-switching vlan members sales
```

7. Configure a trunk interface:

```
user@Access-Switch-A# set interfaces xe-0/1/1 unit 0 family
ethernet-switching port-mode trunk
```

8. Enable GVRP on the trunk interface:

```
[edit]
user@Access-Switch-A# set protocols gvrp interface xe-0/1/1.0
```

Results Check the results of the configuration:

```
interfaces {
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members finance;
        }
      }
    }
  }
  ge-0/0/2 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members lab;
        }
      }
    }
  }
  ge-0/0/3 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members sales;
        }
      }
    }
  }
  xe-0/1/1 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
      }
    }
  }
  ge-0/1/2 {
    unit 0 {
```

```

        family ethernet-switching;
    }
}
ge-0/1/3 {
    unit 0 {
        family ethernet-switching;
    }
}
}
protocols {
    igmp-snooping {
        vlan all;
    }
    lldp {
        interface all;
    }
    lldp-med {
        interface all;
    }
    gvrp {
        interface xe-0/1/1.0;
    }
    rstp;
}
ethernet-switching-options {
    storm-control {
        interface all {
            level 50;
        }
    }
}
}
vlands {
    finance {
        vlan-id 100;
    }
    lab {
        vlan-id 200;
    }
    sales {
        vlan-id 300;
    }
}
}

```

Configuring VLANs and GVRP on Access Switch B

To configure three VLANs on the switch, bind access interfaces for PC4 and PC5 to the VLANs (finance and lab), and enable GVRP on the trunk interface of Access Switch B, perform these tasks:

CLI Quick Configuration To quickly configure Access Switch B to support the three VLANs, bind interfaces for the two PCs to the appropriate VLANs, and enable GVRP on the trunk interface, copy the following commands and paste them into the switch terminal window of Switch B:

```
[edit]
```

```

set vlans finance vlan-id 100
set vlans lab vlan-id 200
set vlans sales vlan-id 300
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members finance
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members lab
set interfaces xe-0/1/0 unit 0 family ethernet-switching port-mode trunk
set protocols gvrp interface xe-0/1/0.0

```

Step-by-Step Procedure To configure Access Switch B to support the three VLANs, bind interfaces for the two PCs to the appropriate VLAN, and enable GVRP on the trunk interface, copy the following commands and paste them into the switch terminal window of Switch B:

1. Configure the finance VLAN:

```

[edit]
user@Access-Switch-B# set vlans finance vlan-id 100

```

2. Configure the lab VLAN:

```

[edit]
user@Access-Switch-B# set vlans lab vlan-id 200

```

3. Configure the sales VLAN:

```

[edit]
user@Access-Switch-B# set vlans sales vlan-id 300

```

4. Configure an Ethernet interface as a member of the finance VLAN:

```

[edit]
user@Access-Switch-B# set interfaces ge-0/0/0 unit 0 family
ethernet-switching vlan members finance

```

5. Configure an Ethernet interface as a member of the lab VLAN:

```

[edit]
user@Access-Switch-B# set interfaces ge-0/0/1 unit 0 family
ethernet-switching vlan members lab

```

6. Configure a trunk interface:

```

user@Access-Switch-B# set interfaces xe-0/1/0 unit 0 family
ethernet-switching port-mode trunk

```

7. Enable GVRP on the trunk interface:

```

[edit]
user@Access-Switch-B# set protocols gvrp xe-0/1/0.0

```



NOTE: As we recommend, default GVRP timers are used in this example. The default values associated with each GVRP timer are: 200 ms for the **join-timer**, 600 ms for the **leave-timer**, and 1000 cs (10000 ms) for the **leaveall-timer**. Modifying timers to inappropriate values may cause an imbalance in the operation of GVRP. Refer to IEEE 802.1D [2004] Clause 12 for more information. The timer values are displayed when you use the **show gvrp** command to verify that GVRP is enabled. For more information on the timers, see **gvrp** and its associated configuration statements.

Results Check the results of the configuration:

```
[edit]
user@Access-Switch-B #show
interfaces {
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members finance;
        }
      }
    }
  }
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members lab;
        }
      }
    }
  }
  ge-0/0/2 {
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/3 {
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/1/0 {
    unit 0 {
      family ethernet-switching;
    }
  }
  xe-0/1/0 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
      }
    }
  }
}
```

```

    }
  }
  ge-0/1/1 {
    unit 0 {
      family ethernet-switching;
    }
  }
  xe-0/1/1 {
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/1/2 {
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/1/3 {
    unit 0 {
      family ethernet-switching;
    }
  }
}
protocols {
  igmp-snooping {
    vlan all;
  }
  lldp {
    interface all;
  }
  lldp-med {
    interface all;
  }
  gvrp {
    interface xe-0/1/0.0;
  }
  rstp;
}
ethernet-switching-options {
  storm-control {
    interface all {
      level 50;
    }
  }
}
}
vlans {
  finance {
    vlan-id 100;
  }
  lab {
    vlan-id 200;
  }
  sales {
    vlan-id 300;
  }
}
}

```

Configuring VLANs and GVRP on the Distribution Switch

CLI Quick Configuration To quickly configure the finance, lab, and sales VLANs on the Distribution Switch and to enable GVRP on the trunk interface of the Distribution Switch, copy the following commands and paste them into the switch terminal window of the Distribution Switch:

```
[edit]
set vlans finance vlan-id 100
set vlans lab vlan-id 200
set vlans sales vlan-id 300
set interfaces xe-0/1/1 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/1/0 unit 0 family ethernet-switching port-mode trunk
set protocols gvrp interface xe-0/1/1.0
set protocols gvrp interface xe-0/1/0.0
```

Step-by-Step Procedure To configure the three VLANs on the Distribution Switch, to configure the trunk interfaces, and to enable GVRP on the trunk interface of the Distribution Switch:

1. Configure the finance VLAN:

```
[edit]
user@Distribution-Switch# set vlans finance vlan-id 100
```

2. Configure the lab VLAN:

```
[edit]
user@Distribution-Switch# set vlans lab vlan-id 200
```

3. Configure the sales VLAN:

```
[edit]
user@Distribution-Switch# set vlans sales vlan-id 300
```

4. Configure the trunk interface to Access Switch A:

```
[edit]
user@Distribution-Switch# set interfaces xe-0/1/1 unit 0 family
ethernet-switching port-mode trunk
```

5. Configure the trunk interface to Access Switch B:

```
[edit]
user@Distribution-Switch# set interfaces xe-0/1/0 unit 0 family
ethernet-switching port-mode trunk
```

6. Enable GVRP on the trunk interface for xe-0/1/1 :

```
[edit]
user@Distribution-Switch# set protocols gvrp interface xe-0/1/1.0
```

7. Enable GVRP on the trunk interface for xe-0/1/0 :

```
[edit]
user@Distribution-Switch# set protocols gvrp interface xe-0/1/0.0
```

Results Display the results of the configuration:

```
[edit]
user@Distribution Switch-D #show
interfaces {
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/2 {
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/3 {
    unit 0 {
      family ethernet-switching;
    }
  }
  xe-0/1/0 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
      }
    }
  }
  ge-0/1/1 {
    unit 0 {
      family ethernet-switching;
    }
  }
  xe-0/1/1 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
      }
    }
  }
  ge-0/1/2 {
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/1/3 {
```

```

        unit 0 {
            family ethernet-switching;
        }
    }
    protocols {
        igmp-snooping {
            vlan all;
        }
        lldp {
            interface all;
        }
        lldp-med {
            interface all;
        }
        gvrp {
            interface xe-0/1/0.0;
            interface xe-0/1/1.0;
        }
        rstp;
    }
    ethernet-switching-options {
        storm-control {
            interface all {
                level 50;
            }
        }
    }
    vlans {
        finance {
            vlan-id 100;
        }
        lab {
            vlan-id 300;
        }
        sales {
            vlan-id 300;
        }
    }
}

```

Verification

To confirm that the configuration is updating VLAN membership, perform these tasks:

- Verifying That GVRP Is Enabled on Access Switch A on page 59
- Verifying That GVRP Is Updating VLAN Membership on Switch A on page 59
- Verifying That GVRP Is Enabled on Access Switch B on page 59
- Verifying That GVRP Is Updating VLAN Membership on Switch B on page 60
- Verifying That GVRP Is Enabled on the Distribution Switch on page 60
- Verifying That GVRP Is Updating VLAN Membership on the Distribution Switch on page 60

Verifying That GVRP Is Enabled on Access Switch A

Purpose Verify that GVRP is enabled on the switch.

Action Show the GVRP configuration, using the `show gvrp` command:

```
user@Access-Switch-A> show gvrp

Global GVRP configuration
  GVRP status : Enabled
  GVRP Timers (ms)
    Join      : 200
    Leave     : 600
    LeaveAll  : 10000
Interface Name      Protocol Status
-----
xe-0/1/1.0         Enabled
```

Meaning The results show that GVRP is enabled on the trunk interface of Switch A and that the default timers are used.

Verifying That GVRP Is Updating VLAN Membership on Switch A

Purpose To verify that GVRP is updating VLAN membership, display the Ethernet switching interfaces and associated VLANs that are active on switch A:

Action List Ethernet switching interfaces on the switch, using the `show ethernet-switching interfaces` command:

```
user@Access-Switch-A> show ethernet-switching interfaces

Interface  State  VLAN members      Blocking
ge-0/0/1.0 up     finance           unblocked
ge-0/0/2.0 up     lab               unblocked
ge-0/0/3.0 up     sales            unblocked
xe-0/1/1.0 up     finance          unblocked
           up     lab              unblocked
```

Meaning GVRP has automatically added `finance` and `lab` as VLAN members on the trunk interface, because they are being advertised by Access Switch B.

Verifying That GVRP Is Enabled on Access Switch B

Purpose Verify that GVRP is enabled on the switch.

Action Show the GVRP configuration:

```
user@Access-Switch-B> show gvrp

Global GVRP configuration
  GVRP status : Enabled
  GVRP Timers (ms)
    Join      : 200
    Leave     : 600
    LeaveAll  : 10000
Interface Name      Protocol Status
```

```

-----
xe-0/1/0.0      Enabled
    
```

Meaning The results show that GVRP is enabled on the trunk interface of Switch B and that the default timers are used.

Verifying That GVRP Is Updating VLAN Membership on Switch B

Purpose To verify that GVRP is updating VLAN membership, display the Ethernet switching interfaces and associated VLANs that are active on switch B:

Action List Ethernet switching interfaces on the switch:

```

user@Access-Switch-B> show ethernet-switching interfaces
Interface  State  VLAN members      Blocking
ge-0/0/0.0 up     finance           unblocked
ge-0/0/1.0 up     lab               unblocked
xe-0/1/1.0 up     finance          unblocked
              lab             unblocked
              sales          unblocked
    
```

Meaning GVRP has automatically added `finance`, `lab`, and `sales` as VLAN members on the trunk interface because they are being advertised by Access Switch A.

Verifying That GVRP Is Enabled on the Distribution Switch

Purpose Verify that GVRP is enabled on the switch.

Action Show the GVRP configuration:

```

user@Distribution-Switch> show gvrp

Global GVRP configuration
  GVRP status : Enabled
  GVRP Timers (ms)
    Join      : 200
    Leave    : 600
    LeaveAll  : 10000
Interface Name      Protocol Status
-----
xe-0/1/0.0          Enabled
xe-0/1/1.0          Enabled
    
```

Verifying That GVRP Is Updating VLAN Membership on the Distribution Switch

Purpose To verify that GVRP is updating VLAN membership on the distribution switch, display the Ethernet switching interfaces and associated VLANs on the Distribution Switch:

Action List the Ethernet switching interfaces on the switch:

```

user@Distribution-Switch> show ethernet-switching interfaces
Interface  State  VLAN members      Blocking
xe-0/1/1.0 up     finance          unblocked
              lab             unblocked
    
```

```

xe-0/1/0.0 up      sales      unblocked
                  finance    unblocked
                  lab       unblocked

```

Meaning The Distribution Switch has two trunk interfaces. Interface `xe-0/1/1.0` connects the Distribution Switch to Access Switch A and is therefore updated to show that it is a member of all the VLANs that are active on Access Switch A. Any traffic for those VLANs will be passed on from the Distribution Switch to Access Switch A, through interface `xe-0/1/1.0`. Interface `xe-0/1/0.0` connects the Distribution Switch to Access Switch B and is updated to show that it is a member of the two VLANs that are active on Access Switch B. Thus, the Distribution Switch sends traffic for `finance` and `lab` to both Access Switch A and Access Switch B. But the Distribution Switch sends traffic for `sales` only to Access Switch A.

- Related Topics**
- Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21
 - Understanding Bridging and VLANs on EX Series Switches on page 3

Example: Configuring Redundant Trunk Links for Faster Recovery

Simplify the convergence configuration in a typical enterprise network by configuring a primary link and a secondary link on trunk ports. If the primary link fails, the secondary link automatically takes over without waiting for normal STP convergence.

This example describes how to create a redundant trunk group:

- Requirements on page 61
- Overview and Topology on page 62
- Configuration on page 63
- Verification on page 64

Requirements

This example uses the following hardware and software components:

- Two EX4200 distribution switches.
- One EX3200 access switch.
- JUNOS Release 9.0 or later for EX Series switches

Before you configure the redundant trunk links network on the access and distribution switches, be sure you have:

- Installed the access switch. See *Installing and Connecting an EX3200 or EX4200 Switch*.
- Installed the two distribution switches. See *Installing and Connecting an EX3200 or EX4200 Switch*.
- Performed the initial switch configuration. See *Connecting and Configuring an EX Series Switch (J-Web Procedure)*.

Overview and Topology

This example shows a simple configuration to illustrate the basic steps for creating a redundant trunk group.

Configuring redundant trunk links places the primary link and the secondary link in a redundant group. However, a primary link need not be configured. If a primary link is not specified, the software compares the two links and selects the link with the highest port number as the active link. For example, if the two interfaces are `ge-0/1/0` and `ge-0/1/1`, the software assigns `ge-0/1/1` as the active link..

Whether a primary link is specified as the active link, or whether it is calculated by the software, traffic is handled in the same manner. Traffic passes through the active link but is blocked on the secondary link. If the active link goes down or is disabled administratively, the secondary link becomes active and begins forwarding traffic. However, there is a difference between the behavior of a primary, active link and an active link that is calculated to be active by the software. If an active link goes down, the secondary link begins forwarding traffic. If the old, active link comes up again, the following occurs:


- If the old, active link was configured as the primary link, then it resumes the role of active link and the other link is blocked. An interface configured as primary continues to carry with it the primary role whenever it becomes active.
- If no primary link was configured, and the active link was calculated by the software when the redundant group was formed, then the old, active link will not preempt the other interface (new active).



NOTE: The JUNOS Software for EX Series switches does not allow an interface to be in a redundant trunk group and in an STP topology at the same time.

Figure 5 on page 63 displays an example topology containing three switches. Switch 1 and Switch 2 make up the distribution layer, and Switch 3 makes up the access layer. Switch 3 is connected to the distribution layer through trunk ports `ge-0/0/9.0` (Link 1) and `ge-0/0/10.0` (Link 2).

Table 5 on page 63 lists the components used in this redundant trunk group.

Figure 5: Topology for Configuring the Redundant Trunk Links**Table 5: Components of the Redundant Trunk Link Topology**

Property	Settings
Switch hardware	<ul style="list-style-type: none"> ■ Switch 1–1 EX4200 distribution switch ■ Switch 2–1 EX4200 distribution switch ■ Switch 3–1 EX3200 access switch
Trunk port interfaces	On Switch 3 (access switch): ge-0/0/9.0 and ge-0/0/10.0
Redundant trunk group	group1

This configuration example creates a redundant trunk group called **group1** on Switch 3. The trunk ports **ge-0/0/9.0** and **ge-0/0/10.0** are the two links in **group1**. The trunk port **ge-0/0/9.0** will be configured administratively as the primary link. The trunk port **ge-0/0/10.0** will be the secondary link.

Configuration

CLI Quick Configuration To quickly configure the redundant trunk group **group1** on Switch 3, copy the following commands and paste them into the switch terminal window:

```
[edit]
set ethernet-switching-options redundant-trunk-group group-name group1
set ethernet-switching-options redundant-trunk-group group-name group1 interface
ge-0/0/9.0 primary
```

```
set ethernet-switching-options redundant-trunk-group group-name group1 interface
ge-0/0/10.0
```

Step-by-Step Procedure Configure the redundant trunk group `group1` on Switch 3 and specify the primary and secondary links.

1. Configure the redundant trunk group `group1`:

```
[edit ethernet-switching-options]
user@switch# set redundant-trunk-group group-name group1
```

2. Configure the trunk port `ge-0/0/9.0` as the primary link and `ge-0/0/10` as the secondary link:

```
[edit ethernet-switching-options]
user@switch# set redundant-trunk-group group-name group1 interface ge-0/0/9.0
primary
user@switch# set redundant-trunk-group group-name group1 interface
ge-0/0/10.0
```

Results Display the results of the configuration:

```
user@switch# show
ethernet-switching-options {
  redundant-trunk-group {
    group-name group1 {
      interface ge-0/0/9.0 primary;
      interface ge-0/0/10.0;
    }
  }
}
```

Verification

Verify that the redundant trunk group `group1` has been created and is operating properly:

- Verifying That the Redundant Group Has Been Created on page 64

Verifying That the Redundant Group Has Been Created

Purpose Verify that the redundant trunk group `group1` has been created on the switch and that trunk ports are members of the redundant trunk group.

Action List all redundant trunk groups configured on the switch:

```
user@switch> show redundant-trunk-group group1
Redundant-trunk-group: group1
Interfaces           : ge-0/0/9.0 (P) , DOWN
                    : ge-0/0/10.0 (A) , UP
Bandwidth            : 1000 Mbps, 1000 Mbps
```

```
Last Time of Flap : 1970-01-01 00:19:12 UTC (00:00:06 ago), Never
#Flaps           : 1, 0
```

Meaning The `show redundant-trunk-group` command lists all redundant trunk groups configured on the switch and which trunk links are members of the group. For this configuration example, the output shows that the redundant trunk group `group1` is configured on the switch. The (P) beside trunk port `ge-0/0/9.0` indicates that it is configured as the primary link. The (A) beside the `ge-0/0/10.0` trunk port indicates that it is the active link.

Related Topics ■ Understanding Redundant Trunk Links on EX Series Switches on page 11

Example: Setting Up Q-in-Q Tunneling on EX Series Switches

Service providers can use Q-in-Q tunneling to transparently pass Layer 2 VLAN traffic from a customer site, through the service provider network, to another customer site without removing or changing the customer VLAN tags or class-of-service (CoS) settings. You can configure Q-in-Q tunneling on EX Series switches.

This example describes how to set up Q-in-Q:

- Requirements on page 65
- Overview and Topology on page 65
- Configuration on page 66
- Verification on page 67

Requirements

This example requires one EX Series switch with JUNOS Release 9.3 or later for EX Series switches.

Before you begin setting up Q-in-Q tunneling, make sure you have created and configured the necessary customer VLANs. See “Configuring VLANs for EX Series Switches (CLI Procedure)” on page 98 or “Configuring VLANs for EX Series Switches (J-Web Procedure)” on page 95.

Overview and Topology

In this service provider network, there are multiple customer VLANs mapped to one service VLAN.

Table 6 on page 65 lists the settings for the example topology.

Table 6: Components of the Topology for Setting Up Q-in-Q Tunneling

Interface	Description
ge-0/0/11.0	Tagged S-VLAN trunk port

Table 6: Components of the Topology for Setting Up Q-in-Q Tunneling (continued)

ge-0/0/12.0	Untagged customer-facing access port
ge-0/0/13.0	Untagged customer-facing access port
ge-0/0/14.0	Tagged S-VLAN trunk port

Configuration

CLI Quick Configuration To quickly create and configure Q-in-Q tunneling, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans qinqvlan vlan-id 4001
set vlans qinqvlan dot1q-tunneling customer-vlans 1-100
set vlans qinqvlan dot1q-tunneling customer-vlans 201-300
set interfaces ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/11 unit 0 family ethernet-switching vlan members 4001
set interfaces ge-0/0/12 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/12 unit 0 family ethernet-switching vlan members 4001
set interfaces ge-0/0/13 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/13 unit 0 family ethernet-switching vlan members 4001
set interfaces ge-0/0/14 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/14 unit 0 family ethernet-switching vlan members 4001
set ethernet-switching-options dot1q-tunneling ether-type 0x9100
```

Step-by-Step Procedure To configure Q-in-Q tunneling:

1. Set the VLAN ID for the S-VLAN:

```
[edit vlans]
user@switch# set qinqvlan vlan-id 4001
```

2. Enable Q-in-Q tunneling and specify the customer VLAN ranges:

```
[edit vlans]
user@switch# set qinqvlan dot1q-tunneling customer-vlans 1-100
user@switch# set qinqvlan dot1q-tunneling customer-vlans 201-300
```

3. Set the port mode and VLAN information for the interfaces:

```
[edit interfaces]
user@switch# set ge-0/0/11 unit 0 family ethernet-switching port-mode
trunk
user@switch# set ge-0/0/11 unit 0 family ethernet-switching vlan members
4001
user@switch# set ge-0/0/12 unit 0 family ethernet-switching port-mode
access
user@switch# set ge-0/0/12 unit 0 family ethernet-switching vlan members
4001
user@switch# set ge-0/0/13 unit 0 family ethernet-switching port-mode
access
user@switch# set ge-0/0/13 unit 0 family ethernet-switching vlan members
4001
```

```

user@switch# set ge-0/0/14 unit 0 family ethernet-switching port-mode
trunk
user@switch# set ge-0/0/14 unit 0 family ethernet-switching vlan members
4001

```

4. Set the Q-in-Q Ethertype value:

```

[edit]
user@switch# set ethernet-switching-options dot1q-tunneling ether-type
0x9100

```

Results Check the results of the configuration:

```

user@switch> show configuration vlans qinqvlan
vlan-id 4001;
  dot1q-tunneling {
    customer-vlans [ 1-100 201-300 ];
  }

```

Verification

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

- Verifying That Q-in-Q Tunneling Was Enabled on page 67

Verifying That Q-in-Q Tunneling Was Enabled

Purpose Verify that Q-in-Q tunneling was properly enabled on the switch.

Action Use the show vlans command:

```

user@switch> show vlans qinqvlan extensive
VLAN: qinqvlan, Created at: Thu Sep 18 07:17:53 2008
802.1Q Tag: 4001, Internal index: 18, Admin State: Enabled, Origin: Static
Dot1q Tunneling Status: Enabled
Customer VLAN ranges:
    1-100
    201-300
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged 4 (Active = 0)
    ge-0/0/11.0, tagged, trunk
    ge-0/0/14.0, tagged, trunk
    ge-0/0/12.0, untagged, access
    ge-0/0/13.0, untagged, access

```

Meaning The output indicates that Q-in-Q tunneling is enabled and that the VLAN is tagged and shows the associated customer VLANs.

Related Topics ■ [Configuring Q-in-Q Tunneling \(CLI Procedure\) on page 106](#)

Example: Configuring a Private VLAN on an EX Series Switch

For security reasons, it is often useful to restrict the flow of broadcast and unknown unicast traffic and to even limit the communication between known hosts. The private VLAN (PVLAN) feature on EX Series switches allow an administrator to split a broadcast domain into multiple isolated broadcast subdomains, essentially putting a VLAN inside a VLAN.

This example describes how to create a private VLAN primary VLAN and secondary VLANs:



NOTE: Configuring a voice over IP (VoIP) VLAN on PVLAN interfaces is not supported.

- [Requirements on page 68](#)
- [Overview and Topology on page 68](#)
- [Configuration on page 69](#)
- [Verification on page 72](#)

Requirements

This example requires one EX Series switch with JUNOS Release 9.3 or later for EX Series switches.

Before you begin configuring a private VLAN, make sure you have created and configured the necessary VLAN. See “Configuring VLANs for EX Series Switches (CLI Procedure)” on page 98 or “Configuring VLANs for EX Series Switches (J-Web Procedure)” on page 95.

Overview and Topology

In a large office with multiple buildings and VLANs, you might need to isolate some workgroups or other endpoints for security reasons or to partition the broadcast domain. This configuration example shows a simple topology to illustrate how to create a private VLAN with one primary VLAN and two community VLANs, one for HR and one for finance, as well as two isolated ports for the mail server and the backup server.

Table 7 on page 68 lists the settings for the example topology.

Table 7: Components of the Topology for Configuring a Private VLAN

Interface	Description
ge-0/0/0.0	Primary VLAN (pvlan) trunk interface

Table 7: Components of the Topology for Configuring a Private VLAN (continued)

Interface	Description
ge-0/0/11.0	User 1, HR Community (hr-comm)
ge-0/0/12.0	User 2, HR Community (hr-comm)
ge-0/0/13.0	User 3, Finance Community (finance-comm)
ge-0/0/14.0	User 4, Finance Community (finance-comm)
ge-0/0/15.0	Mail server, Isolated (isolated)
ge-0/0/16.0	Backup server, Isolated (isolated)
ge-1/0/0.0	Primary VLAN (pvlan) trunk interface

Configuration

CLI Quick Configuration To quickly create and configure a private VLAN, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans pvlan vlan-id 1000
set interfaces ge-0/0/0 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members pvlan
set interfaces ge-1/0/0 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-1/0/0 unit 0 family ethernet-switching vlan members pvlan
set interfaces ge-0/0/11 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/12 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/13 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/14 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/15 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/16 unit 0 family ethernet-switching port-mode access
set vlans pvlan no-local-switching
set vlans pvlan interface ge-0/0/0.0
set vlans pvlan interface ge-1/0/0.0
set vlans hr-comm interface ge-0/0/11.0
set vlans hr-comm interface ge-0/0/12.0
set vlans finance-comm interface ge-0/0/13.0
set vlans finance-comm interface ge-0/0/14.0
set vlans hr-comm primary-vlan pvlan
set vlans finance-comm primary-vlan pvlan
```

Step-by-Step Procedure To configure the private VLAN:

1. Set the VLAN ID for the primary VLAN:

```
[edit vlans]
user@switch# set pvlan vlan-id 1000
```

2. Set the interfaces and port modes:

```
[edit interfaces]
```

```

user@switch# set ge-0/0/0 unit 0 family ethernet-switching port-mode trunk

user@switch# set ge-0/0/0 unit 0 family ethernet-switching vlan members
pvlan

user@switch# set ge-1/0/0 unit 0 family ethernet-switching port-mode trunk

user@switch# set ge-1/0/0 unit 0 family ethernet-switching vlan members
pvlan

user@switch# set ge-0/0/11 unit 0 family ethernet-switching port-mode
access

user@switch# set ge-0/0/12 unit 0 family ethernet-switching port-mode
access

user@switch# set ge-0/0/13 unit 0 family ethernet-switching port-mode
access

user@switch# set ge-0/0/14 unit 0 family ethernet-switching port-mode
access

user@switch# set ge-0/0/15 unit 0 family ethernet-switching port-mode
access

user@switch# set ge-0/0/16 unit 0 family ethernet-switching port-mode
access

```

3. Set the primary VLAN to have no local switching:



NOTE: The primary VLAN must be a tagged VLAN.

