



Junos[®] OS for EX Series Ethernet Switches

Ethernet Switching on EX Series Switches

Release

14.1X53



Published: 2014-12-19

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

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

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

Junos[®] OS for EX Series Ethernet Switches Ethernet Switching on EX Series Switches
Release 14.1X53
Copyright © 2014, Juniper Networks, Inc.
All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

	About the Documentation	xv
	Documentation and Release Notes	xv
	Supported Platforms	xv
	Using the Examples in This Manual	xv
	Merging a Full Example	xvi
	Merging a Snippet	xvi
	Documentation Conventions	xvii
	Documentation Feedback	xix
	Requesting Technical Support	xix
	Self-Help Online Tools and Resources	xix
	Opening a Case with JTAC	xx
Part 1	Overview	
Chapter 1	Bridging and VLANs	3
	Understanding Bridging and VLANs on EX Series Switches	3
	History of VLANs	4
	How Bridging of VLAN Traffic Works	4
	Packets Are Either Tagged or Untagged	5
	Switch Interface Modes—Access, Trunk, or Tagged Access	6
	Access Mode	6
	Trunk Mode	6
	Trunk Mode and Native VLAN	7
	Tagged-Access Mode	7
	Additional Advantages of Using VLANs	8
	Maximum VLANs and VLAN Members Per Switch	8
	A Default VLAN Is Configured on Most Switches	9
	Assigning Traffic to VLANs	10
	Assign VLAN Traffic According to the Interface Port Source	10
	Assign VLAN Traffic According to the Source MAC Address	10
	Forwarding VLAN Traffic	10
	VLANs Communicate with Integrated Routing and Bridging Interfaces or Routed VLAN Interfaces	10
	Understanding Integrated Routing and Bridging Interfaces and Routed VLAN Interfaces on EX Series Switches	12
	When Should I Use an IRB Interface or RVI?	13
	How Does an IRB Interface or RVI Work?	13
	Creating an IRB Interface or RVI	14
	Viewing IRB Interface and RVI Statistics	15
	IRB Interfaces and RVI Functions and Other Technologies	15

Understanding Private VLANs on EX Series Switches	16
Typical Structure and Primary Application of PVLANS	17
Routing Between Isolated and Community VLANs	20
PVLANS Use 802.1Q Tags to Identify Packets	20
PVLANS Use IP Addresses Efficiently	20
PVLANS Use Four Different Ethernet Switch Port Types	20
Understanding PVLAN Traffic Flows Across Multiple Switches	22
Community VLAN Sending Untagged Traffic	22
Isolated VLAN Sending Untagged Traffic	23
PVLAN Tagged Traffic Sent on a Promiscuous Port	24
Understanding Virtual Routing Instances on EX Series Switches	25
Understanding Multiple VLAN Registration Protocol (MVRP) on EX Series Switches	26
How MVRP Updates, Creates, and Deletes VLANs on the Switches	26
MVRP Is Disabled by Default on the Switches	26
MRP Timers Control MVRP Updates	27
MVRP Uses MRP Messages to Transmit Switch and VLAN States	27
Compatibility Issues with Junos OS Releases of MVRP	28
Understanding MAC Notification on EX Series Switches	29
Understanding MAC Address Aging	30
Understanding MAC Address Assignment in an EX Series Switch	31
Understanding Edge Virtual Bridging for Use with VEPA Technology	33
What Is EVB?	33
What Is VEPA?	33
Why Use VEPA Instead of VEB?	33
How Does EVB Work?	33
How Do I Implement EVB?	34
Understanding Ethernet Ring Protection Switching Functionality	34
Acronyms	35
Ring Nodes	35
Ring Node States	35
Failure Detection	36
Logical Ring	36
FDB Flush	36
Traffic Blocking and Forwarding	36
RAPS Message Blocking and Forwarding	36
Dedicated Signaling Control Channel	38
RAPS Message Termination	38
Multiple Rings	38
Node ID	38
Bridge Domains with the Ring Port (MX Series Routers Only)	38
Ethernet Ring Protection Switching Overview	39
Chapter 2 Q-in-Q Tunneling	41
Understanding Q-in-Q Tunneling on EX Series Switches	41
How Q-in-Q Tunneling Works	41
Disabling MAC Address Learning	42

	Mapping C-VLANs to S-VLANs	42
	All-in-One Bundling	43
	Many-to-One Bundling	43
	Mapping a Specific Interface	43
	Routed VLAN Interfaces on Q-in-Q VLANs	44
	Limitations for Q-in-Q Tunneling	44
Chapter 3	Layer 2 Protocol Tunneling	45
	Understanding Layer 2 Protocol Tunneling on EX Series Switches	45
	Layer 2 Protocols Supported by L2PT on EX Series Switches	45
	How L2PT Works	46
	L2PT Basics on EX Series Switches	48
Chapter 4	Redundant Trunk Groups	49
	Understanding Redundant Trunk Links	49
Chapter 5	Proxy ARP	53
	Understanding Proxy ARP on EX Series Switches	53
	What Is ARP?	53
	Proxy ARP Overview	53
	Best Practices for Proxy ARP on EX Series Switches	54
Part 2	Configuration	
Chapter 6	Configuration Examples	57
	Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch	57
	Example: Setting Up Bridging with Multiple VLANs for EX Series Switches	65
	Example: Connecting an Access Switch to a Distribution Switch	72
	Example: Configuring a Private VLAN on a Single EX Series Switch	81
	Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches	88
	Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches	92
	Example: Setting Up Q-in-Q Tunneling on EX Series Switches	103
	Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches	107
	Example: Configuring Redundant Trunk Links for Faster Recovery	112
	Example: Configuring Proxy ARP on an EX Series Switch	117
	Example: Configuring a Private VLAN Spanning Multiple EX Series Switches	119
	Example: Configuring Edge Virtual Bridging for Use with VEPA Technology	134
	Example: Configuring Ethernet Ring Protection Switching on EX Series Switches	140
Chapter 7	Configuration Tasks	157
	Configuring VLANs for EX Series Switches (J-Web Procedure)	158
	Configuring VLANs for EX Series Switches (CLI Procedure)	160
	Why Create a VLAN?	161
	Create a VLAN Using the Minimum Procedure	161
	Create a VLAN Using All of the Options	162
	Configuration Guidelines for VLANs	163
	Configuring Routed VLAN Interfaces (CLI Procedure)	164
	Configuring MAC Table Aging (CLI Procedure)	166

Configuring the Native VLAN Identifier (CLI Procedure)	167
Creating a Series of Tagged VLANs (CLI Procedure)	168
Creating a Private VLAN on a Single EX Series Switch (CLI Procedure)	169
Creating a Private VLAN Spanning Multiple EX Series Switches (CLI Procedure)	171
Configuring a Routed VLAN Interface in a Private VLAN (CLI Procedure)	173
Configuring Virtual Routing Instances (CLI Procedure)	174
Configuring MAC Notification (CLI Procedure)	175
Enabling MAC Notification	175
Disabling MAC Notification	175
Setting the MAC Notification Interval	176
Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)	176
Enabling MVRP	176
Disabling MVRP	177
Disabling Dynamic VLANs	177
Configuring Timer Values	177
Configuring MVRP Registration Mode	178
Using MVRP in a Mixed-Release Network	178
Configuring Q-in-Q Tunneling (CLI Procedure)	180
Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure)	181
Configuring Redundant Trunk Groups (J-Web Procedure)	183
Configuring Redundant Trunk Links for Faster Recovery (CLI Procedure)	185
Configuring Proxy ARP (CLI Procedure)	186
Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)	187
Configuring Edge Virtual Bridging (CLI Procedure)	188
Configuring Ethernet Ring Protection Switching (CLI Procedure)	190
Enabling VLAN Pruning for Broadcast, Multicast, and Unknown Unicast Traffic in an EX Series Virtual Chassis (CLI Procedure)	192
Chapter 8 Configuration Statements	193
[edit ethernet-switching-options] Configuration Statement Hierarchy on EX Series Switches	195
Supported Statements in the [edit ethernet-switching-options] Hierarchy Level	195
Unsupported Statements in the [edit ethernet-switching-options] Hierarchy Level	198
[edit interfaces] Configuration Statement Hierarchy on EX Series Switches	198
[edit protocols] Configuration Statement Hierarchy on EX Series Switches	199
[edit routing-instances] Configuration Statement Hierarchy on EX Series Switches	200
Supported Statements in the [edit routing-instances] Hierarchy Level	201
Unsupported Statements in the [edit routing-instances] Hierarchy Level	204
[edit vlans] Configuration Statement Hierarchy on EX Series Switches	205
Supported Statements in the [edit vlans] Hierarchy Level	205
Unsupported Statements in the [edit vlans] Hierarchy Level	206
add-attribute-length-in-pdu	207
arp (System)	208

arp-on-stp	209
control-channel	210
control-vlan	211
customer-vlans	212
data-channel	213
description (VLANs)	214
disable (MVRP)	214
dot1q-tunneling (Ethernet Switching)	215
dot1q-tunneling (VLANs)	216
drop-threshold	217
east-interface	218
edge-virtual-bridging	219
ethernet-ring	220
ether-type	221
ethernet-switching-options	222
filter (VLANs)	225
group (Redundant Trunk Groups)	226
guard-interval	227
instance-type	228
interface (Redundant Trunk Groups)	230
interface (Routing Instances)	231
interface (VLANs)	232
interface (MVRP)	233
interfaces (Q-in-Q Tunneling)	234
isolation-id	234
join-timer (MVRP)	235
layer2-protocol-tunneling	236
l3-interface (VLANs)	238
l3-interface-ingress-counting	239
leaveall-timer (MVRP)	240
leave-timer (MVRP)	241
mac (Static MAC-Based VLANs)	242
mac-limit (VLANs)	243
mac-lookup-length	245
mac-notification	246
mac-table-aging-time	247
mapping	248
members	249
mvrp	251
native-vlan-id	252
next-hop (Static MAC-Based VLANs)	252
no-dynamic-vlan	253
no-local-switching	253
no-mac-learning (Q-in-Q VLANs)	254
no-mac-learning (Q-in-Q Interfaces)	254
node-id	255
notification-interval	256
port-mode	257
preempt-cutover-timer	258

primary-vlan	259
protection-group	260
proxy-arp	262
pvlan-trunk	263
redundant-trunk-group	264
registration	265
restore-interval	266
ring-protection-link-end	267
ring-protection-link-owner	267
routing-instances	268
shutdown-threshold	269
static (Static MAC-Based VLANs)	270
traceoptions (Ethernet Ring Protection)	271
traceoptions (Edge Virtual Bridging)	273
vlan (802.1Q Tagging)	274
vlan (Static MAC-based VLANs)	275
vlan-id (802.1Q Tagging)	276
vlan-prune	277
vlan-range	278
vlangs	279
vrf-mtu-check	280
vsi-discovery	281
vsi-policy	282
west-interface	283

Part 3

Administration

Chapter 9

Routine Monitoring 287

Verifying That a Series of Tagged VLANs Has Been Created	287
Verifying That Virtual Routing Instances Are Working	289
Verifying That Q-in-Q Tunneling Is Working	290
Verifying Routed VLAN Interface Status and Statistics	291
Verifying That a Private VLAN Is Working	292
Verifying That MVRP Is Working Correctly	297
Verifying That MAC Notification Is Working Properly	299
Verifying That Proxy ARP Is Working Correctly	299
Monitoring Ethernet Switching	300

Chapter 10

Operational Commands 303

clear edge-virtual-bridging	304
clear ethernet-switching layer2-protocol-tunneling error	305
clear ethernet-switching layer2-protocol-tunneling statistics	306
clear ethernet-switching table	307
clear mvrp statistics	309
show edge-virtual-bridging	310
show ethernet-switching interfaces	313
show ethernet-switching layer2-protocol-tunneling interface	317
show ethernet-switching layer2-protocol-tunneling statistics	319
show ethernet-switching layer2-protocol-tunneling vlan	322
show ethernet-switching mac-learning-log	324

show ethernet-switching mac-notification	326
show ethernet-switching statistics aging	328
show ethernet-switching statistics mac-learning	330
show ethernet-switching table	334
show mvrp	339
show mvrp dynamic-vlan-memberships	341
show mvrp statistics	342
show protection-group ethernet-ring aps	345
show protection-group ethernet-ring configuration	347
show protection-group ethernet-ring interface	350
show protection-group ethernet-ring node-state	353
show protection-group ethernet-ring statistics	356
show redundant-trunk-group	359
show system statistics arp	361
show vlans	362

Part 4

Chapter 11

Troubleshooting

Troubleshooting Procedure	375
Troubleshooting Ethernet Switching	375
MAC Address in the Switch's Ethernet Switching Table Is Not Updated After a MAC Address Move	375

List of Figures

Part 1	Overview	
Chapter 1	Bridging and VLANs	3
	Figure 1: An IRB Interface or RVI on a Switch Providing Routing Between Two Access Switches	13
	Figure 2: Creating an IRB Interface or RVI	14
	Figure 3: Private VLAN on a Single EX Switch	18
	Figure 4: PVLAN Spanning Multiple EX Series Switches	19
	Figure 5: Community VLAN Sends Untagged Traffic	22
	Figure 6: Isolated VLAN Sends Untagged Traffic	23
	Figure 7: PVLAN Tagged Traffic Sent on a Promiscuous Port	24
	Figure 8: Protocol Packets from the Network to the Router	36
	Figure 9: Protocol Packets from the Router or Switch to the Network	37
Chapter 3	Layer 2 Protocol Tunneling	45
	Figure 10: L2PT Example	47
Chapter 4	Redundant Trunk Groups	49
	Figure 11: Redundant Trunk Group, Link 1 Active	50
	Figure 12: Redundant Trunk Group, Link 2 Active	50
Part 2	Configuration	
Chapter 6	Configuration Examples	57
	Figure 13: Topology for Configuration	73
	Figure 14: Topology of a Private VLAN on a Single EX Series Switch	83
	Figure 15: MVRP Configured on Two Access Switches and One Distribution Switch for Automatic VLAN Administration	94
	Figure 16: L2PT Topology	108
	Figure 17: Topology for Configuring the Redundant Trunk Links	114
	Figure 18: PVLAN Topology Spanning Multiple Switches	121
	Figure 19: Topology	135
	Figure 20: Ethernet Ring Protection Switching Example	142

List of Tables

	About the Documentation	xv
	Table 1: Notice Icons	xvii
	Table 2: Text and Syntax Conventions	xvii
Part 1	Overview	
Chapter 1	Bridging and VLANs	3
	Table 3: Tracking IRB Interface and RVI Usage	15
	Table 4: When VLANs in a PVLAN Need 802.1Q Tags	20
	Table 5: PVLAN Ports and Layer 2 Connectivity	21
	Table 6: Junos OS MVRP Versions and Inclusion of Extra Byte in PDU	28
	Table 7: MVRP Environments and Description of Required Actions	28
Chapter 3	Layer 2 Protocol Tunneling	45
	Table 8: Protocol Destination MAC Addresses	47
Part 2	Configuration	
Chapter 6	Configuration Examples	57
	Table 9: Components of the Basic Bridging Configuration Topology	59
	Table 10: Components of the Multiple VLAN Topology	66
	Table 11: Components of the Topology for Connecting an Access Switch to a Distribution Switch	73
	Table 12: Components of the Topology for Configuring a PVLAN	82
	Table 13: Components of the Network Topology	94
	Table 14: Components of the Topology for Setting Up Q-in-Q Tunneling	104
	Table 15: Components of the Redundant Trunk Link Topology	114
	Table 16: Components of Switch 1 in the Topology for Configuring a PVLAN Spanning Multiple EX Series Switches	122
	Table 17: Components of Switch 2 in the Topology for Configuring a PVLAN Spanning Multiple EX Series Switches	122
	Table 18: Components of Switch 3 in the Topology for Configuring a PVLAN Spanning Multiple EX Series Switches	123
	Table 19: Components of the Topology for Configuring EVB	136
	Table 20: Components to Configure for This Example	142
Chapter 7	Configuration Tasks	157
	Table 21: VLAN Configuration Details	158
	Table 22: RTG Configuration Fields	184
Chapter 8	Configuration Statements	193

Table 23: Unsupported [edit routing-instances] Configuration Statements on EX Series Switches	205
Table 24: Unsupported [edit vlans] Configuration Statements on EX Series Switches	206

Part 3

Chapter 9

Administration

Routine Monitoring	287
-------------------------------------	------------

Table 25: Ethernet Switching Output Fields	300
--	-----

Chapter 10

Operational Commands	303
---------------------------------------	------------

Table 26: show edge-virtual-bridging Output Field Descriptions	310
--	-----

Table 27: show ethernet-switching interfaces Output Fields	314
--	-----

Table 28: show ethernet-switching layer2-protocol-tunneling interface Output Fields	317
---	-----

Table 29: show ethernet-switching layer2-protocol-tunneling statistics Output Fields	320
--	-----

Table 30: show ethernet-switching layer2-protocol-tunneling vlan Output Fields	322
--	-----

Table 31: show ethernet-switching mac-learning-log Output Fields	324
--	-----

Table 32: show ethernet-switching mac-notification Output Fields	326
--	-----

Table 33: show ethernet-switching statistics aging Output Fields	328
--	-----

Table 34: show ethernet-switching statistics mac-learning Output Fields	331
---	-----

Table 35: show ethernet-switching table Output Fields	335
---	-----

Table 36: show mvrp Output Fields	339
---	-----

Table 37: show mvrp dynamic-vlan-memberships Output Fields	341
--	-----

Table 38: show mvrp statistics Output Fields	342
--	-----

Table 39: show protection-group ethernet-ring aps Output Fields	345
---	-----

Table 40: show protection-group ethernet-ring configuration Output Fields	347
---	-----

Table 41: MX Series Routers show protection-group ethernet-ring interface Output Fields	350
---	-----

Table 42: show protection-group ethernet-ring node-state Output Fields	353
--	-----

Table 43: show protection-group ethernet-ring statistics Output Fields	356
--	-----

Table 44: show redundant-trunk-group Output Fields	359
--	-----

Table 45: show vlans Output Fields	363
--	-----

About the Documentation

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

Documentation and Release Notes

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

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

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

Supported Platforms

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

- EX Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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


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

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

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

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

Documentation Conventions

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

Table 1: Notice Icons







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

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

Table 2: Text and Syntax Conventions

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

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

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

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

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

Documentation Feedback

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

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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

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

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

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

- [Bridging and VLANs on page 3](#)
- [Q-in-Q Tunneling on page 41](#)
- [Layer 2 Protocol Tunneling on page 45](#)
- [Redundant Trunk Groups on page 49](#)
- [Proxy ARP on page 53](#)

CHAPTER 1

Bridging and VLANs

- [Understanding Bridging and VLANs on EX Series Switches on page 3](#)
- [Understanding Integrated Routing and Bridging Interfaces and Routed VLAN Interfaces on EX Series Switches on page 12](#)
- [Understanding Private VLANs on EX Series Switches on page 16](#)
- [Understanding PVLAN Traffic Flows Across Multiple Switches on page 22](#)
- [Understanding Virtual Routing Instances on EX Series Switches on page 25](#)
- [Understanding Multiple VLAN Registration Protocol \(MVRP\) on EX Series Switches on page 26](#)
- [Understanding MAC Notification on EX Series Switches on page 29](#)
- [Understanding MAC Address Aging on page 30](#)
- [Understanding MAC Address Assignment in an EX Series Switch on page 31](#)
- [Understanding Edge Virtual Bridging for Use with VEPA Technology on page 33](#)
- [Understanding Ethernet Ring Protection Switching Functionality on page 34](#)
- [Ethernet Ring Protection Switching Overview on page 39](#)

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:

- [History of VLANs on page 4](#)
- [How Bridging of VLAN Traffic Works on page 4](#)
- [Packets Are Either Tagged or Untagged on page 5](#)
- [Switch Interface Modes—Access, Trunk, or Tagged Access on page 6](#)
- [Additional Advantages of Using VLANs on page 8](#)
- [Maximum VLANs and VLAN Members Per Switch on page 8](#)
- [A Default VLAN Is Configured on Most Switches on page 9](#)
- [Assigning Traffic to VLANs on page 10](#)

- [Forwarding VLAN Traffic on page 10](#)
- [VLANs Communicate with Integrated Routing and Bridging Interfaces or Routed VLAN Interfaces on page 10](#)

History of VLANs

Ethernet LANs were originally designed for small, simple networks that primarily carried text. However, over time, the type of data carried by LANs grew to include voice, graphics, and video. This more complex data, when combined with the ever-increasing speed of transmission, eventually became too much of a load for the original Ethernet LAN design. Multiple packet collisions were significantly slowing down the larger LANs.

The IEEE 802.1D-2004 standard helped evolve Ethernet LANs to cope with the higher data and transmission requirements by defining the concept of *transparent bridging* (generally called simply *bridging*). Bridging divides a single physical LAN (now called a single *broadcast domain*) into two or more virtual LANs, or VLANs. Each VLAN is a collection of some of the LAN nodes grouped together to form individual broadcast domains.

When VLANs are grouped logically by function or organization, a significant percentage of data traffic stays within the VLAN. This relieves the load on the LAN because all traffic no longer has to be forwarded to all nodes on the LAN. A VLAN first transmits packets within the VLAN, thereby reducing the number of packets transmitted on the entire LAN. Because packets whose origin and destination are in the same VLAN are forwarded only within the local VLAN, packets that are not destined for the local VLAN are the only ones forwarded to other broadcast domains. This way, bridging and VLANs limit the amount of traffic flowing across the entire LAN by reducing the possible number of collisions and packet retransmissions within VLANs and on the LAN as a whole.

How Bridging of VLAN Traffic Works

Because the objective of the IEEE 802.1D-2004 standard was to reduce traffic and therefore reduce potential transmission collisions for Ethernet, a system was implemented to reuse information. Instead of having a switch go through a location process every time a frame is sent to a node, the transparent bridging protocol allows a switch to record the location of known nodes. When packets are sent to nodes, those destination node locations are stored in address-lookup tables called *Ethernet switching tables*. Before sending a packet, a switch using bridging first consults the switching tables to see if that node has already been located. If the location of a node is known, the frame is sent directly to that node.

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

- Learning
- Forwarding
- Flooding
- Filtering
- Aging

The key bridging mechanism used by LANs and VLANs is *learning*. When a switch is first connected to an Ethernet LAN or VLAN, it has no information about other nodes on the network. As packets are sent, the switch learns the embedded MAC addresses of the sending nodes and stores them in the Ethernet switching table, along with two other pieces of information—the interface (or port) on which the traffic was received on the destination node and the time the address was learned.

Learning allows switches to then do *forwarding*. By consulting the Ethernet switching table to see whether the table already contains the frame's destination MAC address, switches save time and resources when forwarding packets to the known MAC addresses. If the Ethernet switching table does not contain an entry for an address, the switch uses flooding to learn that address.

Flooding finds a particular destination MAC address without using the Ethernet switching table. When traffic originates on the switch and the Ethernet switching table does not yet contain the destination MAC address, the switch first floods the traffic to all other interfaces within the VLAN. When the destination node receives the flooded traffic, it can send an acknowledgment packet back to the switch, allowing it to learn the MAC address of the node and 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—it learns 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.

To keep entries in the Ethernet switching table current, the switch uses a fifth bridging mechanism, *aging*. Aging is the reason that the Ethernet switching table entries include timestamps. 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 eventually flushes unavailable network nodes out of the Ethernet switching table.

Packets Are Either Tagged or Untagged

When an Ethernet LAN is divided into VLANs, each VLAN is identified by a unique 802.1Q ID. The VLAN IDs 1 through 4094 can be assigned to VLANs, while VLAN IDs 0 and 4095 are reserved by Junos OS and cannot be assigned.

Ethernet packets include a tag protocol identifier (TPID) EtherType field, which identifies the protocol being transported. When a device within a VLAN generates a packet, this field includes a value of 0x8100, which indicates that the packet is a VLAN-tagged packet. The packet also has a VLAN ID field that includes the unique 802.1Q ID, which identifies the VLAN to which the packet belongs.

In addition to the TPID EtherType value of 0x8100, EX Series switches that run Junos OS that does not support the Enhanced Layer 2 Software (ELS) configuration style also

support values of 0x88a8 (Provider Bridging and Shortest Path Bridging) and 0x9100 (Q-inQ).

For a simple network that has only a single VLAN, all packets include a default 802.1Q tag, which is the only VLAN membership that does not mark the packet as tagged. These packets are untagged packets.

Switch Interface Modes—Access, Trunk, or Tagged Access

Ports, or interfaces, on a switch operate in one of three modes:

- Access mode
- Trunk mode
- Tagged-access mode

Access Mode

An interface in access mode connects a switch to a single network device, such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. Access interfaces accept only untagged packets.

By default, when you boot an EX Series switch that runs Junos OS that does not support ELS and use the factory default configuration, or when you boot such a switch and do not explicitly configure a port mode, all interfaces on the switch are in access mode and accept only untagged packets from the VLAN named **default**. You can optionally configure another VLAN and use that VLAN instead of **default**.

On an EX Series switch that runs Junos OS that supports ELS, the VLAN named **default** is not supported. Therefore, on such switches, you must explicitly configure at least one VLAN, even if your network is simple and you want only one broadcast domain to exist. After you assign an interface to a VLAN, the interface functions in access mode.

For EX Series switches that run either type of software, you can also configure a trunk port or interface to accept untagged packets from a user-configured VLAN. For details about this concept (native VLAN), see [“Trunk Mode and Native VLAN” on page 7](#).

Trunk Mode

Trunk mode interfaces are generally used to connect switches to one another. Traffic sent between switches can then consist of packets from multiple VLANs, with those packets multiplexed so that they can be sent over the same physical connection. Trunk interfaces usually accept only tagged packets and use the VLAN ID tag to determine both the packets' VLAN origin and VLAN destination.

On an EX Series switch that runs software that does not support ELS, an untagged packet is not recognized on a trunk port unless you configure additional settings on that port.

On an EX Series switch that runs Junos OS that supports ELS, a trunk port recognizes untagged control packets for protocols such as the Link Aggregation Control Protocol (LACP) and the Link Layer Discovery Protocol (LLDP). However, the trunk port does not recognize untagged data packets unless you configure additional settings on that port.

In the rare case where you want untagged packets to be recognized by a trunk port on EX Series switches that run either type of software, you must configure the single VLAN on a trunk port as a *native VLAN*. For more information about native VLANs, see [“Trunk Mode and Native VLAN” on page 7](#).

Trunk Mode and Native VLAN

On an EX Series switch that runs Junos OS that does not support ELS, a trunk port does not recognize packets that do not include VLAN tags, which are also known as untagged packets. On an EX Series switch that runs Junos OS that supports ELS, a trunk port recognizes untagged control packets, but it does not recognize untagged data packets. With native VLAN configured, untagged packets that a trunk port normally does not recognize are sent over the trunk interface. In a situation where packets pass from a device, such as an IP phone or printer, to a switch in access mode, and you want those packets sent from the switch over a trunk port, use native VLAN mode. Create a native VLAN by configuring a VLAN ID for it, and specify that the trunk port is a member of the native VLAN.

The switch's trunk port will then treat those packets differently than the other tagged packets. For example, if a trunk port has three VLANs, 10, 20, and 30, assigned to it with VLAN 10 being the native VLAN, packets on VLAN 10 that leave the trunk port on the other end have no 802.1Q header (tag).

There is another native VLAN option for EX Series switches that do not support ELS. You can have the switch add and remove tags for untagged packets. To do this, you first configure the single VLAN as a native VLAN on a port attached to a device on the edge. Then, assign a VLAN ID tag to the single native VLAN on the port connected to a device. Last, add the VLAN ID to the trunk port. Now, when the switch receives the untagged packet, it adds the ID you specified and sends and receives the tagged packets on the trunk port configured to accept that VLAN.

Tagged-Access Mode

Only EX Series switches that run Junos OS that does not use the ELS configuration style support tagged-access mode.

Tagged-access mode accommodates cloud computing scenarios, specifically deployments including servers that adhere to the edge virtual bridging (EVB) standard (IEEE 803.1Qbg). See [“Understanding Edge Virtual Bridging for Use with VEPA Technology” on page 33](#).

Because several virtual computers can be included on one physical server, the packets generated by one server can contain an aggregation of VLAN packets from different virtual machines on that server. To accommodate this situation, tagged-access mode reflects packets back to the physical server on the same downstream port when the destination address of the packet was learned on that downstream port. Packets are also reflected back to the physical server on the downstream port when the destination has not yet been learned. Therefore, the third interface mode, tagged access, has some characteristics of access mode and some characteristics of trunk mode:

- Like access mode, tagged-access mode connects the switch to an access layer device. Unlike access mode, tagged-access mode is capable of accepting VLAN tagged packets.

- Like trunk mode, tagged-access mode accepts VLAN tagged packets from multiple VLANs. Unlike trunk port interfaces, which are connected at the core/distribution layer, tagged-access port interfaces connect devices at the access layer.

Like trunk mode, tagged-access mode also supports native VLAN.



NOTE: Control packets are never reflected back on the downstream port.

Additional Advantages of Using VLANs

In addition to reducing traffic and thereby speeding up the network, VLANs have the following advantages:

- VLANs provide segmentation services traditionally provided by routers in LAN configurations, thereby reducing hardware equipment costs.
- Packets coupled to a VLAN can be reliably identified and sorted into different domains. You can contain broadcasts within parts of the network, thereby freeing up network resources. For example, when a DHCP server is plugged into a switch and starts broadcasting its presence, you can prevent some hosts from accessing it by using VLANs to split up the network.
- For security issues, VLANs provide granular control of the network because each VLAN is identified by a single IP subnetwork. All packets passing in and out of a VLAN are consistently tagged with the VLAN ID of that VLAN, thereby providing easy identification, because a VLAN ID on a packet cannot be altered. (For an EX Series switch that runs Junos OS that does not support ELS, we recommend that you avoid using 1 as a VLAN ID, because that ID is a default value.)
- VLANs react quickly to host relocation—this is also due to the persistent VLAN tag on packets.
- On an Ethernet LAN, all network nodes must be physically connected to the same network. In VLANs, the physical location of nodes is not important—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 physical location.

Maximum VLANs and VLAN Members Per Switch

The number of VLANs supported per switch varies for each switch. Use the configuration-mode command **set vlans *vlan-name* *vlan-id* ?** to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because you have to assign a specific ID number when you create a VLAN—you could overwrite one of the numbers, but you cannot exceed the limit.

You can, however, exceed the recommended VLAN member maximum for a switch.

On an EX Series switch that runs Junos OS that does not support the ELS configuration style, the maximum number of VLAN members allowed on the switch is eight times the maximum number of VLANs that the switch supports (vmember limit = vlan max * 8). If the configuration of the switch exceeds the recommended VLAN member maximum,

a warning message appears when you commit the configuration. If you commit the configuration despite the warning, the commit succeeds, but there is a risk of the Ethernet switching process (eswd) failing as a result of memory allocation failure.

On an EX Series switch that runs Junos OS that supports ELS, the maximum number of VLAN members allowed on the switch is 24 times the maximum number of VLANs that the switch supports ($\text{vmember limit} = \text{vlan max} * 24$). If the configuration of the switch exceeds the recommended VLAN member maximum, a warning message appears in the system log (syslog).

A Default VLAN Is Configured on Most Switches



NOTE: EX Series switches that run Junos OS with the ELS configuration style do not support a default VLAN.

Some EX Series switches that run Junos OS that does not support the ELS configuration style are preconfigured with a VLAN named **default** that does not tag packets and operates only with untagged packets. On these switches, each interface already belongs to the VLAN named **default** and all traffic uses this VLAN until you configure more VLANs and assign traffic to those VLANs.

The following EX Series switches that run Junos OS that does not support the ELS are not preconfigured to belong to **default** or any other VLAN:.

- Modular switches, such as the EX8200 switches and EX6200 switches
- Switches that are part of a Virtual Chassis

The reason that these switches are not preconfigured is that the physical configuration in both situations is flexible. There is no way of knowing which line cards have been inserted in either the EX8200 switch or EX6200 switch. There is also no way of knowing which switches are included in the Virtual Chassis. Switch interfaces in these two cases must first be defined as Ethernet switching interfaces. After an interface is defined as an Ethernet switching interface, the default VLAN appears in the output from the ? help and other commands.



NOTE: When a Juniper Networks EX4500 Ethernet Switch, EX4200 Ethernet Switch, or EX3300 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.

Assigning Traffic to VLANs

You can assign traffic on any switch to a particular VLAN by referencing either the interface port of the traffic or the MAC addresses of devices sending traffic.

Assign VLAN Traffic According to the Interface Port Source

This method is most commonly used to assign traffic to VLANs. In this case, you specify that all traffic received on a particular switch interface is assigned to a specific VLAN. You configure this VLAN assignment when you configure the switch, by using either the VLAN number (called a VLAN ID) or by using the VLAN name, which the switch then translates into a numeric VLAN ID. This method is referred to simply as creating a VLAN because it is the most commonly used method.

Assign VLAN Traffic According to the Source MAC Address

In this case, all traffic received from a specific MAC address is forwarded to a specific egress interface (next hop) on the switch. MAC-based VLANs are either static (named MAC addresses configured one at a time) or dynamic (configured using a RADIUS server).

To configure a static MAC-based VLAN on an EX Series switch that supports ELS, see [“Adding a Static MAC Address Entry to the Ethernet Switching Table \(CLI Procedure\)” on page 187](#). To configure a static MAC-based VLAN on an EX Series switch that does not support ELS, see [“Adding a Static MAC Address Entry to the Ethernet Switching Table \(CLI Procedure\)” on page 187](#).

For information about using 802.1X authentication to authenticate end devices and allow access to dynamic VLANs configured on a RADIUS server, see *Understanding Dynamic VLANs for 802.1X on EX Series Switches*. You can optionally implement this feature to offload the manual assignment of VLAN traffic to automated RADIUS server databases.

Forwarding VLAN Traffic

To pass traffic within a VLAN, the switch uses Layer 2 forwarding protocols, including IEEE 802.1Q spanning-tree protocols and Multiple VLAN Registration Protocol (MVRP).

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.

To pass traffic from a single device on an access port to a switch and then pass those packets on a trunk port, use the native mode configuration previously discussed under [“Trunk Mode” on page 6](#).

VLANs Communicate with Integrated Routing and Bridging Interfaces or Routed VLAN Interfaces

Traditionally, switches sent traffic to hosts that were part of the same broadcast domain (VLAN) but routers were needed to route traffic from one broadcast domain to another. Also, only routers performed other Layer 3 functions such as traffic engineering.