```

[edit vlans]
user@switch# set pvlan no-local-switching

```

4. Add the trunk interfaces to the primary VLAN:

```

[edit vlans]
user@switch# set pvlan interface ge-0/0/0.0

user@switch# set pvlan interface ge-1/0/0.0

```

5. For each secondary VLAN, configure access interfaces:



NOTE: The secondary VLANs must be untagged VLANs.

```

[edit vlans]
user@switch# set hr-comm interface ge-0/0/11.0

```

```

user@switch# set hr-comm interface ge-0/0/12.0

user@switch# set finance-comm interface ge-0/0/13.0

user@switch# set finance-comm interface ge-0/0/14.0

```

6. For each community VLAN, set the primary VLAN:

```

[edit vlans]
user@switch# set hr-comm primary-vlan pvlan

user@switch# set finance-comm primary-vlan pvlan

```

7. Add each isolated interface to the primary VLAN:

```

[edit vlans]
user@switch# set pvlan interface ge-0/0/15.0

user@switch# set pvlan interface ge-0/0/16.0

```

Results Check the results of the configuration:

```

user@switch> show configuration vlans
finance-comm {
  interface {
    ge-0/0/13.0;
    ge-0/0/14.0;
  }
  primary-vlan pvlan;
}
hr-comm {
  interface {
    ge-0/0/11.0;
    ge-0/0/12.0;
  }
  primary-vlan pvlan;
}
pvlan {
  vlan-id 1000;
  interface {
    ge-0/0/15.0;
    ge-0/0/16.0;
    ge-0/0/0.0;
    ge-1/0/0.0;
  }
  no-local-switching;
}

```

Verification

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

- Verifying the Private VLAN and Secondary VLANs Were Created on page 72

Verifying the Private VLAN and Secondary VLANs Were Created

Purpose Verify that the primary VLAN and secondary VLANs were properly created on the switch.

Action Use the show vlans command:

```

user@switch> show vlans pvlan extensive
VLAN: pvlan, Created at: Tue Sep 16 17:59:47 2008
802.1Q Tag: 1000, Internal index: 18, Admin State: Enabled, Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged 6 (Active = 0)
    ge-0/0/0.0, tagged, trunk
    ge-0/0/11.0, untagged, access
    ge-0/0/12.0, untagged, access
    ge-0/0/13.0, untagged, access
    ge-0/0/14.0, untagged, access
    ge-0/0/15.0, untagged, access
    ge-0/0/16.0, untagged, access
    ge-1/0/0.0, tagged, trunk
Secondary VLANs: Isolated 2, Community 2
  Isolated VLANs :
    __pvlan_pvlan_ge-0/0/15.0__
    __pvlan_pvlan_ge-0/0/16.0__
  Community VLANs :
    finance-comm
    hr-comm
user@switch> show vlans hr-comm extensive
VLAN: hr-comm, Created at: Tue Sep 16 17:59:47 2008
Internal index: 22, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged 2 (Active = 0)
    ge-0/0/0.0, tagged, trunk
    ge-0/0/11.0, untagged, access
    ge-0/0/12.0, untagged, access
    ge-1/0/0.0, tagged, trunk
user@switch> show vlans finance-comm extensive
VLAN: finance-comm, Created at: Tue Sep 16 17:59:47 2008
Internal index: 21, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged 2 (Active = 0)
    ge-0/0/0.0, tagged, trunk
    ge-0/0/13.0, untagged, access
    ge-0/0/14.0, untagged, access
    ge-1/0/0.0, tagged, trunk
user@switch> show vlans __pvlan_pvlan_ge-0/0/15.0__ extensive
VLAN: __pvlan_pvlan_ge-0/0/15.0__, Created at: Tue Sep 16 17:59:47 2008
Internal index: 19, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: pvlan

```

```

Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged 1 (Active = 0)
    ge-0/0/0.0, tagged, trunk
    ge-0/0/15.0, untagged, access
    ge-1/0/0.0, tagged, trunk
user@switch> show vlans __pvlan_pvlan_ge-0/0/16.0__ extensive
VLAN: __pvlan_pvlan_ge-0/0/16.0__, Created at: Tue Sep 16 17:59:47 2008
Internal index: 20, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: pvlan
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged 1 (Active = 0)
    ge-0/0/0.0, tagged, trunk
    ge-0/0/16.0, untagged, access
    ge-1/0/0.0, tagged, trunk

```

Meaning The output shows that the primary VLAN was created and identifies the interfaces and secondary VLANs associated with it.

Related Topics ■ Creating a Private VLAN (CLI Procedure) on page 105

Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches

Virtual routing instances allow each EX Series switch to have multiple routing tables on a device. With virtual routing instances, you can segment your network to isolate traffic without setting up additional devices.

This example describes how to create virtual routing instances:

- Requirements on page 73
- Overview and Topology on page 73
- Configuration on page 74
- Verification on page 75

Requirements

This example uses the following hardware and software components:

- One EX Series switch
- JUNOS Release 9.2 or later for EX Series switches

Before you create the virtual routing instances, make sure you have:

- Configured the necessary VLANs. See “Configuring VLANs for EX Series Switches (CLI Procedure)” on page 98 or “Configuring VLANs for EX Series Switches (J-Web Procedure)” on page 95.

Overview and Topology

In a large office, you may need multiple VLANs to properly manage your traffic. This configuration example shows a simple topology to illustrate how to connect a single

EX Series switch with a virtual routing instance for each of two VLANs, enabling traffic to pass between those VLANs.

In the example topology, the LAN is segmented into two VLANs, each associated with an interface and a routing instance on the EX Series switch.

Configuration

CLI Quick Configuration To quickly create and configure virtual routing instances, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/0/3 vlan-tagging
set interfaces ge-0/0/3 unit 0 vlan-id 1030 family inet address 103.1.1.1/24
set interfaces ge-0/0/3 unit 1 vlan-id 1031 family inet address 103.1.1.1/24
set routing-instances r1 instance-type virtual-router
set routing-instances r1 interface ge-0/0/1.0
set routing-instances r1 interface ge-0/0/3.0
set routing-instances r2 instance-type virtual-router
set routing-instances r2 interface ge-0/0/2.0
set routing-instances r2 interface ge-0/0/3.1
```

Step-by-Step Procedure To configure virtual routing instances:

1. Create a VLAN-tagged interface:

```
[edit]
user@switch# set interfaces ge-0/0/3 vlan-tagging
```

2. Create two subinterfaces, on the interface, one for each routing instance:

```
[edit]
user@switch# set interfaces ge-0/0/3 unit 0 vlan-id 1030 family inet
address 103.1.1.1/24

user@switch# set interfaces ge-0/0/3 unit 1 vlan-id 1031 family inet
address 103.1.1.1/24
```

3. Create two virtual routers:

```
[edit]
user@switch# set routing-instances r1 instance-type virtual-router
user@switch# set routing-instances r2 instance-type virtual-router
```

4. Set the interfaces for the virtual routers:

```
[edit]
user@switch# set routing-instances r1 interface ge-0/0/1.0

user@switch# set routing-instances r1 interface ge-0/0/3.0

user@switch# set routing-instances r2 interface ge-0/0/2.0

user@switch# set routing-instances r2 interface ge-0/0/3.1
```

Results Check the results of the configuration:

```

user@switch> show configuration
interfaces {
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/2 {
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/3 {
    vlan-tagging;
    unit 0 {
      vlan-id 1030;
      family inet {
        address 103.1.1.1/24;
      }
    }
    unit 1 {
      vlan-id 1031;
      family inet {
        address 103.1.1.1/24;
      }
    }
  }
}
routing-instances {
  r1 {
    instance-type virtual-router;
    interface ge-0/0/1.0;
    interface ge-0/0/3.0;
  }
  r2 {
    instance-type virtual-router;
    interface ge-0/0/2.0;
    interface ge-0/0/3.1;
  }
}

```

Verification

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

- Verifying That the Routing Instances Were Created on page 75

Verifying That the Routing Instances Were Created

Purpose Verify that the virtual routing instances were properly created on the switch.

Action Use the show route instance command:

```

user@switch> show route instance
Instance          Type
Primary RIB
master            forwarding
inet.0           3/0/0
r1                virtual-router
r1.inet.0        1/0/0
r2                virtual-router
r2.inet.0        1/0/0

```

Meaning Each routing instance created is displayed, along with its type, information about whether it is active or not, and its primary routing table.

Related Topics ■ Configuring Virtual Routing Instances (CLI Procedure) on page 105

Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches

As a network expands and the number of clients and VLANs increases, VLAN administration becomes complex and the task of efficiently configuring VLANs on multiple EX Series switches becomes increasingly difficult. To automate VLAN administration, you can enable Multiple VLAN Registration Protocol (MVRP) on the network.

MVRP can also be used to dynamically create VLANs, further simplifying the network overhead required to statically configure VLANs.



NOTE: Only trunk interfaces can be enabled for MVRP.

This example describes how to use MVRP to automate administration of VLAN membership changes within your network and how to use MVRP to dynamically create VLANs:

- Requirements on page 76
- Overview and Topology on page 77
- Configuring VLANs and MVRP on Access Switch A on page 79
- Configuring VLANs and MVRP on Access Switch B on page 82
- Configuring VLANs and MVRP on Distribution Switch C on page 84
- Verification on page 85

Requirements

This example uses the following hardware and software components:

- Two EX Series access switches

- One EX Series distribution switch
- JUNOS Release 10.0 or later for EX Series switches

Overview and Topology

MVRP is used to manage dynamic VLAN registration in a LAN. It can also be used to dynamically create VLANs.

This example uses MVRP to dynamically create VLANs on the switching network. You can disable dynamic VLAN creation and create VLANs statically, if desired. Enabling MVRP on the trunk interface of each switch in your switching network ensures that the active VLAN information for the switches in the network is propagated to each switch through the trunk interfaces, assuming dynamic VLAN creation is enabled for MVRP.

MVRP ensures that the VLAN membership information on the trunk interface is updated as the switch's access interfaces become active or inactive in the configured VLANs in a static or dynamic VLAN creation setup.

You do not need to explicitly bind a VLAN to the trunk interface. When MVRP is enabled, the trunk interface advertises all the VLANs that are active (bound to access interfaces) on that switch. An MVRP-enabled trunk interface does not advertise VLANs that have been configured on the switch but that are not currently bound to an access interface. Thus, MVRP provides the benefit of reducing network overhead—by limiting the scope of broadcast, unknown unicast, and multicast (BUM) traffic to interested devices only.

When VLAN access interfaces become active or inactive, MVRP ensures that the updated information is advertised on the trunk interface. Thus, in this example, distribution Switch C does not forward traffic to inactive VLANs.

This example shows a network with three VLANs: **finance**, **sales**, and **lab**.

Access Switch A has been configured to support all three VLANs and all three VLANs are active, bound to interfaces that are connected to personal computers:

- `ge-0/0/1`—Connects PC1 as a member of **finance**, VLAN ID 100
- `ge-0/0/2`—Connects PC2 as a member of **lab**, VLAN ID 200
- `ge-0/0/3`—Connects PC3 as a member of **sales**, VLAN ID 300

Access Switch B has also been configured to support three VLANs. However, currently only two VLANs are active, bound to interfaces that are connected to personal computers:

- `ge-0/0/0`—Connects PC4 as a member of **finance**, VLAN ID 100
- `ge-0/0/1`—Connects PC5 as a member of **lab**, VLAN ID 200

Distribution Switch C learns the VLANs dynamically using MVRP through the connection to the access switches. Distribution Switch C has two trunk interfaces:

- `xe-0/1/1`—Connects the switch to access Switch A.

- xe-0/1/0—Connects the switch to access Switch B.

Figure 1 shows MVRP configured on two access switches and one distribution switch.

Figure 6: MVRP Configured on Two Access Switches and One Distribution Switch for Automatic VLAN Administration




Table 1 explains the components of the example topology.

Table 8: Components of the Network Topology

Property	Settings
Switch hardware	<ul style="list-style-type: none"> ■ Access Switch A ■ Access Switch B ■ Distribution Switch C
VLAN names and tag IDs	finance, tag 100 lab, tag 200 sales, tag 300

Table 8: Components of the Network Topology (continued)

Interfaces	<p>Access Switch A interfaces:</p> <ul style="list-style-type: none"> ■ <code>ge-0/0/1</code>—Connects PC1 to access Switch A. ■ <code>ge-0/0/2</code>—Connects PC2 to access Switch A. ■ <code>ge-0/0/3</code>—Connects PC3 to access Switch A. ■ <code>xe-0/1/1</code>—Connects access Switch A to distribution Switch C (trunk). <p>Access Switch B interfaces:</p> <ul style="list-style-type: none"> ■ <code>ge-0/0/0</code>—Connects PC4 to access Switch B. ■ <code>ge-0/0/1</code>—Connects PC5 to access Switch B. ■ <code>xe-0/1/0</code>—Connects access Switch B to distribution Switch C. (trunk) <p>Distribution Switch C interfaces:</p> <ul style="list-style-type: none"> ■ <code>xe-0/1/1</code>—Connects distribution Switch C to access Switch A. (trunk) ■ <code>xe-0/1/0</code>—Connects distribution Switch C to access Switch B. (trunk)
------------	---

Configuring VLANs and MVRP on Access Switch A

To configure VLANs on the switch, bind access interfaces to the VLANs, and enable MVRP on the trunk interface of access Switch A, perform these tasks:

CLI Quick Configuration To quickly configure access Switch A for MVRP, copy the following commands and paste them into the switch terminal window of Switch A:

```
[edit]
set vlans finance vlan-id 100
set vlans lab vlan-id 200
set vlans sales vlan-id 300
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members finance
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members lab
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members sales
set interfaces xe-0/1/1 unit 0 family ethernet-switching port-mode trunk
set protocols mvrp interface xe-0/1/1.0
```



NOTE: As we recommend as a best practice, default MVRP timers are used in this example. The default values associated with each MVRP timer are: 200 ms for the join timer, 1000 ms for the leave timer, and 10000 ms for the leaveall timer. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

Step-by-Step Procedure To configure access Switch A for MVRP:

1. Configure the finance VLAN:

```
[edit]
user@Access-Switch-A# set vlans finance vlan-id 100
```

2. Configure the lab VLAN:

```
[edit]
user@Access-Switch-A# set vlans lab vlan-id 200
```

3. Configure the sales VLAN:

```
[edit]
user@Access-Switch-A# set vlans sales vlan-id 300
```

4. Configure an Ethernet interface as a member of the finance VLAN:

```
[edit]
user@Access-Switch-A# set interfaces ge-0/0/1 unit 0 family
  ethernet-switching vlan members finance
```

5. Configure an Ethernet interface as a member of the lab VLAN:

```
[edit]
user@Access-Switch-A# set interfaces ge-0/0/2 unit 0 family
  ethernet-switching vlan members lab
```

6. Configure an Ethernet interface as a member of the sales VLAN:

```
[edit]
user@Access-Switch-A# set interfaces ge-0/0/3 unit 0 family
  ethernet-switching vlan members sales
```

7. Configure a trunk interface:

```
[edit]
user@Access-Switch-A# set interfaces xe-0/1/1 unit 0 family
  ethernet-switching port-mode trunk
```

8. Enable MVRP on the trunk interface:

```
[edit]
user@Access-Switch-A# set protocols mvrp interface xe-0/1/1.0
```

Results Check the results of the configuration:

```
[edit]
user@switch# show
interfaces {
```

```
ge-0/0/1 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members finance;
      }
    }
  }
}
ge-0/0/2 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members lab;
      }
    }
  }
}
ge-0/0/3 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members sales;
      }
    }
  }
}
xe-0/1/1 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
    }
  }
}

protocols {
  mvrp {
    interface xe-0/1/1.0;
  }
}

vlans {
  finance {
    vlan-id 100;
  }
  lab {
    vlan-id 200;
  }
  sales {
    vlan-id 300;
  }
}
```

Configuring VLANs and MVRP on Access Switch B

To configure three VLANs on the switch, bind access interfaces for PC4 and PC5 to the VLANs, and enable MVRP on the trunk interface of access Switch B, perform these tasks:

CLI Quick Configuration To quickly configure Access Switch B for MVRP, copy the following commands and paste them into the switch terminal window of Switch B:

```
[edit]
set vlans finance vlan-id 100
set vlans lab vlan-id 200
set vlans sales vlan-id 300
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members finance
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members lab
set interfaces xe-0/1/0 unit 0 family ethernet-switching port-mode trunk
set protocols mvrp interface xe-0/1/0.0
```

Step-by-Step Procedure To configure access Switch B for MVRP:

1. Configure the finance VLAN:

```
[edit]
user@Access-Switch-B# set vlans finance vlan-id 100
```

2. Configure the lab VLAN:

```
[edit]
user@Access-Switch-B# set vlans lab vlan-id 200
```

3. Configure the sales VLAN:

```
[edit]
user@Access-Switch-B# set vlans sales vlan-id 300
```

4. Configure an Ethernet interface as a member of the finance VLAN:

```
[edit]
user@Access-Switch-B# set interfaces ge-0/0/0 unit 0 family
ethernet-switching vlan members finance
```

5. Configure an Ethernet interface as a member of the lab VLAN:

```
[edit]
user@Access-Switch-B# set interfaces ge-0/0/1 unit 0 family
ethernet-switching vlan members lab
```

6. Configure a trunk interface:

```
user@Access-Switch-B# set interfaces xe-0/1/0 unit 0 family
ethernet-switching port-mode trunk
```

7. Enable MVRP on the trunk interface:

```
[edit]
user@Access-Switch-B# set protocols mvrp xe-0/1/0.0
```



NOTE: As we recommend as a best practice, default MVRP timers are used in this example. The default values associated with each MVRP timer are: 200 ms for the join timer, 1000 ms for the leave timer, and 10000 ms for the leaveall timer. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

Results Check the results of the configuration:

```
[edit]
user@Access-Switch-B# show
interfaces {
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members finance;
        }
      }
    }
  }
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members lab;
        }
      }
    }
  }
  xe-0/1/0 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
      }
    }
  }
}

protocols {
  mvrp {
```

```

        interface xe-0/1/0.0;
    }
}
vlans {
    finance {
        vlan-id 100;
    }
    lab {
        vlan-id 200;
    }
    sales {
        vlan-id 300;
    }
}
}

```

Configuring VLANs and MVRP on Distribution Switch C

CLI Quick Configuration To quickly configure distribution Switch C for MVRP, copy the following commands and paste them into the switch terminal window of distribution Switch C:

```

[edit]
set interfaces xe-0/1/1 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/1/0 unit 0 family ethernet-switching port-mode trunk
set protocols mvrp interface xe-0/1/1.0
set protocols mvrp interface xe-0/1/0.0

```

Step-by-Step Procedure To configure distribution Switch C for MVRP:

1. Configure the trunk interface to access Switch A:

```

[edit]
user@Distribution-Switch-C# set interfaces xe-0/1/1 unit 0 family
ethernet-switching port-mode trunk

```

2. Configure the trunk interface to access Switch B:

```

[edit]
user@Distribution-Switch-C# set interfaces xe-0/1/0 unit 0 family
ethernet-switching port-mode trunk

```

3. Enable MVRP on the trunk interface for xe-0/1/1 :

```

[edit]
user@Distribution-Switch-C# set protocols mvrp interface xe-0/1/1.0

```

4. Enable MVRP on the trunk interface for xe-0/1/0 :

```

[edit]
user@Distribution-Switch-C# set protocols mvrp interface xe-0/1/0.0

```

Results Check the results of the configuration:

```
[edit]
user@Distribution Switch-D# show
interfaces {
  xe-0/1/0 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
      }
    }
  }
  xe-0/1/1 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
      }
    }
  }
}
protocols {
  mvrp {
    interface xe-0/1/0.0;
    interface xe-0/1/1.0;
  }
}
```

Verification

To confirm that the configuration is updating VLAN membership, perform these tasks:

- Verifying That MVRP Is Enabled on Access Switch A on page 85
- Verifying That MVRP Is Updating VLAN Membership on Access Switch A on page 86
- Verifying That MVRP Is Enabled on Access Switch B on page 86
- Verifying That MVRP Is Updating VLAN Membership on Access Switch B on page 87
- Verifying That MVRP Is Enabled on Distribution Switch C on page 87
- Verifying That MVRP Is Updating VLAN Membership on Distribution Switch C on page 87

Verifying That MVRP Is Enabled on Access Switch A

Purpose Verify that MVRP is enabled on the switch.

Action Show the MVRP configuration:

```
user@Access-Switch-A> show mvrp
MVRP configuration
MVRP status : Enabled
MVRP dynamic VLAN creation : Enabled

MVRP timers (ms):
Interface      Join   Leave   LeaveAll
```

```

-----
all          200  1000  10000
xe-0/1/1.0  200  1000  10000

Interface    Status      Registration Mode
-----
all          Disabled   Normal
xe-0/1/1.0  Enabled    Normal
    
```

Meaning The results show that MVRP is enabled on the trunk interface of Switch A and that the default timers are used.

Verifying That MVRP Is Updating VLAN Membership on Access Switch A

Purpose Verify that MVRP is updating VLAN membership by displaying the Ethernet switching interfaces and associated VLANs that are active on Switch A.

Action List Ethernet switching interfaces on the switch:

```

user@Access-Switch-A> show ethernet-switching interfaces
Interface  State  VLAN members  Blocking
ge-0/0/1.0 up     finance       unblocked
ge-0/0/2.0 up     lab           unblocked
ge-0/0/3.0 up     sales        unblocked
xe-0/1/1.0 up     finance       unblocked
           lab           unblocked
    
```

Meaning MVRP has automatically added finance and lab as VLAN members on the trunk interface because they are being advertised by access Switch B.

Verifying That MVRP Is Enabled on Access Switch B

Purpose Verify that MVRP is enabled on the switch.

Action Show the MVRP configuration:

```

user@Access-Switch-B> show mvrp

MVRP configuration
MVRP status           : Enabled
MVRP dynamic VLAN creation : Enabled

MVRP timers (ms):
Interface    Join  Leave  LeaveAll
-----
all          200  1000  10000
xe-0/1/0.0  200  1000  10000

Interface    Status      Registration Mode
-----
all          Disabled   Normal
xe-0/1/0.0  Enabled    Normal
    
```

Meaning The results show that MVRP is enabled on the trunk interface of Switch B and that the default timers are used.

Verifying That MVRP Is Updating VLAN Membership on Access Switch B

Purpose Verify that MVRP is updating VLAN membership by displaying the Ethernet switching interfaces and associated VLANs that are active on Switch B.

Action List Ethernet switching interfaces on the switch:

```
user@Access-Switch-B> show ethernet-switching interfaces
Interface  State  VLAN members  Blocking
ge-0/0/0.0 up     finance        unblocked
ge-0/0/1.0 up     lab            unblocked
xe-0/1/1.0 up     finance        unblocked
           lab            unblocked
           sales        unblocked
```

Meaning MVRP has automatically added finance, lab, and sales as VLAN members on the trunk interface because they are being advertised by access Switch A.

Verifying That MVRP Is Enabled on Distribution Switch C

Purpose Verify that MVRP is enabled on the switch.

Action Show the MVRP configuration:

```
user@Distribution-Switch-C> show mvrp

MVRP configuration
MVRP status           : Enabled
MVRP dynamic VLAN creation : Enabled

MVRP timers (ms):
Interface              Join   Leave   LeaveAll
-----
all                    200   1000   10000
xe-0/0/1.0             200   1000   10000
xe-0/1/1.0             200   1000   10000

Interface              Status      Registration Mode
-----
all                    Disabled   Normal
xe-0/0/1.0             Enabled    Normal
xe-0/1/1.0             Enabled    Normal
```

Verifying That MVRP Is Updating VLAN Membership on Distribution Switch C

Purpose Verify that MVRP is updating VLAN membership on distribution Switch C by displaying the Ethernet switching interfaces and associated VLANs on distribution Switch C.

Action List the Ethernet switching interfaces on the switch:

```
user@Distribution-Switch-C> show ethernet-switching interfaces
Interface  State  VLAN members  Blocking
xe-0/1/1.0 up     __mvrp_100__  unblocked
           __mvrp_200__  unblocked
```

```

xe-0/1/0.0 up      __mvrp_300__      unblocked
                  __mvrp_100__      unblocked
                  __mvrp_200__      unblocked

```

List the VLANs that were created dynamically using MVRP on the switch:

```

user@Distribution-Switch-C> show mvrp dynamic-vlan-memberships
VLAN Name          Interfaces
-----
__mvrp_100__       xe-0/1/1.0
                   xe-0/1/0.0
__mvrp_200__       xe-0/1/1.0
                   xe-0/1/0.0
__mvrp_300__       xe-0/1/1.0

```

Meaning Distribution Switch C has two trunk interfaces. Interface `xe-0/1/1.0` connects distribution Switch C to Access Switch A and is therefore updated to show that it is a member of all the VLANs that are active on Switch A. Any traffic for those VLANs will be passed on from distribution Switch C to Switch A, through interface `xe-0/1/1.0`. Interface `xe-0/1/0.0` connects distribution Switch C to Switch B and is updated to show that it is a member of the two VLANs that are active on Switch B. Thus, distribution Switch C sends traffic for **finance** and **lab** to both Switch A and Switch B. But distribution Switch C sends traffic for **sales** only to Switch A.

Distribution Switch C also has three dynamic VLANs created using MVRP: `mvrp_100`, `mvrp_200`, and `mvrp_300`. The dynamically created VLANs `mvrp_100` and `mvrp_200` are active on interfaces `xe-0/1/1.0` and `xe-0/1/1.0`, and dynamically created VLAN `mvrp_300` is active on interface `xe-0/1/1.0`.