EX Series switches that run Junos OS that supports the ELS configuration style perform inter-VLAN routing functions using an integrated routing and bridging (IRB) interface

named `irb`, while EX Series switches that run Junos OS that does not support ELS perform these functions using a routed VLAN interface (RVI) named `vlan`. These interfaces detect both MAC addresses and IP addresses and route data to Layer 3 interfaces, thereby frequently eliminating the need to have both a switch and a router.

**Related
Documentation**

- [Understanding Private VLANs on EX Series Switches on page 16](#)
- [Understanding Layer 2 Protocol Tunneling on EX Series Switches on page 45](#)
- [Understanding Multiple VLAN Registration Protocol \(MVRP\) on EX Series Switches on page 26](#)
- [Understanding Integrated Routing and Bridging Interfaces and Routed VLAN Interfaces on EX Series Switches on page 12](#)
- [Understanding Edge Virtual Bridging for Use with VEPA Technology on page 33](#)
- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 57](#)
- *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*
- [Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 65](#)
- [Example: Connecting an Access Switch to a Distribution Switch on page 72](#)
- *Example: Connecting Access Switches to a Distribution Switch*

Understanding Integrated Routing and Bridging Interfaces and Routed VLAN Interfaces on EX Series Switches

Virtual LANs (VLANs), by definition, divide a LAN's broadcast environment into isolated virtual broadcast domains, thereby limiting the amount of traffic flowing across the entire LAN and reducing the possible number of collisions and packet retransmissions within the LAN. For example, you might want to create a VLAN that includes the employees in a department and the resources that they use often, such as printers, servers, and so on.

Of course, you also want to allow these employees to communicate with people and resources in other VLANs. To forward packets between VLANs, you traditionally needed a router that connected the VLANs. However, you can also accomplish this forwarding with a switch by configuring one of the following features:

- On Juniper Networks EX Series Ethernet Switches that run Juniper Networks Junos operating system (Junos OS) that supports the Enhanced Layer 2 Software (ELS) configuration style, configure an integrated routing and bridging (IRB) interface.
- On EX Series switches that run Junos OS that does not support ELS, configure a routed VLAN interface (RVI).



NOTE: IRB interfaces and RVIs provide the same functionality. Where the functionality for both features is the same, this topic uses the term *these interfaces* to refer collectively to both IRB interfaces and RVIs. Where differences exist between the two features, this topic calls out the IRB interfaces and RVIs separately.


Configuring a switch to route traffic between VLANs reduces complexity and eliminates the costs associated with purchasing, installing, managing, powering, and cooling a router.

These interfaces route only VLAN traffic and work by logically dividing a switch into multiple virtual routing instances, thereby isolating VLAN traffic traveling across the network into virtual segments. These interfaces allow switches to recognize which packets are being sent to another VLAN's MAC addresses—then, packets are bridged (switched) whenever the destination is within the same VLAN and are routed through these interfaces only when necessary. Whenever packets can be switched instead of routed, several layers of processing are eliminated. The switches rely on their Layer 3 capabilities to provide this basic routing between VLANs:

- Two VLANs on the same switch
- Two VLANs on different switches (routing is provided by an intermediary third switch.)

Figure 1 on page 13 illustrates a switch routing VLAN traffic between two access layer switches using one of these interfaces.

Figure 1: An IRB Interface or RVI on a Switch Providing Routing Between Two Access Switches



This topic describes:

- [When Should I Use an IRB Interface or RVI? on page 13](#)
- [How Does an IRB Interface or RVI Work? on page 13](#)
- [Creating an IRB Interface or RVI on page 14](#)
- [Viewing IRB Interface and RVI Statistics on page 15](#)
- [IRB Interfaces and RVI Functions and Other Technologies on page 15](#)

When Should I Use an IRB Interface or RVI?

Configure an IRB interface or an RVI for a VLAN if you need to:

- Allow traffic to be routed between VLANs.
- Provide Layer 3 IP connectivity to the switch.
- Monitor individual VLANs for billing purposes. Service providers often need to monitor traffic for this purpose, but this capability can be useful for enterprises where various groups share the cost of the network.

How Does an IRB Interface or RVI Work?

For an IRB interface, the switch provides the name `irb`, and for an RVI, the switch provides the name `vlan`. Like all Layer 3 interfaces, these interfaces require a logical unit number with an IP address assigned to it. In fact, to be useful, the implementation of these interfaces in an enterprise with multiple VLANs requires at least two logical units and two IP addresses—you must create units with addresses in each of the subnets associated with the VLANs between which you want traffic to be routed. That is, if you have two

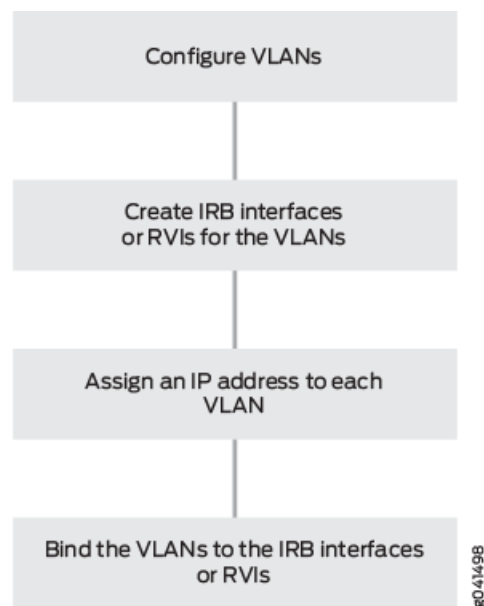
VLANs (for example, VLAN **red** and VLAN **blue**) with corresponding subnets, your interfaces must have a logical unit with an address in the subnet for **red** and a logical unit with an address in the subnet for **blue**. The switch automatically creates direct routes to these subnets and uses these routes to forward traffic between VLANs.

The interface on the switch detects both MAC addresses and IP addresses, then routes data to other Layer 3 interfaces on routers or other switches. These interfaces detect both IPv4 and IPv6 unicast and multicast virtual routing and forwarding (VRF) traffic. Each logical interface can belong to only one routing instance and is further subdivided into logical interfaces, each with a logical interface number appended as a suffix to the names `irb` and `vlan`—for example, `irb.10` and `vlan.10`.

Creating an IRB Interface or RVI

There are four basic steps in creating an IRB interface or RVI as shown in [Figure 2 on page 14](#).

Figure 2: Creating an IRB Interface or RVI



The following explanations correspond to the four steps for creating a VLAN, as depicted in [Figure 2 on page 14](#).

- **Configure VLANs**—Virtual LANs are groups of hosts that communicate as if they were attached to the same broadcast stream. VLANs are created with software and do not require a physical router to forward traffic. VLANs are Layer 2 constructs.
- **Create IRB interfaces or RVIs for the VLANs**—The switch's IRB interfaces and RVIs use Layer 3 logical interfaces (unlike routers, which can use either physical or logical interfaces).

- Assign an IP address to each VLAN—An IRB interface or RVI cannot be activated unless it is associated with a physical interface.
- Bind the VLANs to the logical interfaces—There is a one-to-one mapping between a VLAN and an IRB interface or RVI, which means that only one of these interfaces can be mapped to a VLAN.

For specific instructions for creating an IRB interface, see *Configuring Integrated Routing and Bridging Interfaces (CLI Procedure)*, and for an RVI, see “[Configuring Routed VLAN Interfaces \(CLI Procedure\)](#)” on page 164.

Viewing IRB Interface and RVI Statistics

Some switches automatically track IRB interface and RVI traffic statistics. Other switches allow you to configure tracking. [Table 3 on page 15](#) illustrates the IRB interface- and RVI-tracking capability on various switches.

Table 3: Tracking IRB Interface and RVI Usage

Switch	Input (ingress)	Output (Egress)
EX4300	Automatic	Automatic
EX3200, EX4200	Automatic	–
EX8200	Configurable	Automatic
EX2200, EX3300, EX4500, EX6200	–	–

You can view input (ingress) and output (egress) totals with the following commands:

- For IRB interfaces, use the **show interfaces irb extensive** command. Look at the input and output values in the Transit Statistics field for IRB interface activity values.
- For RVI, use the **show interfaces vlan extensive** command. Look at the input and output values in the Logical Interface Transit Statistics field for RVI activity values.

IRB Interfaces and RVI Functions and Other Technologies

IRB interfaces and RVIs are similar to switch virtual interfaces (SVIs) and bridge-group virtual interfaces (BVI), which are supported on other vendors’ devices. They can also be combined with other functions:

- 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. For more information about VRF, see “[Understanding Virtual Routing Instances on EX Series Switches](#)” on page 25.
- For redundancy, you can combine an IRB interface or RVI with implementations of the Virtual Router Redundancy Protocol (VRRP) in both bridging and virtual private LAN service (VPLS) environments. For more information about VRRP, see *Understanding VRRP on EX Series Switches*.

**Related
Documentation**

- [Understanding Bridging and VLANs on EX Series Switches on page 3](#)

Understanding Private VLANs on EX Series Switches

VLANs limit broadcasts to specified users. Private VLANs (PVLANS) take this concept a step further by limiting communication within a VLAN. PVLANS accomplish this by restricting traffic flows through their member switch ports (which are called *private ports*) so that these ports communicate only with a specified uplink trunk port or with specified ports within the same VLAN. The uplink trunk port or link aggregation group (LAG) is usually connected to a router, firewall, server, or provider network. Each PVLAN typically contains many private ports that communicate only with a single uplink port, thereby preventing the ports from communicating with each other. PVLANS provide Layer 2 isolation between ports within the same VLAN, splitting a broadcast domain into multiple isolated broadcast subdomains and essentially putting secondary VLANs inside a primary VLAN.

Just like regular VLANs, PVLANS are isolated on Layer 2 and require one of the following options to route Layer 3 traffic among the secondary VLANs:

- A promiscuous port connection with a router
- A routed VLAN interface (RVI), which can be configured only on an EX8200 switch or EX8200 Virtual Chassis in the PVLAN



NOTE: To route Layer 3 traffic among secondary VLANs, a PVLAN needs only one of the options mentioned above. If you use an RVI, you can still implement a promiscuous port connection to a router with the promiscuous port set up to handle only traffic that enters and exits the PVLAN.

PVLANS are useful for restricting the flow of broadcast and unknown unicast traffic and for limiting the communication between known hosts. Service providers use PVLANS to keep their customers isolated from each other. Another typical use for a PVLAN is to provide per-room Internet access in a hotel.



NOTE: You can configure a PVLAN to span switches that support PVLANS.

This topic explains the following concepts regarding PVLANS on EX Series switches:

- [Typical Structure and Primary Application of PVLANS on page 17](#)
- [Routing Between Isolated and Community VLANs on page 20](#)
- [PVLANS Use 802.1Q Tags to Identify Packets on page 20](#)
- [PVLANS Use IP Addresses Efficiently on page 20](#)
- [PVLANS Use Four Different Ethernet Switch Port Types on page 20](#)

Typical Structure and Primary Application of PVLANS

The configured PVLAN is the *primary* domain (primary VLAN). Within the PVLAN, you configure *secondary* VLANs, which become subdomains nested within the primary domain. A PVLAN can be configured on a single switch or can be configured to span multiple switches.

Following are the types of domains, interfaces, and ports that you configure within a PVLAN:

- **Primary VLAN**—The primary VLAN of the PVLAN is defined with an 802.1Q tag (VLAN ID) for the complete PVLAN. The primary PVLAN can contain multiple secondary VLANs (one isolated VLAN and multiple community VLANs).
- **Isolated VLAN**—The isolated VLAN is a secondary VLAN nested within the primary VLAN. A primary VLAN can contain only one isolated VLAN. An interface within an isolated VLAN (isolated interface) can forward packets only to a promiscuous port or the PVLAN trunk port. An isolated interface cannot forward packets to another isolated interface; nor can an isolated interface receive packets from another isolated interface. If a customer device needs to have access *only* to a router, the device must be attached to an isolated trunk port.
- **Community VLAN**—A community VLAN is a secondary VLAN nested within the primary VLAN. You can configure multiple community VLANs within a single PVLAN. An interface within a specific community VLAN can establish Layer 2 communications with any other interface that belongs to the same community VLAN. An interface within a community VLAN can also communicate with a promiscuous port or the PVLAN trunk port.
- **Interswitch isolated VLAN**—An interswitch isolated VLAN is a secondary VLAN nested within the primary VLAN. This VLAN is used to forward isolated VLAN traffic from one switch to another through a PVLAN trunk port.
- **Promiscuous port**—A promiscuous port has Layer 2 communications with all the interfaces that are in the PVLAN, regardless of whether the interface belongs to an isolated VLAN or a community VLAN. A promiscuous port is a member of the primary VLAN, but is not included within one of the secondary subdomains. Layer 3 gateways, DHCP servers, and other trusted devices that need to communicate with endpoint devices are typically connected to a promiscuous port.
- **RVI**—On an EX8200 switch or EX8200 Virtual Chassis in the PVLAN, you can optionally configure one RVI for the primary VLAN. When configured, this RVI routes Layer 3 packets received by isolated and community VLAN interfaces.
- **PVLAN trunk link**—The PVLAN trunk link, which is also known as the interswitch link, is required only when a PVLAN is configured to span multiple switches. The PVLAN trunk link connects the multiple switches that compose the PVLAN.

Figure 3 on page 18 shows a PVLAN on a single switch, where the primary VLAN (VLAN 100) contains two community VLANs (VLAN 300 and VLAN 400) and one isolated VLAN (VLAN 50).

Figure 3: Private VLAN on a Single EX Switch





Figure 4 on page 19 shows a PVLAN spanning multiple switches, where the primary VLAN (VLAN 100) contains two community VLANs (VLAN 300 and VLAN 400) and one isolated VLAN (VLAN 200). It also shows that Switches 1 and 2 are connected through an interswitch link (PVLAN trunk link).

Figure 4: PVLAN Spanning Multiple EX Series Switches



Also, the PVLANS shown in [Figure 3 on page 18](#) and [Figure 4 on page 19](#) use a promiscuous port connected to a router as the means to route Layer 3 traffic among the community and isolated VLANs. Instead of using the promiscuous port connected to a router, you can configure an RVI on the switch in [Figure 3 on page 18](#) or one of the switches shown in [Figure 4 on page 19](#), provided that it is an EX200 switch or EX200 Virtual Chassis.

For information about configuring PVLANS on a single switch and on multiple switches, see “[Creating a Private VLAN on a Single EX Series Switch \(CLI Procedure\)](#)” on [page 169](#) and “[Creating a Private VLAN Spanning Multiple EX Series Switches \(CLI Procedure\)](#)” on [page 171](#), respectively. For information about configuring an RVI, see “[Configuring a Routed VLAN Interface in a Private VLAN \(CLI Procedure\)](#)” on [page 173](#).

Routing Between Isolated and Community VLANs

To route Layer 3 traffic between isolated and community VLANs, you must either connect a router to a promiscuous port, as shown in [Figure 3 on page 18](#) and [Figure 4 on page 19](#), or on an EX8200 switch or EX8200 Virtual Chassis in the PVLAN domain, configure an RVI.

If you choose the RVI option, you must configure one RVI for the primary VLAN on one EX8200 switch or EX8200 Virtual Chassis in the PVLAN domain. This RVI serves the entire PVLAN domain regardless of whether the domain includes one or more switches. After you configure the RVI, Layer 3 packets received by the secondary VLAN interfaces are mapped to and routed by the RVI.

When setting up the RVI, you must also enable proxy Address Resolution Protocol (ARP) so that the RVI can handle ARP requests received by the secondary VLAN interfaces.

PVLANS Use 802.1Q Tags to Identify Packets

When packets are marked with a customer-specific 802.1Q tag, that tag identifies ownership of the packets for any switch or router in the network. Sometimes, 802.1Q tags are needed within PVLANS to keep track of packets from different subdomains. [Table 4 on page 20](#) indicates when an 802.1Q tag is needed on the primary VLAN or on secondary VLANs.

Table 4: When VLANs in a PVLAN Need 802.1Q Tags

	On a Single Switch	On Multiple Switches
Primary VLAN	Specify an 802.1Q tag by setting a VLAN ID.	Specify an 802.1Q tag by setting a VLAN ID.
Secondary VLAN	No 802.1Q tag needed.	Specify an 802.1Q tag for each community VLAN by setting a VLAN ID. Specify an 802.1Q tag for an isolation VLAN ID by setting an isolation ID.

PVLANS Use IP Addresses Efficiently

PVLANS provide IP address conservation and efficient allocation of IP addresses. In a typical network, VLANs usually correspond to a single IP subnet. In PVLANS, the hosts in all secondary VLANs belong to the same IP subnet because the subnet is allocated to the primary VLAN. Hosts within the secondary VLAN are assigned IP addresses based on IP subnets associated with the primary VLAN, and their IP subnet masking information reflects that of the primary VLAN subnet.

PVLANS Use Four Different Ethernet Switch Port Types

[Table 5 on page 21](#) summarizes whether or not Layer 2 connectivity exists between the different types of ports within a PVLAN.

Table 5: PVLAN Ports and Layer 2 Connectivity

Port Type To: → From: ↓	Promiscuous	Community	Isolated	PVLAN Trunk	RVI
Promiscuous	Yes	Yes	Yes	Yes	Yes
Community	Yes	Yes—same community only	No	Yes	Yes
Isolated	Yes	No	No	Yes <i>NOTE: This communication is unidirectional.</i>	Yes
PVLAN trunk	Yes	Yes—same community only	Yes <i>NOTE: This communication is unidirectional.</i>	Yes	Yes
RVI	Yes	Yes	Yes	Yes	Yes

As noted in [Table 5 on page 21](#), Layer 2 communication between an isolated port and a PVLAN trunk port is unidirectional. That is, an isolated port can only send packets to a PVLAN trunk port, and a PVLAN trunk port can only receive packets from an isolated port. Conversely, a PVLAN trunk port cannot send packets to an isolated port, and an isolated port cannot receive packets from a PVLAN trunk port.



NOTE: If you enable `no-mac-learning` on a primary VLAN, all isolated VLANs (or the interswitch isolated VLAN) in the PVLAN inherit that setting. However, if you want to disable MAC address learning on any community VLANs, you must configure `no-mac-learning` on each of those VLANs.

Related Documentation

- [Understanding Bridging and VLANs on EX Series Switches on page 3](#)
- [Example: Configuring a Private VLAN on a Single EX Series Switch on page 81](#)
- [Example: Configuring a Private VLAN Spanning Multiple EX Series Switches on page 119](#)

Understanding PVLAN Traffic Flows Across Multiple Switches

This topic illustrates and explains three different traffic flows on a sample multiswitch network configured with a private VLAN (PVLAN). PVLANS restrict traffic flows through their member switch ports (which are called “private ports”) so that they communicate only with a specific uplink trunk port or with specified ports within the same VLAN.


This topic describes:

- [Community VLAN Sending Untagged Traffic on page 22](#)
- [Isolated VLAN Sending Untagged Traffic on page 23](#)
- [PVLAN Tagged Traffic Sent on a Promiscuous Port on page 24](#)

Community VLAN Sending Untagged Traffic

In this scenario, a VLAN in Community-1 of Switch 1 at interface ge-0/0/0 sends untagged traffic. The arrows in [Figure 5 on page 22](#) represent this traffic flow.

Figure 5: Community VLAN Sends Untagged Traffic



In this scenario, the following activity takes place on Switch 1:

- Community-1 VLAN on interface ge-0/0/0: Learning
- pvlan100 on interface ge-0/0/0: Replication
- Community-1 VLAN on interface ge-0/0/12: Receives traffic
- PVLAN trunk port: Traffic exits from ge-1/0/2 and from ae0 with tag 10
- Community-2: Interface receives no traffic
- Isolated VLANs: Interfaces receive no traffic

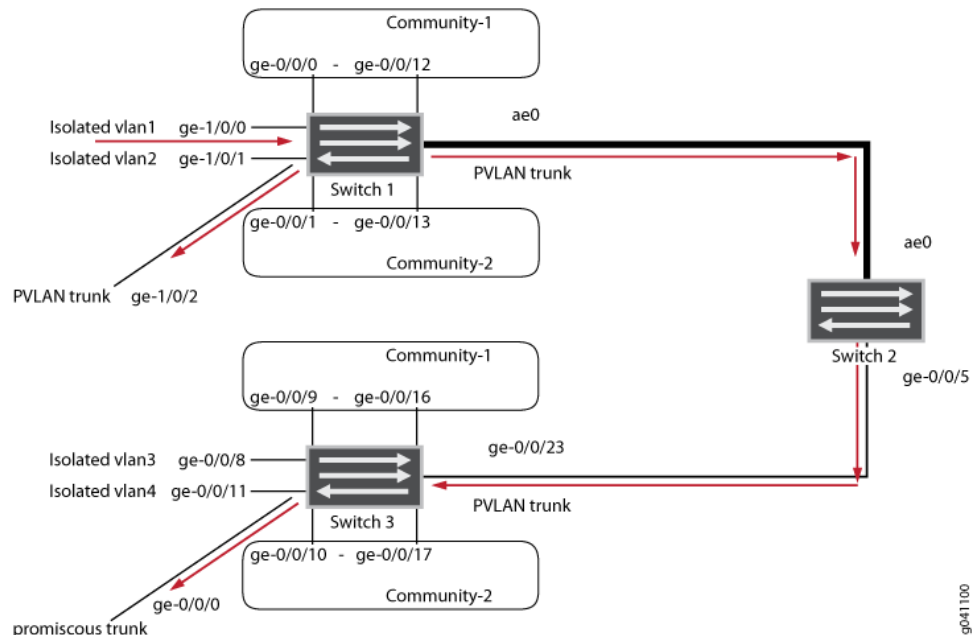
In this scenario, this activity takes place on Switch 3:

- Community-1 VLAN on interface ge-0/0/23 (PVLAN trunk): Learning
- pvlan100 on interface ge-0/0/23: Replication
- Community-1 VLAN on interface ge-0/0/9 and ge-0/0/16: Receives traffic
- Promiscuous trunk port: Traffic exits from ge-0/0/0 with tag 100
- Community-2: Interface receives no traffic
- Isolated VLANs: Interfaces receive no traffic

Isolated VLAN Sending Untagged Traffic

In this scenario, isolated VLAN1 on Switch 1 at interface ge-1/0/0 sends untagged traffic. The arrows in [Figure 6 on page 23](#) represent this traffic flow.

Figure 6: Isolated VLAN Sends Untagged Traffic



In this scenario, the following activity takes place on Switch 1:

- Isolated VLAN1 on interface ge-1/0/0: Learning
- pvlan100 on interface ge-1/0/0: Replication
- Traffic exits from pvlan-trunk ge-1/0/2 and ae0 with tag 50
- Community-1 and Community-2: Interfaces receive no traffic
- Isolated VLANs: Interfaces receive no traffic

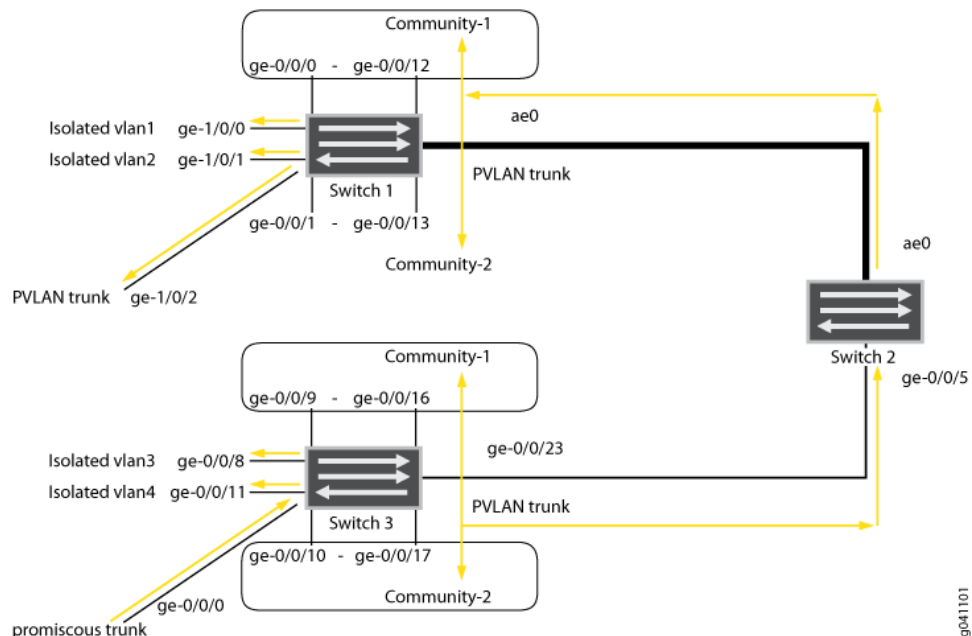
In this scenario, this activity takes place on Switch 3:

- VLAN on interface ge-0/0/23 (PVLAN trunk port): Learning
- pvlan100 on interface ge0/0/23: Replication
- Promiscuous trunk port: Traffic exits from ge-0/0/0 with tag 100
- Community-1 and Community-2: Interfaces receive no traffic
- Isolated VLANs: Receive no traffic

PVLAN Tagged Traffic Sent on a Promiscuous Port

In this scenario, PVLAN tagged traffic is sent on a promiscuous port. The arrows in [Figure 7 on page 24](#) represent this traffic flow.

Figure 7: PVLAN Tagged Traffic Sent on a Promiscuous Port



In this scenario, the following activity takes place on Switch 1:

- pvlan100 VLAN on interface ae0 (PVLAN trunk): Learning
- Community-1, Community-2, and all isolated VLANs on interface ae0: Replication
- VLAN on interface ae0: Replication
- Traffic exits from pvlan-trunk ge-1/0/2 with tag 100
- Community-1 and Community-2: Interfaces receive traffic
- Isolated VLANs: Receive traffic

In this scenario, this activity takes place on Switch 3:

- pvlan100 on interface ge-0/0/0: Learning
- Community-1, Community-2 and all isolated VLANs on interface ge-0/0/0: Replication

- VLAN on interface ge-0/0/0: Replication
- Community-1 and Community-2: Interfaces receive traffic
- Isolated VLANs: Receive traffic

Related Documentation

- [Understanding Private VLANs on EX Series Switches on page 16](#)
- [Example: Configuring a Private VLAN on a Single EX Series Switch on page 81](#)
- [Example: Configuring a Private VLAN Spanning Multiple EX Series Switches on page 119](#)
- [Creating a Private VLAN on a Single EX Series Switch \(CLI Procedure\) on page 169](#)
- [Creating a Private VLAN Spanning Multiple EX Series Switches \(CLI Procedure\) on page 171](#)
- *Understanding Private VLANs*
- *Creating a Private VLAN on a Single Switch*
- *Creating a Private VLAN Spanning Multiple Switches*
- *Example: Configuring a Private VLAN on a Single Switch*
- *Example: Configuring a Private VLAN Spanning Multiple Switches*

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.

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.

EX Series switches support IPv4 and IPv6 unicast and multicast VRF traffic. See *EX Series Switch Software Features Overview* for details on VRF support by switch per Junos OS release.

Related Documentation

- *Understanding Layer 3 Subinterfaces*
- [Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 88](#)
- [Configuring Virtual Routing Instances \(CLI Procedure\) on page 174](#)

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

Multiple VLAN Registration Protocol (MVRP) is a Layer 2 messaging protocol that manages the addition, deletion, and renaming of active virtual LANs, thereby reducing network administrators' time spent on these tasks. Use MVRP on Juniper Networks EX Series Ethernet Switches to dynamically register and unregister active VLANs on trunk interfaces. Using MVRP means that you do not have to manually register VLANs on all connections—that is, you do not need to explicitly bind a VLAN to each trunk interface. With MVRP, you configure a VLAN on one switch interface and the VLAN configuration is distributed through all active switches in the domain.

MVRP is an application protocol of the Multiple Registration Protocol (MRP) and is defined in the IEEE 802.1ak standard. MRP and MVRP replace Generic Attribute Registration Protocol (GARP) and GARP VLAN Registration Protocol (GVRP) and overcome GARP and GVRP limitations.

This topic describes:

- [How MVRP Updates, Creates, and Deletes VLANs on the Switches on page 26](#)
- [MVRP Is Disabled by Default on the Switches on page 26](#)
- [MRP Timers Control MVRP Updates on page 27](#)
- [MVRP Uses MRP Messages to Transmit Switch and VLAN States on page 27](#)
- [Compatibility Issues with Junos OS Releases of MVRP on page 28](#)

How MVRP Updates, Creates, and Deletes VLANs on the Switches

When any MVRP-member VLAN is changed, that VLAN sends a protocol data unit (PDU) to all other MVRP-member active VLANs. The PDU informs the other VLANs which switches and interfaces currently belong to the sending VLAN. This way, all MVRP-member VLANs are always updated with the current VLAN state of all other MVRP-member VLANs. Timers dictate when PDUs can be sent and when switches receiving MVRP PDUs can update their MVRP VLAN information.

In addition to sending PDU updates, MVRP dynamically creates VLANs on member interfaces when a new VLAN is added to any one interface. This way, VLANs created on one member switch are propagated to other member switches as part of the MVRP message exchange process.

To keep VLAN membership information current, MVRP removes switches and interfaces when they become unavailable. Pruning VLAN information has these benefits:

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

MVRP Is Disabled by Default on the Switches

MVRP is disabled by default on the switches and, when enabled, affects only trunk interfaces. Once you enable MVRP, all VLAN interfaces on the switch belong to MVRP

(the default **normal** registration mode) and those interfaces accept PDU messages and send their own PDU messages. To prevent one or more interfaces from participating in MVRP, you can specifically configure an interface to **forbidden** registration mode instead of the default **normal** mode.

VLAN updating, dynamic VLAN configuration through MVRP, and VLAN pruning are all active on trunk interfaces when MVRP is enabled.

MRP Timers Control MVRP Updates

MVRP registration and updates are controlled by timers that are part of the MRP. The timers define when MVRP PDUs can be sent and when MVRP information can be updated on a switch.

The timers are set on a per-interface basis, and on EX Series switches that use Juniper Networks Junos operating system (Junos OS) with support for the Enhanced Layer 2 Software (ELS) configuration style, the timers are also set on a per-switch basis.

On an EX Series switch that uses Junos OS with support for ELS, if the timer value set on an interface level is different from the value set on a switch level, the value on the interface level takes precedence.

The following MRP 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: Unless there is a compelling reason to change the timer settings, leave the default settings in place. Modifying timers to inappropriate values can cause an imbalance in the operation of MVRP.

MVRP Uses MRP Messages to Transmit Switch and VLAN States

MVRP uses MRP messages to register and declare MVRP states for a switch or VLAN and to inform the switching network that a switch or VLAN is leaving MVRP. These messages are communicated as part of the PDU sent by any switch interface to the other switches in the network.

The following MRP messages are communicated for MVRP:

- Empty—MVRP information is not declared and no VLAN is registered.
- In—MVRP information is not declared but a VLAN is registered.
- JoinEmpty—MVRP information is declared but no VLAN is registered.
- JoinIn—MVRP information is declared and a VLAN is registered.

- Leave—MVRP information that was previously declared is withdrawn.
- LeaveAll—Unregister all VLANs on the switch. VLANs must re-register to participate in MVRP.
- New—The MVRP information is new and a VLAN might not be registered yet.

Compatibility Issues with Junos OS Releases of MVRP

Except in Junos OS Releases 11.2 and earlier, MVRP has conformed with IEEE standard 802.1ak and IEEE Draft 802.1Q regarding the inclusion of an extra byte in the protocol data units (PDUs) sent and received by MVRP. [Table 6 on page 28](#) outlines the MVRP versions and whether or not each version includes the extra byte in the PDU.

[Table 6 on page 28](#) also labels each MVRP version with a scenario number, which is used throughout the remainder of this discussion for brevity.

Table 6: Junos OS MVRP Versions and Inclusion of Extra Byte in PDU

MVRP in Junos OS Releases 11.2 and Earlier For EX Series Switches That Do Not Support Enhanced Layer 2 Software (ELS) Configuration Style	MVRP in Junos OS Releases 11.3 and Later For EX Series Switches That Do Not Support ELS	MVRP in Junos OS Releases 13.2 and Later For EX Series Switches With Support For ELS
Scenario 1	Scenario 2	Scenario 3
Includes extra byte in the PDU	By default, does not include extra byte in the PDU	By default, includes extra byte in the PDU

As a result of the non-conformance of Releases 11.2 and earlier and changes in the standards with regard to the extra byte, a compatibility issue exists between some of the Junos OS versions of MVRP. This issue can result in some versions of MVRP not recognizing PDUs without the extra byte.

To address this compatibility issue, the MVRP versions described in scenarios 2 and 3 include the ability to control whether or not the PDU includes the extra byte. Before using these controls, however, you must understand each scenario that applies to your environment and plan carefully so that you do not inadvertently create an additional compatibility issue between the MVRP versions in scenarios 2 and 3.

[Table 7 on page 28](#) provides a summary of environments that include the various MVRP scenarios and whether or not a particular environment requires you to take action.

Table 7: MVRP Environments and Description of Required Actions

Environment	Action Required?	Action Description
Includes MVRP versions in scenario 1 only	No	—
Includes MVRP versions in scenario 2 only	No	—
Includes MVRP versions in scenario 3 only	No	—

Table 7: MVRP Environments and Description of Required Actions (*continued*)

Environment	Action Required?	Action Description
Includes MVRP versions in scenarios 1 and 2. MVRP version in scenario 2 is in its default state.	Yes	On switches running MVRP version in scenario 2, use the add-attribute-length-in-pdu statement. For more information, see “Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)” on page 176 .
Includes MVRP versions in scenarios 1 and 3. MVRP version in scenario 3 is in its default state.	No	—
Includes MVRP versions in scenarios 2 and 3, and both versions are in their respective default states	Yes	Do one of the following: <ul style="list-style-type: none"> On switches running MVRP version in scenario 2, use the add-attribute-length-in-pdu statement. For more information, see “Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)” on page 176. On switches running MVRP version in scenario 3, use the no-attribute-length-in-pdu statement. For more information, see Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure).

You can determine whether the switches in your network are running incompatible versions of MVRP by issuing the **show mvrp statistics** command. For more information on diagnosing and correcting this MVRP compatibility situation, see [“Configuring Multiple VLAN Registration Protocol \(MVRP\) \(CLI Procedure\)” on page 176](#) or [Configuring Multiple VLAN Registration Protocol \(MVRP\) \(CLI Procedure\)](#).