- Related Topics**
- Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 109
 - Understanding Multiple VLAN Registration Protocol (MVRP) on EX Series Switches on page 17

Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches

Layer 2 protocol tunneling (L2PT) allows you to send Layer 2 protocol data units (PDUs) across a service provider network and deliver them to EX Series switches that are not part of the local broadcast domain. This feature is useful when you want to run Layer 2 protocols on a network that includes switches located at remote sites that are connected across a service provider network.

This example describes how to configure L2PT:

- Requirements on page 89
- Overview and Topology on page 89
- Configuration on page 91
- Verification on page 92

Requirements

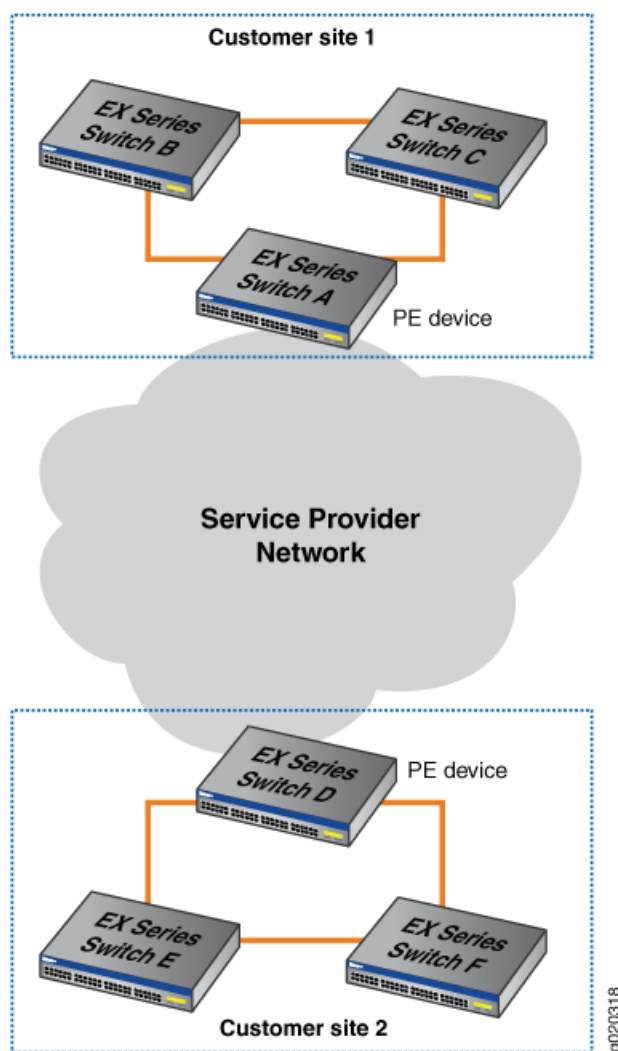
This example uses the following hardware and software components:

- Six EX Series switches, with three each at two customer sites, with one of the switches at each site designated as the provider edge (PE) device
- JUNOS Release 10.0 or later for EX Series switches

Overview and Topology

L2PT allows you to send Layer 2 PDUs across a service provider network and deliver them to EX Series switches that are not part of the local broadcast domain.

Figure 7 on page 90 shows a customer network that includes two sites that are connected across a service provider network. Site 1 contains three switches connected in a Layer 2 network, with Switch A designated as a provider edge (PE) device in the service provider network. Site 2 contains a Layer 2 network with a similar topology to that of Site 1, with Switch D designated as a PE device.

Figure 7: L2PT Topology

When you enable L2PT on a VLAN, Q-in-Q tunneling is also (and must be) enabled. Q-in-Q tunneling ensures that Switches A, B, C, D, E, and F are part of the same broadcast domain.

This example uses STP as the Layer 2 protocol being tunneled, but you could substitute any of the supported protocols for STP. You can also use the `all` keyword to enable L2PT for all supported Layer 2 protocols.

Tunneled Layer 2 PDUs do not normally arrive at high rate. If the tunneled Layer 2 PDUs do arrive at high rate, there might be a problem in the network. Typically, you would want to shut down the interface that is receiving a high rate of tunneled Layer 2 PDUs so that problem can be isolated. However, if you do not want to completely shut down the interface, you can configure the switch to drop tunneled Layer 2 PDUs that exceed a certain threshold.

The **drop-threshold** configuration statement allows you to specify the maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the switch begins dropping the Layer 2 PDUs. The drop threshold must be less than or equal to the shutdown threshold. If the drop threshold is greater than the shutdown threshold and you try to commit the configuration, the commit will fail.

The **shutdown-threshold** configuration statement allows you to specify the maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the specified interface is disabled. The shutdown threshold must be greater than or equal to the drop threshold. You can specify a drop threshold without specifying a shutdown threshold, and you can specify a shutdown threshold without specifying a drop threshold. If you do not specify these thresholds, then no thresholds are enforced. As a result, the switch tunnels all Layer 2 PDUs regardless of the speed at which they are received, although the number of packets tunneled per second might be limited by other factors.

In this example, we will configure both a drop threshold and a shutdown threshold to show how this is done.

If L2PT-encapsulated packets are received on an access interface, the switch reacts as it does when there is a loop between the service provider network and the customer network and shuts down (disables) the access interface.

Once an interface is disabled, you must explicitly reenable it using the **clear ethernet-switching layer2-protocol-tunneling error** command or else the interface will remain disabled.

Configuration

To configure L2PT, perform these tasks:

CLI Quick Configuration To quickly configure L2PT, copy the following commands and paste them into the switch terminal window of each PE device (in Figure 7 on page 90, Switch A and Switch D are the PE devices):

```
[edit]
set vlans customer-1 dot1q-tunneling
set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling stp
set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling stp drop-threshold
50
set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling stp
shutdown-threshold 100
```

Step-by-Step Procedure To configure L2PT, perform these tasks on each PE device (in Figure 7 on page , Switch A and Switch D are the PE devices):

1. Enable Q-in-Q tunneling on VLAN customer-1:

```
[edit]
user@switch# set vlans customer-1 dot1q-tunneling
```

2. Enable L2PT for STP on VLAN customer-1:

```
[edit]
user@switch# set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling stp
```

3. Configure the drop threshold as 50:

```
[edit]
user@switch# set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling
stp drop-threshold 50
```

4. Configure the shutdown threshold as 100:

```
[edit]
user@switch# set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling
stp shutdown-threshold 100
```

Results Check the results of the configuration:

```
[edit]
user@switch# show vlans customer-1 dot1q-tunneling
layer2-protocol-tunneling {
  stp {
    drop-threshold 50;
    shutdown-threshold 100;
  }
}
```

Verification

To verify that L2PT is working correctly, perform this task:

- Verify That L2PT Is Working Correctly on page 92

Verify That L2PT Is Working Correctly

Purpose Verify that Q-in-Q tunneling and L2PT are enabled.

Action Check to see that Q-in-Q tunneling and L2PT are enabled on each PE device (Switch A and Switch D are the PE devices):

```
user@switchA> show vlans extensive customer-1
VLAN: customer-1, Created at: Thu Jun 25 05:07:38 2009
802.1Q Tag: 100, Internal index: 4, Admin State: Enabled, Origin: Static
Dot1q Tunneling status: Enabled
Layer2 Protocol Tunneling status: Enabled
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 0 (Active = 0), Untagged 3 (Active = 0)
ge-0/0/7.0, untagged, access
```

```

ge-0/0/8.0, untagged, access
ge-0/0/9.0, untagged, access

```

Check to see that L2PT is tunneling STP on VLAN `customer-1` and that `drop-threshold` and `shutdown-threshold` have been configured:

```

user@switchA> show ethernet-switching layer2-protocol-tunneling vlan customer-1

```

```

Layer2 Protocol Tunneling VLAN information:
VLAN          Protocol      Drop          Shutdown
              Threshold    Threshold
customer-1    stp           50            100

```

Check the state of the interfaces on which L2PT has been enabled, including what kind of operation (encapsulation or decapsulation) they are performing:

```

user@switchA> show ethernet-switching layer2-protocol-tunneling interface

```

```

Layer2 Protocol Tunneling information:
Interface      Operation      State          Description
ge-0/0/0.0     Encapsulation  Shutdown      Shutdown threshold exceeded
ge-0/0/1.0     Decapsulation  Shutdown      Loop detected
ge-0/0/2.0     Decapsulation  Active

```

Meaning The `show vlans extensive customer-1` command shows that Q-in-Q tunneling and L2PT have been enabled. The `show ethernet-switching layer2-protocol-tunneling vlan customer-1` command shows that L2PT is tunneling the STP protocol on VLAN `customer-1`, the drop threshold is set to 50, and the shutdown threshold is set to 100. The `show ethernet-switching layer2-protocol-tunneling interface` command shows the type of operation being performed on each interface, the state of each interface and, if the state is `Shutdown`, the reason why the interface is shut down.

- Related Topics**
- Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) on page 112
 - Understanding Layer 2 Protocol Tunneling on EX Series Switches on page 19

Chapter 3

Configuring Bridging and VLANs

- Configuring VLANs for EX Series Switches (J-Web Procedure) on page 95
- Configuring VLANs for EX Series Switches (CLI Procedure) on page 98
- Configuring Routed VLAN Interfaces (CLI Procedure) on page 99
- Configuring MAC Table Aging (CLI Procedure) on page 101
- Configuring the Native VLAN Identifier (CLI Procedure) on page 101
- Creating a Series of Tagged VLANs (CLI Procedure) on page 103
- Configuring Virtual Routing Instances (CLI Procedure) on page 105
- Creating a Private VLAN (CLI Procedure) on page 105
- Configuring Q-in-Q Tunneling (CLI Procedure) on page 106
- Configuring GVRP (J-Web Procedure) on page 107
- Configuring Redundant Trunk Groups (J-Web Procedure) on page 108
- Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 109
- Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) on page 112

Configuring VLANs for EX Series Switches (J-Web Procedure)

You can use the VLAN configuration page to add a new VLAN or to edit or delete an existing VLAN.

To access the VLAN configuration page:

1. From the **Configure** menu, select **Switching > VLAN**.

The VLAN configuration page displays a list of existing VLANs. If you select a specific VLAN, the specific VLAN details are displayed in the Details section.

2. Click one:
 - **Add**—creates a VLAN.
 - **Edit**—edits an existing VLAN configuration.
 - **Delete**—deletes an existing VLAN.



NOTE: If you delete a VLAN, the VLAN configuration for all the associated interfaces is also deleted.

When you are adding or editing a VLAN, enter information as described in Table 9 on page 96.

Table 9: VLAN Configuration Details

Field	Function	Your Action
General tab		
VLAN Name	Specifies a unique name for the VLAN.	Enter a name.
VLAN Id/Range	Specifies the identifier or range for the VLAN.	Select one: <ul style="list-style-type: none"> ■ VLAN ID—Type a unique identification number from 1 through 4094. If no value is specified, it defaults to 0. ■ VLAN Range—Type a number range to create VLANs with IDs corresponding to the range. For example, the range 2–3 will create two VLANs with the IDs 2 and 3.
Description	Describes the VLAN.	Enter a brief description for the VLAN.
MAC-Table-Aging-Time	Specifies the maximum time that an entry can remain in the forwarding table before it 'ages out'.	Type the number of seconds from 60 through 1000000.
Input filter	Specifies the VLAN firewall filter that is applied to incoming packets.	To apply an input firewall filter, select the firewall filter from the list.
Output filter	Specifies the VLAN firewall filter that is applied to outgoing packets.	To apply an output firewall filter, select the firewall filter from the list.
Ports tab		
Ports	Specifies the ports (interfaces) to be associated with this VLAN for data traffic. You can also remove the port association.	Click one: <ul style="list-style-type: none"> ■ Add—Select the ports from the available list. ■ Remove—Select the port that you do not want associated with the VLAN.
IP address tab		

Table 9: VLAN Configuration Details (continued)

Field	Function	Your Action
IPv4 address	Specifies IPv4 address options for the VLAN.	<p>Select IPv4 address to enable the IPv4 address options.</p> <p>To configure IPv4:</p> <ol style="list-style-type: none"> 1. Enter the IP address. 2. Enter the subnet mask—for example, 255.255.255.0. You can also specify the address prefix. 3. To apply an input firewall filter to an interface, select the firewall filter from the list. 4. To apply an output firewall filter to an interface, select the firewall filter from the list. 5. Click the ARP/MAC Details button. Enter the static IP address and MAC address in the window that is displayed.
IPv6 address	Specifies IPv6 address options for the VLAN.	<p>Select IPv6 address to enable the IPv6 address options.</p> <p>To configure IPv6:</p> <ol style="list-style-type: none"> 1. Enter the IP address—for example: 2001:ab8:85a3::8a2e:370:7334. 2. Specify the subnet mask.
Voip tab		
Ports	<p>Specifies the ports to be associated with this VLAN for voice traffic. You can also remove the port association.</p> <p>NOTE: VoIP is not supported on EX8200 switches.</p>	<p>Click one:</p> <ul style="list-style-type: none"> ■ Add—Select the ports from the available list. ■ Remove—Select the port that you do not want associated with the VLAN.

- Related Topics**
- Configuring VLANs for EX Series Switches (CLI Procedure) on page 98
 - Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21
 - Understanding Bridging and VLANs on EX Series Switches on page 3
 - Configuring Routed VLAN Interfaces (CLI Procedure) on page 99

Configuring VLANs for EX Series Switches (CLI Procedure)

EX Series switches use VLANs to make logical groupings of network nodes with their own broadcast domains. You can use VLANs to limit the traffic flowing across the entire LAN and reduce collisions and packet retransmissions.

For each endpoint on the VLAN, configure the following VLAN parameters on the corresponding interface:

1. Set the description of the VLAN:

```
[edit interfaces ge-chassis/pic/port unit 0]
user@switch# set description vlan-description
```

2. Set the unique name of the VLAN:

```
[edit interfaces ge-chassis/pic/port unit 0]
user@switch# set family ethernet-switching vlan members vlan-name
```

3. Create the subnet for the VLAN:

```
[edit interfaces]
user@switch# set vlan unit 0 family inet address ip-address
```

4. Configure the VLAN tag ID or VLAN ID range for the VLAN:

```
[edit vlans]
user@switch# set vlan-name vlan-id vlan-id-number
```

or

```
[edit vlans]
user@switch# set vlan-name vlan-range vlan-id-low-vlan-id-high
```

5. To specify the maximum time that an entry can remain in the forwarding table before it ages out:

```
[edit vlans]
user@switch# set vlan-name mac-table-aging-time time
```

6. To specify a VLAN firewall filter to be applied to incoming or outgoing packets:

```
[edit vlans]
user@switch# set vlan-name filter (input | output) filter-name
```

- Related Topics**
- Configuring VLANs for EX Series Switches (J-Web Procedure) on page 95
 - Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21

- Configuring Routed VLAN Interfaces (CLI Procedure) on page 99
- Creating a Series of Tagged VLANs (CLI Procedure) on page 103
- Understanding Bridging and VLANs on EX Series Switches on page 3

Configuring Routed VLAN Interfaces (CLI Procedure)

Routed VLAN interfaces (RVIs) enable the EX Series switch to recognize which packets are being sent to local addresses so that they are bridged (switched) whenever possible and are routed only when needed. Whenever packets can be switched instead of routed, several layers of processing are eliminated. Switching also reduces the number of address lookups.

An interface named `vlan` functions as the logical router, on which you can configure a Layer 3 logical interface for each VLAN. For redundancy, an RVI can be combined with implementations of the Virtual Router Redundancy Protocol (VRRP) in both bridging and VPLS environments.

Jumbo frames of up to 9216 bytes are supported on an RVI. 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.

To configure the routed VLAN interface (RVI):

1. Create a Layer 2 VLAN by assigning it a name (for example, `support`) and a VLAN ID (for example, `111`).

```
[edit]
user@switch# set vlans support vlan-id 111
```

2. Assign an interface (for example, `ge-0/0/18`) to the VLAN (`support`) by naming the VLAN as a trunk member on the logical interface, thereby making the interface part of the VLAN's broadcast domain.

```
[edit]
user@switch# set interfaces ge-0/0/18 unit 0 family ethernet-switching
vlan members support
```

3. Create a logical Layer 3 RVI (`vlan.111`) on a subnet for the VLAN's broadcast domain.

```
[edit]
```

```
user@switch# set interfaces vlan unit 111 family inet address
111.111.111.1/24
```

4. Link the Layer 2 VLAN to the logical Layer 3 interface.

```
[edit]
user@switch# set vlans support 13-interface vlan.111
```



NOTE: Layer 3 interfaces on trunk ports allow the interface to transfer traffic between multiple VLANs. Within a VLAN, traffic is bridged, while across VLANs, traffic is routed.

You can display the configuration settings:

```
user@switch> show interfaces vlan terse
Interface           Admin Link Proto   Local           Remote
vlan
vlan.111             up    up    inet    111.111.111.1/24
```

```
user@switch> show vlans
Name      Tag  Interfaces
default
employee-vlan  20  ge-1/0/0.0, ge-1/0/1.0, ge-1/0/2.0
marketing  40  ge-1/0/10.0, ge-1/0/20.0, ge-1/0/30.0
support   111  ge-0/0/18.0
mgmt
bme0.32769, bme0.32771*
```

```
user@switch> show ethernet-switching table
Ethernet-switching table: 1 entries, 0 learned
VLAN      MAC address      Type      Age Interfaces
support   00:19:e2:50:95:a0 Static      - Router
```

- Related Topics**
- Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 28
 - Example: Connecting an Access Switch to a Distribution Switch on page 36
 - Example: Configuring IP Directed Broadcast on an EX Series Switch
 - Understanding Bridging and VLANs on EX Series Switches on page 3

Configuring MAC Table Aging (CLI Procedure)

The aging process ensures that the EX Series switch tracks only active nodes on the network and that it is able to flush out network nodes that are no longer available.

To manage MAC entries more efficiently, you can configure an entry's aging time, which is the maximum time that an entry can remain in the Ethernet Switching table before it “ages out”.

To configure how long entries remain in the Ethernet Switching table before expiring, using the CLI (here, the VLAN is `employee-vlan`):

```
[edit vlans employee-vlan]
user@switch# set mac-table-aging-time 200
```

- Related Topics**
- Understanding Bridging and VLANs on EX Series Switches on page 3
 - Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21
 - Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 28
 - Example: Connecting an Access Switch to a Distribution Switch on page 36

Configuring the Native VLAN Identifier (CLI Procedure)

EX Series switches support receiving and forwarding routed or bridged Ethernet frames with 802.1Q VLAN tags. The logical interface on which untagged packets are to be received must be configured with the same native VLAN ID as that configured on the physical interface.

To configure the native VLAN ID using the CLI:

1. Configure the port mode so that the interface is in multiple VLANs and can multiplex traffic between different VLANs. Trunk interfaces typically connect to other switches and to routers on the LAN. Configure the port mode as `trunk`:

```
[edit interfaces ge-0/0/3 unit 0 family ethernet-switching]
user@switch# set port-mode trunk
```

2. Configure the native VLAN ID:

```
[edit interfaces ge-0/0/3 unit 0 family ethernet-switching]
user@switch# set native-vlan-id 3
```

- Related Topics**
- Understanding Bridging and VLANs on EX Series Switches on page 3
 - Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 28

- Example: Connecting an Access Switch to a Distribution Switch on page 36
- Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21

Creating a Series of Tagged VLANs (CLI Procedure)

To identify which VLAN traffic belongs to, all frames on an Ethernet VLAN are identified by a tag, as defined in the IEEE 802.1Q standard. These frames are *tagged* and are encapsulated with 802.1Q tags. For a simple network that has only a single VLAN, all traffic has the same 802.1Q tag.

Instead of configuring VLANs and 802.1Q tags one at a time for a trunk interface, you can configure a VLAN range to create a series of tagged VLANs.

When an Ethernet LAN is divided into VLANs, each VLAN is identified by a unique 802.1Q tag. The tag is applied to all frames so that the network nodes receiving the frames know which VLAN the frames belong to. Trunk ports, which multiplex traffic among a number of VLANs, use the tag to determine the origin of frames and where to forward them.

For example, you could configure the VLAN `employee` and specify a tag range of `10-12`. This creates the following VLANs and tags:

- VLAN `employee-10`, tag `10`
- VLAN `employee-11`, tag `11`
- VLAN `employee-12`, tag `12`

Creating tagged VLANs in a series has the following limitations:

- Layer 3 interfaces do not support this feature.
- Because an access interface can only support one VLAN member, access interfaces also do not support this feature.
- Voice over IP (VoIP) configurations do not support a range of tagged VLANs.

To configure a series of tagged VLANs using the CLI (here, the VLAN is `employee`):

1. Configure the series (here, a VLAN series from 120 through 130):

```
[edit]
user@switch# set vlans employee vlan-range 120-130
```

2. Associate a series of tagged VLANs when you configure an interface in one of two ways:

- Include the name of the series:

```
[edit interfaces]
user@switch# set interfaces ge-0/0/22.0 family ethernet-switching vlan
members employee
```

- Include the VLAN range:

```
[edit interfaces]
user@switch# set interfaces ge-0/0/22.0 family ethernet-switching vlan
members 120-130
```

Associating a series of tagged VLANs to an interface by name or by VLAN range have the same result: VLANs `__employee_120__` through `__employee_130__` are created.



NOTE: When a series of VLANs are created using the `vlan-range` command, the VLAN names are prefixed and suffixed with a double underscore.

-
- Related Topics**
- Verifying That a Series of Tagged VLANs Has Been Created on page 115
 - Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21
 - Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 28
 - Example: Connecting an Access Switch to a Distribution Switch on page 36
 - Understanding Bridging and VLANs on EX Series Switches on page 3

Configuring Virtual Routing Instances (CLI Procedure)

Use virtual routing and forwarding (VRF) to divide an EX Series switch into multiple virtual routing instances. VRF allows you to isolate VLAN traffic without using multiple devices to segment your network.

Before you begin to configure these multiple virtual routing instances, make sure to set up your VLANs. See “Configuring VLANs for EX Series Switches (CLI Procedure)” on page 98 or “Configuring VLANs for EX Series Switches (J-Web Procedure)” on page 95.

To configure virtual routing instances:

1. Create a routing instance:

```
[edit routing-instances]
user@switch# set routing-instance-name instance-type virtual-router
```

2. Bind each routing instance to the corresponding interfaces:

```
[edit routing-instances]
user@switch# set routing-instance-name interface
ge-chassis/slot/port.logical-unit-number
```

3. Create each of the logical interfaces bound to each routing instance:

```
[edit interfaces]
user@switch# set ge-chassis/slot/port unit logical-unit-number family inet
address ip-address
```

4. Enable VLAN tagging on each interface:

```
[edit interfaces]
user@switch# set ge-chassis/slot/port vlan-tagging
```

- Related Topics**
- Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 73
 - Verifying That Virtual Routing Instances Are Working on page 117
 - Understanding Virtual Routing Instances on EX Series Switches on page 10

Creating a Private VLAN (CLI Procedure)

The private VLAN (PVLAN) feature on EX Series switches allows an administrator to split a broadcast domain into multiple isolated broadcast subdomains, essentially putting a VLAN inside a VLAN.

Before you begin, make sure you set up your VLANs. See “Configuring VLANs for EX Series Switches (CLI Procedure)” on page 98 or “Configuring VLANs for EX Series Switches (J-Web Procedure)” on page 95.



NOTE: Configuring a voice over IP (VoIP) VLAN on PVLAN interfaces is not supported.

To configure private VLANs:

1. Set the primary VLAN to have no local switching:



NOTE: The primary VLAN must be a tagged VLAN.

[edit vlans]

```
user@switch# set primary-vlan-name no-local-switching
```

2. For each community VLAN, configure access interfaces:



NOTE: The secondary VLANs must be untagged VLANs.

[edit vlans]

```
user@switch# set community-vlan-name interface ge-chassis/slot/port
```

3. For each community VLAN, set the primary VLAN:

[edit vlans]

```
user@switch# set community-vlan-name primary-vlan primary-vlan-name
```

4. For each isolated VLAN, add the interface to the primary VLAN:

[edit vlans]

```
user@switch# set primary-vlan-name interface ge-chassis/slot/port
```

- Related Topics**
- Example: Configuring a Private VLAN on an EX Series Switch on page 68
 - Verifying That a Private VLAN Is Working on page 118
 - Understanding Private VLANs on EX Series Switches on page 9

Configuring Q-in-Q Tunneling (CLI Procedure)

Q-in-Q tunneling allows service providers on Ethernet access networks to segregate or bundle customer traffic into different VLANs by adding another layer of 802.1Q tags. You can configure Q-in-Q tunneling on EX Series switches.

Before you begin configuring Q-in-Q tunneling, make sure you set up your VLANs. See “Configuring VLANs for EX Series Switches (CLI Procedure)” on page 98 or “Configuring VLANs for EX Series Switches (J-Web Procedure)” on page 95.

To configure Q-in-Q tunneling:

1. Enable Q-in-Q tunneling on the S-VLAN:

```
[edit vlans]
user@switch# set s-vlan-name dot1q-tunneling
```

2. Set the allowed C-VLANs on the S-VLAN (optional). Here, the C-VLANs are identified by VLAN range:

```
[edit vlans]
user@switch# set s-vlan-name dot1q-tunneling customer-vlans range
```

3. Change the global Ethertype value (optional):

```
[edit]
user@switch# set ethernet-switching-options dot1q-tunneling ether-type ether-type-value
```

4. Disable MAC address learning on the S-VLAN (optional):

```
[edit vlans]
user@switch# set s-vlan-name no-mac-learning
```

- Related Topics**
- Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 65
 - Verifying That Q-in-Q Tunneling Is Working on page 118
 - Understanding Q-in-Q Tunneling on EX Series Switches on page 13

Configuring GVRP (J-Web Procedure)

As a network expands and the number of clients and VLANs increases, VLAN administration becomes complex, and the task of efficiently configuring VLANs on multiple EX Series switches becomes increasingly difficult. To automate VLAN administration, you can enable GARP VLAN Registration Protocol (GVRP) on the network.

GVRP learns VLANs on a particular 802.1Q trunk port and adds the corresponding trunk interface to the VLAN if the advertised VLAN is preconfigured or existing already on the switch. For example, a VLAN named “sales” is advertised to trunk interface 1 on the GVRP-enabled switch. The switch adds trunk interface 1 to the sales VLAN if the sales VLAN already exists on the switch.

As individual interfaces become active and send requests to join a VLAN, the VLAN configuration is updated and propagated among the switches. Limiting the VLAN configuration to active participants reduces the network overhead. GVRP also provides the benefit of pruning VLANs to limit the scope of broadcast, unknown unicast, and multicast (BUM) traffic to interested network devices only.

To configure GVRP using the J-Web interface:

1. Select **Configure > Switching > GVRP**. Interfaces on which GVRP has been enabled are listed.
2. To enable GVRP on an interface, click **Add**. Click the arrow key to move the interface from the **Interface Out of GVRP** list to the **Interface under GVRP** list, and click **OK**.
3. To modify GVRP timers, click **Global Settings**. When you are modifying GVRP Timer settings for the interface, enter information as described in Table 10 on page 108.
4. Click **OK** to apply changes to the configuration or click **Cancel** to cancel without saving changes.

To disable an interface, select the interface and click **Disable Port**.

Table 10: GVRP Timer Settings

Field	Function	Your Action
Join Timer	Specifies the maximum number of milliseconds the interfaces wait before sending VLAN advertisements.	Type a number.
Leave Timer	Specifies the number of milliseconds an interface waits after receiving a leave message before the interface leaves the VLAN specified in the message.	Type a number.
Leave All Timer	Specifies the interval in milliseconds at which Leave All messages are sent on interfaces. Leave All messages help to maintain current GVRP VLAN membership information in the network.	Type a number.
Disable GVRP	Disables GVRP on all interfaces.	To disable GVRP, select the check box. To enable GVRP, clear the check box.

- Related Topics**
- Example: Configure Automatic VLAN Administration Using GVRP on page 46
 - Monitoring GVRP on page 121

Configuring Redundant Trunk Groups (J-Web Procedure)

A redundant trunk link provides a simple solution for network recovery when a trunk interface goes down. Traffic is routed to another trunk interface, keeping network convergence time to a minimum. You can configure redundant trunk groups (RTGs) with a primary link and a secondary link on trunk interfaces, or configure dynamic selection of the active interface. If the primary link fails, the secondary link automatically takes over without waiting for normal STP convergence. An RTG can be created only if the following conditions are satisfied:

- A minimum of two trunk interfaces that are not part of any RTG are available.

- All the selected trunk interfaces to be added to the RTG have the same VLAN configuration.
- The selected trunk interfaces are not part of a spanning-tree configuration.

To configure an RTG using the J-Web interface:

1. From the **Configure** menu, select **Switching > RTG**.

The RTG Configuration page displays a list of existing RTGs. If you select a specific RTG, the details of the selected RTG are displayed in the Details of group section.

2. Click one:
 - **Add**—Creates an RTG.
 - **Edit**—Modifies an RTG.
 - **Delete**—Deletes an RTG.

When you are adding or editing an RTG, enter information as described in Table 11 on page 109.

3. Click **OK** to apply changes to the configuration or click **Cancel** to cancel without saving changes.

Table 11: RTG Configuration Fields

Field	Function	Your Action
Group Name	Specifies a unique name for the RTG.	Enter a name.
Member Interface 1	Specifies a logical interface containing multiple trunk interfaces.	Select a trunk interface from the list.
Member Interface 2	Specifies a trunk interface containing multiple VLANs.	Select a trunk interface from the list.
Select Primary Interface	Enables you to specify one of the interfaces in the RTG as the primary link. The interface without this option is the secondary link in the RTG.	<ol style="list-style-type: none"> 1. Select the option button. 2. Select the primary interface.
Dynamically select my active interface	Specifies that the system dynamically selects the active interface.	Select the option button.

- Related Topics**
- Example: Configuring Redundant Trunk Links for Faster Recovery on page 61
 - Understanding Redundant Trunk Links on EX Series Switches on page 11

Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)

Multiple VLAN Registration Protocol (MVRP) is used to manage dynamic VLAN registration in a LAN. You can use MVRP on EX Series switches.

MVRP is disabled by default on EX Series switches.

To enable MVRP or set MVRP options, follow these instructions:

- Enabling MVRP on page 110
- Disabling MVRP on page 110
- Disabling Dynamic VLANs on page 110
- Configuring Timer Values on page 111
- Configuring MVRP Registration Mode on page 112

Enabling MVRP

MVRP can only be enabled on trunk interfaces.

To enable MVRP on all trunk interfaces on the switch:

```
[edit protocols mvrp]
user@switch# set interface all
```

To enable MVRP on a specific trunk interface:

```
[edit protocols mvrp]
user@switch# set interface xe-0/0/1.0
```

Disabling MVRP

MVRP is disabled by default. You only need to perform this procedure if you have previously enabled MVRP.

To disable MVRP on all trunk interfaces on the switch:

```
[edit protocols mvrp]
user@switch# set disable
```

To disable MVRP on a specific trunk interface:

```
[edit protocols mvrp]
user@switch# set disable interface xe-0/0/1.0
```

Disabling Dynamic VLANs

Dynamic VLANs can be created on interfaces participating in MVRP by default. Dynamic VLANs are VLANs created on one switch that are propagated to other switches dynamically; in this case, using MVRP.

Dynamic VLAN creation through MVRP cannot be disabled per switch interface. To disable dynamic VLAN creation for interfaces participating in MVRP, you must disable it for all interfaces on the switch.

To disable dynamic VLAN creation:

```
[edit protocols mvrp]
```

```
user@switch# set no-dynamic-vlan
```

Configuring Timer Values

The timers in MVRP define the amount of time an interface waits to join or leave MVRP or to send or process the MVRP information for the switch after receiving an MVRP PDU. The join timer controls the amount of time the switch waits to accept a registration request, the leave timer controls the period of time that the switch waits in the Leave state before changing to the unregistered state, and the leaveall timer controls the frequency with which the LeaveAll messages are communicated.

The default MVRP timer values are 200 ms for the join timer, 1000 ms for the leave timer, and 10000 ms for the leaveall timer.



BEST PRACTICE: Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

To set the join timer for all interfaces on the switch:

```
[edit protocols mvrp]
user@switch# set interface all join-timer 300
```

To set the join timer for a specific interface:

```
[edit protocols mvrp]
user@switch# set interface xe-0/0/1.0 300
```

To set the leave timer for all interfaces on the switch:

```
[edit protocols mvrp]
user@switch# set interface all leave-timer 1200
```

To set the leave timer for a specific interface:

```
[edit protocols mvrp]
user@switch# set interface xe-0/0/1.0 leave-timer 1200
```

To set the leaveall timer for all interfaces on the switch:

```
[edit protocols mvrp]
user@switch# set interface all leaveall-timer 12000
```

To set the leaveall timer for a specific interface:

```
[edit protocols mvrp]
user@switch# set interface xe-0/0/1.0 leaveall-timer 12000
```

Configuring MVRP Registration Mode

The default MVRP registration mode for any interface participating in MVRP is normal. An interface in normal registration mode participates in MVRP when MVRP is enabled on the switch.

An interface in forbidden registration mode does not participate in MVRP even if MVRP is enabled on the switch.

To set all interfaces to forbidden registration mode:

```
[edit protocols mvrp]
user@switch# set interface all registration forbidden
```

To set one interface to forbidden registration mode:

```
[edit protocols mvrp]
user@switch# set interface xe-0/0/1.0 registration forbidden
```

To set all interfaces to normal registration mode:

```
[edit protocols mvrp]
user@switch# set interface all registration normal
```

To set one interface to normal registration mode:

```
[edit protocols mvrp]
user@switch# set interface xe-0/0/1.0 registration normal
```

- Related Topics**
- Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 76
 - Verifying That MVRP Is Working Correctly on page 122

Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure)

Layer 2 protocol tunneling (L2PT) allows you to send Layer 2 protocol data units (PDUs) across a service provider network and deliver them to EX Series switches at a remote location. This feature is useful when you have a network that includes remote sites that are connected across a service provider network and you want to run Layer 2 protocols on switches connected across the service provider network.

Tunneled Layer 2 PDUs do not normally arrive at high rate. If the tunneled Layer 2 PDUs do arrive at high rate, there might be a problem in the network. Typically, you would want to shut down the interface that is receiving a high rate of tunneled Layer 2 PDUs so that the problem can be isolated. You do so using the `shutdown-threshold` statement. However, if you do not want to completely shut down the interface, you can configure the switch to drop tunneled Layer 2 PDUs that exceed a certain threshold using the `drop-threshold` statement.

There are no default settings for `drop-threshold` and `shutdown-threshold`. If you do not specify these thresholds, then no thresholds are enforced. As a result, the switch

tunnels all Layer 2 PDUs regardless of the speed at which they are received, although the number of packets tunneled per second might be limited by other factors.

You can specify a drop threshold value without specifying a shutdown threshold value, and you can specify a shutdown threshold value without specifying a drop threshold value. If you specify both threshold values, then the drop threshold value must be less than or equal to the shutdown threshold value. If the drop threshold value is greater than the shutdown threshold value and you try to commit the configuration, the commit will fail.



NOTE: If the switch receives untagged Layer 2 control PDUs to be tunneled, then you must configure the switch to map untagged (native) packets to an L2PT-enabled VLAN. Otherwise, the untagged Layer 2 control PDU packets are discarded. For more information, see “Understanding Q-in-Q Tunneling on EX Series Switches” on page 13 and “Configuring Q-in-Q Tunneling (CLI Procedure)” on page 106.

To configure L2PT on an EX Series switch:

1. Because L2PT operates under the Q-in-Q tunneling configuration, you must enable Q-in-Q tunneling before you can configure L2PT. Enable Q-in-Q tunneling on VLAN customer-1:

```
[edit]
user@switch# set vlans customer-1 dot1q-tunneling
```

2. Enable L2PT for the Layer 2 protocol you want to tunnel, on the VLAN:

- To enable L2PT for a specific protocol (here, STP):

```
[edit]
user@switch# set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling stp
```

- To enable L2PT for all supported protocols:

```
[edit]
user@switch# set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling
all
```

3. (Optional) Configure the drop threshold:
-



NOTE: If you also configure the shutdown threshold, ensure that you configure the drop threshold value to be less than or equal to the shutdown threshold value. If the drop threshold value is greater than the shutdown threshold value and you try to commit the configuration changes, the commit will fail.

```
[edit]
```

```
user@switch# set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling
stp drop-threshold 50
```

4. (Optional) Configure the shutdown threshold:



NOTE: If you also configure the drop threshold, ensure that you configure the shutdown threshold value to be greater than or equal to the drop threshold value. If the shutdown threshold value is less than the drop threshold value and you try to commit the configuration changes, the commit will fail.

```
[edit]
user@switch# set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling
stp shutdown-threshold 100
```



NOTE: Once an interface is disabled, you must explicitly reenable it using the `clear ethernet-switching layer2-protocol-tunneling error` command. Otherwise, the interface remains disabled.

- Related Topics**
- Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 88
 - Understanding Layer 2 Protocol Tunneling on EX Series Switches on page 19

Chapter 4

Verifying Bridging and VLAN Configuration

- Verifying That a Series of Tagged VLANs Has Been Created on page 115
- Verifying That Virtual Routing Instances Are Working on page 117
- Verifying That Q-in-Q Tunneling Is Working on page 118
- Verifying That a Private VLAN Is Working on page 118
- Monitoring Ethernet Switching on page 120
- Monitoring GVRP on page 121
- Verifying That MVRP Is Working Correctly on page 122

Verifying That a Series of Tagged VLANs Has Been Created

Purpose Verify that a series of tagged VLANs is created on the switch.

Action Display the VLANs in the ascending order of their VLAN ID:

```
user@switch> show vlans sort-by tag
```

Name	Tag	Interfaces
__employee_120__	120	ge-0/0/22.0*
__employee_121__	121	ge-0/0/22.0*
__employee_122__	122	ge-0/0/22.0*
__employee_123__	123	ge-0/0/22.0*
__employee_124__	124	ge-0/0/22.0*
__employee_125__	125	ge-0/0/22.0*
__employee_126__	126	ge-0/0/22.0*
__employee_127__	127	ge-0/0/22.0*
__employee_128__	128	ge-0/0/22.0*
__employee_129__	129	ge-0/0/22.0*
__employee_130__	130	ge-0/0/22.0*

Display the VLANs by the alphabetical order of the VLAN name:

```
user@switch> show vlans sort-by name
```

Name	Tag	Interfaces
__employee_120__	120	ge-0/0/22.0*
__employee_121__	121	ge-0/0/22.0*
__employee_122__	122	ge-0/0/22.0*
__employee_123__	123	ge-0/0/22.0*
__employee_124__	124	ge-0/0/22.0*
__employee_125__	125	ge-0/0/22.0*
__employee_126__	126	ge-0/0/22.0*
__employee_127__	127	ge-0/0/22.0*
__employee_128__	128	ge-0/0/22.0*
__employee_129__	129	ge-0/0/22.0*
__employee_130__	130	ge-0/0/22.0*

Display the VLANs by specifying the VLAN-range name (here, the VLAN-range name is `employee`):

```
user@switch> show vlans employee
```

Name	Tag	Interfaces
__employee_120__	120	ge-0/0/22.0*
__employee_121__	121	ge-0/0/22.0*
__employee_122__	122	ge-0/0/22.0*
__employee_123__	123	ge-0/0/22.0*
__employee_124__	124	ge-0/0/22.0*
__employee_125__	125	ge-0/0/22.0*
__employee_126__	126	ge-0/0/22.0*
__employee_127__	127	ge-0/0/22.0*
__employee_128__	128	ge-0/0/22.0*
__employee_129__	129	ge-0/0/22.0*
__employee_130__	130	ge-0/0/22.0*

Meaning The sample output shows the VLANs configured on the switch. The series of tagged VLANs is displayed: `__employee_120__` through `__employee_130__`. Each of the tagged VLANs is configured on the trunk interface `ge-0/0/22.0`. The asterisk (*) beside the interface name indicates that the interface is UP.

When a series of VLANs is created using the `vlan-range` statement, the VLAN names are prefixed and suffixed with a double underscore.

Related Topics ■ Creating a Series of Tagged VLANs (CLI Procedure) on page 103

Verifying That Virtual Routing Instances Are Working

Purpose After creating a virtual routing instance, make sure it is set up properly.

Action 1. Use the `show route instance` command to list all of the routing instances and their properties:

```
user@switch> show route instance

Instance          Type
      Primary RIB
Active/holddown/hidden
master            forwarding
      inet.0                                3/0/0

__juniper_private1__ forwarding
      __juniper_private1__.inet.0          1/0/3

__juniper_private2__ forwarding

instance1         forwarding

r1                virtual-router
      r1.inet.0                            1/0/0

r2                virtual-router
      r2.inet.0                            1/0/0
```

2. Use the `show route forwarding-table` command to view the forwarding table information for each routing instance:

```
user@switch> show route forwarding-table

Routing table: r1.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0      0                  rjct  539  2
0.0.0.0/32       perm  0      0                  dscd  537  1
103.1.1.0/24     ifdn  0      0                  rslv  579  1
ge-0/0/3.0
103.1.1.0/32     iddn  0 103.1.1.0         recv  577  1
ge-0/0/3.0
103.1.1.1/32     user  0      0                  rjct  539  2
103.1.1.1/32     intf  0 103.1.1.1         locl  578  2
103.1.1.1/32     iddn  0 103.1.1.1         locl  578  2
103.1.1.255/32  iddn  0 103.1.1.255       bcst  576  1
ge-0/0/3.0
224.0.0.0/4      perm  0      0                  mdsc  538  1
224.0.0.1/32    perm  0 224.0.0.1         mcst  534  1
255.255.255.255/32 perm  0      0                  bcst  535  1
```

Meaning The output confirms that the virtual routing instances are created and the links are up and displays the routing table information.

- Related Topics**
- Configuring Virtual Routing Instances (CLI Procedure) on page 105
 - Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 73

Verifying That Q-in-Q Tunneling Is Working

Purpose After creating a Q-in-Q VLAN, verify that it is set up properly.

- Action**
1. Use the `show configuration vlans` command to determine if you successfully created the primary and secondary VLAN configurations:

```
user@switch> show configuration vlans

svlan {
  vlan-id 300;
  dot1q-tunneling {
    customer-vlans [ 101-200 ];
  }
}
```

2. Use the `show vlans` command to view VLAN information and link status:

```
user@switch> show vlans s-vlan-name extensive

VLAN: svlan, Created at: Thu Oct 23 16:53:20 2008
802.1Q Tag: 300, Internal index: 2, Admin State: Enabled, Origin: Static
Dot1q Tunneling Status: Enabled
Customer VLAN ranges:
    101-200
Protocol: Port Mode
Number of interfaces: Tagged 1 (Active = 0), Untagged 1 (Active = 0)
    ge-0/0/1, tagged, trunk
    ge-0/0/2, untagged, access
```

Meaning The output confirms that Q-in-Q tunneling is enabled and that the VLAN is tagged, and lists the customer VLANs that are associated with the tagged VLAN.

- Related Topics**
- Configuring Q-in-Q Tunneling (CLI Procedure) on page 106
 - Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 65

Verifying That a Private VLAN Is Working

Purpose After creating and configuring private VLANs, verify they are set up properly.

- Action**
1. Use the `show configuration vlans` command to determine if you successfully created the primary and secondary VLAN configurations:

```
user@switch> show configuration vlans
```

```
community1 {
  interface {
    interface a;
    interface b;
  }
  primary-vlan pvlan;
}
community2 {
  interface {
    interface d;
    interface e;
  }
  primary-vlan pvlan;
}
pvlan {
  vlan-id 1000;
  interface {
    isolated1;
    isolated2;
    trunk1;
    trunk2;
  }
  no-local-switching;
}
```

2. Use the `show vlans` command to view VLAN information and link status:

```
user@switch> show vlans pvlan extensive
```

```
VLAN: pvlan, Created at: time
802.1Q Tag: vlan-id, Internal index: index-number, Admin State: Enabled,
Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged 6 (Active = 0)
  trunk1, tagged, trunk
  interface a, untagged, access
  interface b, untagged, access
  interface c, untagged, access
  interface d, untagged, access
  interface e, untagged, access
  interface f, untagged, access
  trunk2, tagged, trunk
Secondary VLANs: Isolated 2, Community 2
Isolated VLANs :
  __pvlan_pvlan_isolated1__
  __pvlan_pvlan_isolated2__
Community VLANs :
  community1
  community2
```

3. Use the `show ethernet-switching table vlan` command to view logs for MAC learning on the VLANs:

```
user@switch> vlan pvlan extensive
```

```

pvlan, *
  Interface(s): trunk1
  Interface(s): interface a
  Interface(s): interface b
  Interface(s): interface c
  Interface(s): interface d
  Interface(s): interface e
  Interface(s): interface f
  Interface(s): trunk2
  Type: Flood
  Nexthop index: 1344
    
```

Meaning The output shows that the primary and secondary VLANs were created and associated and displays MAC learning information.

- Related Topics**
- Creating a Private VLAN (CLI Procedure) on page 105
 - Example: Configuring a Private VLAN on an EX Series Switch on page 68

Monitoring Ethernet Switching

Purpose Use the monitoring feature to view details that the EX Series switch maintains in its Ethernet switching table. These are details about the nodes on the LAN such as VLAN name, VLAN ID, member interfaces, MAC addresses, and so on.

Action To display Ethernet switching details in the J-Web interface, select **Monitor > Switching > Ethernet Switching**.

To view Ethernet switching details in the CLI, enter the following commands:

- **show ethernet-switching table**
- **show vlans**
- **show ethernet-switching interfaces**

Meaning Table 12 on page 120 summarizes the Ethernet switching output fields.

Table 12: Ethernet Switching Output Fields

Field	Value
Ethernet Switching Table Information	
MAC Table Count	The number of entries added to the Ethernet switching table.
MAC Table Learned	The number of dynamically learned MAC addresses in the Ethernet switching table.
Ethernet Switching Table Information	
VLAN	The VLAN name.

Table 12: Ethernet Switching Output Fields (continued)

Field	Value
MAC Address	The MAC address associated with the VLAN. If a VLAN range has been configured for a VLAN, the output displays the MAC addresses for the entire series of VLANs that were created with that name.
Type	The type of MAC address. Values are: <ul style="list-style-type: none"> ■ static—The MAC address is manually created. ■ learn—The MAC address is learned dynamically from a packet's source MAC address. ■ flood—The MAC address is unknown and flooded to all members.
Age	The time remaining before the entry ages out and is removed from the Ethernet switching table.
MAC Learning Log	
VLAN-Name	The VLAN name.
MAC Address	The learned MAC address associated with the VLAN ID.
Time	Timestamp for the time at which when the MAC address was added or deleted from the MAC learning log.
State	Operating state of the interface. Values are Up and Down.

- Related Topics**
- Configuring MAC Table Aging (CLI Procedure) on page 101
 - Understanding Bridging and VLANs on EX Series Switches on page 3

Monitoring GVRP

- Purpose** Use the monitoring feature to view information about the GVRP configuration on the EX Series switch.
- Action** To monitor GVRP in the J-Web interface, select **Monitor > Switching > GVRP**.
- To monitor GVRP in the CLI, enter the following command:
- **show gvrp**
- Meaning** Table 13 on page 121 summarizes the GVRP output fields.

Table 13: Summary of GVRP Output Fields

Field	Value
Global GVRP Configuration	
GVRP Status	Displays whether GVRP is enabled or disabled.

Table 13: Summary of GVRP Output Fields (continued)

Field	Value
GVRP Timers	<ul style="list-style-type: none"> ■ Join—The number of milliseconds the interfaces must wait before sending VLAN advertisements. ■ Leave—The number of milliseconds an interface must wait after receiving a Leave message to remove the interface from the VLAN specified in the message. ■ Leave All—The interval in milliseconds at which Leave All messages are sent on interfaces. Leave All messages maintain current GVRP VLAN membership information in the network.
GVRP Interface Details	
Interface Name	The interface on which GVRP is configured.
Protocol Status	Displays whether GVRP is enabled or disabled on the interface.

- Related Topics**
- Configuring GVRP (J-Web Procedure) on page 107
 - Example: Configure Automatic VLAN Administration Using GVRP on page 46

Verifying That MVRP Is Working Correctly

Purpose After configuring your EX Series switch to participate in MVRP, verify that the configuration is properly set and that MVRP messages are being sent and received on your switch.

- Action**
1. Confirm that MVRP is enabled on your switch.

```
user@switch> show mvrp
```

```
Global MVRP configuration
```

```
MVRP status : Enabled
MVRP dynamic vlan creation: Enabled
MVRP Timers (ms):
Interface      Join   Leave  LeaveAll
-----
all            200   600    10000
xe-0/1/1.0     200   600    10000
```

```
Interface based configuration:
```

```
Interface      Status      Registration  Dynamic VLAN Creation
-----
all            Disabled   Fixed         Enabled
xe-0/1/1.0     Enabled    Normal        Enabled
```

2. Confirm that MVRP messages are being sent and received on your switch.

```
user@switch> show mvrp statistics interface xe-0/1/1.0
```

```
MVRP statistics
MRPDU received : 3342
```

```

Invalid PDU received      : 0
New received             : 2
Join Empty received      : 1116
Join In received         : 2219
Empty received           : 2
In received              : 2
Leave received            : 1
LeaveAll received         : 1117
MRPDU transmitted        : 3280
MRPDU transmit failures  : 0
New transmitted           : 0
Join Empty transmitted    : 1114
Join In transmitted      : 2163
Empty transmitted        : 1
In transmitted           : 1
Leave transmitted         : 1
LeaveAll transmitted      : 1111

```

Meaning The output of `show mvrp` shows that interface `xe-0/1/1.0` is enabled for MVRP participation as shown in the status in the **Interface based configuration** field.

The output for `show mvrp statistics interface xe-0/1/1.0` confirms that MVRP messages are being transmitted and received on the interface.

- Related Topics**
- Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 76
 - Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 109

Chapter 5

Troubleshooting Bridging and VLAN Configuration

- Troubleshooting Ethernet Switching on page 125

Troubleshooting Ethernet Switching

Troubleshooting issues for Ethernet switching on EX Series switches:

- MAC Address in the Switch's Ethernet Switching Table Is Not Updated After a MAC Address Move on page 125

MAC Address in the Switch's Ethernet Switching Table Is Not Updated After a MAC Address Move

Problem Sometimes a MAC address entry in the switch's Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch. Typically, the switch does not wait for a MAC address expiration when a MAC move operation occurs. As soon as the switch detects the MAC address on the new interface, it immediately updates the table. Many network devices send a gratuitous ARP packet when switching an IP address from one device to another. The switch updates its ARP cache table after receipt of such gratuitous ARP messages, and then it also updates its Ethernet switching table. However, sometimes silent devices, such as SYSLOG servers or SNMP Trap receivers that receive UDP traffic but do not return acknowledgement (ACK) messages to the traffic source, do not send gratuitous ARP packets when a device moves. If such a move occurs when the system administrator is not available to explicitly clear the affected interfaces by issuing the `clear ethernet-switching table` command, the entry for the moved device in the Ethernet switching table is not updated.

Solution Set up the switch to handle unattended MAC address switchovers.