Related Documentation

- [Understanding Bridging and VLANs on EX Series Switches on page 3](#)
- [Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 92](#)
- [Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches](#)

Understanding MAC Notification on EX Series Switches

Juniper Networks EX Series Switches track clients on a network by storing Media Access Control (MAC) addresses in the Ethernet switching table on the switch. When switches learn or unlearn a MAC address, SNMP notifications can be sent to the network management system at regular intervals to record the addition or removal of the MAC address. This process is known as MAC notification.

The MAC Notification MIB controls MAC notification for the network management system. For general information on the MAC Notification MIB, see the [Junos OS Network Management Configuration Guide](#).

The MAC notification interval defines how often these SNMP notifications are sent to the network management system. The MAC notification interval works by tracking all of the MAC address additions or removals on the switch over a period of time and then sending all of the tracked MAC address additions or removals to the network management

server at the end of the interval. For instance, if the MAC notification interval is set to 10, all of the MAC address addition and removal SNMP notifications are sent to the network management system every 10 seconds.

Enabling MAC notification allows users to monitor the addition and removal of MAC addresses from the Ethernet switching table remotely using a network management system. The advantage of setting a high MAC notification interval is that the amount of network traffic is reduced because updates are sent less frequently. The advantage of setting a low MAC notification interval is that the network management system is better synchronized with the switch.

MAC notification is disabled by default. When MAC notification is enabled, the default MAC notification interval is 30 seconds.

**Related
Documentation**

- [Configuring MAC Notification \(CLI Procedure\) on page 175](#)
- [Configuring SNMP \(J-Web Procedure\)](#)

Understanding MAC Address Aging

Juniper Networks EX Series Ethernet Switches store MAC addresses in the Ethernet switching table, also called the *MAC table*. When the aging time for a MAC address in the table expires, the address is removed.

If your switch runs Juniper Networks Junos operating system (Junos OS) for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style, you can configure the MAC table aging time on all VLANs on the switch. If your switch runs Junos OS that does not support ELS, you can configure the MAC table aging time on all VLANs on the switch or on specified VLANs, as well as configure aging time to be unlimited, either on all VLANs or on specified VLANs, so that MAC addresses never age out of the 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 on which the traffic was received and the time when the address was learned.

When the switch receives traffic on an interface, it searches the Ethernet switching table for the MAC address of the destination. If the MAC address is not found, the traffic is flooded out all of the other interfaces associated with the VLAN. For example, if traffic is received on an interface that is associated with VLAN v-10 and there is no entry in the Ethernet switching table for VLAN v-10 (the Ethernet switching table is organized by VLAN), then the traffic is flooded to all access and trunk interfaces that are members of VLAN v-10.

Flooding allows the switch to learn about destinations that are not yet in its Ethernet switching table. If a particular destination MAC address is not in the Ethernet switching table, the switch floods the traffic to all interfaces except the interface on which it was received. When the destination node receives the flooded traffic, it sends an acknowledgment packet back to the switch, allowing the switch to learn the MAC address of the node and to add the address to its Ethernet switching table.

The switch uses a mechanism called aging to keep 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 that is in its Ethernet switching table, it updates the timestamp of that MAC address. A timer on the switch periodically checks the timestamp, and if the MAC address of a node is older than the value set, the switch removes that MAC address from the Ethernet switching table. This aging process ensures that the switch tracks only active MAC addresses on the network and that it is able to flush out from the Ethernet switching table MAC addresses that are no longer available.

You configure how long MAC addresses remain in the Ethernet switching table by:

- (On switches that run Junos OS with support for the ELS configuration style) Using the **global-mac-table-aging-time** statement in the **[edit protocols l2-learning]** hierarchy.
- (On switches that run Junos OS that does not support ELS) Using the **mac-table-aging-time** statement in either the **[edit ethernet-switching-options]** or the **[edit vlans]** hierarchy, depending on whether you want to configure it for the entire switch or only for specific VLANs.

For example, in a topology with EX switches that run Junos OS that does not support ELS, if you have a printer VLAN, you might choose to configure the aging time for that VLAN to be considerably longer than for other VLANs so that MAC addresses of printers on this VLAN age out less frequently. Because the MAC addresses remain in the table, even if a printer has been idle for some time before traffic arrives for it, the switch still finds the MAC address and does not need to flood the traffic to all other interfaces.

Similarly, in a data center environment where the list of servers connected to the switch is fairly stable, you might choose to increase MAC address aging time, or even set it to unlimited, to increase the efficiency of the utilization of network bandwidth by reducing flooding.

**Related
Documentation**

- [Configuring MAC Table Aging \(CLI Procedure\) on page 166](#)
- [Configuring MAC Table Aging \(CLI Procedure\)](#)
- [Controlling Authentication Session Timeouts \(CLI Procedure\)](#)

Understanding MAC Address Assignment in an EX Series Switch

This topic describes MAC address assignment for interfaces on standalone Juniper Networks EX Series Ethernet Switches. For information regarding MAC address assignments in a Virtual Chassis, see *Understanding MAC Address Assignment on a Virtual Chassis*.

MAC addresses are used to identify network devices at Layer 2. Because all Layer 2 traffic decisions are based on an interface's MAC address, understanding MAC address assignment is important to understanding how network traffic is forwarded and received by the switch. For additional information on how a network uses MAC addresses to forward and receive traffic, see ["Understanding Bridging and VLANs on EX Series Switches" on page 3](#).

A MAC address comprises six groups of two hexadecimal digits, with each group separated from the next group by a colon—for instance, aa:bb:cc:dd:ee:00. The first five groups of hexadecimal digits are derived from the switch and are the same for all interfaces on the switch.

The assignment of a unique MAC address to each network interface helps ensure that functions that require MAC address differentiation—such as redundant trunk groups (RTGs), Link Aggregation Control Protocol (LACP), and general monitoring functions—can properly function.

On switches that use line cards, this MAC addressing scheme differentiates the Layer 2 interfaces on different line cards in the switch.

For EX Series switches, the first five groups of hexadecimal digits are determined when the switch is manufactured. The switch then assigns a unique MAC address to each interface by assigning a unique identifier as the last group of hexadecimal digits. The assignment depends on how the interface is configured. The switch uses a different pattern to distinguish between an interface that is configured as any of a routed VLAN interface (RVI), a virtual management Ethernet (VME) interface, or an aggregated Ethernet interface or is not configured as any of an RVI, a VME, or as an aggregated Ethernet interface.

For aggregated Ethernet interfaces, the MAC address assignment remains constant regardless of whether the configuration of the interface is Layer 2 or Layer 3.



NOTE: In Junos OS Release 11.3 and later releases through Release 12.1, the MAC address assignment for aggregated Ethernet interfaces changes if the interface is changed from Layer 2 to Layer 3 or the reverse. Starting with Junos Release 12.2, the MAC address assignment for aggregated Ethernet interfaces remains constant regardless of whether the interface is Layer 2 or Layer 3.



NOTE: Prior to Junos OS Release 11.3, MAC addresses for Layer 2 interfaces could be shared between interfaces and RVIs on different line cards in the same switch. However, if you upgrade from Junos OS Release 11.2 or earlier to Junos OS Release 11.3 or later on a switch that supports line cards, the MAC addresses of these interfaces will change.

MAC addresses are assigned to interfaces automatically—no user configuration is possible or required. You can view MAC addresses assigned to interfaces using the **show interfaces** command.

Related Documentation

- *EX Series Switches Interfaces Overview*

Understanding Edge Virtual Bridging for Use with VEPA Technology

Servers using virtual Ethernet port aggregator (VEPA) do not send packets directly from one virtual machine (VM) to another. Instead, the packets are sent to virtual bridges on an adjacent switch for processing. EX Series switches use edge virtual bridging (EVB) as a virtual bridge to return the packets on the same interface that delivered the packets.

- [What Is EVB? on page 33](#)
- [What Is VEPA? on page 33](#)
- [Why Use VEPA Instead of VEB? on page 33](#)
- [How Does EVB Work? on page 33](#)
- [How Do I Implement EVB? on page 34](#)

What Is EVB?

EVB is a software capability on a switch running Junos OS that allows multiple virtual machines to communicate with each other and with external hosts in the Ethernet network environment.

What Is VEPA?

VEPA is a software capability on a server that collaborates with an adjacent, external switch to provide bridging support between multiple virtual machines and external networks. The VEPA collaborates with the adjacent switch by forwarding all VM-originated frames to the adjacent switch for frame processing and frame relay (including hairpin forwarding) and by steering and replicating frames received from the VEPA uplink to the appropriate destinations.

Why Use VEPA Instead of VEB?

Even though virtual machines are capable of sending packets directly to one another with a technology called virtual Ethernet bridging (VEB), you typically want to use physical switches for switching because VEB uses expensive server hardware to accomplish the task. Instead of using VEB, you can install VEPA on a server to offload switching functionality to an adjacent, less expensive physical switch. Additional advantages of using VEPA include:

- VEPA reduces complexity and allows higher performance at the server.
- VEPA takes advantage of the physical switch's security and tracking features.
- VEPA provides visibility of inter-virtual-machine traffic to network management tools designed for an adjacent bridge.
- VEPA reduces the amount of network configuration required by server administrators, and as a consequence, reduces work for the network administrator.

How Does EVB Work?

EVB uses two protocols, Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP) and Edge Control Protocol (ECP), to program policies for each individual

virtual switch instance—specifically, EVB maintains the following information for each VSI instance:

- VLAN ID
- VSI type
- VSI type version
- MAC address of the server

VDP is used by the VEPA server to propagate VSI information to the switch. This allows the switch to program policies on individual VSIs and supports virtual machine migration by implementing logic to preassociate a VSI with a particular interface.

ECP is a Link Layer Discovery Protocol (LLDP)-like transport layer that allows multiple upper layer protocols to send and receive protocol data units (PDUs). ECP improves upon LLDP by implementing sequencing, retransmission and an ack mechanism, while at the same time remaining lightweight enough to be implemented on a single-hop network. ECP is implemented in an EVB configuration when you configure LLDP on interfaces that you have configured for EVB. That is, you configure LLDP, not ECP.

How Do I Implement EVB?

You can configure EVB on a switch when that switch is adjacent to a server that includes VEPA technology. In general, this is what you do to implement EVB:

- The network manager creates a set of VSI types. Each VSI type is represented by a VSI type ID and a VSI version--the network manager can deploy one or more VSI versions at any given time.
- The VM manager configures VSI (which is a virtual station interface for a VM that is represented by a MAC address and VLAN ID pair) . To accomplish this, the VM manager queries available VSI type IDs (VTIDs) and creates a VSI instance consisting of a VSI Instance ID and the chosen VTID. This instance is known as VTDB and contains a VSI manager ID, a VSI type ID, a VSI version, and a VSI instance ID.

Related Documentation

- [Understanding Bridging and VLANs on EX Series Switches on page 3](#)
- [Example: Configuring Edge Virtual Bridging for Use with VEPA Technology on page 134](#)

Understanding Ethernet Ring Protection Switching Functionality

- [Acronyms on page 35](#)
- [Ring Nodes on page 35](#)
- [Ring Node States on page 35](#)
- [Failure Detection on page 36](#)
- [Logical Ring on page 36](#)
- [FDB Flush on page 36](#)
- [Traffic Blocking and Forwarding on page 36](#)

- [RAPS Message Blocking and Forwarding on page 36](#)
- [Dedicated Signaling Control Channel on page 38](#)
- [RAPS Message Termination on page 38](#)
- [Multiple Rings on page 38](#)
- [Node ID on page 38](#)
- [Bridge Domains with the Ring Port \(MX Series Routers Only\) on page 38](#)

Acronyms

The following acronyms are used in the discussion about Ethernet ring protection switching:

- MA—Maintenance association
- MEP—Maintenance association end point
- OAM—Operations, administration, and management (Ethernet ring protection switching uses connectivity fault management daemon)
- FDB—MAC forwarding database
- STP—Spanning Tree Protocol
- RAPS—Ring automatic protection switching
- WTR—Wait to restore
- RPL—Ring protection link

Ring Nodes

Multiple nodes are used to form a ring. There are two different node types:

- Normal node—The node has no special role on the ring.
- RPL owner node—The node owns the RPL and blocks or unblocks traffic over the RPL. This node also initiates the RAPS message.

Ring Node States

There are three different states for each node of a specific ring:

- init—Not a participant of a specific ring.
- idle—No failure on the ring; the node is performing normally. For a normal node, traffic is unblocked on both ring ports. For the RPL owner, traffic is blocked on the ring port that connects to the RPL and unblocked on the other ring port.
- protection—A failure occurred on the ring. For a normal node, traffic is blocked on the ring port that connects to the failing link and unblocked on working ring ports. For the RPL owner, traffic is unblocked on both ring ports if they connect to non-failure links.

There can be only one RPL owner for each ring. The user configuration must guarantee this, because the APS protocol cannot check this.

Failure Detection

Ethernet ring operation depends on quick and accurate failure detection. The failure condition *signal failure (SF)* is supported. For SF detection, an Ethernet continuity check MEP must be configured for each ring link. For fast protection switching, a 10-ms transmission period for this MEP group is supported. OAM monitors the MEP group's MA and reports SF or SF clear events to the Ethernet ring control module. For this MEP group, the action profile must be configured to update the interface device IFF_LINKDOWN flag. OAM updates the IFF_LINKDOWN flag to notify the Ethernet ring control module.

Logical Ring

This feature currently supports only the physical ring, which means that two adjacent nodes of a ring must be physically connected and the ring must operate on the physical interface, not the VLAN.

FDB Flush

When ring protection switching occurs, normally an *FDB flush* is executed. The Ethernet ring control module uses the same mechanism as the STP to trigger the FDB flush. The Ethernet ring control module controls the ring port physical interface's default STP index to execute the FDB flush.

Traffic Blocking and Forwarding

Ethernet ring control uses the same mechanism as the STP to control forwarding or discarding of user traffic. The Ethernet ring control module sets the ring port physical interface default STP index state to forwarding or discarding in order to control user traffic.

RAPS Message Blocking and Forwarding

The router or switch treats the ring automatic protection switching (RAPS) message the same as it treats user traffic for forwarding RAPS messages between two ring ports. The ring port physical interface default STP index state also controls forwarding RAPS messages between the two ring ports. Other than forwarding RAPS messages between the two ring ports, as shown in [Figure 8 on page 36](#), the system also needs to forward the RAPS message between the CPU (Ethernet ring control module) and the ring port. This type of forwarding does not depend on the ring port physical interfaces' STP index state. The RAPS message is always sent by the router or switch through the ring ports, as shown in [Figure 9 on page 37](#). A RAPS message received from a discarding ring port is sent to the Ethernet ring control module, but is not sent to the other ring port.

Figure 8: Protocol Packets from the Network to the Router





Figure 9: Protocol Packets from the Router or Switch to the Network

Juniper Networks switches and Juniper Networks routers use different methods to achieve these routes.

The switches use forwarding database entries to direct the RAPS messages. The forwarding database entry (keyed by the RAPS multicast address and VLAN) has a composite next hop associated with it—the composite next hop associates the two ring interfaces with the forwarding database entry and uses the split horizon feature to prevent sending the packet out on the interface that it is received on. This is an example of the forwarding database entry relating to the RAPS multicast MAC (a result of the **show ethernet-switching table detail** command):

```
VLAN: v1, Tag: 101, MAC: 01:19:a7:00:00:01, Interface: ERP
Interfaces:      ge-0/0/9.0, ge-0/0/3.0
Type: Static
Action: Mirror
Nexthop index: 1333
```

The routers use an implicit filter to achieve ERP routes. Each implicit filter binds to a bridge domain. Therefore, the east ring port control channel and the west ring port control channel of a particular ring instance must be configured to the same bridge domain. For each ring port control channel, a filter term is generated to control RAPS message forwarding. The filter number is the same as the number of bridge domains that contain the ring control channels. If a bridge domain contains control channels from multiple rings, the filter related to this bridge domain will have multiple terms and each term will relate to a control channel. The filter has command parts and control-channel related parts, as follows:

- Common terms:
 - term 1: if [Ethernet type is not OAM Ethernet type (0x8902)]
 { accept packet }
 - term 2: if [source MAC address belongs to this bridge]
 { drop packet, our packet loop through the ring and come back to home }
 - term 3: if [destination is the RAPS PDU multicast address(0x01,0x19,0xa7,0x00,0x00,0x01) AND[ring port STP status is DISCARDING]]
 { send to CPU }
 - Control channel related terms:
 - if [destination is the RAPS PDU multicast address(0x01,0x19,0xa7,0x00,0x00,0x01) AND[ring port STP status is FORWARDING] AND [Incoming interface IFL equal to control channel IFL]]
 { send packet to CPU and send to the other ring port }
- default term: accept packet.

Dedicated Signaling Control Channel

For each ring port, a dedicated signaling control channel with a dedicated VLAN ID must be configured. In Ethernet ring configuration, only this control logical interface is configured and the underlying physical interface is the physical ring port. Each ring requires that two control physical interfaces be configured. These two logical interfaces must be configured in a bridge domain for routers (or the same VLAN for switches) in order to forward RAPS protocol data units (PDUs) between the two ring control physical interfaces. If the router control channel logical interface is not a trunk port, only control logical interfaces will be configured in ring port configuration. If this router control channel logical interface is a trunk port, in addition to the control channel logical interfaces, a dedicated VLAN ID must be configured for routers. For switches, always specify either a VLAN name or VLAN ID for all links.

RAPS Message Termination

The RAPS message starts from the originating node, travels through the entire ring, and terminates in the originating node unless a failure is present in the ring. The originating node must drop the RAPS message if the source MAC address in the RAPS message belongs to itself. The source MAC address is the node's node ID.

Multiple Rings

The Ethernet ring control module supports multiple rings in each node (two logical interfaces are part of each ring). However, interconnection of multiple rings is not supported in this release. The interconnection of two rings means that two rings may share the same link or share the same node.

Node ID

For each node in the ring, a unique *node ID* identifies each node. The node ID is the node's MAC address.

For routers only, you can configure this node ID when configuring the ring on the node or automatically select an ID such as STP. In most cases, you will not configure this and the router will select a node ID, like STP does. It should be the manufacturing MAC address. The ring node ID should not be changed, even if you change the manufacturing MAC address. Any MAC address can be used if you make sure each node in the ring has a different node ID. The node ID on switches is selected automatically and is not configurable.

Bridge Domains with the Ring Port (MX Series Routers Only)

On the routers, the protection group is seen as an abstract logical port that can be configured to any bridge domain. Therefore, if you configure one ring port or its logical interface in a bridge domain, you must configure the other related ring port or its logical interface to the same bridge domain. The bridge domain that includes the ring port acts as any other bridge domain and supports the IRB Layer 3 interface.

- Related Documentation**
- [Ethernet Ring Protection Switching Overview on page 39](#)
 - [Configuring Ethernet Ring Protection Switching](#)

- [Example: Ethernet Ring Protection Switching Configuration on MX Routers](#)
- [Ethernet Interfaces Feature Guide for Routing Devices](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches on page 140](#)
- [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 190](#)

Ethernet Ring Protection Switching Overview

Ethernet ring protection switching (ERPS) helps achieve high reliability and network stability. Links in the ring will never form loops that fatally affect the network operation and services availability. The basic idea of an Ethernet ring is to use one specific link to protect the whole ring. This special link is called a *ring protection link (RPL)*. If no failure happens in other links of the ring, the RPL blocks the traffic and is not used. The RPL is controlled by a special node called an *RPL owner*. There is only one RPL owner in a ring. The RPL owner is responsible for blocking traffic over the RPL. Under ring failure conditions, the RPL owner is responsible for unblocking traffic over the RPL. A ring failure results in protection switching of the RPL traffic. An automatic protection switching (APS) protocol is used to coordinate the protection actions over the ring. Protection switching blocks traffic on the failed link and unblocks the traffic on the RPL. When the failure clears, revertive protection switching blocks traffic over the RPL and unblocks traffic on the link on which the failure is cleared.

The following standards provide detailed information on Ethernet ring protection switching:

- IEEE 802.1Q - 1998
- IEEE 802.1D - 2004
- IEEE 802.1Q - 2003
- Draft ITU-T Recommendation G.8032/Y.1344, *Ethernet Ring protection switching*
- ITU-T Y.1731, *OAM functions and mechanisms for Ethernet-based networks*

For additional information on configuring Ethernet ring protection switching on EX Series switches, see [“Example: Configuring Ethernet Ring Protection Switching on EX Series Switches” on page 140](#).

For additional information on configuring Ethernet ring protection switching on MX Series routers, see the *Layer 2 Configuration Guide* for a complete example of Ethernet rings and information about STP loop avoidance and prevention.

Related Documentation

- [Understanding Ethernet Ring Protection Switching Functionality on page 34](#)
- [Configuring Ethernet Ring Protection Switching](#)
- [Example: Ethernet Ring Protection Switching Configuration on MX Routers](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches on page 140](#)

- *Ethernet Interfaces Feature Guide for Routing Devices*

CHAPTER 2

Q-in-Q Tunneling

- [Understanding Q-in-Q Tunneling on EX Series Switches on page 41](#)

Understanding Q-in-Q Tunneling on EX Series Switches



NOTE: This topic is not applicable to Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Understanding Q-in-Q Tunneling on EX Series Switches*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

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 operating system (Junos OS) implementation of Q-in-Q tunneling supports the IEEE 802.1ad standard.

This topic describes:

- [How Q-in-Q Tunneling Works on page 41](#)
- [Disabling MAC Address Learning on page 42](#)
- [Mapping C-VLANs to S-VLANs on page 42](#)
- [Routed VLAN Interfaces on Q-in-Q VLANs on page 44](#)
- [Limitations for Q-in-Q Tunneling on page 44](#)

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 43](#)
- [Many-to-One Bundling on page 43](#)
- [Mapping a Specific Interface on page 43](#)

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.

Q-in-Q tunneling supports only two VLAN tags.

Related Documentation

- [Understanding Bridging and VLANs on EX Series Switches on page 3](#)
- [Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 103](#)
- [Configuring Q-in-Q Tunneling \(CLI Procedure\) on page 180](#)

CHAPTER 3

Layer 2 Protocol Tunneling

- [Understanding Layer 2 Protocol Tunneling on EX Series Switches on page 45](#)

Understanding Layer 2 Protocol Tunneling on EX Series Switches

Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud 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 45](#)
- [How L2PT Works on page 46](#)
- [L2PT Basics on EX Series Switches on page 48](#)

Layer 2 Protocols Supported by L2PT on EX Series Switches

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

- 802.1X authentication
- 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM)



NOTE: If you enable L2PT for untagged OAM LFM (Operation, Administration, and Maintenance of link fault management) packets, do not configure link fault management (LFM) on the corresponding access interface.

- Cisco Discovery Protocol (CDP)
- Ethernet local management interface (E-LMI)
- MVRP VLAN Registration Protocol (MVRP)
- Link Aggregation Control Protocol (LACP)



NOTE: If you enable L2PT for untagged LACP packets, do not configure Link Aggregation Control Protocol (LACP) on the corresponding access interface.

- Link Layer Discovery Protocol (LLDP)
- Multiple MAC Registration Protocol (MMRP)
- Multiple VLAN Registration Protocol (MVRP)
- Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP)
- Unidirectional Link Detection (UDLD)
- VLAN Spanning Tree Protocol (VSTP)
- VLAN Trunking Protocol (VTP)




NOTE: CDP, UDLD, and VTP cannot be configured on EX Series switches. L2PT does, however, tunnel CDP, UDLD, 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. This process is illustrated in [Figure 10 on page 47](#).

Figure 10: L2PT Example



L2PT supports tunneling of STP, LLDP, CDP and VTP control PDUs across the service provider network. The PE device identifies the Layer 2 control protocols by their encapsulated MAC address. The destination MAC address used by different protocols is listed in [Table 8 on page 47](#):

Table 8: Protocol Destination MAC Addresses

Protocol	Ethernet Encapsulation	MAC Address
802.1X	Ether-II	01:80:C2:00:00:03
802.3ah	Ether-II	01:80:C2:00:00:02
Cisco Discovery Protocol (CDP)	SNAP	01:00:0C:CC:CC:CC
Ethernet local management interface (E-LMI)	Ether-II	01:80:C2:00:00:07
MVRP VLAN Registration Protocol (MVRP)	Ether-II	01:80C2:00:00:21
Link Aggregation Control Protocol (LACP)	Ether-II	01:80:C2:00:00:02
Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP)	SNAP	01:80:C2:00:00:21
Link Layer Discovery Protocol (LLDP)	Ether-II	01:80:0C:00:00:0E
Multiple MAC Registration Protocol (MMRP)	Ether-II	01:80:C2:00:00:0E
Unidirectional Link Detection (UDLD)	SNAP	01:00:0C:CC:CC:CC
VLAN Spanning Tree Protocol (VSTP)	SNAP	01:00:0C:CC:CC:CD
VLAN Trunking Protocol (VTP)	SNAP	01:00:0C:CC:CC:CC

When a PE device receives a Layer 2 control PDU from any of the customer PE devices, it changes the destination MAC address to 01:00:0C:CD:CD:D0. The modified packet is then sent to the provider network. All devices on the provider network treat these packets as multicast Ethernet packets and deliver them to all PE devices for the customer. The egress PE devices receive all the control PDUs with the same MAC address (01:00:0C:CD:CD:D0). Then they identify the packet type by doing deeper packet inspection and replace the destination MAC address 01:00:0C:CD:CD:D0 with the appropriate destination address. The modified PDUs are sent out to the customer PE devices, thus ensuring the Layer 2 control PDUs are delivered, in their original state, across the provider network. The L2PT protocol is valid for all types of packets (untagged, tagged, and Q-in-Q tagged).

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**. L2PT PDUs are flooded to all trunk and access ports within a given S-VLAN.



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 or priority-tagged Layer 2 control PDUs to be tunneled, then you must configure the switch to map untagged and priority-tagged packets to an L2PT-enabled VLAN. For more information on assigning untagged and priority-tagged packets to VLANs, see [“Understanding Q-in-Q Tunneling on EX Series Switches” on page 41](#) and [“Configuring Q-in-Q Tunneling \(CLI Procedure\)” on page 180](#).

Related Documentation

- [Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 107](#)
- [Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 103](#)

CHAPTER 4

Redundant Trunk Groups

- [Understanding Redundant Trunk Links on page 49](#)

Understanding Redundant Trunk Links

In a typical enterprise network composed of distribution and access layers, a redundant trunk link provides a simple solution for network recovery when a trunk port on a switch goes down. In that case, traffic is routed to another trunk port, keeping network convergence time to a minimum.

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 spanning-tree protocol 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 *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 switches on the secondary link.

Rapid Spanning Tree Protocol (RSTP) is enabled by default on the switches to create a loop-free topology, but an interface is not allowed to be in both a redundant trunk group and in a spanning-tree protocol topology at the same time. You must disable RSTP on an interface if a redundant trunk group is configured on that interface. For example, in [Figure 11 on page 50](#), in addition to disabling RSTP on the Switch 3 interfaces, you must also disable RSTP on the Switch 1 and Switch 2 interfaces connected to Switch 3. Spanning-tree protocols can, however, continue operating on other interfaces on those switches—for example on the link between Switch 1 and Switch 2.

[Figure 11 on page 50](#) 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 11: Redundant Trunk Group, Link 1 Active





Figure 12 on page 50 illustrates how the redundant trunk link topology works when the primary link goes down.

Figure 12: Redundant Trunk Group, Link 2 Active



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

- Related Documentation**
- [Example: Configuring Redundant Trunk Links for Faster Recovery on page 112](#)
 - *Example: Configuring Redundant Trunk Links for Faster Recovery*

CHAPTER 5

Proxy ARP

- [Understanding Proxy ARP on EX Series Switches on page 53](#)

Understanding Proxy ARP on EX Series Switches

You can configure proxy Address Resolution Protocol (ARP) on your Juniper Networks EX Series Ethernet Switch to enable the switch to respond to ARP queries for network addresses by offering its own Ethernet media access control (MAC) address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

Proxy ARP is useful in situations where hosts are on different physical networks and you do not want to use subnet masking. Because ARP broadcasts are not propagated between hosts on different physical networks, hosts will not receive a response to their ARP request if the destination is on a different subnet. Enabling the switch to act as an ARP proxy allows the hosts to transparently communicate with each other through the switch. Proxy ARP can help hosts on a subnet reach remote subnets without your having to configure routing or a default gateway.

- [What Is ARP? on page 53](#)
- [Proxy ARP Overview on page 53](#)
- [Best Practices for Proxy ARP on EX Series Switches on page 54](#)

What Is ARP?

Ethernet LANs use ARP to map Ethernet MAC addresses to IP addresses. Each device maintains a cache containing a mapping of MAC addresses to IP addresses. The switch maintains this mapping in a cache that it consults when forwarding packets to network devices. If the ARP cache does not contain an entry for the destination device, the host (the DHCP client) broadcasts an ARP request for that device's address and stores the response in the cache.

Proxy ARP Overview

When proxy ARP is enabled, if the switch receives an ARP request for which it has a route to the target (destination) IP address, the switch responds by sending a proxy ARP reply packet containing its own MAC address. The host that sent the ARP request then sends its packets to the switch, which forwards them to the intended host.



NOTE: For security reasons, the source address in an ARP request must be on the same subnet as the interface on which the ARP request is received.

You can configure proxy ARP for each interface. You can also configure proxy ARP for an integrated routing and bridging (IRB) interface named `irb` or a routed VLAN interface (RVI) named `vlan`. (On EX Series switches that use Juniper Networks Junos operating system (Junos OS) with support for the Enhanced Layer 2 Software (ELS) configuration style, the feature is known as an IRB interface. On EX Series switches that use Junos OS that does not support ELS, the feature is known as an RVI.)

EX Series switches support two modes of proxy ARP, restricted and unrestricted. Both modes require that the switch have an active route to the destination address of the ARP request.

- **Restricted**—The switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are on the same subnet. In this mode, hosts on the same subnet communicate without proxy ARP. We recommend that you use this mode on the switch.
- **Unrestricted**—The switch responds to all ARP requests for which it has a route to the destination. This is the default mode (because it is the default mode in Juniper Networks Junos operating system (Junos OS) configurations other than those on the switch). We recommend using restricted mode on the switch.

Best Practices for Proxy ARP on EX Series Switches

We recommend these best practices for configuring proxy ARP on the switches:

- Set proxy ARP on the interfaces that you want, including IRB interfaces or RVIs, to restricted mode.
- If you set proxy ARP to unrestricted, disable gratuitous ARP requests on each interface enabled for proxy ARP.

Related Documentation

- [Example: Configuring Proxy ARP on an EX Series Switch on page 117](#)
- [Configuring Proxy ARP \(CLI Procedure\)](#)

PART 2

Configuration

- [Configuration Examples on page 57](#)
- [Configuration Tasks on page 157](#)
- [Configuration Statements on page 193](#)

CHAPTER 6

Configuration Examples

- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 57](#)
- [Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 65](#)
- [Example: Connecting an Access Switch to a Distribution Switch on page 72](#)
- [Example: Configuring a Private VLAN on a Single EX Series Switch on page 81](#)
- [Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 88](#)
- [Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 92](#)
- [Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 103](#)
- [Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 107](#)
- [Example: Configuring Redundant Trunk Links for Faster Recovery on page 112](#)
- [Example: Configuring Proxy ARP on an EX Series Switch on page 117](#)
- [Example: Configuring a Private VLAN Spanning Multiple EX Series Switches on page 119](#)
- [Example: Configuring Edge Virtual Bridging for Use with VEPA Technology on page 134](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches on page 140](#)

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



NOTE: This example uses Junos OS for EX Series switches that does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*

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 58](#)
- [Overview and Topology on page 58](#)
- [Configuration on page 59](#)
- [Verification on page 63](#)

Requirements

This example uses the following software and hardware components:

- Junos OS 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 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 9 on page 59](#) details the topology used in this configuration example.

Table 9: Components of the Basic Bridging Configuration Topology

Property	Settings
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:

```
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 63](#)
- [Verifying That Interfaces Are Associated with the Proper VLANs on page 64](#)

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 OS creates VLANs on logical interfaces, not directly on physical interfaces.

Related Documentation

- [Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 65](#)
- [Example: Connecting an Access Switch to a Distribution Switch on page 72](#)
- [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 65](#)
- [Overview and Topology on page 65](#)
- [Configuration on page 66](#)
- [Verification on page 70](#)

Requirements

This example uses the following hardware and software components:

- One EX4200-48P Virtual Chassis switch
- Junos OS 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 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 10 on page 66](#) explains the components of the example topology.

Table 10: 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
user@switch# set support l3-interface vlan.1
```

Results 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;
        l3-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;
        l3-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

To 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 70](#)
- [Verifying That Traffic Is Being Routed Between the Two VLANs on page 71](#)
- [Verifying That Traffic Is Being Switched Between the Two VLANs on page 71](#)

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-0/0/24.0, 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 Documentation

- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 57](#)
- [Example: Connecting an Access Switch to a Distribution Switch on page 72](#)
- [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 72](#)
- [Overview and Topology on page 72](#)
- [Configuring the Access Switch on page 74](#)
- [Configuring the Distribution Switch on page 78](#)
- [Verification on page 80](#)

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 OS 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 the installation instructions for your 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 13 on page 73](#) shows one EX4200 switch that is connected to the three access switches.

Figure 13: Topology for Configuration




Table 11 on page 73 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 11: 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)
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

Table 11: Components of the Topology for Connecting an Access Switch to a Distribution Switch (*continued*)

Property	Settings
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 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
```

**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-id 100
user@access-switch# set l3-interface vlan.0
```

5. Configure the support VLAN:

```
[edit vlans support]
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 l3-interface vlan.0
user@access-switch# set support l3-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;
 }
}
vpls {
 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.

## 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 (802.1Q Tagging) 100
user@distribution-switch# set l3-interface (VLANs) vlan.0
```

The reason that the VLAN configuration for this distribution switch includes the statement `set l3-interface vlan.0` is that the VLAN is being configured for an attached router. The access switch VLAN configuration did not include this statement because the access switch is not monitoring IP addresses, but is instead passing them to the distribution switch for interpretation.

5. Configure the support VLAN:

```
[edit vlans support]
user@distribution-switch# set vlan-description "Support VLAN"
user@distribution-switch# set vlan-id (802.1Q Tagging) 200
user@distribution-switch# set l3-interface (VLANs) vlan.1
```

The reason that the VLAN configuration for this distribution switch includes the statement **set l3-interface vlan.1** is that the VLAN is being configured for an attached router. The access switch VLAN configuration did not include this statement because the access switch is not monitoring IP addresses, but is instead passing them to the distribution switch for interpretation.