1. Reduce the system-wide ARP aging timer. (By default, the ARP aging timer is set at 20 minutes. In JUNOS Release 9.4 and later, the range of the ARP aging timer is from 1 through 240 minutes.)

```
[edit system arp]
user@switch# set aging-timer 3
```

2. Set the MAC aging timer to the same value as the ARP timer. (By default, the MAC aging timer is set to 300 seconds. The range is 15 to 1,000,000 seconds.)

```
[edit vlans]  
user@switch# set vlans sales mac-table-aging-time 180
```

The ARP entry and the MAC address entry for the moved device expire within the times specified by the aging timer values. After the entries expire, the switch sends a new ARP message to the IP address of the device. The device responds to the ARP, thereby refreshing the entries in the switch's ARP cache table and Ethernet switching table

- Related Topics**
- arp
 - mac-table-aging-time

Chapter 6

Configuration Statements for Bridging and VLANs

- [edit ethernet-switching-options] Configuration Statement Hierarchy on page 127
- [edit interfaces] Configuration Statement Hierarchy on page 129
- [edit protocols] Configuration Statement Hierarchy on page 130
- [edit routing-instances] Configuration Statement Hierarchy on page 136
- [edit vlans] Configuration Statement Hierarchy on page 136

[edit ethernet-switching-options] Configuration Statement Hierarchy

```
ethernet-switching-options {
  analyzer {
    name {
      loss-priority priority;
      ratio number;
      input {
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
        egress {
          interface (all | interface-name);
        }
      }
      output {
        interface interface-name;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdu-block {
    disable-timeout timeout;
    interface (all | [interface-name]);
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100) ;
  }
  interfaces interface-name {
    no-mac-learning;
  }
  mac-table-aging-time seconds;
}
```

```

port-error-disable {
  disable-timeout timeout;
}
redundant-trunk-group {
  group-name name {
    interface interface-name <primary>;
  }
}
secure-access-port {
  dhcp-snooping-file {
    location local_pathname | remote_URL;
    timeout seconds;
    write-interval seconds;
  }
  interface (all | interface-name) {
    allowed-mac {
      mac-address-list;
    }
    (dhcp-trusted | no-dhcp-trusted );
    mac-limit limit action action;
    no-allowed-mac-log;
    static-ip ip-address {
      vlan vlan-name;
      mac mac-address;
    }
  }
}
vlan (all | vlan-name) {
  (arp-inspection | no-arp-inspection );
  dhcp-option82 {
    circuit-id {
      prefix hostname;
      use-interface-description;
      use-vlan-id;
    }
    remote-id {
      prefix hostname | mac | none;
      use-interface-description;
      use-string string;
    }
    vendor-id [string];
  }
  (examine-dhcp | no-examine-dhcp );
  (ip-source-guard | no-ip-source-guard);
  mac-move-limit limit action action;
}
}
storm-control {
  interface (all | interface-name) {
    bandwidth bandwidth;
    no-broadcast;
    no-unknown-unicast;
  }
}
traceoptions {
  file filename <files number> <no-stamp> <replace> <size size> <world-readable
  | no-world-readable>;
}

```

```

    flag flag <disable>;
  }
  unknown-unicast-forwarding {
    vlan (all | vlan-name) {
      interface interface-name;
    }
  }
  voip {
    interface (all | [interface-name | access-ports]) {
      vlan vlan-name ;
      forwarding-class <assured-forwarding | best-effort | expedited-forwarding |
        network-control>;
    }
  }
}

```

- Related Topics**
- Understanding Port Mirroring on EX Series Switches
 - Port Security for EX Series Switches Overview
 - Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches
 - Understanding Redundant Trunk Links on EX Series Switches on page 11
 - Understanding Storm Control on EX Series Switches
 - Understanding 802.1X and VoIP on EX Series Switches
 - Understanding Q-in-Q Tunneling on EX Series Switches on page 13
 - Understanding Unknown Unicast Forwarding on EX Series Switches

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

[edit protocols] Configuration Statement Hierarchy

```

protocols {
  connections {
    remote-interface-switch connection-name {
      interface interface-name.unit-number;
      transmit-lsp label-switched-path;
      receive-lsp label-switched-path;
    }
  }
  dot1x {
    authenticator {
      authentication-profile-name profile-name;
      interface (all | [ interface-names ]) {
        disable;
        guest-vlan ( vlan-id | vlan-name);
        mac-radius <restrict>;
        maximum-requests number;
        no-reauthentication;
        quiet-period seconds;
        reauthentication {
          interval seconds;
        }
      }
      retries number;
      server-fail (deny | permit | use-cache | vlan-id | vlan-name);
      server-reject-vlan ( vlan-id | vlan-name);
      server-timeout seconds;
      supplicant (multiple | single | single-secure);
      supplicant-timeout seconds;
      transmit-period seconds;
    }
  }
}

```

```

    }
    static mac-address {
        interface interface-name;
        vlan-assignment (vlan-id |vlan-name);
    }
}
gvrp {
    <enable | disable>;
    interface (all | [interface-name]) {
        disable;
    }
    join-timer milliseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
}
igmp-snooping {
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
        flag flag (detail | disable | receive | send);
    }
    vlan (vlan-id | vlan-number) {
        data-forwarding {
            source {
                groups group-prefix;
            }
            receiver {
                source-vlans vlan-list;
                install ;
            }
        }
        disable {
            interface interface-name
        }
        immediate-leave;
        interface interface-name {
            group-limit limit;
            multicast-router-interface;
            static {
                group ip-address;
            }
        }
        proxy ;
        query-interval seconds;
        query-last-member-interval seconds;
        query-response-interval seconds;
        robust-count number;
    }
}
lldp {
    disable;
    advertisement-interval seconds;
    hold-multiplier number;
    interface (all | interface-name) {
        disable;
    }
}

```

```

traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>
    <match regex>;
  flag flag (detail | disable | receive | send);
}
}
lldp-med {
  disable;
  fast-start number;
  interface (all | interface-name) {
    disable;
    location {
      elin number;
      civic-based {
        what number;
        country-code code;
        ca-type {
          number {
            ca-value value;
          }
        }
      }
    }
  }
}
mpls {
  interface ( all | interface-name );
  label-switched-path lsp-name to remote-provider-edge-switch;
  path destination {
    <address | hostname> <strict | loose>
  }
}
mstp {
  disable;
  bpdu-block-on-edge;
  bridge-priority priority;
  configuration-name name;
  forward-delay seconds;
  hello-time seconds;
  interface (all | interface-name) {
    disable;
    bpdu-timeout-action {
      block;
      alarm;
    }
    cost cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
  }
  max-age seconds;
  max-hops hops;
  msti msti-id {
    vlan (vlan-id | vlan-name);
    interface interface-name {
      disable;

```

```

        cost cost;
        edge;
        mode mode;
        priority priority;
    }
}
revision-level revision-level;
traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
    no-world-readable>;
    flag flag;
}
}
mvrp {
    disable
    interface (all | interface-name) {
        disable;
        join-timer milliseconds;
        leave-timer milliseconds;
        leaveall-timer milliseconds;
        registration (forbidden | normal);
    }
    no-dynamic-vlan;
    traceoptions {
        file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
        flag flag;
    }
}
}
oam {
    ethernet{
        link-fault-management {
            action-profile profile-name;
            action {
                syslog;
                link-down;
            }
            event {
                link-adjacency-loss;
                link-event-rate;
                frame-error count;
                frame-period count;
                frame-period-summary count;
                symbol-period count;
            }
        }
        interface interface-name {
            link-discovery (active | passive);
            pdu-interval interval;
            event-thresholds threshold-value;
            remote-loopback;
            event-thresholds {
                frame-error count;
                frame-period count;
                frame-period-summary count;
                symbol-period count;
            }
        }
    }
}
}

```

```

    }
    negotiation-options {
        allow-remote-loopback;
        no-allow-link-events;
    }
}
}
rstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
            block;
            alarm;
        }
        cost cost;
        edge;
        mode mode;
        no-root-port;
        priority priority;
    }
    max-age seconds;
}
}
traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
    no-world-readable>;
    flag flag;
}
}
}
sflow {
    collector {
        ip-address;
        udp-port port-number;
    }
    disable;
    interfaces interface-name {
        disable;
        polling-interval seconds;
        sample-rate number;
    }
    polling-interval seconds;
    sample-rate number;
}
}
stp {
    disable;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
            block;

```

```

        alarm;
    }
    cost cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
}
max-age seconds;
}
traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
    no-world-readable>;
    flag flag;
}
vstp {
    bpdu-block-on-edge;
    disable;
    force-version stp;
    vlan (all | vlan-id | vlan-name) {
        bridge-priority priority;
        forward-delay seconds;
        hello-time seconds;
        interface (all | interface-name) {
            bpdu-timeout-action {
                alarm;
                block;
            }
            cost cost;
            disable;
            edge;
            mode mode;
            no-root-port;
            priority priority;
        }
        max-age seconds;
        traceoptions {
            file filename <files number > <size size> <no-stamp | world-readable |
            no-world-readable>;
            flag flag;
        }
    }
}
}
}

```

- Related Topics**
- 802.1X for EX Series Switches Overview
 - Example: Configure Automatic VLAN Administration Using GVRP on page 46
 - Understanding MAC RADIUS Authentication on EX Series Switches
 - Understanding Server Fail Fallback and 802.1X Authentication on EX Series Switches
 - IGMP Snooping on EX Series Switches Overview
 - Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches
 - Understanding MSTP for EX Series Switches

- Understanding RSTP for EX Series Switches
- Understanding STP for EX Series Switches
- Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch
- Understanding VSTP for EX Series Switches
- Understanding Ethernet OAM Link Fault Management for an EX Series Switch

[edit routing-instances] Configuration Statement Hierarchy

```
routing-instances routing-instance-name {
  instance-type virtual-router
  interface interface-name
}
```

- Related Topics**
- Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 73
 - Configuring Virtual Routing Instances (CLI Procedure) on page 105

[edit vlans] Configuration Statement Hierarchy

```
vlans {
  vlan-name {
    description text-description;
    dot1q-tunneling {
      customer-vlans (id | native | range);
      layer2-protocol-tunneling all | protocol-name {
        drop-threshold number;
        shutdown-threshold number;
      }
    }
  }
  filter input filter-name;
  filter output filter-name;
  interface interface-name {
    mapping (native (push | swap) | policy | tag (push | swap));
  }
  l3-interface vlan.logical-interface-number;
  mac-limit number;
  mac-table-aging-time seconds;
  no-local-switching;
  no-mac-learning;
  primary-vlan vlan-name;
  vlan-id number;
  vlan-range vlan-id-low-vlan-id-high;
}
}
```

- Related Topics**
- Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21
 - Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 28
 - Example: Configure Automatic VLAN Administration Using GVRP on page 46
 - Example: Connecting an Access Switch to a Distribution Switch on page 36
 - Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 65
 - Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 88
 - Creating a Private VLAN (CLI Procedure) on page 105
 - Understanding Q-in-Q Tunneling on EX Series Switches on page 13

arp

Syntax arp {
 aging-timer *minutes*;
}

Hierarchy Level [edit system]

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

Description Set the time interval between ARP updates.

Options aging-timer *minutes*—Time interval in minutes between ARP updates. In environments where the number of ARP entries to update is high, increasing the time between updates can improve system performance.

Range: 5 to 240 minutes

Default: 20 minutes

Required Privilege Level system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Related Topics ■ For more information about ARP updates, see the *JUNOS Software System Basics Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos96/index.html>.

bridge-priority

Syntax	bridge-priority <i>priority</i> ;
Hierarchy Level	[edit protocols mstp], [edit protocols mstp msti <i>msti-id</i>], [edit protocols rstp], [edit protocols stp], [edit protocols vstp vlan <i>vlan-id</i>]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches. Statement updated in JUNOS Release 9.4 for EX Series switches to add VSTP support.
Description	Configure the bridge priority. The bridge priority determines which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.
Default	32,768
Options	<i>priority</i> —Bridge priority. It can be set only in increments of 4096. Range: 0 through 61,440 Default: 32,768
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 spanning-tree bridge ■ show spanning-tree interface ■ Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches ■ Understanding MSTP for EX Series Switches ■ Understanding VSTP for EX Series Switches


customer-vlans

Syntax	customer-vlans (<i>id</i> native <i>range</i>);
Hierarchy Level	[edit vlans <i>vlan-name</i> dot1q-tunneling]
Release Information	Statement introduced in JUNOS Release 9.3 for EX Series switches. Option native introduced in JUNOS Release 9.6 for EX Series switches.
Description	Limit the set of accepted C-VLAN tags to a range or to discrete values.
Options	<i>id</i> —Numeric identifier for a VLAN. <i>native</i> —Accepts untagged and priority-tagged packets from access interfaces and assigns the configured S-VLAN to the packet. <i>range</i> —Range of numeric identifiers for VLANs.
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">■ dot1q-tunneling■ ether-type■ Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 65■ Configuring Q-in-Q Tunneling (CLI Procedure) on page 106■ Understanding Q-in-Q Tunneling on EX Series Switches on page 13

description

Syntax	<code>description text-description;</code>
Hierarchy Level	[edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Provide a textual description of the VLAN. The text has no effect on the operation of the VLAN or switch.
Options	<i>text-description</i> —Text to describe the interface. It can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. If the text includes spaces, enclose the entire text in quotation marks.
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 ■ Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21 ■ Understanding Bridging and VLANs on EX Series Switches on page 3

disable

Syntax	<code>disable;</code>
Hierarchy Level	[edit protocols <i>gvrp</i>], [edit protocols <i>gvrp</i> interface [<i>interface-name</i>]]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
	NOTE: As of JUNOS Release 9.2, GVRP can be enabled only on trunk interfaces.
Description	Disable the GVRP configuration on the interface.
Default	If you do not configure GVRP, it is disabled. You can use this command to disable a prior configuration of GVRP on a specified interface.
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 gvrp ■ Example: Configure Automatic VLAN Administration Using GVRP on page 46

disable (MVRP)

Syntax	disable;
Hierarchy Level	[edit protocols mvrp], [edit protocols mvrp interface(all <i>interface-name</i>)]
Release Information	Statement introduced in JUNOS Release 10.0 for EX Series switches.
Description	Disable the MVRP configuration on the interface.
Default	MVRP is disabled by default.
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"> ■ Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 109

dot1q-tunneling (Ethernet Switching)

Syntax	dot1q-tunneling { ether-type (0x8100 0x88a8 0x9100); }
Hierarchy Level	[edit ethernet-switching-options]
Release Information	Statement introduced in JUNOS Release 9.3 for EX Series switches. The remaining statement is explained separately.
Description	Set a global value for the Ethertype for Q-in-Q tunneling.
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"> ■ dot1q-tunneling ■ Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 65 ■ Configuring Q-in-Q Tunneling (CLI Procedure) on page 106

dot1q-tunneling (VLANs)

Syntax dot1q-tunneling {
 customer-vlans (*id* | *native* | *range*);
 layer2-protocol-tunneling all | *protocol-name* {
 drop-threshold *number*;
 shutdown-threshold *number*;
 }
 }

Hierarchy Level [edit vlans *vlan-name*]

Release Information Statement introduced in JUNOS Release 9.3 for EX Series switches.
 Option *native* introduced in JUNOS Release 9.6 for EX Series switches.
 Options *layer2-protocol-tunneling*, *drop-threshold*, and *shutdown-threshold* introduced in JUNOS Release 10.0 for EX Series switches.

Description Enable Q-in-Q tunneling on the specified VLAN.



NOTE: The VLAN on which you enable Q-in-Q tunneling must be a tagged VLAN.

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**
- dot1q-tunneling
 - Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 65
 - Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 88
 - Configuring Q-in-Q Tunneling (CLI Procedure) on page 106
 - Understanding Q-in-Q Tunneling on EX Series Switches on page 13

drop-threshold

Syntax drop-threshold *number*;

Hierarchy Level [edit vlans *vlan-name* dot1q-tunneling layer2-protocol-tunneling all | *protocol-name*]

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

Description Specify the maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the switch begins dropping the Layer 2 PDUs. The drop threshold value must be less than or equal to the shutdown threshold value.



NOTE: If the drop threshold value is greater than the shutdown threshold value and you try to commit the configuration, the commit will fail.

You can specify a drop threshold value without specifying a shutdown threshold value.

Default No drop threshold is specified.

Options *number*—Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the switch begins dropping the Layer 2 PDUs.
Range: 1 through 1000

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

Related Topics

- shutdown-threshold
- Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 88
- Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) on page 112

ether-type

Syntax	ether-type (0x8100 0x88a8 0x9100)
Hierarchy Level	[edit ethernet-switching-options dot1q-tunneling]
Release Information	Statement introduced in JUNOS Release 9.3 for EX Series switches.
Description	Configure a global value for the Ethertype. Only one Ethertype value is supported at a time. The Ethertype value appears in the Ethernet type field of the packet. It specifies the protocol being transported in the Ethernet frame.
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">■ dot1q-tunneling■ Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 65■ Configuring Q-in-Q Tunneling (CLI Procedure) on page 106

ethernet-switching-options

```

Syntax ethernet-switching-options {
  analyzer {
    name {
      loss-priority priority;
      ratio number;
      input {
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
        egress {
          interface (all | interface-name);
        }
      }
      output {
        interface interface-name;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdu-block {
    disable-timeout timeout;
    interface (all | [interface-name]);
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100)
  }
  interfaces interface-name {
    no-mac-learning;
  }
  mac-table-aging-time seconds;
  port-error-disable {
    disable-timeout timeout;
  }
  redundant-trunk-group {
    group-name name {
      interface interface-name <primary>;
      interface interface-name;
    }
  }
  secure-access-port {
    dhcp-snooping-file {
      location local_pathname | remote_URL;
      timeout seconds;
      write-interval seconds;
    }
    interface (all | interface-name) {
      allowed-mac {
        mac-address-list;
      }
      (dhcp-trusted | no-dhcp-trusted);
      mac-limit limit action action;
    }
  }
}

```

```

    no-allowed-mac-log;
    static-ip ip-address {
        vlan vlan-name;
        mac mac-address;
    }
}
vlan (all | vlan-name) {
    (arp-inspection | no-arp-inspection);
    dhcp-option82 {
        circuit-id {
            prefix hostname;
            use-interface-description;
            use-vlan-id;
        }
        remote-id {
            prefix hostname | mac | none;
            use-interface-description;
            use-string string;
        }
        vendor-id [string];
    }
    (examine-dhcp | no-examine-dhcp);
    (ip-source-guard | no-ip-source-guard);
    mac-move-limit limit action action;
}
}
storm-control {
    interface (all | interface-name) {
        level level;
        no-broadcast;
        no-unknown-unicast;
    }
}
traceoptions {
    file filename <files number> <no-stamp> <replace> <size size> <world-readable |
    no-world-readable>;
    flag flag <disable>;
}
unknown-unicast-forwarding {
    vlan (all | vlan-name) {
        interface interface-name;
    }
}
}
voip {
    interface (all | [interface-name | access-ports]) {
        vlan vlan-name ;
        forwarding-class <assured-forwarding | best-effort | expedited-forwarding |
        network-control>;
    }
}
}
}

```

Hierarchy Level [edit]

Release Information Statement introduced in JUNOS Release 9.0 for EX Series switches.
 Support for storm control and BPDU protection added in JUNOS Release 9.1 for EX Series switches.
 Option `ip-source-guard` added in JUNOS Release 9.2 for EX Series switches.
 Options `dhcp-option82`, `dot1q-tunneling`, and `no-allowed-mac-log` added in JUNOS Release 9.3 for EX Series switches.
 Options `dhcp-snooping-file` and `mac-table-aging-time` introduced in JUNOS Release 9.4 for EX Series switches.
 Options `interfaces` and `no-mac-learning` introduced in JUNOS Release 9.5 for EX Series switches.
 Options `port-error-disable` and `disable-timeout` introduced in JUNOS Release 9.6 for EX Series switches.

Description Configure Ethernet switching options.

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**
- Understanding Port Mirroring on EX Series Switches
 - Port Security for EX Series Switches Overview
 - Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches
 - Understanding Redundant Trunk Links on EX Series Switches on page 11
 - Understanding Storm Control on EX Series Switches
 - Understanding 802.1X and VoIP on EX Series Switches
 - Understanding Q-in-Q Tunneling on EX Series Switches on page 13
 - Understanding Unknown Unicast Forwarding on EX Series Switches

filter

Syntax	filter (input output) <i>filter-name</i> ;
Hierarchy Level	[edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Apply a firewall filter to traffic coming into or exiting from the VLAN.
Default	All incoming traffic is accepted unmodified to the VLAN, and all outgoing traffic is sent unmodified from the VLAN.
Options	<p><i>filter-name</i> —Name of a firewall filter defined in a filter statement.</p> <ul style="list-style-type: none"> ■ input—Apply a firewall filter to VLAN ingress traffic. ■ output—Apply a firewall filter to VLAN egress traffic.
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 Firewall Filters (CLI Procedure) ■ Configuring Firewall Filters (J-Web Procedure) ■ Firewall Filters for EX Series Switches Overview

group-name

Syntax group-name *name* {
 interface *interface-name* <primary>;
 interface *interface-name*;
}

Hierarchy Level [edit ethernet-switching-options redundant-trunk-group]

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

Description Create a redundant trunk group.

Options *name*—The name of the redundant trunk group. The group name must start with a letter and can consist of letters, numbers, dashes, and underscores.

The remaining options 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 Redundant Trunk Links for Faster Recovery on page 61
 - Understanding Redundant Trunk Links on EX Series Switches on page 11

gvrp

Syntax gvrp {
 interface [*interface-name*] {
 disable;
 }
 join-timer *milliseconds*;
 leave-timer *milliseconds*;
 leaveall-timer *milliseconds*;
 }

Hierarchy Level [edit protocols]

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



NOTE: As of JUNOS Release 9.2, GVRP can be enabled only on trunk interfaces.

Description When GVRP is configured on a trunk interface, it ensures that the VLAN membership information on the trunk interface is updated as the switch's access interfaces become active or inactive in the configured VLANs.

The statements are explained separately.

Default GVRP is disabled by default.

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

- Related Topics**
- show gvrp
 - Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 28
 - Example: Configure Automatic VLAN Administration Using GVRP on page 46

instance-type

Syntax	instance-type virtual-router
Hierarchy Level	[edit routing-instances]
Release Information	Statement introduced in JUNOS Release 9.2 for EX Series switches.
Description	Specify the type of routing instance.
Options	virtual-router—A logical entity.
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: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 73 ■ Configuring Virtual Routing Instances (CLI Procedure) on page 105

interface

Syntax	interface (all [<i>interface-name</i>]) { <enable disable>; }
Hierarchy Level	[edit protocols gvrp]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Configure GARP VLAN Registration Protocol (GVRP) for one or more interfaces.
Default	By default, GVRP is disabled.
Options	all—All interfaces. <i>interface-name</i> —The list of interfaces to be configured for GVRP. 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	<ul style="list-style-type: none"> ■ show gvrp ■ Example: Configure Automatic VLAN Administration Using GVRP on page 46

interface (MVRP)

Syntax interface (all | *interface-name*) {
 disable;
 join-timer *milliseconds*;
 leave-timer *milliseconds*;
 leaveall-timer *milliseconds*;
 registration (forbidden | normal);
 }

Hierarchy Level [edit protocols mvrp]

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

Description Specify interfaces on which to configure Multiple VLAN Registration Protocol (MVRP).

Default By default, MVRP is disabled.

Options all—All interfaces on the switch.

interface-name—Names of interface to be configured for MVRP.

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 Automatic VLAN Administration Using MVRP on EX Series Switches on page 76
 - Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 109

interface

Syntax	<code>interface <i>interface-name</i> <primary>; interface <i>interface-name</i>;</code>
Hierarchy Level	[edit ethernet-switching-options redundant-trunk-group group-name <i>name</i>]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Configure a primary link and secondary link on trunk ports. If the primary link fails, the secondary link automatically takes over as the primary link without waiting for normal STP convergence.
Options	<p><code>interface <i>interface-name</i></code>—A logical interface or an aggregated interface containing multiple ports.</p> <p><code>primary</code>—(Optional) Specify one of the interfaces in the redundant group as the primary link. The interface without this option is the secondary link in the redundant group. If a link is not specified as primary, the software compares the two links and selects the link with the highest port number as the active link. For example, if the two interfaces are <code>ge-0/1/0</code> and <code>ge-0/1/1</code>, the software assigns <code>ge-0/1/1</code> as the active link.</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 Redundant Trunk Links for Faster Recovery on page 61 ■ Understanding Redundant Trunk Links on EX Series Switches on page 11

interface

Syntax	<code>interface <i>interface-name</i>;</code>
Hierarchy Level	[edit routing-instances]
Release Information	Statement introduced in JUNOS Release 9.2 for EX Series switches.
Description	For virtual routing instances, configure an interface.
Options	<code><i>interface-name</i></code> —Name of a Gigabit Ethernet interface.
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: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 73 ■ Configuring Virtual Routing Instances (CLI Procedure) on page 105 ■ Understanding Virtual Routing Instances on EX Series Switches on page 10

interface

Syntax `interface interface-name {
 mapping (native (push | swap) | policy | tag (push | swap));
 }`

Hierarchy Level [edit vlans *vlan-name*]

Release Information Statement introduced in JUNOS Release 9.3 for EX Series switches.
 Option `mapping` introduced in JUNOS Release 9.6 for EX Series switches.
 Option `swap` introduced in JUNOS Release 10.0 for EX Series switches.

Description For a specific VLAN, configure an interface.

Options *interface-name*—Name of a Gigabit Ethernet interface.

The remaining statement is 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: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21
 - Configuring VLANs for EX Series Switches (CLI Procedure) on page 98
 - Understanding Bridging and VLANs on EX Series Switches on page 3
 - Understanding Q-in-Q Tunneling on EX Series Switches on page 13

interfaces

Syntax	<code>interfaces <i>interface-name</i> { no-mac-learning; }</code>
Hierarchy Level	[edit ethernet-switching-options]
Release Information	Statement introduced in JUNOS Release 9.5 for EX Series switches.
Description	Configure settings for interfaces that have been assigned to family <code>ethernet-switching</code> .
Options	<i>interface-name</i> --Name of an interface that is configured for family <code>ethernet-switching</code> . The remaining statement is explained separately.
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none"> ■ Understanding Q-in-Q Tunneling on EX Series Switches on page 13

join-timer

Syntax	<code>join-timer <i>milliseconds</i>;</code>
Hierarchy Level	[edit protocols gvrp]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	For GARP VLAN Registration Protocol (GVRP), configure the maximum number of milliseconds interfaces must wait before sending VLAN advertisements.
Default	20 milliseconds
Options	<i>milliseconds</i> —Number of milliseconds. Default: 20 milliseconds
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none"> ■ <code>show gvrp</code> ■ Example: Configure Automatic VLAN Administration Using GVRP on page 46

join-timer (MVRP)

Syntax	join-timer <i>milliseconds</i> ;
Hierarchy Level	[edit protocols mvrp interface (all <i>interface-name</i>)]
Release Information	Statement introduced in JUNOS Release 10.0 for EX Series switches.
Description	<p>Configure the maximum number of milliseconds interfaces must wait before sending Multiple VLAN Registration Protocol (MVRP) protocol data units (PDUs).</p> <p>Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.</p>
Default	200 milliseconds
Options	<i>milliseconds</i> —Number of milliseconds that the interface must wait before sending MVRP PDUs.
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"> ■ leave-timer ■ leaveall-timer ■ Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 76 ■ Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 109

I3-interface

Syntax	I3-interface vlan. <i>logical-interface-number</i> ;
Hierarchy Level	[edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Associate a Layer 3 interface with the VLAN. Configure Layer 3 interfaces on trunk ports to allow the interface to transfer traffic between multiple VLANs. Within a VLAN, traffic is bridged, while across VLANs, traffic is routed.
Default	No Layer 3 (routing) interface is associated with the VLAN.
Options	<i>vlan.logical-interface-number</i> —Number of the logical interface defined with a <code>set interfaces vlan unit</code> command. For the logical interface number, use the same number you configure in the <code>unit</code> statement.
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">■ <code>show ethernet-switching interfaces</code>■ <code>show vlans</code>■ Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21■ Example: Connecting an Access Switch to a Distribution Switch on page 36■ Configuring Routed VLAN Interfaces (CLI Procedure) on page 99■ Understanding Bridging and VLANs on EX Series Switches on page 3

layer2-protocol-tunneling

Syntax layer2-protocol-tunneling all | *protocol-name* {
 drop-threshold *number*;
 shutdown-threshold *number*;
 }

Hierarchy Level [edit vlans *vlan-name* dot1q-tunneling]

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

Description Enable Layer 2 protocol tunneling (L2PT) on the VLAN.

The remaining statements are explained separately.

Default L2PT is not enabled.

Options all—Enable all supported Layer 2 protocols.

protocol-name—Name of the Layer 2 protocol. Values are:

- cdp—Cisco Discovery Protocol
- gvrp—GARP VLAN Registration Protocol
- llpd—Link Layer Discovery Protocol
- mvrp—Multiple VLAN Registration Protocol
- stp—Spanning Tree Protocol, Rapid Spanning Tree Protocol, and Multiple Spanning Tree Protocol
- vstp—VLAN Spanning Tree Protocol
- vtp—VLAN Trunking Protocol

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

- Related Topics**
- show ethernet-switching layer2-protocol-tunneling interface
 - show ethernet-switching layer2-protocol-tunneling statistics
 - show ethernet-switching layer2-protocol-tunneling vlan
 - Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 88
 - Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) on page 112

leave-timer

Syntax	leave-timer <i>milliseconds</i> ;
Hierarchy Level	[edit protocols gvrp]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	For GARP VLAN Registration Protocol (GVRP), configure the number of milliseconds an interface waits after receiving a leave message before the interface leaves the VLAN specified in the message. If the interface receives a join message before the timer expires, the software keeps the interface in the VLAN.
Default	60 centiseconds
Options	<i>milliseconds</i> —Number of milliseconds. At a minimum, the leave timer interval should be twice the join timer interval. Default: 60 centiseconds
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 gvrp■ Example: Configure Automatic VLAN Administration Using GVRP on page 46

leave-timer (MVRP)

Syntax	leave-timer <i>milliseconds</i> ;
Hierarchy Level	[edit protocols mvrp interface (all <i>interface-name</i>)]
Release Information	Statement introduced in JUNOS Release 10.0 for EX Series switches.
Description	<p>For Multiple VLAN Registration Protocol (MVRP), configure the number of milliseconds the switch retains a VLAN in the Leave state before the VLAN is unregistered. If the interface receives a join message before this timer expires, the VLAN remains registered.</p> <p>Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.</p>
Default	1000 milliseconds
Options	<i>milliseconds</i> —Number of milliseconds that the switch retains a VLAN in the Leave state before the VLAN is unregistered. At a minimum, set the leave-timer interval at twice the join-timer interval.
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"> ■ join-timer ■ leaveall-timer ■ Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 76 ■ Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 109

leaveall-timer

Syntax	leaveall-timer <i>milliseconds</i> ;
Hierarchy Level	[edit protocols gvrp]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	For GARP VLAN Registration Protocol (GVRP), configure the interval at which Leave All messages are sent on the interfaces. Leave All messages maintain current GVRP VLAN membership information in the network. A Leave All message instructs the port to change the GVRP state for all its VLANs to a leaving state and remove them unless a Join message is received before the leave timer expires.
Default	1000 centiseconds
Options	<i>milliseconds</i> —Number of milliseconds. Range: 5 times <i>leave-timer</i> value Default: 1000 centiseconds
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 gvrp■ Example: Configure Automatic VLAN Administration Using GVRP on page 46

leaveall-timer (MVRP)

Syntax	leaveall-timer <i>milliseconds</i> ;
Hierarchy Level	[edit protocols mvrp interface (all <i>interface-name</i>)]
Release Information	Statement introduced in JUNOS Release 10.0 for EX Series switches.
Description	<p>For Multiple VLAN Registration Protocol (MVRP), configure the interval at which the LeaveAll state operates on the interface.</p> <p>Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.</p>
Default	10000 milliseconds
Options	<i>milliseconds</i> —Number of milliseconds between the sending of Leave All messages.
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"> ■ join-timer ■ leave-timer ■ Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 76 ■ Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 109

mac-limit

Syntax `mac-limit number;`

Hierarchy Level `[edit vlans vlan-name]`

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

Description Configure the number of MAC addresses allowed on a VLAN.

Default MAC limit is disabled.

Options *number*—Maximum number of MAC addresses.

Range: 1 through 32768

Required Privilege Level routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Related Topics ■ `show vlans`

- Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21
- Configuring MAC Table Aging (CLI Procedure) on page 101
- Understanding Bridging and VLANs on EX Series Switches on page 3

mac-table-aging-time

Syntax	mac-table-aging-time <i>seconds</i> ;
Hierarchy Level	[edit ethernet-switching-options], [edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches. Statement updated in JUNOS Release 9.4 for EX Series switches to include [edit ethernet-switching-options] hierarchy level.
Description	Define how long entries remain in the Ethernet switching table before expiring: <ul style="list-style-type: none"> ■ If you specify this statement at the [edit ethernet-switching-options] hierarchy level, it applies to all VLANs on the switch. ■ If you specify this statement at the [edit vlans] hierarchy level, it applies to the specified VLAN.
Default	Entries remain in the Ethernet switching table for 300 seconds
Options	<i>seconds</i> —Time that entries remain in the Ethernet switching table before being removed. <ul style="list-style-type: none"> ■ Range—60 through 1,000,000 seconds ■ Default—300 seconds
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 ethernet-switching statistics aging ■ Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21 ■ Configuring MAC Table Aging (CLI Procedure) on page 101 ■ Configuring VLANs for EX Series Switches (CLI Procedure) on page 98 ■ Understanding Bridging and VLANs on EX Series Switches on page 3

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 on page 21
 - Example: Connecting an Access Switch to a Distribution Switch on page 36
 - Configuring Gigabit Ethernet Interfaces (CLI Procedure)
 - Configuring Gigabit Ethernet Interfaces (J-Web Procedure)
 - Creating a Series of Tagged VLANs (CLI Procedure) on page 103
 - Understanding Bridging and VLANs on EX Series Switches on page 3

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

mvrp

Syntax

```
mvrp {
  disable
  interface (all | interface-name) {
    disable;
    join-timer milliseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
    registration (forbidden | normal);
  }
  no-dynamic-vlan;
  traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
}
```

Hierarchy Level [edit protocols]

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

Description Configure Multiple VLAN Registration Protocol (MVRP) on a trunk interface to ensure that the VLAN membership information on the trunk interface is updated as the switch's access interfaces become active or inactive in the configured VLANs.

The remaining statements are explained separately.

Default MVRP is disabled by default.

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

- Related Topics**
- Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 76
 - Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 109

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)■ <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 on page 3

no-dynamic-vlan

Syntax	no-dynamic-vlan;
Hierarchy Level	[edit protocols mvrp]
Release Information	Statement introduced in JUNOS Release 10.0 for EX Series switches.
Description	<p>Disable the dynamic creation of VLANs using Multiple VLAN Registration Protocol (MVRP) for interfaces participating in MVRP.</p> <p>Dynamic VLAN configuration can be enabled on an interface independent of MVRP. The MVRP dynamic VLAN configuration setting does not override the interface configuration dynamic VLAN configuration setting. If dynamic VLAN creation is disabled on the interface in the interface configuration, no dynamic VLANs are created on the interface, including dynamic VLANs created using MVRP.</p> <p>This option can only be applied globally; it cannot be applied per interface.</p>
Default	If MVRP is enabled, the dynamic creation of VLANs as a result of MVRP protocol exchange messages is enabled.
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"> ■ Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 109

no-local-switching

Syntax	no-local-switching
Hierarchy Level	[edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in JUNOS Release 9.3 for EX Series switches.
Description	Specify that access ports in this VLAN domain do not forward packets to each other. You use this statement with primary VLANs and isolated secondary VLANs.
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 a Private VLAN on an EX Series Switch on page 68 ■ Creating a Private VLAN (CLI Procedure) on page 105

no-mac-learning

Syntax no-mac-learning;

Hierarchy Level [edit vlans *vlan-name*]

Release Information Statement introduced in JUNOS Release 9.5 for EX Series switches.

Description Disables MAC address learning for the specified VLAN.

Options There are no options to this statement.

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

- Related Topics**
- Configuring Q-in-Q Tunneling (CLI Procedure) on page 106
 - Understanding Q-in-Q Tunneling on EX Series Switches on page 13

no-mac-learning

Syntax	no-mac-learning;
Hierarchy Level	[edit ethernet-switching-options interfaces <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 9.5 for EX Series switches.
Description	Disable MAC address learning for the specified interface. Disabling MAC address learning on an interface disables learning for all the VLANs of which that interface is a member.
Options	There are no options to this statement.
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"> ■ Understanding Q-in-Q Tunneling on EX Series Switches on page 13

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	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) ■ Configuring Gigabit Ethernet Interfaces (J-Web Procedure) ■ <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 on page 36

primary-vlan

Syntax primary-vlan *vlan-name*

Hierarchy Level [edit vlans *vlan-name*]

Release Information Statement introduced in JUNOS Release 9.3 for EX Series switches. Statement updated with enhanced ? (CLI completion feature) functionality in JUNOS Release 9.5 for EX Series switches.

Description Configure the primary VLAN for this community VLAN. The primary VLAN must be tagged, and the community VLAN must be untagged.



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

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

- Related Topics**
- Example: Configuring a Private VLAN on an EX Series Switch on page 68
 - Creating a Private VLAN (CLI Procedure) on page 105

redundant-trunk-group

Syntax `redundant-trunk-group {
 group-name name {
 interface interface-name <primary>;
 interface interface-name;
 }
}`

Hierarchy Level [edit ethernet-switching-options]

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

Description Configure a primary link and secondary link on trunk ports. If the primary link fails, the secondary link automatically takes over without waiting for normal STP convergence.

Options The 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 Redundant Trunk Links for Faster Recovery on page 61
- Understanding Redundant Trunk Links on EX Series Switches on page 11

registration

Syntax	registration (forbidden normal);
Hierarchy Level	[edit protocols mvrp interface (all <i>interface-name</i>)]
Release Information	Statement introduced in JUNOS Release 10.0 for EX Series switches.
Description	Specifies the Multiple VLAN Registration Protocol (MVRP) registration mode for the interface if MVRP is enabled.
Default	normal
Options	<p>forbidden—The interface or interfaces do not register and do not participate in MVRP.</p> <p>normal—The interface or interfaces accept MVRP messages and participate in MVRP.</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"> ■ Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 109

routing-instances

Syntax	<pre>routing-instances <i>routing-instance-name</i> { instance-type virtual-router; interface <i>interface-name</i>; }</pre>
Hierarchy Level	[edit]
Release Information	Statement introduced in JUNOS Release 9.2 for EX Series switches.
Description	Configure a virtual routing entity.
Options	<p><i>routing-instance-name</i>—Name for this routing instance.</p> <p>The remaining statements are explained separately.</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: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 73 ■ Configuring Virtual Routing Instances (CLI Procedure) on page 105

shutdown-threshold

Syntax	shutdown-threshold <i>number</i> ;
Hierarchy Level	[edit vlans <i>vlan-name</i> dot1q-tunneling layer2-protocol-tunneling all <i>protocol-name</i>]
Release Information	Statement introduced in JUNOS Release 10.0 for EX Series switches.
Description	<p>Specify the maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the interface is disabled. Once an interface is disabled, you must explicitly reenable it using the <code>clear ethernet-switching layer2-protocol-tunneling error</code> command. Otherwise, the interface remains disabled.</p> <p>The shutdown threshold value must be greater than or equal to the drop threshold value. If the shutdown threshold value is less than the drop threshold value, the drop threshold value has no effect.</p> <p>You can specify a shutdown threshold value without specifying a drop threshold value.</p>
Default	No shutdown threshold is specified.
Options	<p><i>number</i>—Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the interface is disabled.</p> <p>Range: 1 through 1000</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"> ■ drop-threshold ■ Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 88 ■ Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) on page 112

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 on page 28
 - Configuring Routed VLAN Interfaces (CLI Procedure) on page 99
 - Understanding Bridging and VLANs on EX Series Switches on page 3
 - *JUNOS Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>

vlan-id

Syntax	<code>vlan-id number;</code>
Hierarchy Level	[edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Configure an 802.1Q tag to apply to all traffic that originates on the VLAN.
Default	If you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of 0.
Options	<i>number</i> —VLAN tag identifier. Range: 0 through 4093.
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: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 28 ■ Understanding Bridging and VLANs on EX Series Switches on page 3

vlan-range

Syntax	<code>vlan-range vlan-id-low-vlan-id-high;</code>
Hierarchy Level	[edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in JUNOS Release 9.2 for EX Series switches.
Description	Configure multiple VLANs. Each VLAN is assigned a VLAN ID number from the range.
Default	None.
Options	<i>vlan-id-low-vlan-id-high</i> —Specify the first and last VLAN ID number for the group of VLANs.
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"> ■ Configuring VLANs for EX Series Switches (CLI Procedure) on page 98 ■ Configuring VLANs for EX Series Switches (J-Web Procedure) on page 95 ■ Configuring Routed VLAN Interfaces (CLI Procedure) on page 99 ■ Understanding Bridging and VLANs on EX Series Switches on page 3

vlan

```

Syntax  vlan {
            vlan-name {
                description text-description;
                dot1q-tunneling {
                    customer-vlans (id | range)
                    layer2-protocol-tunneling all | protocol-name {
                        drop-threshold number;
                        shutdown-threshold number;
                    }
                }
            }
            filter input filter-name;
            filter output filter-name;
            interface interface-name {
                mapping (native (push | swap) | policy | tag (push | swap));
            }
            l3-interface vlan.logical-interface-number;
            mac-limit number;
            mac-table-aging-time seconds;
            no-local-switching;
            no-mac-learning;
            primary-vlan vlan-name;
            vlan-id number;
            vlan-range vlan-id-low-vlan-id-high;
        }
    }

```

Hierarchy Level [edit]

Release Information Statement introduced in JUNOS Release 9.0 for EX Series switches. Options `dot1q-tunneling`, `no-local-switching`, and `primary-vlan` introduced in JUNOS Release 9.3 for EX Series switches. Periods (.) in VLAN names introduced in JUNOS Release 9.4 for EX Series switches. Option `no-mac-learning` introduced in JUNOS Release 9.5 for EX Series switches. Option `mapping` introduced in JUNOS Release 9.6 for EX Series switches. Option `swap` introduced in JUNOS Release 10.0 for EX Series switches.

Description Configure VLAN properties on EX Series switches. The following configuration guidelines apply:

- Only private VLAN (PVLAN) firewall filters can be used when the VLAN is enabled for Q-in-Q tunneling.
- An S-VLAN tag is added to the packet if the VLAN is dot1q-tunneled and the packet is arriving from an access interface.
- You cannot use a firewall filter to assign a routed VLAN interface (RVI) to a VLAN.
- VLAN assignments performed using a firewall filter override all other VLAN assignments.

Default If you use the default factory configuration, all switch interfaces become part of the VLAN default.

Options *vlan-name*—Name of the VLAN. The name can contain letters, numbers, hyphens (-), and periods (.) and can be up to 255 characters long.

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**
- Configuring VLANs for EX Series Switches (CLI Procedure) on page 98
 - Configuring VLANs for EX Series Switches (J-Web Procedure) on page 95
 - Configuring Q-in-Q Tunneling (CLI Procedure) on page 106
 - Creating a Series of Tagged VLANs (CLI Procedure) on page 103
 - Configuring Routed VLAN Interfaces (CLI Procedure) on page 99
 - Understanding Q-in-Q Tunneling on EX Series Switches on page 13
 - Understanding Bridging and VLANs on EX Series Switches on page 3

Chapter 7

Operational Mode Commands for Bridging and VLANs

clear ethernet-switching layer2-protocol-tunneling error

Syntax	clear ethernet-switching layer2-protocol-tunneling error <interface <i>interface-name</i> >
Release Information	Command introduced in JUNOS Release 10.0 for EX Series switches.
Description	Clear Layer 2 protocol tunneling (L2PT) errors on one or more interfaces. If an interface has been disabled because the amount of Layer 2 protocol traffic exceeded the <code>shutdown-threshold</code> or because the switch has detected an error in the network topology or configuration, use this command to reenable the interface.
Options	none—Clears L2PT errors on all interfaces. interface <i>interface-name</i> —(Optional) Clear L2PT errors on the specified interface.
Required Privilege Level	view
Related Topics	<ul style="list-style-type: none"> ■ Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 88 ■ Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) on page 112
List of Sample Output	clear ethernet-switching layer2-protocol-tunneling error on page 182 clear ethernet-switching layer2-protocol-tunneling error interface ge-0/1/1.0 on page 182
clear ethernet-switching layer2-protocol-tunneling error	user@switch> clear ethernet-switching layer2-protocol-tunneling error
clear ethernet-switching layer2-protocol-tunneling error interface ge-0/1/1.0	user@switch> clear ethernet-switching layer2-protocol-tunneling error interface ge-0/1/1.0

clear ethernet-switching layer2-protocol-tunneling statistics

Syntax	clear ethernet-switching layer2-protocol-tunneling statistics <interface <i>interface-name</i> > <vlan <i>vlan-name</i> >
Release Information	Command introduced in JUNOS Release 10.0 for EX Series switches.
Description	Clear Layer 2 protocol tunneling (L2PT) statistics on one or more interfaces or VLANs.
Options	none—Clear L2PT statistics on all interfaces and VLANs. interface <i>interface-name</i> —(Optional) Clear L2PT statistics on the specified interface. vlan <i>vlan-name</i> —(Optional) Clear L2PT statistics on the specified VLAN.
Required Privilege Level	view
Related Topics	<ul style="list-style-type: none"> ■ show ethernet-switching layer2-protocol-tunneling statistics ■ Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 88 ■ Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) on page 112
List of Sample Output	clear ethernet-switching layer2-protocol-tunneling statistics on page 183 clear ethernet-switching layer2-protocol-tunneling error interface ge-0/1/1.0 on page 183 clear ethernet-switching layer2-protocol-tunneling error vlan v2 on page 183
clear ethernet-switching layer2-protocol-tunneling statistics	user@switch> clear ethernet-switching layer2-protocol-tunneling statistics
clear ethernet-switching layer2-protocol-tunneling error interface ge-0/1/1.0	user@switch> clear ethernet-switching layer2-protocol-tunneling statistics interface ge-0/1/1.0
clear ethernet-switching layer2-protocol-tunneling error vlan v2	user@switch> clear ethernet-switching layer2-protocol-tunneling statistics vlan v2

clear gvrp statistics

Syntax clear gvrp statistics

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

Description Clear GARP VLAN Registration Protocol (GVRP) statistics.

Required Privilege Level clear

- Related Topics**
- show spanning-tree statistics
 - Example: Configure Automatic VLAN Administration Using GVRP on page 46

List of Sample Output clear gvrp statistics on page 184

clear gvrp statistics user@switch> **clear gvrp statistics**

clear mvrp statistics

Syntax	clear mvrp statistics <interface <i>interface-name</i> >
Release Information	Command introduced in JUNOS Release 10.0 for EX Series switches.
Description	Clear Multiple VLAN Registration Protocol (MVRP) statistics.
Options	<p>none—Clear all MVRP statistics.</p> <p>interface <i>interface-name</i>—Clear the MVRP statistics on the specified interface.</p>
Required Privilege Level	clear
Related Topics	<ul style="list-style-type: none"> ■ show mvrp statistics ■ Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 76
List of Sample Output	<p>clear mvrp statistics on page 185</p> <p>clear mvrp statistics interface ge-0/0/1.0 on page 185</p>
Output Fields	When you enter this command, you are provided feedback on the status of your request.
clear mvrp statistics	user@switch> clear mvrp statistics
clear mvrp statistics interface ge-0/0/1.0	user@switch> clear mvrp statistics interface ge-0/0/1.0

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 187</p> <p>show ethernet-switching interfaces ge-0/0/15 brief on page 188</p> <p>show ethernet-switching interfaces ge-0/0/2 detail (Blocked by RTG rtggroup) on page 188</p> <p>show ethernet-switching interfaces ge-0/0/15 detail (Blocked by STP) on page 188</p> <p>show ethernet-switching interfaces ge-0/0/17 detail (Disabled by bpdu-control) on page 188</p> <p>show ethernet-switching interfaces detail (C-VLAN to S-VLAN Mapping) on page 188</p>
Output Fields	Table 14 on page 186 lists the output fields for the show ethernet-switching interfaces command. Output fields are listed in the approximate order in which they appear.

Table 14: 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 14: 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
detail (Blocked by RTG Interface: ge-0/0/2.0, Index: 65, State: up, Port mode: Access
rtggroup) VLAN membership:
                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
detail (Blocked by STP) Interface: ge-0/0/15.0, Index: 70, State: up, Port mode: Trunk
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
detail (Disabled by Interface: ge-0/0/17.0, Index: 71, State: down, Port mode: Trunk
bpdu-control) VLAN membership:
                vlan100, 802.1Q Tag: 100, tagged, msti-id: 1, Disabled by bpdu-control
                vlan200, 802.1Q Tag: 200, tagged, msti-id: 2, Disabled by bpdu-control
Number of MACs learned on IFL: 0

```

```

show ethernet-switching 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 ethernet-switching layer2-protocol-tunneling interface

Syntax	show ethernet-switching-layer2-protocol-tunneling interface <interface-name>
Release Information	Command introduced in JUNOS Release 10.0 for EX Series switches.
Description	Display information about Layer 2 protocol tunneling (L2PT) on interfaces that have been configured for L2PT.
Options	none—Display L2PT information about all interfaces on which L2PT is enabled. <i>interface-name</i> —(Optional) Display L2PT information for the specified interface.
Required Privilege Level	view
Related Topics	<ul style="list-style-type: none"> ■ show ethernet-switching layer2-protocol-tunneling statistics ■ show ethernet-switching layer2-protocol-tunneling vlan ■ Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) on page 112
List of Sample Output	show ethernet-switching layer2-protocol-tunneling interface on page 189 show ethernet-switching layer2-protocol-tunneling interface ge-0/0/0.0 on page 190
Output Fields	Table 15 on page 189 lists the output fields for the show ethernet-switching layer2-protocol-tunneling interface command. Output fields are listed in the approximate order in which they appear.

Table 15: show ethernet-switching layer2-protocol-tunneling interface Output Fields

Field Name	Field Description
Interface	Name of an interface on the switch.
Operation	Type of operation being performed on the interface. Values are Encapsulation and Decapsulation.
State	State of the interface. Values are active and shutdown.
Description	If the interface state is shutdown, displays why the interface is shut down. If the description says Loop detected, it means that the interface is an access interface that has received L2PT-enabled PDUs. Access interfaces should not receive L2PT-enabled PDUs. This scenario might mean that there is a loop in the network.

```

show ethernet-switching user@switch> show ethernet-switching layer2-protocol-tunneling interface
layer2-protocol-tunneling
interface
Layer2 Protocol Tunneling information:
Interface      Operation      State      Description
ge-0/0/0.0    Encapsulation  Shutdown   Shutdown threshold exceeded
ge-0/0/1.0    Decapsulation  Shutdown   Loop detected
ge-0/0/2.0    Decapsulation  Active

```

```
user@switch> show ethernet-switching layer2-protocol-tunneling interface ge-0/0/0.0
show ethernet-switching
layer2-protocol-tunneling
interface ge-0/0/0.0
Layer2 Protocol Tunneling information:
Interface      Operation      State      Description
ge-0/0/0.0     Encapsulation  Shutdown   Shutdown threshold exceeded
```

show ethernet-switching layer2-protocol-tunneling statistics

Syntax show ethernet-switching-layer2-protocol-tunneling statistics
 <interface *interface-name*>
 <vlan *vlan-name*>

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

Description Display Layer 2 protocol tunneling (L2PT) statistics for Layer 2 PDU packets received by the switch.



NOTE: The show ethernet-switching-layer2-protocol-tunneling statistics command does not display L2PT statistics for Layer 2 PDU packets transmitted from the switch.

Options none—Display L2PT statistics for all interfaces on which you enabled L2PT.

<interface *interface-name*>—(Optional) Display L2PT statistics for the specified interface.

<vlan *vlan-name*>—(Optional) Display L2PT statistics for the specified VLAN.

Required Privilege Level view

- Related Topics**
- clear ethernet-switching layer2-protocol-tunneling statistics
 - show ethernet-switching layer2-protocol-tunneling interface
 - show ethernet-switching layer2-protocol-tunneling vlan
 - show vlans
 - Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 88
 - Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) on page 112

List of Sample Output show ethernet-switching layer2-protocol-tunneling statistics on page 192
 show ethernet-switching layer2-protocol-tunneling statistics interface ge-0/0/0.0 on page 192
 show ethernet-switching layer2-protocol-tunneling statistics vlan v2 on page 192

Output Fields Table 16 on page 191 lists the output fields for the show ethernet-switching layer2-protocol-tunneling statistics command. Output fields are listed in the approximate order in which they appear.

Table 16: show ethernet-switching layer2-protocol-tunneling statistics Output Fields

VLAN	Field Description
VLAN	Name of a VLAN on which L2PT has been configured.