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 80](#)
- [Verifying the VLAN Members and Interfaces on the Distribution Switch on page 80](#)

### 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,<br>ge-0/0/6.0, ge-0/0/7.0, ge-0/0/8.0*, ge-0/0/9.0,<br><br>ge-0/0/10.0, ge-0/0/11.0*, ge-0/0/12.0, ge-0/0/13.0,<br>ge-0/0/14.0, ge-0/0/15.0, ge-0/0/16.0, ge-0/0/17.0,<br>ge-0/0/18.0, ge-0/0/19.0*, ge-0/0/21.0, ge-0/0/23.0,<br>ge-0/0/25.0, ge-0/0/27.0*, ge-0/0/28.0, ge-0/0/29.0,<br>ge-0/0/30.0, ge-0/0/31.0*, ge-0/0/32.0, ge-0/0/33.0,<br>ge-0/0/34.0, ge-0/0/35.0*, ge-0/0/36.0, ge-0/0/37.0,<br>ge-0/0/38.0, ge-0/0/39.0*, ge-0/0/40.0, ge-0/0/41.0,<br>ge-0/0/42.0, ge-0/0/43.0*, ge-0/0/45.0, ge-0/0/47.0,<br>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,<br>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,<br>ge-0/0/5.0, ge-0/0/6.0, ge-0/0/7.0*, ge-0/0/8.0,<br><br>ge-0/0/9.0, ge-0/0/10.0*, ge-0/0/11.0, ge-0/0/12.0,<br>ge-0/0/13.0, ge-0/0/14.0, ge-0/0/15.0, ge-0/0/16.0,<br>ge-0/0/17.0, ge-0/0/18.0*, ge-0/0/19.0, ge-0/0/20.0,<br>ge-0/0/21.0, ge-0/0/22.0*, ge-0/0/23.0, ge-0/1/1.0*,<br>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 Documentation**
- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 57](#)
  - [Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 65](#)
  - [Understanding Bridging and VLANs on EX Series Switches on page 3](#)

## Example: Configuring a Private VLAN on a Single 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 allows 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 PVLAN on a single EX Series switch:



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

- [Requirements on page 82](#)
- [Overview and Topology on page 82](#)
- [Configuration on page 83](#)
- [Verification on page 86](#)

## Requirements

This example uses the following hardware and software components:

- One EX Series switch
- Junos OS Release 9.3 or later for EX Series switches

Before you begin configuring a PVLAN, make sure you have created and configured the necessary VLANs. See [“Configuring VLANs for EX Series Switches \(CLI Procedure\)”](#) on page 160.

## 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 PVLAN with one primary VLAN and two community VLANs, one for HR and one for finance, as well as two isolated ports—one for the mail server and the other for the backup server.


[Table 12 on page 82](#) lists the settings for the example topology.

**Table 12: Components of the Topology for Configuring a PVLAN**

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

[Figure 14 on page 83](#) shows the topology for this example.

Figure 14: Topology of a Private VLAN on a Single EX Series Switch



## Configuration

To configure a PVLAN, perform these tasks:

### CLI Quick Configuration

To quickly create and configure a PVLAN, 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**
- ```

set vlans pvlan interface ge-0/0/15.0
set vlans pvlan interface ge-0/0/16.0

```
- To configure the PVLAN:
- Set the VLAN ID for the primary VLAN:


```

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

```
 - 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

```
 - 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

```
- 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

```
  - For each secondary VLAN, configure access interfaces:



**NOTE:** We recommend that the secondary VLANs be untagged VLANs. It does not impair functioning if you tag the secondary VLANs. However, the tags are not used when a secondary VLAN is configured on a single switch.

- ```

[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

```
- 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

```
 - 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:

```
[edit]
user@switch# show
interfaces {
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members pvlan;
        }
      }
    }
  }
  ge-1/0/0 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members pvlan;
        }
      }
    }
  }
  ge-0/0/11 {
    unit 0 {
      family ethernet-switching {
        port-mode access;
      }
    }
  }
  ge-0/0/12 {
    unit 0 {
      family ethernet-switching {
        port-mode access;
        vlan {
          members pvlan;
        }
      }
    }
  }
  ge-0/0/13 {
    unit 0 {
      family ethernet-switching {
        port-mode access;
      }
    }
  }
  ge-0/0/14 {
    unit 0 {
```

```
        family ethernet-switching {
            port-mode access;
        }
    }
}
vllans {
    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 That the Private VLAN and Secondary VLANs Were Created on page 86](#)

Verifying That 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 Documentation

- [Example: Configuring a Private VLAN Spanning Multiple EX Series Switches on page 119](#)
- [Creating a Private VLAN on a Single EX Series Switch \(CLI Procedure\) on page 169](#)

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

Virtual routing instances enable an EX Series switch to have multiple routing tables. With virtual routing instances, you can segment your network to isolate traffic without setting up additional devices.

In a large office, you might need multiple VLANs to properly manage your traffic. This example describes how to create a virtual routing instance associated with each VLAN for this large office:

- [Requirements on page 88](#)
- [Overview and Topology on page 88](#)
- [Configuration on page 89](#)
- [Verification on page 91](#)

Requirements

This example uses the following hardware and software components:

- One EX Series switch
- Junos OS 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 160](#), [Configuring VLANs for EX Series Switches \(CLI Procedure\)](#), or [“Configuring VLANs for EX Series Switches \(J-Web Procedure\)” on page 158](#).

Overview and Topology

This configuration example shows a simple large office topology wherein a LAN is segmented into two VLANs, each of which is associated with an interface and a virtual routing instance.

For example, VLAN 1031 is associated with interface ge-0/0/3.1 and virtual routing instance **r1**. VLAN 1032 is associated with interface ge-0/0/3.2 and virtual routing instance **r2**.

This example also includes interfaces ge-0/0/1.0 and ge-0/0/2.0. Although these interfaces are not part of either VLAN, interface ge-0/0/1.0 is associated with virtual routing instance **r1**, while interface ge-0/0/2.0 is associated with virtual routing instance **r2**.

This example also shows how to use policy statements to import routes from virtual routing instance **r1** to **r2** and from virtual routing instance **r2** to **r1**.



NOTE: On EX Series switches, importing directly connected routes from one virtual routing instance to another is not supported. (A directly connected route is a route that is created by specifying an IP address on one of the switch interfaces or for a VLAN.) For more information, including a workaround for this situation, see [KB23027 - Exchanging \(leaking\) directly connected routes across routing instances is not supported](#).

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 1 vlan-id 1031 family inet address 10.3.1.1/24
set interfaces ge-0/0/3 unit 2 vlan-id 1032 family inet address 10.4.1.1/24
set interfaces ge-0/0/1 unit 0 family inet address 10.1.1.1/24
set interfaces ge-0/0/2 unit 0 family inet address 10.2.1.1/24
set routing-instances r1 instance-type virtual-router
set routing-instances r1 interface ge-0/0/1.0
set routing-instances r1 interface ge-0/0/3.1
set routing-instances r1 routing-options instance-import import-from-r2
set routing-instances r2 instance-type virtual-router
set routing-instances r2 interface ge-0/0/2.0
set routing-instances r2 interface ge-0/0/3.2
set routing-instances r2 routing-options instance-import import-from-r1
set policy-options policy-statement import-from-r1 term 1 from instance r1
set policy-options policy-statement import-from-r1 term 1 then accept
set policy-options policy-statement import-from-r2 term 1 from instance r2
set policy-options policy-statement import-from-r2 term 1 then accept
```

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 one or more logical interfaces on the interfaces to be included in each virtual routing instance:


```
[edit]
user@switch# set interfaces ge-0/0/3 unit 1 vlan-id 1031 family inet address 10.3.1.1/24
user@switch# set interfaces ge-0/0/3 unit 2 vlan-id 1032 family inet address 10.4.1.1/24
user@switch# set interfaces ge-0/0/1 unit 0 family inet address 10.1.1.1/24
user@switch# set interfaces ge-0/0/2 unit 0 family inet address 10.2.1.1/24
```
3. Create two virtual routing instances:


```
[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 routing instances:


```
[edit]
user@switch# set routing-instances r1 interface ge-0/0/1.0
user@switch# set routing-instances r1 interface ge-0/0/3.1
user@switch# set routing-instances r2 interface ge-0/0/2.0
user@switch# set routing-instances r2 interface ge-0/0/3.2
```

5. Apply a policy to routes being imported into each of the virtual routing instances:

```
[edit]
```

```
user@switch# set routing-instances r1 routing-options instance-import import-from-r2
```

```
user@switch# set routing-instances r2 routing-options instance-import import-from-r1
```

6. Create a policy for importing routes from virtual routing instance r1 to r2 and another policy for importing routes from virtual routing instance r2 to r1:

```
[edit]
```

```
user@switch# set policy-options policy-statement import-from-r1 term 1 from instance r1
```

```
user@switch# set policy-options policy-statement import-from-r1 term 1 then accept
```

```
user@switch# set policy-options policy-statement import-from-r2 term 1 from instance  
r2
```

```
user@switch# set policy-options policy-statement import-from-r2 term 1 then accept
```

Results Check the results of the configuration:

```
user@switch> show configuration
```

```
interfaces {  
    ge-0/0/1 {  
        unit 0 {  
            family inet {  
                address 10.1.1.1/24;  
            }  
        }  
    }  
    ge-0/0/2 {  
        unit 0 {  
            family inet {  
                address 10.2.1.1/24;  
            }  
        }  
    }  
    ge-1/0/3 {  
        vlan-tagging;  
        unit 1 {  
            vlan-id 1031;  
            family inet {  
                address 10.3.1.1/24;  
            }  
        }  
        unit 2 {  
            vlan-id 1032;  
            family inet {  
                address 10.4.1.1/24;  
            }  
        }  
    }  
}  
policy-options {  
    policy-statement import-from-r1 {  
        term 1 {  
            from instance r1;  
            then accept;  
        }  
    }  
    policy-statement import-from-r2 {  
        term 1 {  
            from instance r2;  
            then accept;  
        }  
    }  
}
```

```

    }
  }
  routing-instances {
    r1 {
      instance-type virtual-router;
      interface ge-0/0/1.0;
      interface ge-0/0/3.1;
      routing-options {
        instance-import import-from-r2;
      }
    }
    r2 {
      instance-type virtual-router;
      interface ge-0/0/2.0;
      interface ge-0/0/3.2;
      routing-options {
        instance-import import-from-r1;
      }
    }
  }
}

```

Verification

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

- [Verifying That the Routing Instances Were Created on page 91](#)

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      6/0/0
  iso.0       1/0/0
  inet6.0     2/0/0
...
r1            virtual-router
  r1.inet.0   7/0/0
r2            virtual-router
  r2.inet.0   7/0/0

```

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 Documentation

- [Configuring Virtual Routing Instances \(CLI Procedure\) on page 174](#)

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



NOTE: This example uses Junos OS for EX Series switches that does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

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 also dynamically creates 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 92](#)
- [Overview and Topology on page 92](#)
- [Configuring VLANs and MVRP on Access Switch A on page 95](#)
- [Configuring VLANs and MVRP on Access Switch B on page 97](#)
- [Configuring VLANs and MVRP on Distribution Switch C on page 99](#)
- [Verification on page 100](#)

Requirements

This example uses the following hardware and software components:

- Two EX Series access switches
- One EX Series distribution switch
- Junos OS 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.



NOTE: This example shows a network with three VLANs: **finance**, **sales**, and **lab**. All three VLANs are running the same version of Junos OS. If switches in this network were running a mix of Junos OS releases that included Release 11.3, additional configuration would be necessary—see [“Configuring Multiple VLAN Registration Protocol \(MVRP\) \(CLI Procedure\)”](#) on page 176 for details.

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 15 on page 94](#) shows MVRP configured on two access switches and one distribution switch.

Figure 15: MVRP Configured on Two Access Switches and One Distribution Switch for Automatic VLAN Administration




Table 13 on page 94 explains the components of the example topology.

Table 13: Components of the Network Topology

Settings	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 13: Components of the Network Topology (*continued*)

Settings	Settings
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 C (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 C. (trunk) <p>Distribution Switch C interfaces:</p> <ul style="list-style-type: none"> • xe-0/1/1—Connects distribution Switch C to access Switch A. (trunk) • xe-0/1/0—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 recommended 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 on Switch A:

```
[edit]
user@Access-Switch-A# 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 {
        members sales;
      }
    }
  }
  xe-0/1/1 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
      }
    }
  }
}
protocols {
  mvrp {
    interface xe-0/1/1.0;
  }
}
vlands {
  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 for Switch B:

```
[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;
    }
}

vlands {
    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 for Switch C:

```
[edit]
user@Distribution Switch-C# 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 100](#)
- [Verifying That MVRP Is Updating VLAN Membership on Access Switch A on page 101](#)
- [Verifying That MVRP Is Enabled on Access Switch B on page 101](#)
- [Verifying That MVRP Is Updating VLAN Membership on Access Switch B on page 102](#)
- [Verifying That MVRP Is Enabled on Distribution Switch C on page 102](#)
- [Verifying That MVRP Is Updating VLAN Membership on Distribution Switch C on page 102](#)

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  Tag  Tagging  Blocking
ge-0/0/1.0 up     finance       100  untagged  unblocked
ge-0/0/2.0 up     lab           200  untagged  unblocked
ge-0/0/3.0 up     sales         300  untagged  unblocked
xe-0/1/1.0 up     finance       100  untagged  unblocked
           up     lab           200  untagged  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	Tag	Tagging	Blocking
ge-0/0/0.0	up	finance	100	untagged	unblocked
ge-0/0/1.0	up	lab	200	untagged	unblocked
xe-0/1/1.0	up	finance	100	untagged	unblocked
		lab	200	untagged	unblocked
		sales	300	untagged	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	Tag	Tagging	Blocking
xe-0/1/1.0	up	__mvrp_100__			unblocked

```

xe-0/1/0.0 up    __mvrp_200__    unblocked
                  __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
```

```

MVRP dynamic vlans for routing instance 'default-switch'
(s) static vlan, (f) fixed registration

```

VLAN ID	Interfaces
100	xe-0/1/1.0
	xe-0/1/0.0
200	xe-0/1/1.0
	xe-0/1/0.0
300	xe-0/1/1.0

Note that this scenario does not have any fixed registration, which is typical when MVRP is enabled.

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 Documentation

- [Configuring Multiple VLAN Registration Protocol \(MVRP\) \(CLI Procedure\) on page 176](#)
- [Understanding Multiple VLAN Registration Protocol \(MVRP\) on EX Series Switches on page 26](#)

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 104](#)
- [Overview and Topology on page 104](#)

- [Configuration on page 104](#)
- [Verification on page 106](#)

Requirements

This example requires one EX Series switch with Junos OS 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 160 or “[Configuring VLANs for EX Series Switches \(J-Web Procedure\)](#)” on page 158.

Overview and Topology

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

[Table 14 on page 104](#) lists the settings for the example topology.

Table 14: Components of the Topology for Setting Up Q-in-Q Tunneling

Interface	Description
ge-0/0/11.0	Tagged S-VLAN trunk port
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 (VLAN Tagging and Layer 3 Subinterfaces) 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 ];
  }
}
```

```
user@switch> show configuration interfaces
```

```
ge-0/0/11 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan members 4001;
    }
  }
}
ge-0/0/12 {
  unit 0 {
    family ethernet-switching {
      port-mode access;
      vlan members 4001;
    }
  }
}
ge-0/0/13 {
  unit 0 {
    family ethernet-switching {
      port-mode access;
      vlan members 4001;
    }
  }
}
```

```
    }  
  }  
}  
ge-0/0/14 {  
  unit 0 {  
    family ethernet-switching {  
      port-mode trunk;  
      vlan members 4001;  
    }  
  }  
}  
user@switch> show ethernet-switching-options  
dot1q-tunneling {  
  ether-type 0x9100;  
}
```

Verification

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

- [Verifying That Q-in-Q Tunneling Was Enabled on page 106](#)

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 Documentation • [Configuring Q-in-Q Tunneling \(CLI Procedure\) on page 180](#)

Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches



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

Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud 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.



NOTE: L2PT and VLAN translation configured with the `mapping` statement cannot both be configured on the same VLAN. However, L2PT can be configured on one VLAN on a switch while VLAN translation can be configured on a different VLAN that has no L2PT.