Table 16: show ethernet-switching layer2-protocol-tunneling statistics Output Fields (continued)

VLAN	Field Description
Interface	Name of an interface on which L2PT has been configured.
Protocol	Name of a protocol for which L2PT has been enabled. Values are all, cdp, gvrp, lldp, mvrp, stp, vstp, and vtp.
Operation	Type of operation being performed on the interface. Values are Encapsulation and Decapsulation.
Packets	Number of packets that have been encapsulated or decapsulated.
Drops	Number of packets that have exceeded the drop threshold and have been dropped.
Shutdowns	Number of times that packets have exceeded the shutdown threshold and the interface has been shut down.

```

show ethernet-switching layer2-protocol-tunneling statistics
user@switch> show ethernet-switching layer2-protocol-tunneling statistics
Layer2 Protocol Tunneling Statistics:

```

VLAN	Interface	Protocol	Operation	Packets	Drops	Shutdowns
v1	ge-0/0/0.0	mvrp	Encapsulation	0	0	0
v1	ge-0/0/1.0	mvrp	Decapsulation	0	0	0
v1	ge-0/0/2.0	mvrp	Decapsulation	60634	0	0
v2	ge-0/0/0.0	cdp	Encapsulation	0	0	0
v2	ge-0/0/0.0	gvrp	Encapsulation	0	0	0
v2	ge-0/0/0.0	lldp	Encapsulation	0	0	0

```

show ethernet-switching layer2-protocol-tunneling statistics interface ge-0/0/0.0
user@switch> show ethernet-switching layer2-protocol-tunneling statistics interface ge-0/0/0.0
Layer2 Protocol Tunneling Statistics:

```

VLAN	Interface	Protocol	Operation	Packets	Drops	Shutdowns
v1	ge-0/0/0.0	mvrp	Encapsulation	0	0	0
v2	ge-0/0/0.0	cdp	Encapsulation	0	0	0
v2	ge-0/0/0.0	gvrp	Encapsulation	0	0	0
v2	ge-0/0/0.0	lldp	Encapsulation	0	0	0
v2	ge-0/0/0.0	mvrp	Encapsulation	0	0	0
v2	ge-0/0/0.0	stp	Encapsulation	0	0	0
v2	ge-0/0/0.0	vtp	Encapsulation	0	0	0
v2	ge-0/0/0.0	vstp	Encapsulation	0	0	0

```

show ethernet-switching layer2-protocol-tunneling statistics vlan v2
user@switch> show ethernet-switching layer2-protocol-tunneling statistics vlan v2
Layer2 Protocol Tunneling Statistics:

```

VLAN	Interface	Protocol	Operation	Packets	Drops	Shutdowns
v2	ge-0/0/0.0	cdp	Encapsulation	0	0	0
v2	ge-0/0/0.0	gvrp	Encapsulation	0	0	0
v2	ge-0/0/0.0	lldp	Encapsulation	0	0	0
v2	ge-0/0/0.0	mvrp	Encapsulation	0	0	0
v2	ge-0/0/0.0	stp	Encapsulation	0	0	0
v2	ge-0/0/0.0	vtp	Encapsulation	0	0	0
v2	ge-0/0/0.0	vstp	Encapsulation	0	0	0
v2	ge-0/0/1.0	cdp	Decapsulation	0	0	0
v2	ge-0/0/1.0	gvrp	Decapsulation	0	0	0
v2	ge-0/0/1.0	lldp	Decapsulation	0	0	0
v2	ge-0/0/1.0	mvrp	Decapsulation	0	0	0

```
v2    ge-0/0/1.0  stp    Decapsulation  0      0      0
v2    ge-0/0/1.0  vtp    Decapsulation  0      0      0
```

show ethernet-switching layer2-protocol-tunneling vlan

Syntax	show ethernet-switching-layer2-protocol-tunneling vlan <vlan-name>
Release Information	Command introduced in JUNOS Release 10.0 for EX Series switches.
Description	Display information about Layer 2 protocol tunneling (L2PT) on VLANs that have been configured for L2PT.
Options	<p>none—Display information about L2PT for the VLANs on which you have configured L2PT.</p> <p>vlan-name—(Optional) Display information about L2PT for the specified VLAN.</p>
Required Privilege Level	view
Related Topics	<ul style="list-style-type: none"> ■ show ethernet-switching layer2-protocol-tunneling interface ■ show ethernet-switching layer2-protocol-tunneling statistics ■ show vlans ■ Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 88 ■ Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) on page 112
List of Sample Output	<p>show ethernet-switching layer2-protocol-tunneling vlan on page 195</p> <p>show ethernet-switching layer2-protocol-tunneling vlan v2 on page 195</p>
Output Fields	Table 17 on page 194 lists the output fields for the show ethernet-switching layer2-protocol-tunneling vlan command. Output fields are listed in the approximate order in which they appear.

Table 17: show ethernet-switching layer2-protocol-tunneling vlan Output Fields

Field Name	Field Description
VLAN	Name of the VLAN on which L2PT has been configured.
Protocol	Name of a protocol for which L2PT has been enabled. Values are all cdp, gvrp, lldp, mvrp, stp, vstp, and vtp.
Drop Threshold	Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the VLAN before the switch begins dropping the Layer 2 PDUs.
Shutdown Threshold	Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the VLAN before the interface is disabled.

show ethernet-switching layer2-protocol-tunneling vlan user@switch> **show ethernet-switching layer2-protocol-tunneling vlan**

```
Layer2 Protocol Tunneling VLAN information:
VLAN          Protocol      Drop          Shutdown
              Threshold    Threshold
v1            mvrp         100          200
v2            cdp          0            0
v2            cdp          0            0
v2            gvrp         0            0
```

show ethernet-switching layer2-protocol-tunneling vlan v2 user@switch> **show ethernet-switching layer2-protocol-tunneling vlan v2**

```
Layer2 Protocol Tunneling VLAN information:
VLAN          Protocol      Drop          Shutdown
              Threshold    Threshold
v2            cdp          0            0
v2            cdp          0            0
v2            gvrp         0            0
```

show ethernet-switching mac-learning-log

Syntax	show ethernet-switching mac-learning-log
Release Information	Command introduced in JUNOS Release 9.0 for EX Series switches.
Description	Displays the event log of learned MAC addresses.
Required Privilege Level	view
Related Topics	<ul style="list-style-type: none"> ■ show ethernet-switching table ■ show ethernet-switching interfaces ■ Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21 ■ Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 28 ■ Example: Configure Automatic VLAN Administration Using GVRP on page 46 ■ Example: Connecting an Access Switch to a Distribution Switch on page 36
List of Sample Output	show ethernet-switching mac-learning-log on page 196
Output Fields	Table 18 on page 196 lists the output fields for the show ethernet-switching mac-learning-log command. Output fields are listed in the approximate order in which they appear.

Table 18: show ethernet-switching mac-learning-log Output Fields

Field Name	Field Description
Date and Time	Timestamp when the MAC address was added or deleted from the log.
vlan_name	VLAN name. A value defined by the user for all user-configured VLANs.
MAC	Learned MAC address.
Deleted Added	MAC address deleted or added to the MAC learning log.
Blocking	The forwarding state of the interface: <ul style="list-style-type: none"> ■ blocked—Traffic is not being forwarded on the interface. ■ unblocked—Traffic is forwarded on the interface.

```

show ethernet-switching mac-learning-log
user@switch> show ethernet-switching mac-learning-log
Mon Feb 25 08:07:05 2008
  vlan_name v1 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
  vlan_name v9 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
  vlan_name HR_vlan mac 00:00:00:00:00:00 was deleted

```

```
Mon Feb 25 08:07:05 2008
  vlan_name v3 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
  vlan_name v12 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
  vlan_name v13 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
  vlan_name sales_vlan mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
  vlan_name employee1 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
  vlan_name employee2 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
  vlan_name v3 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
  vlan_name HR_vlan mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
  vlan_name employee2 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
  vlan_name employee1 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
  vlan_name employee2 mac 00:00:05:00:00:05 was learned
Mon Feb 25 08:07:05 2008
  vlan_name employee1 mac 00:30:48:90:54:89 was learned
Mon Feb 25 08:07:05 2008
  vlan_name HR_vlan mac 00:00:5e:00:01:00 was learned
Mon Feb 25 08:07:05 2008
  vlan_name sales_vlan mac 00:00:5e:00:01:08 was learned
[output truncated]
```

show ethernet-switching statistics aging

Syntax	show ethernet-switching statistics aging
Release Information	Command introduced in JUNOS Release 9.4 for EX Series switches.
Description	Display media access control (MAC) aging statistics.
Options	none—(Optional) Display MAC aging statistics. brief detail—(Optional) Display the specified level of output.
Required Privilege Level	view
Related Topics	<ul style="list-style-type: none"> ■ show ethernet-switching statistics mac-learning ■ Configuring MAC Table Aging (CLI Procedure) on page 101
List of Sample Output	show ethernet-switching statistics aging on page 198
Output Fields	Table 19 on page 198 lists the output fields for the show ethernet-switching statistics aging command. Output fields are listed in the approximate order in which they appear.

Table 19: show ethernet-switching statistics aging Output Fields

Field Name	Field Description	Level of Output
Total age messages received	Total number of aging messages received from the hardware.	All levels
Immediate aging	Aging message indicating that the entry should be removed immediately.	All levels
MAC address seen	Aging message indicating that the MAC address has been detected by hardware and that the aging timer should be stopped.	All levels
MAC address not seen	Aging message indicating that the MAC address has not been detected by the hardware and that the aging timer should be started.	All levels
Error age messages	The received aging message contains the following errors: <ul style="list-style-type: none"> ■ Invalid VLAN—The VLAN of the packet does not exist. ■ No such entry—The MAC address and VLAN pair provided by the aging message does not exist. ■ Static entry—An unsuccessful attempt was made to age out a static MAC entry. 	All levels

```

show ethernet-switching statistics aging user@switch> show ethernet-switching statistics aging
Total age messages received: 0
Immediate aging: 0, MAC address seen: 0, MAC address not seen: 0
  
```

```
Error age messages: 0  
  Invalid VLAN: 0, No such entry: 0, Static entry: 0
```

show ethernet-switching statistics mac-learning

Syntax	show ethernet-switching statistics mac-learning
Release Information	Command introduced in JUNOS Release 9.4 for EX Series switches.
Description	Display media access control (MAC) learning statistics.
Options	<p>none—(Optional) Display MAC learning statistics for all interfaces.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i> —(Optional) Display MAC learning statistics for the specified interface.</p>
Required Privilege Level	view
Related Topics	<ul style="list-style-type: none"> ■ show ethernet-switching statistics aging ■ show ethernet-switching mac-learning-log ■ show ethernet-switching table ■ show ethernet-switching interfaces ■ Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21 ■ Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 28 ■ Example: Configure Automatic VLAN Administration Using GVRP on page 46
List of Sample Output	<p>show ethernet-switching statistics mac-learning on page 201</p> <p>show ethernet-switching statistics mac-learning detail on page 201</p> <p>show ethernet-switching statistics mac-learning interface ge-0/0/1 on page 202</p>
Output Fields	Table 20 on page 200 lists the output fields for the <code>show ethernet-switching statistics mac-learning</code> command. Output fields are listed in the approximate order in which they appear.

Table 20: show ethernet-switching statistics mac-learning Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface for which statistics are being reported.	All levels
Learning message from local packets	MAC learning message generated due to packets coming in on the management interface.	All levels
Learning message from transit packets	MAC learning message generated due to packets coming in on network interfaces.	All levels

Table 20: show ethernet-switching statistics mac-learning Output Fields (continued)

Field Name	Field Description	Level of Output
Learning message with error	<p>MAC learning messages received with errors:</p> <ul style="list-style-type: none"> ■ Invalid VLAN—The VLAN of the packet does not exist. ■ Invalid MAC—The MAC address is either NULL or a multicast MAC address. ■ Security violation—The MAC address is not an allowed MAC address. ■ Interface down—The MAC address is learned on an interface that is down. ■ Incorrect membership—The MAC address is learned on an interface that is not a member of the VLAN. ■ Interface limit—The number of MAC addresses learned on the interface has exceeded the limit. ■ MAC move limit—This MAC address has moved among multiple interfaces too many times in a given interval. ■ VLAN limit—The number of MAC addresses learned on the VLAN has exceeded the limit. ■ Invalid VLAN index—The VLAN of the packet, while configured, does not yet exist in the kernel. ■ Interface not learning—The MAC address is learned on an interface that does not yet allow learning—for example, the interface is blocked. ■ No nexthop—The MAC address is learned on an interface that does not have a unicast next hop. ■ MAC learning disabled—The MAC address is learned on an interface on which MAC learning has been disabled. ■ Others—The message contains some other error. 	All levels

```

show ethernet-switching statistics mac-learning
user@switch> show ethernet-switching statistics mac-learning
Learning stats: 0 learn msg rcvd, 0 error
Interface      Local pkts      Transit pkts      Error
ge-0/0/0.0     0                0                 0
ge-0/0/1.0     0                0                 0
ge-0/0/2.0     0                0                 0
ge-0/0/3.0     0                0                 0

show ethernet-switching statistics mac-learning detail
user@switch> show ethernet-switching statistics mac-learning detail
Learning stats: 0 learn msg rcvd, 0 error

Interface: ge-0/0/0.0
Learning message from local packets: 0
Learning message from transit packets: 1
Learning message with error: 0
  Invalid VLAN: 0      Invalid MAC: 0
  Security violation: 0  Interface down: 0
  Incorrect membership: 0 Interface limit: 0
  MAC move limit: 0    VLAN limit: 0
  Invalid VLAN index: 0 Interface not learning: 0
  No nexthop: 0       MAC learning disabled: 0
  Others: 0

Interface: ge-0/0/1.0
Learning message from local packets: 0
Learning message from transit packets: 2

```

```

Learning message with error:      0
  Invalid VLAN:                   0   Invalid MAC:                   0
  Security violation:             0   Interface down:               0
  Incorrect membership:          0   Interface limit:              0
  MAC move limit:                0   VLAN limit:                   0
  Invalid VLAN index:            0   Interface not learning:       0
  No nexthop:                    0   MAC learning disabled:        0
  Others:                          0
    
```

```

show ethernet-switching user@switch> show ethernet-switching statistics mac-learning interface ge-0/0/1
statistics mac-learning Interface      Local pkts      Transit pkts      Error
interface ge-0/0/1      ge-0/0/1.0      0                1                1
    
```

show ethernet-switching table

Syntax	show ethernet-switching table <brief detail extensive summary> <interface <i>interface-name</i> > <management-vlan> <vlan (<i>vlan-name</i>)>
Release Information	Command introduced in JUNOS Release 9.0 for EX Series switches. Options <i>summary</i> , <i>management-vlan</i> , and <i>vlan <i>vlan-name</i></i> introduced in JUNOS Release 9.6 for EX Series switches.
Description	Displays the Ethernet switching table.
Options	<p><i>none</i>—(Optional) Display brief information about the Ethernet switching table.</p> <p><i>brief detail extensive summary</i>—(Optional) Display the specified level of output.</p> <p><i>management-vlan</i>—(Optional) Display the Ethernet switching table for a management VLAN.</p> <p><i>interface-name</i>—(Optional) Display the Ethernet switching table for a specific interface.</p> <p><i>vlan <i>vlan-name</i></i>—(Optional) Display the Ethernet switching table for a specific VLAN.</p>
Required Privilege Level	view
Related Topics	<ul style="list-style-type: none"> ■ Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21 ■ Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 28 ■ Example: Configure Automatic VLAN Administration Using GVRP on page 46 ■ Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 65
List of Sample Output	<p>show ethernet-switching table on page 204</p> <p>show ethernet-switching table brief on page 204</p> <p>show ethernet-switching table detail on page 205</p> <p>show ethernet-switching table extensive on page 206</p> <p>show ethernet-switching table interface ge-0/0/1 on page 208</p>
Output Fields	Table 21 on page 203 lists the output fields for the <code>show ethernet-switching table</code> command. Output fields are listed in the approximate order in which they appear.

Table 21: show ethernet-switching table Output Fields

Field Name	Field Description	Level of Output
VLAN	The name of a VLAN.	All levels
MAC address	The MAC address associated with the VLAN.	All levels

Table 21: show ethernet-switching table Output Fields (continued)

Field Name	Field Description	Level of Output
Type	The type of MAC address. Values are: <ul style="list-style-type: none"> ■ static—The MAC address is manually created. ■ learn—The MAC address is learned dynamically from a packet's source MAC address. ■ flood—The MAC address is unknown and flooded to all members. 	All levels
Age	The time remaining before the entry ages out and is removed from the Ethernet switching table.	All levels
Interfaces	Interface associated with learned MAC addresses or All-members (flood entry).	All levels
Learned	For learned entries, the time which the entry was added to the Ethernet-switching table.	detail, extensive

```

show ethernet-switching table user@switch> show ethernet-switching table
Ethernet-switching table: 57 entries, 17 learned
VLAN      MAC address      Type      Age Interfaces
F2         *                Flood     - All-members
F2         00:00:05:00:00:03 Learn     0 ge-0/0/44.0
F2         00:19:e2:50:7d:e0 Static    - Router
Linux      *                Flood     - All-members
Linux      00:19:e2:50:7d:e0 Static    - Router
Linux      00:30:48:90:54:89 Learn     0 ge-0/0/47.0
T1         *                Flood     - All-members
T1         00:00:05:00:00:01 Learn     0 ge-0/0/46.0
T1         00:00:5e:00:01:00 Static    - Router
T1         00:19:e2:50:63:e0 Learn     0 ge-0/0/46.0
T1         00:19:e2:50:7d:e0 Static    - Router
T10        *                Flood     - All-members
T10        00:00:5e:00:01:09 Static    - Router
T10        00:19:e2:50:63:e0 Learn     0 ge-0/0/46.0
T10        00:19:e2:50:7d:e0 Static    - Router
T111       *                Flood     - All-members
T111       00:19:e2:50:63:e0 Learn     0 ge-0/0/15.0
T111       00:19:e2:50:7d:e0 Static    - Router
T111       00:19:e2:50:ac:00 Learn     0 ge-0/0/15.0
T2         *                Flood     - All-members
T2         00:00:5e:00:01:01 Static    - Router
T2         00:19:e2:50:63:e0 Learn     0 ge-0/0/46.0
T2         00:19:e2:50:7d:e0 Static    - Router
T3         *                Flood     - All-members
T3         00:00:5e:00:01:02 Static    - Router
T3         00:19:e2:50:63:e0 Learn     0 ge-0/0/46.0
T3         00:19:e2:50:7d:e0 Static    - Router
T4         *                Flood     - All-members
T4         00:00:5e:00:01:03 Static    - Router
T4         00:19:e2:50:63:e0 Learn     0 ge-0/0/46.0
[output truncated]
    
```

```

show ethernet-switching table brief user@switch> show ethernet-switching table brief
Ethernet-switching table: 57 entries, 17 learned
VLAN      MAC address      Type      Age Interfaces
F2         *                Flood     - All-members
    
```

```

F2          00:00:05:00:00:03 Learn      0 ge-0/0/44.0
F2          00:19:e2:50:7d:e0 Static    - Router
Linux       *                   Flood    - All-members
Linux       00:19:e2:50:7d:e0 Static    - Router
Linux       00:30:48:90:54:89 Learn    0 ge-0/0/47.0
T1          *                   Flood    - All-members
T1          00:00:05:00:00:01 Learn    0 ge-0/0/46.0
T1          00:00:5e:00:01:00 Static    - Router
T1          00:19:e2:50:63:e0 Learn    0 ge-0/0/46.0
T1          00:19:e2:50:7d:e0 Static    - Router
T10         *                   Flood    - All-members
T10         00:00:5e:00:01:09 Static    - Router
T10         00:19:e2:50:63:e0 Learn    0 ge-0/0/46.0
T10         00:19:e2:50:7d:e0 Static    - Router
T111        *                   Flood    - All-members
T111        00:19:e2:50:63:e0 Learn    0 ge-0/0/15.0
T111        00:19:e2:50:7d:e0 Static    - Router
T111        00:19:e2:50:ac:00 Learn    0 ge-0/0/15.0
T2          *                   Flood    - All-members
T2          00:00:5e:00:01:01 Static    - Router
T2          00:19:e2:50:63:e0 Learn    0 ge-0/0/46.0
T2          00:19:e2:50:7d:e0 Static    - Router
T3          *                   Flood    - All-members
T3          00:00:5e:00:01:02 Static    - Router
T3          00:19:e2:50:63:e0 Learn    0 ge-0/0/46.0
T3          00:19:e2:50:7d:e0 Static    - Router
T4          *                   Flood    - All-members
T4          00:00:5e:00:01:03 Static    - Router
T4          00:19:e2:50:63:e0 Learn    0 ge-0/0/46.0

```

[output truncated]

**show ethernet-switching
table detail**

```

user@switch> show ethernet-switching table detail
Ethernet-switching table: 57 entries, 17 learned

```

```

F2, *
  Interface(s): ge-0/0/44.0
  Type: Flood
  Nexthop index: 0

F2, 00:00:05:00:00:03
  Interface(s): ge-0/0/44.0
  Type: Learn, Age: 0, Learned: 2:03:09
  Nexthop index: 0

F2, 00:19:e2:50:7d:e0
  Interface(s): Router
  Type: Static
  Nexthop index: 0

Linux, *
  Interface(s): ge-0/0/47.0
  Type: Flood
  Nexthop index: 0

Linux, 00:19:e2:50:7d:e0
  Interface(s): Router
  Type: Static
  Nexthop index: 0

Linux, 00:30:48:90:54:89
  Interface(s): ge-0/0/47.0
  Type: Learn, Age: 0, Learned: 2:03:08

```

```

    Nexthop index: 0

T1, *
  Interface(s): ge-0/0/46.0
  Type: Flood
  Nexthop index: 0

T1, 00:00:05:00:00:01
  Interface(s): ge-0/0/46.0
  Type: Learn, Age: 0, Learned: 2:03:07
  Nexthop index: 0

T1, 00:00:5e:00:01:00
  Interface(s): Router
  Type: Static
  Nexthop index: 0

T1, 00:19:e2:50:63:e0
  Interface(s): ge-0/0/46.0
  Type: Learn, Age: 0, Learned: 2:03:07
  Nexthop index: 0

T1, 00:19:e2:50:7d:e0
  Interface(s): Router
  Type: Static
  Nexthop index: 0

T10, *
  Interface(s): ge-0/0/46.0
  Type: Flood
  Nexthop index: 0

T10, 00:00:5e:00:01:09
  Interface(s): Router
  Type: Static
  Nexthop index: 0

T10, 00:19:e2:50:63:e0
  Interface(s): ge-0/0/46.0
  Type: Learn, Age: 0, Learned: 2:03:08
  Nexthop index: 0

T10, 00:19:e2:50:7d:e0
  Interface(s): Router
  Type: Static
  Nexthop index: 0

T111, *
  Interface(s): ge-0/0/15.0
  Type: Flood
  Nexthop index: 0
[output truncated]

```

show ethernet-switching table extensive user@switch> **show ethernet-switching table extensive**
 Ethernet-switching table: 57 entries, 17 learned

```

F2, *
  Interface(s): ge-0/0/44.0
  Type: Flood
  Nexthop index: 0

F2, 00:00:05:00:00:03

```

```

Interface(s): ge-0/0/44.0
Type: Learn, Age: 0, Learned: 2:03:09
Nexthop index: 0

```

```

F2, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

```

```

Linux, *
Interface(s): ge-0/0/47.0
Type: Flood
Nexthop index: 0

```

```

Linux, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

```

```

Linux, 00:30:48:90:54:89
Interface(s): ge-0/0/47.0
Type: Learn, Age: 0, Learned: 2:03:08
Nexthop index: 0

```

```

T1, *
Interface(s): ge-0/0/46.0
Type: Flood
Nexthop index: 0

```

```

T1, 00:00:05:00:00:01
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
Nexthop index: 0

```

```

T1, 00:00:5e:00:01:00
Interface(s): Router
Type: Static
Nexthop index: 0

```

```

T1, 00:19:e2:50:63:e0
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
Nexthop index: 0

```

```

T1, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

```

```

T10, *
Interface(s): ge-0/0/46.0
Type: Flood
Nexthop index: 0

```

```

T10, 00:00:5e:00:01:09
Interface(s): Router
Type: Static
Nexthop index: 0

```

```

T10, 00:19:e2:50:63:e0
Interface(s): ge-0/0/46.0

```

```

Type: Learn, Age: 0, Learned: 2:03:08
Nexthop index: 0

T10, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

T111, *
Interface(s): ge-0/0/15.0
Type: Flood
Nexthop index: 0
[output truncated]

```

**show ethernet-switching
table interface ge-0/0/1**

```

user@switch> show ethernet-switching table interface ge-0/0/1
Ethernet-switching table: 1 unicast entries
VLAN      MAC address      Type      Age Interfaces
V1        *                Flood     - All-members
V1        00:00:05:00:00:05 Learn     0 ge-0/0/1.0

```

show gvrp

Syntax	show gvrp
Release Information	Command introduced in JUNOS Release 9.0 for EX Series switches.
Description	Display GARP VLAN Registration Protocol (GVRP) information.
Options	none—Displays all GVRP configuration attributes. interface <i>interface-name</i> —(Optional) Displays GVRP statistics for a specific interface only.
Required Privilege Level	view
Related Topics	<ul style="list-style-type: none"> ■ show gvrp statistics ■ Example: Configure Automatic VLAN Administration Using GVRP on page 46
List of Sample Output	show gvrp on page 209
Output Fields	Table 22 on page 209 lists the output fields for the show gvrp command. Output fields are listed in the approximate order in which they appear.

Table 22: show gvrp Output Fields

Field Name	Field Description
Global GVRP Configuration	Displays global GVRP information: <ul style="list-style-type: none"> ■ GVRP status—Displays whether GVRP is enabled or disabled. ■ Join—The maximum number of milliseconds the interfaces must wait before sending VLAN advertisements. ■ Leave— The number of milliseconds an interface must wait after receiving a Leave message to remove the interface from the VLAN specified in the message. ■ Leaveall—The interval at which Leave All messages are sent on interfaces. Leave all messages maintain current GVRP VLAN membership information in the network.
Interface based configuration	Displays interface-specific GVRP information: <ul style="list-style-type: none"> ■ Interface—The interface on which GVRP is configured.. ■ GVRP status—Displays whether GVRP is enabled or disabled.

```

show gvrp user@switch> show gvrp

Global GVRP configuration
GVRP status      : Enabled
GVRP timers (ms)
  Join           : 40
  Leave          : 120
  Leaveall       : 2000
  
```

```
Interface based configuration:
Interface  GVRP status
-----  -----
ge-0/0/0.0 Enabled
```

show gvrp statistics

Syntax	show gvrp statistics
Release Information	Command introduced in JUNOS Release 9.0 for EX Series switches.
Description	Display Generic VLAN Registration Protocol (GVRP) statistics in the form of GARP Information Propagation (GIP) messages.
Required Privilege Level	clear
Related Topics	<ul style="list-style-type: none"> ■ show gvrp ■ Example: Configure Automatic VLAN Administration Using GVRP on page 46
List of Sample Output	show gvrp statistics on page 212
Output Fields	Table 23 on page 211 lists the output fields for the show gvrp statistics command. Output fields are listed in the approximate order in which they appear.

Table 23: show gvrp statistics Output Fields

Field Name	Field Description
Join Empty received	Number of GIP Join Empty messages received on the switch.
Join In received	Number of GIP Join In messages received on the switch.
Empty received	Number of GIP Empty messages received on the switch.
Leave In received	Number of GIP Leave In messages received on the switch.
Leave Empty received	Number of GIP Leave Empty messages received on the switch.
Leave All received	Number of GIP Leave All messages received on the switch.
Join Empty transmitted	Number of GIP Join Empty messages sent from the switch.
Join In transmitted	Number of GIP Join In messages sent from the switch.
Empty transmitted	Number of GIP Empty messages sent from the switch.
Leave In transmitted	Number of GIP Leave In messages sent from the switch.
Leave Empty transmitted	Number of GIP Leave Empty messages sent from the switch.
Leave All transmitted	Number of GIP Leave All messages sent from the switch.

```
show gvrp statistics user@switch> show gvrp statistics  
GVRP statistics  
Join Empty received      : 0  
Join In received        : 12  
Empty received          : 0  
Leave In received        : 0  
Leave Empty received     : 0  
Leave All received       : 0  
Join Empty transmitted  : 0  
Join In transmitted     : 48  
Empty transmitted       : 4  
Leave In transmitted     : 0  
Leave Empty transmitted  : 0  
Leave All transmitted    : 4
```

show mvrp

Syntax	show mvrp
Release Information	Command introduced in JUNOS Release 10.0 for EX Series switches.
Description	Display Multiple VLAN Registration Protocol (MVRP) configuration information.
Required Privilege Level	view
Related Topics	<ul style="list-style-type: none"> ■ show mvrp statistics ■ Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 76 ■ Verifying That MVRP Is Working Correctly on page 122
List of Sample Output	show mvrp on page 214
Output Fields	Table 24 on page 213 lists the output fields for the show mvrp command. Output fields are listed in the approximate order in which they appear.

Table 24: show mvrp Output Fields

Field Name	Field Description
Global MVRP configuration	Displays global MVRP information: <ul style="list-style-type: none"> ■ MVRP status—Displays whether MVRP is Enabled or Disabled. ■ MVRP dynamic vlan creation—Displays whether global MVRP dynamic VLAN creation is Dnabled or Disabled.
MVRP Timers (ms)	Displays MVRP timer information: <ul style="list-style-type: none"> ■ Interface—The interface on which MVRP is configured. ■ Join—The maximum number of milliseconds the interfaces must wait before sending VLAN advertisements. ■ Leave—The number of milliseconds an interface must wait after receiving a Leave message to remove the interface from the VLAN specified in the message. ■ LeaveAll—The interval at which LeaveAll messages are sent on interfaces. LeaveAll messages maintain current MVRP VLAN membership information in the network.
Interface based configuration	Displays interface-specific MVRP information: <ul style="list-style-type: none"> ■ Interface—The interface on which MVRP is configured. ■ Status—Displays whether MVRP is Enabled or Disabled. ■ Registration—Displays whether registration for the interface is Forbidden or Normal. ■ Dynamic VLAN Creation—Displays whether interface dynamic VLAN creation is Enabled or Disabled.

show mvrp user@switch> **show mvrp**

Global MVRP configuration

MVRP status : Enabled

MVRP dynamic vlan creation: Enabled

MVRP Timers (ms):

Interface	Join	Leave	LeaveAll
all	200	600	10000
xe-0/1/1.0	200	600	10000

Interface based configuration:

Interface	Status	Registration	Dynamic VLAN Creation
all	Disabled	Normal	Enabled
xe-0/1/1.0	Enabled	Normal	Enabled

show mvrp dynamic-vlan-memberships

Syntax	show mvrp dynamic-vlan-memberships
Release Information	Command introduced in JUNOS Release 10.0 for EX Series switches.
Description	Display all VLANs that have been created dynamically using Multiple VLAN Registration Protocol (MVRP) on the switch.
Required Privilege Level	clear
Related Topics	<ul style="list-style-type: none"> ■ show mvrp ■ show mvrp statistics ■ Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 76 ■ Verifying That MVRP Is Working Correctly on page 122
List of Sample Output	show mvrp dynamic-vlan-memberships on page 215
Output Fields	Table 25 on page 215 lists the output fields for the show mvrp dynamic-vlan-memberships command. Output fields are listed in the approximate order in which they appear.

Table 25: show mvrp dynamic-vlan-memberships Output Fields

Field Name	Field Description
VLAN Name	The name of the dynamically created VLAN.
Interfaces	The interface or interfaces that are bound to the dynamically created VLAN.

```

show mvrp      user@switch> show mvrp dynamic-vlan-memberships
dynamic-vlan-memberships
VLAN Name          Interfaces
-----
__mvrp_100__       xe-0/1/1.0
                   xe-0/1/0.0
__mvrp_200__       xe-0/1/1.0
                   xe-0/1/0.0
__mvrp_300__       xe-0/1/1.0

```

show mvrp statistics

Syntax	show mvrp statistics <interface <i>interface-name</i> >
Release Information	Command introduced in JUNOS Release 10.0 for EX Series switches.
Description	Display Multiple VLAN Registration Protocol (MVRP) statistics in the form of Multiple Registration Protocol data unit (MRPDU) messages.
Options	<p>none—Show MVRP statistics for all interfaces on the switch.</p> <p>interface <i>interface-name</i>—Show MVRP statistics for the specified interface.</p>
Required Privilege Level	clear
Related Topics	<ul style="list-style-type: none"> ■ show mvrp ■ clear mvrp statistics ■ Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 76 ■ Verifying That MVRP Is Working Correctly on page 122
List of Sample Output	show mvrp statistics interface xe-0/1/1.0 on page 217
Output Fields	Table 25 on page 215 lists the output fields for the show mvrp statistics command. Output fields are listed in the approximate order in which they appear.

Table 26: show mvrp statistics Output Fields

Field Name	Field Description
MRPDU received	Number of MRPDU messages received on the switch.
Invalid PDU received	Number of invalid MRPDU messages received on the switch.
New received	Number of new messages received on the switch.
Join Empty received	Number of MRP Join Empty messages received on the switch.
Join In received	Number of MRP Join In messages received on the switch.
Empty received	Number of MRP Empty messages received on the switch.
In received	Number of MRP In messages received on the switch.
Leave received	Number of MRP Leave messages received on the switch.
LeaveAll received	Number of LeaveAll messages received on the switch.
MRPDU transmitted	Number of MRPDU messages transmitted from the switch.

Table 26: show mvrp statistics Output Fields (continued)

Field Name	Field Description
MRPDU transmit failures	Number of MRPDU transmit failures from the switch.
New transmitted	Number of new messages transmitted from the switch.
Join Empty transmitted	Number of Join Empty messages sent from the switch.
Join In transmitted	Number of MRP Join In messages sent from the switch.
Empty transmitted	Number of MRP Empty messages sent from the switch.
In transmitted	Number of MRP In messages sent from the switch.
Leave transmitted	Number of MRP Leave Empty messages sent from the switch.
LeaveAll transmitted	Number of MRP LeaveAll messages sent from the switch.

```

show mvrp statistics user@switch> show mvrp statistics interface xe-0/1/1.0
interface xe-0/1/1.0 MVRP statistics
MRPDU received           : 3342
Invalid PDU received    : 0
New received             : 2
Join Empty received     : 1116
Join In received        : 2219
Empty received          : 2
In received              : 2
Leave received           : 1
LeaveAll received        : 1117
MRPDU transmitted       : 3280
MRPDU transmit failures : 0
New transmitted         : 0
Join Empty transmitted  : 1114
Join In transmitted     : 2163
Empty transmitted       : 1
In transmitted          : 1
Leave transmitted        : 1
LeaveAll transmitted     : 1111

```

show redundant-trunk-group

Syntax	show redundant-trunk-group <group-name <i>group-name</i> >
Release Information	Command introduced in JUNOS Release 9.0 for EX Series switches.
Description	Display information about redundant trunk groups.
Options	group-name <i>group-name</i> —Display information about the specified redundant trunk group.
Required Privilege Level	view
Related Topics	<ul style="list-style-type: none"> ■ Example: Configuring Redundant Trunk Links for Faster Recovery on page 61 ■ Understanding Redundant Trunk Links on EX Series Switches on page 11
List of Sample Output	show redundant-trunk-group group-name Group1 on page 218
Output Fields	Table 27 on page 218 lists the output fields for the show redundant-trunk-group command. Output fields are listed in the approximate order in which they appear.

Table 27: show redundant-trunk-group Output Fields

Field Name	Field Description
Group Name	Name of the redundant trunk port group.
Interface	Name of an interface belonging to the trunk port group. <ul style="list-style-type: none"> ■ (P) denotes a primary interface. ■ (A) denotes an active interface. ■ Lack of (A) denotes a blocking interface.
State	Operating state of the interface: UP or DOWN.
Last Time of Flap	Date and time at which the advertised link became unavailable, and then, available again.
# Flaps	Total number of flaps since the last switch reboot.

```

show user@switch> show redundant-trunk-group group-name Group1
redundant-trunk-group show redundant-trunk-group group-name Group1
group-name Group1
Group Name Interface State Last Time of Flap # Flaps
Group1 ge-0/0/45.0 (P) UP Fri Jan 2 04:10:58 0
ge-0/0/47.0 UP Fri Jan 2 04:10:58 0

```

show vlans

Syntax show vlans
 <brief | detail | extensive>
 <dot1q-tunneling>
 <management-vlan>
 <sort-by (name | tag)>
 <summary>
 <vlan-name>
 <vlan-range-name>

Release Information Command introduced in JUNOS Release 9.0 for EX Series switches.
 Command modified in JUNOS Release 9.2 for EX Series switches to display support for MAC-based VLANs and **sort-by (name | tag)** and **vlan-range-name** options.
 Command modified in JUNOS Release 9.4 for EX Series switches to display whether MAC learning is disabled and to include the **management-vlan** option.
 Command modified in JUNOS Release 9.5 for EX Series switches to include the **summary** option.
 Command modified in JUNOS Release 10.0 for EX Series switches to display output for the MAC aging timer and MVRP dynamic VLANs.

Description Display information about VLANs configured on bridged Ethernet interfaces. For interfaces configured to support a VoIP VLAN and a data VLAN, the **show vlans** command displays both tagged and untagged membership for those VLANs.



NOTE: When a series of VLANs is created with the **vlan-range** statement, such VLAN names are prefixed and suffixed with a double underscore. For example, a series of VLANs using the VLAN range 1–3 and the base VLAN name **marketing** are displayed as **__marketing_1__**, **__marketing_2__**, and **__marketing_3__**.



NOTE: To display an 802.1X supplicant successfully authenticated in multiple-supplicant mode with dynamic VLAN movement, use the **show vlans vlan-name extensive** operational mode command, where **vlan-name** is the dynamic VLAN.

Options none—Display information for all VLANs. VLAN information is displayed by VLAN name in ascending order.

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

dot1q-tunneling—(Optional) Display VLANs with the Q-in-Q tunneling feature enabled.

management-vlan—(Optional) Display management VLANs.

sort-by (name | tag)—(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.

summary—(Optional) Display the total number of VLANs and counts of VLANs by type—for example, the number of dynamic, 802.1Q, and Q-in-Q tunneled VLANs.

vlan-name—(Optional) Display information for the specified VLAN.

vlan-range-name—(Optional) Display information for the specified VLAN range. To see information for all members of the VLAN range, specify the base VLAN name—for example, **employee** for a VLAN range that includes **__employee_1__** through **__employee_10__**.

Required Privilege Level view

- Related Topics**
- show ethernet-switching interfaces
 - Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 21
 - Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 28
 - Example: Configure Automatic VLAN Administration Using GVRP on page 46
 - Example: Configuring a Private VLAN on an EX Series Switch on page 68
 - Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 65
 - Understanding Bridging and VLANs on EX Series Switches on page 3

- List of Sample Output**
- show vlans on page 222
 - show vlans brief on page 223
 - show vlans detail on page 223
 - show vlans extensive (MAC-based) on page 224
 - show vlans extensive (Port-based) on page 224
 - show vlans sort-by tag on page 225
 - show vlans sort-by name on page 226
 - show vlans employee (vlan-range-name) on page 226
 - show vlans summary on page 227

Output Fields Table 28 on page 220 lists the output fields for the **show vlans** command. Output fields are listed in the approximate order in which they appear.

Table 28: show vlans Output Fields

Field Name	Field Description	Level of Output
Name	Name of a VLAN.	none, brief
Tag	The 802.1Q tag applied to this VLAN. If none is displayed, no tag is applied.	All levels
Interfaces	Interface associated with learned MAC addresses or all-members (flood entry). An asterisk (*) beside the interface indicates that the interface is UP.	All levels
Address	The IP address.	none, brief
Ports Active / Total	The number of interfaces associated with a VLAN. The Active column indicates interfaces that are UP, and the Total column indicates interfaces that are active and inactive.	brief

Table 28: show vlans Output Fields (continued)

Field Name	Field Description	Level of Output
VLAN	Name of a VLAN.	detail, extensive
Admin state	Indicates whether the physical link is operational and can pass packets.	detail, extensive
Dot1q Tunneling Status	Indicates whether Q-in-Q tunneling is enabled.	detail, extensive
MAC learning Status	Indicates whether MAC learning is disabled.	detail, extensive
Description	A description for the VLAN.	detail,extensive
Primary IP	Primary IP address associated with a VLAN.	detail
Number of interfaces	The number of interfaces associated with a VLAN. Both the total number of interfaces and the number of active interfaces associated with a VLAN are displayed.	detail, extensive
STP	The spanning tree associated with a VLAN.	detail, extensive
RTG	The redundant trunk group associated with a VLAN.	detail, extensive
Tagged interfaces	The tagged interfaces to which a VLAN is associated.	detail, extensive
Untagged interfaces	The untagged interfaces to which a VLAN is associated.	detail, extensive
Customer VLAN Ranges	Lists the customer VLAN (C-VLAN) ranges associated with this service VLAN (S-VLAN).	extensive
Private VLAN Mode	The private VLAN mode for this VLAN. Values are Primary, Isolated, and Community.	extensive
Primary VLAN	The primary VLAN tag for this secondary VLAN.	extensive
Internal Index	VLAN index internal to JUNOS Software.	extensive
Origin	The manner in which the VLAN was created. Values are static and learn.	extensive
Protocol	Port-based VLAN or MAC-based VLAN. MAC-based protocol is displayed when VLAN assignment is done either statically or dynamically through 802.1X.	extensive
Mac aging time	The MAC aging timer.	extensive
IP addresses	IP address associated with a VLAN.	extensive
Number of MAC entries	For MAC-based VLANs created either statically or dynamically, the MAC addresses associated with an interface.	extensive
Secondary VLANs	The secondary VLANs associated with a primary VLAN.	extensive
Isolated VLANs	The isolated VLANs associated with a primary VLAN.	extensive
Community VLANs	The community VLANs associated with a primary VLAN.	extensive

Table 28: show vlans Output Fields (continued)

Field Name	Field Description	Level of Output
VLANs summary	<p>VLAN counts:</p> <ul style="list-style-type: none"> ■ Total—Total number of VLANs on the switch. ■ Configured VLANs—Number of VLANs that are based on user-configured settings. ■ Internal VLANs—Number of VLANs created by the system with no explicit configuration or protocol—for example, the default VLAN and the VLAN created when a trunk interface is not configured with native VLAN membership. ■ Temporary VLANs—Number of VLANs from the previous configuration that the system retains for a limited time after restart. Temporary VLANs are converted into one of the other types of VLAN, or are removed from the system if the current configuration does not require them. 	All levels
Dot1q VLANs summary	<p>802.1Q VLAN counts:</p> <ul style="list-style-type: none"> ■ Total—Total number of 802.1Q VLANs on the switch. ■ Tagged VLANs—Number of tagged 802.1Q VLANs. ■ Untagged VLANs—Number of untagged 802.1Q VLANs. ■ Private VLAN—Counts of the following kinds of 802.1Q private VLANs (PVLANS): <ul style="list-style-type: none"> ■ Primary VLANs—Number of primary forwarding private VLANs. ■ Community VLANs—Number of secondary transporting and forwarding private VLANs. ■ Isolated VLANs—Number of secondary receiving and forwarding private VLANs. 	All levels
Dot1q Tunneled VLANs summary	<p>Q-in-Q VLAN counts:</p> <ul style="list-style-type: none"> ■ Total—Total number of Q-in-Q VLANs on the switch. ■ Private VLAN—Counts of primary, community, and isolated Q-in-Q private VLANs (PVLANS). 	All levels
Dynamic VLANs	<p>Counts of VLANs assigned or created dynamically by a protocol:</p> <ul style="list-style-type: none"> ■ Total—Total number of dynamic VLANs on the switch. ■ Dot1x—Number of 802.1X VLANs authenticated and assigned when the switch learns the MAC address of a supplicant host from a packet's source MAC address. ■ MVRP—Number of VLANs created by the Multiple VLAN Registration Protocol (MVRP). 	All levels

```

show vlans user@switch> show vlans
Name      Tag      Interfaces
default  None
          ge-0/0/34.0, ge-0/0/33.0, ge-0/0/32.0, ge-0/0/31.0,
          ge-0/0/30.0, ge-0/0/29.0, ge-0/0/28.0, ge-0/0/27.0,
          ge-0/0/26.0, ge-0/0/25.0, ge-0/0/19.0, ge-0/0/18.0,
          ge-0/0/17.0, ge-0/0/16.0, ge-0/0/15.0, ge-0/0/14.0,
    
```

```

                                ge-0/0/13.0, ge-0/0/11.0, ge-0/0/9.0, ge-0/0/8.0,
                                ge-0/0/3.0, ge-0/0/2.0, ge-0/0/1.0
v0001          1
v0002          2          ge-0/0/24.0, ge-0/0/23.0, ge-0/0/22.0, ge-0/0/21.0
v0003          3          None
v0004          4          None
v0005          5          None

```

show vlans brief user@switch> **show vlans brief**

Name	Tag	Address	Ports Active/Total
default	None		0/23
v0001	1		0/4
v0002	2		0/0
v0003	3		0/0
v0004	4		0/0
v0005	5		0/0
v0006	6		0/0
v0007	7		0/0
v0008	8		0/0
v0009	9		0/0
v0010	10		0/2
v0011	11		0/0
v0012	12		0/0
v0013	13		0/0
v0014	14		0/0
v0015	15		0/0
v0016	16		0/0

show vlans detail user@switch> **show vlans detail**

```

VLAN: default, Tag: Untagged, Admin state: Enabled
  Description: None
  Primary IP: None, Number of interfaces: 23 (Active = 0)
  STP: None, RTG: None
  Untagged interfaces: ge-0/0/34.0, ge-0/0/33.0, ge-0/0/32.0, ge-0/0/31.0,
  ge-0/0/30.0, ge-0/0/29.0, ge-0/0/28.0, ge-0/0/27.0, ge-0/0/26.0,
  ge-0/0/25.0, ge-0/0/19.0, ge-0/0/18.0, ge-0/0/17.0, ge-0/0/16.0,
  ge-0/0/15.0, ge-0/0/14.0, ge-0/0/13.0, ge-0/0/11.0, ge-0/0/9.0, ge-0/0/8.0,
  ge-0/0/3.0, ge-0/0/2.0, ge-0/0/1.0,
  Tagged interfaces: None

VLAN: v0001, Tag: 802.1Q Tag 1, Admin state: Enabled
  Description: None
  Primary IP: None, Number of interfaces: 4 (Active = 0)
  Dot1q Tunneling Status: Enabled
  STP: None, RTG: None
  Untagged interfaces: None
  Tagged interfaces: ge-0/0/24.0, ge-0/0/23.0, ge-0/0/22.0, ge-0/0/21.0,

VLAN: v0002, Tag: 802.1Q Tag 2, Admin state: Enabled
  Description: None
  Primary IP: None, Number of interfaces: 0 (Active = 0)
  STP: None, RTG: None
  Untagged interfaces: None
  Tagged interfaces: None

```

```
VLAN: v0003, Tag: 802.1Q Tag 3, Admin state: Enabled
Description: None
Primary IP: None, Number of interfaces: 0 (Active = 0)
STP: None, RTG: None
Untagged interfaces: None
Tagged interfaces: None
```

```
VLAN: vlan4000, 802.1Q Tag: Untagged, Admin State: Enabled
MAC learning Status: Disabled
Number of interfaces: 0 (Active = 0)
```

**show vlans extensive
(MAC-based)**

```
user@switch> show vlans extensive
VLAN: default, Created at: Thu May 15 13:43:09 2008
Internal index: 3, Admin State: Enabled, Origin: Static
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 0 (Active = 0), Untagged 2 (Active = 2)
    ge-0/0/0.0*, untagged, access
    ge-0/0/14.0*, untagged, access
```

```
VLAN: vlan_dyn, Created at: Thu May 15 13:43:09 2008
Internal index: 4, Admin State: Enabled, Origin: Static
Protocol: Port Mode
Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)
Protocol: MAC Based
Number of MAC entries: 6
    ge-0/0/0.0*
        00:00:00:00:00:02 (untagged)
        00:00:00:00:00:03 (untagged)
        00:00:00:00:00:04 (untagged)
        00:00:00:00:00:05 (untagged)
        00:00:00:00:00:06 (untagged)
        00:00:00:00:00:07 (untagged)
```

**show vlans extensive
(Port-based)**

```
user@switch> show vlans extensive
VLAN: default, created at Mon Feb 4 12:13:47 2008
Tag: None, Internal index: 0, Admin state: Enabled, Origin: static
Description: None
Dot1q Tunneling Status: Enabled
Customer VLAN ranges:
    1-4100
Private VLAN Mode: Primary
Protocol: Port based, Layer 3 interface: None
IP addresses: None
STP: None, RTG: None.
Number of interfaces: Tagged 0 (Active = 0), Untagged 23 (Active = 0)
    ge-0/0/34.0 (untagged, access)
    ge-0/0/33.0 (untagged, access)
    ge-0/0/32.0 (untagged, access)
    ge-0/0/31.0 (untagged, access)
    ge-0/0/30.0 (untagged, access)
    ge-0/0/29.0 (untagged, access)
    ge-0/0/28.0 (untagged, access)
    ge-0/0/27.0 (untagged, access)
    ge-0/0/26.0 (untagged, access)
    ge-0/0/25.0 (untagged, access)
    ge-0/0/19.0 (untagged, access)
    ge-0/0/18.0 (untagged, access)
    ge-0/0/17.0 (untagged, access)
    ge-0/0/16.0 (untagged, access)
    ge-0/0/15.0 (untagged, access)
    ge-0/0/14.0 (untagged, access)
```

```

ge-0/0/13.0 (untagged, access)
ge-0/0/11.0 (untagged, access)
ge-0/0/9.0 (untagged, access)
ge-0/0/8.0 (untagged, access)
ge-0/0/3.0 (untagged, access)
ge-0/0/2.0 (untagged, access)
ge-0/0/1.0 (untagged, access)

```

Secondary VLANs: Isolated 1, Community 1

```

Isolated VLANs :
  __vlan_pvlan_ge-0/0/3.0__
Community VLANs :
  comm1

```

VLAN: v0001, created at Mon Feb 4 12:13:47 2008

```

Tag: 1, Internal index: 1, Admin state: Enabled, Origin: static
Description: None
Protocol: Port based, Layer 3 interface: None
IP addresses: None
STP: None, RTG: None.
Number of interfaces: Tagged 4 (Active = 0), Untagged 0 (Active = 0)
  ge-0/0/24.0 (tagged, trunk)
  ge-0/0/23.0 (tagged, trunk)
  ge-0/0/22.0 (tagged, trunk)
  ge-0/0/21.0 (tagged, trunk)

```

VLAN: v0002, created at Mon Feb 4 12:13:47 2008

```

Tag: 2, Internal index: 2, Admin state: Enabled, Origin: static
Description: None
Protocol: Port based, Layer 3 interface: None
IP addresses: None
STP: None, RTG: None.
Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)
None

```

VLAN: v0003, created at Mon Feb 4 12:13:47 2008

```

Tag: 3, Internal index: 3, Admin state: Enabled, Origin: static
Description: None
Protocol: Port based, Layer 3 interface: None
IP addresses: None
STP: None, RTG: None.
Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)
None

```

show vlans sort-by tag user@switch> **show vlans sort-by tag**

Name	Tag	Interfaces
default		None
__vlan-x_1__	1	None
__vlan-x_2__	2	None
__vlan-x_3__	3	None
__vlan-x_4__	4	None
__vlan-x_5__	5	None
__vlan-x_6__	6	None
__vlan-x_7__	7	None

```

__vlan-x_8__ 8
None
__vlan-x_9__ 9
None
__vlan-x_10__ 10
None
__vlan-x_11__ 11
None
__vlan-x_12__ 12
None
__vlan-x_13__ 13
None
__vlan-x_14__ 14
None
__vlan-x_15__ 15
None
__vlan-x_16__ 16
None
__vlan-x_17__ 17
None
__vlan-x_18__ 18
None
__vlan-x_19__ 19
None
__vlan-x_20__ 20
None

```

show vlans sort-by name user@switch> **show vlans sort-by name**

```

Name          Tag  Interfaces
__employee_120__ 120  ge-0/0/22.0*
__employee_121__ 121  ge-0/0/22.0*
__employee_122__ 122  ge-0/0/22.0*
__employee_123__ 123  ge-0/0/22.0*
__employee_124__ 124  ge-0/0/22.0*
__employee_125__ 125  ge-0/0/22.0*
__employee_126__ 126  ge-0/0/22.0*
__employee_127__ 127  ge-0/0/22.0*
__employee_128__ 128  ge-0/0/22.0*
__employee_129__ 129  ge-0/0/22.0*
__employee_130__ 130  ge-0/0/22.0*

```

show vlans employee user@switch> **show vlans employee**

(vlan-range-name)

```

Name          Tag  Interfaces
__employee_120__ 120  ge-0/0/22.0*
__employee_121__ 121

```

```

__employee_122__ 122      ge-0/0/22.0*
__employee_123__ 123      ge-0/0/22.0*
__employee_124__ 124      ge-0/0/22.0*
__employee_125__ 125      ge-0/0/22.0*
__employee_126__ 126      ge-0/0/22.0*
__employee_127__ 127      ge-0/0/22.0*
__employee_128__ 128      ge-0/0/22.0*
__employee_129__ 129      ge-0/0/22.0*
__employee_130__ 130      ge-0/0/22.0*

```

show vlans summary

```

user@switch> show vlans summary
VLANs summary:
  Total: 8,   Configured VLANs: 5
  Internal VLANs: 1,   Temporary VLANs: 0

Dot1q VLANs summary:
  Total: 8,   Tagged VLANs: 2,   Untagged VLANs: 6
  Private VLAN:
    Primary VLANs: 2,   Community VLANs: 2,   Isolated VLANs: 3

Dot1q Tunneled VLANs summary:
  Total: 0
  Private VLAN:
    Primary VLANs: 0,   Community VLANs: 0,   Isolated VLANs: 0

Dynamic VLANs:
  Total: 2,   Dot1x: 2,   MVRP: 0

```