This example describes how to configure L2PT:

- [Requirements on page 107](#)
- [Overview and Topology on page 107](#)
- [Configuration on page 109](#)
- [Verification on page 110](#)

Requirements

This example uses the following hardware and software components:

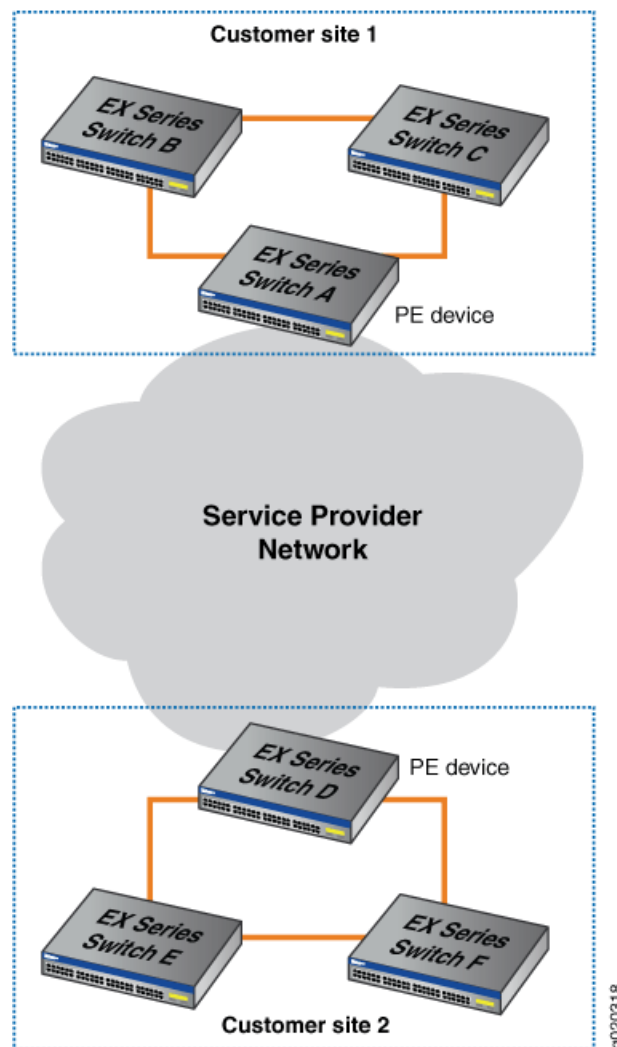
- Six EX Series switches, with three each at two customer sites, with one of the switches at each site designated as the provider edge (PE) device
- Junos OS Release 10.0 or later for EX Series switches

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 16 on page 108](#) 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 16: 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 a high rate. If the tunneled Layer 2 PDUs do arrive at a 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. Alternately, 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 16 on page 108](#), 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 16 on page 108](#), 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 110](#)

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 STP 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 Documentation

- [Configuring Layer 2 Protocol Tunneling on EX Series Switches \(CLI Procedure\) on page 181](#)
- [Understanding Layer 2 Protocol Tunneling on EX Series Switches on page 45](#)

Example: Configuring Redundant Trunk Links for Faster Recovery



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

You can manage network convergence by configuring both a primary link and a secondary link on a switch; this is called a redundant trunk group (RTG). If the primary link in a redundant trunk group fails, it passes its known MAC address locations to the secondary link, which automatically takes over after one minute.

This example describes how to create a redundant trunk group with a primary and a secondary link:

- [Requirements on page 112](#)
- [Overview and Topology on page 112](#)
- [Disabling RSTP on Switches 1 and 2 on page 114](#)
- [Configuring Redundant Trunk Links on Switch 3 on page 115](#)
- [Verification on page 116](#)

Requirements

This example uses the following hardware and software components:

- Two EX Series distribution switches
- One EX Series access switch
- Junos OS Release 10.4 or later for EX Series switches

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

- Configured interfaces **ge-0/0/9** and **ge-0/0/10** on the access switch, Switch 3, as trunk interfaces. See *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*.
- Configured one trunk interface on each distribution switch, Switch 1 and Switch 2.
- Connected the three switches as shown in the topology for this example (see [Figure 17 on page 114](#)).

Overview and Topology

In a typical enterprise network composed of distribution and access layers, a redundant trunk link provides a simple solution for trunk interface network recovery. When a trunk interface fails, data traffic is routed to another trunk interface after one minute, thereby keeping network convergence time to a minimum.

This example shows the configuration of a redundant trunk group that includes one primary link (and its interface) and one unspecified link (and its interface) that serves as the secondary link.

A second type of redundant trunk group, not illustrated in the example, consists of two unspecified links (and their interfaces); in this case, neither of the links is primary. In this second case, the software selects an active link by comparing the port numbers of the two links and activating the link with the higher port number. For example, if the two link interfaces use interfaces **ge-0/1/0** and **ge-0/1/1**, the software activates **ge-0/1/1**. (In the interface names, the final number is the port number.)

The two links in a redundant trunk group generally operate the same way, whether they are configured as primary/unspecified or unspecified/unspecified. Data traffic initially passes through the active link but is blocked on the inactive link. While data traffic is blocked on the secondary link, note that Layer 2 control traffic is still permitted if the link is active. For example, an LLDP session can be run between two switches on the secondary link. If the active link either goes down or is disabled administratively, it broadcasts a list of its known MAC addresses for data traffic; the other link immediately picks up and adds the MAC addresses to its address table, becomes active, and begins forwarding traffic.

The one difference in operation between the two types of redundant trunk groups occurs when a primary link is active, goes down, is replaced by the secondary link, and then reactivates. When a primary link is re-enabled like this while the secondary link is active, the primary link waits 2 minutes (you can change the length of time using the preempt cutover timer to accommodate your network) and then takes over as the active link. In other words, the primary link has priority and is always activated if it is available. This differs from the behavior of two unspecified links, which act as equals. Because the unspecified links are equal, the active link remains active until it either goes down or is disabled administratively; this is the only time that the other unspecified link learns the MAC addresses and immediately becomes active.

The example given here illustrates a primary/unspecified configuration for a redundant trunk group because that configuration gives you more control and is more commonly used.



NOTE: Rapid Spanning Tree Protocol (RSTP) is enabled by default on EX Series switches to create a loop-free topology, but an interface is not allowed to be in both a redundant trunk group and in a spanning-tree protocol topology at the same time. You will need to disable RSTP on the two distribution switches in the example, Switch 1 and Switch 2. Spanning-tree protocols can, however, continue operating in other parts of the network—for example, between the distribution switches and also in links between distribution switches and the enterprise core.

Figure 17 on page 114 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 interfaces **ge-0/0/9.0** (Link 1) and **ge-0/0/10.0** (Link 2).

Table 15 on page 114 lists the components used in this redundant trunk group.

Because RSTP and RTGs cannot operate simultaneously on a switch, you disable RSTP on Switch 1 and Switch 2 in the first configuration task, and you disable RSTP on Switch 3 in the second task.

The second configuration task creates a redundant trunk group called **example 1** on Switch 3. The trunk interfaces **ge-0/0/9.0** and **ge-0/0/10.0** are the two links configured in the second configuration task. You configure the trunk interface **ge-0/0/9.0** as the primary link. You configure the trunk interface **ge-0/0/10.0** as an unspecified link, which becomes the secondary link by default.

Figure 17: Topology for Configuring the Redundant Trunk Links




Table 15: Components of the Redundant Trunk Link Topology

Property	Settings
Switch hardware	<ul style="list-style-type: none"> Switch 1–1 EX Series distribution switch Switch 2–1 EX Series distribution switch Switch 3–1 EX Series access switch
Trunk interfaces	On Switch 3 (access switch): ge-0/0/9.0 and ge-0/0/10.0
Redundant trunk group	example1

Disabling RSTP on Switches 1 and 2

To disable RSTP on Switch 1 and Switch 2, perform this task on each switch:

CLI Quick Configuration To quickly disable RSTP on Switch 1 and Switch 2, copy the following command and paste it into each switch terminal window:

```
[edit]
set protocols rstp disable
```

Step-by-Step Procedure To disable RSTP on Switch 1 and Switch 2:

1. Disable RSTP on Switch 1 and Switch 2:

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

Results Check the results of the configuration:

```
[edit]
user@switch# show
protocols {
  rstp {
    disable;
  }
}
```

Configuring Redundant Trunk Links on Switch 3

To configure redundant trunk links on Switch 3, perform this task:

CLI Quick Configuration To quickly configure the redundant trunk group **example1** on Switch 3, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols rstp disable
set ethernet-switching-options redundant-trunk-group group example1 interface ge-0/0/9.0 primary
set ethernet-switching-options redundant-trunk-group group example1 interface ge-0/0/10.0
set redundant-trunk-group group example1 preempt-cutover-timer 60
```

Step-by-Step Procedure Configure the redundant trunk group **example1** on Switch 3.

1. Turn off RSTP:

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

2. Name the redundant trunk group **example1** while configuring trunk interface **ge-0/0/9.0** as the primary link and **ge-0/0/10** as an unspecified link to serve as the secondary link:

```
[edit ethernet-switching-options]
user@switch# set redundant-trunk-group group example1 interface ge-0/0/9.0 primary
user@switch# set redundant-trunk-group group example1 interface ge-0/0/10.0
```

3. (Optional) Change the length of time (from the default 120 seconds) that a re-enabled primary link waits to take over for an active secondary link:

```
[edit ethernet-switching-options]
user@switch# set redundant-trunk-group group example1 preempt-cutover-timer 60
```

Results Check the results of the configuration:

```
[edit]
user@switch# show
ethernet-switching-options
  redundant-trunk-group {
    group example1 {
      preempt-cutover-timer 60;
      interface ge-0/0/9.0 {
        primary;
      }
      interface ge-0/0/10.0;
    }
  }
protocols {
  rstp {
    disable;
  }
}
```

Verification

To confirm that the configuration is set up correctly, perform this task:

- [Verifying That a Redundant Trunk Group Was Created on page 116](#)

Verifying That a Redundant Trunk Group Was Created

Purpose Verify that the redundant trunk group **example1** has been created on Switch 1 and that trunk interfaces are members of the redundant trunk group.

Action List all redundant trunk groups configured on the switch:

```
user@switch> show redundant-trunk-group
```

Group name	Interface	State	Time of last flap	Flap count
example1	ge-0/0/9.0	Up/Pri	Never	0
	ge-0/0/10.0	Up	Never	0

Meaning The **show redundant-trunk-group** command lists all redundant trunk groups configured on the switch, both links' interface addresses, and the links' current states (up or down for an unspecified link, and up or down and primary for a primary link). For this configuration example, the output shows that the redundant trunk group **example1** is configured on the switch. The **(Up)** beside the interfaces indicates that both link cables are physically connected. The **(Pri)** beside trunk interface **ge-0/0/9.0** indicates that it is configured as the primary link.

Related Documentation

- [Understanding Redundant Trunk Links on page 49](#)

Example: Configuring Proxy ARP on an EX Series Switch

You can configure proxy Address Resolution Protocol (ARP) on your EX Series switch to enable the switch to respond to ARP queries for network addresses by offering its own MAC address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

This example shows how to configure proxy ARP on an access switch:

- [Requirements on page 117](#)
- [Overview and Topology on page 117](#)
- [Configuration on page 117](#)
- [Verification on page 118](#)

Requirements

This example uses the following hardware and software components:

- One EX Series switch
- Junos OS Release 10.0 or later for EX Series switches

Overview and Topology

This example shows the configuration of proxy ARP on an interface of an EX Series switch using restricted mode. In restricted mode, the switch does not act as a proxy for hosts on the same subnet.

The topology for this example consists of one EX Series switch. When a host wants to communicate with a host that is not already in its ARP table, it broadcasts an ARP request for the MAC address of the destination host:

- When proxy ARP is not enabled, a host that shares the same IP address replies directly to the ARP request, providing its MAC address, and future transmissions are sent directly to the destination host MAC address.
- When proxy ARP is enabled, the switch responds to ARP requests, providing the switch's MAC address—even when the destination IP address is the same as the source IP address. Thus, communications must be sent through the switch and then routed through the switch to the appropriate destination.

Configuration

To configure proxy ARP, perform the following tasks:

CLI Quick Configuration

To quickly configure proxy ARP on an interface, copy the following command and paste it into the switch terminal window:

```
[edit]
set interfaces ge-0/0/3 unit 0 proxy-arp restricted
```

Step-by-Step Procedure You configure proxy ARP on individual interfaces.

1. To configure proxy ARP on an interface:

```
[edit interfaces]
user@switch# set ge-0/0/3 unit 0 proxy-arp restricted
```



BEST PRACTICE: We recommend that you configure proxy ARP in restricted mode. In restricted mode, the switch does not act as a proxy if the source and target IP addresses are on the same subnet. If you use unrestricted mode, disable gratuitous ARP requests on the interface to avoid a situation wherein the switch's response to a gratuitous ARP request appears to the host to be an indication of an IP conflict.

```
[edit interfaces]
user@switch# set ge-0/0/3 no-gratuitous-arp-request
```

Results Display the results of the configuration:

```
user@switch> show configuration
interfaces {
  ge-0/0/3 {
    unit 0 {
      proxy-arp restricted;
      family ethernet-switching;
    }
  }
}
```

Verification

To verify that the switch is sending proxy ARP messages, perform these tasks:

- [Verifying That the Switch Is Sending Proxy ARP Messages on page 118](#)

Verifying That the Switch Is Sending Proxy ARP Messages

Purpose Verify that the switch is sending proxy ARP messages.

Action List the system statistics for ARP messages:

```
user@switch> show system statistics arp
arp:
  90060 datagrams received
  34 ARP requests received
  610 ARP replies received
  2 resolution request received
  0 unrestricted proxy requests
  0 restricted proxy requests
  0 received proxy requests
  0 unrestricted proxy requests not proxied
  0 restricted proxy requests not proxied
  0 datagrams with bogus interface
  0 datagrams with incorrect length
  0 datagrams for non-IP protocol
```

```

0 datagrams with unsupported op code
0 datagrams with bad protocol address length
0 datagrams with bad hardware address length
0 datagrams with multicast source address
0 datagrams with multicast target address
0 datagrams with my own hardware address
0 datagrams for an address not on the interface
0 datagrams with a broadcast source address
294 datagrams with source address duplicate to mine
89113 datagrams which were not for me
0 packets discarded waiting for resolution
0 packets sent after waiting for resolution
309 ARP requests sent
35 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
0 replies for from unnumbered interfaces
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor

```

Meaning The statistics show that two proxy ARP requests were received. The **unrestricted proxy requests not proxied** and **restricted proxy requests not proxied** fields indicate that all the unproxied ARP requests received have been proxied by the switch.

Related Documentation

- [Configuring Proxy ARP \(CLI Procedure\) on page 186](#)
- [Configuring Proxy ARP \(CLI Procedure\)](#)
- [Understanding Proxy ARP on EX Series Switches on page 53](#)

Example: Configuring a Private VLAN Spanning Multiple EX Series Switches

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 allows an administrator to split a broadcast domain into multiple isolated broadcast subdomains, essentially putting a VLAN inside a VLAN. A PVLAN can span multiple switches.

This example describes how to create a PVLAN spanning multiple EX Series switches. The example creates one primary PVLAN, containing multiple secondary VLANs:



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

- [Requirements on page 120](#)
- [Overview and Topology on page 120](#)
- [Configuring a PVLAN on Switch 1 on page 123](#)

- [Configuring a PVLAN on Switch 2 on page 125](#)
- [Configuring a PVLAN on Switch 3 on page 128](#)
- [Verification on page 130](#)

Requirements

This example uses the following hardware and software components:

- Three EX Series switches
- Junos OS Release 10.4 or later for EX Series switches

Before you begin configuring a PVLAN, make sure you have created and configured the necessary VLANs. See [“Configuring VLANs for EX Series Switches \(CLI Procedure\)” on page 160](#).

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 how to create a PVLAN spanning multiple EX Series switches, with one primary VLAN containing two community VLANs (one for HR and one for Finance), and an Interswitch isolated VLAN (for the mail server, the backup server, and the CVS server). The PVLAN comprises three switches, two access switches and one distribution switch. The PVLAN is connected to a router through a promiscuous port, which is configured on the distribution switch.



.....

NOTE: The isolated ports on Switch 1 and on Switch 2 do not have Layer 2 connectivity with each other even though they are included within the same domain. See [“Understanding Private VLANs on EX Series Switches” on page 16](#).

.....

[Figure 18 on page 121](#) shows the topology for this example—two access switches connecting to a distribution switch, which has a connection (through a promiscuous port) to the router.

Figure 18: PVLAN Topology Spanning Multiple Switches




Table 16 on page 122, Table 17 on page 122, and Table 18 on page 123 list the settings for the example topology.

Table 16: Components of Switch 1 in the Topology for Configuring a PVLAN Spanning Multiple EX Series Switches

Property	Settings
VLAN names and tag IDs	primary-vlan , tag 100 isolation-id , tag 50 finance-comm , tag 300 hr-comm , tag 400
PVLAN trunk interfaces	ge-0/0/0.0 , Connects Switch 1 to Switch 3 ge-0/0/5.0 , Connects Switch 1 to Switch 2
Interfaces in VLAN isolation	ge-0/0/15.0 , Mail server ge-0/0/16.0 , Backup server
Interfaces in VLAN finance-com	ge-0/0/11.0 ge-0/0/12.0
Interfaces in VLAN hr-comm	ge-0/0/13.0 ge-0/0/14.0

Table 17: Components of Switch 2 in the Topology for Configuring a PVLAN Spanning Multiple EX Series Switches

Property	Settings
VLAN names and tag IDs	primary-vlan , tag 100 isolation-id , tag 50 finance-comm , tag 300 hr-comm , tag 400
PVLAN trunk interfaces	ge-0/0/0.0 , Connects Switch 2 to Switch 3 ge-0/0/5.0 , Connects Switch 2 to Switch 1
Interfaces in VLAN isolation	ge-0/0/17.0 , CVS server
Interfaces in VLAN finance-com	ge-0/0/11.0 ge-0/0/12.0
Interfaces in VLAN hr-comm	ge-0/0/13.0 ge-0/0/14.0

Table 18: Components of Switch 3 in the Topology for Configuring a PVLAN Spanning Multiple EX Series Switches

Property	Settings
VLAN names and tag IDs	primary-vlan , tag 100 isolation-id , tag 50 finance-comm , tag 300 hr-comm , tag 400
PVLAN trunk interfaces	ge-0/0/0.0 , Connects Switch 3 to Switch 1 ge-0/0/1.0 , Connects Switch 3 to Switch 2
Promiscuous port	ge-0/0/2 , Connects the PVLAN to the router NOTE: You must configure the trunk port that connects the PVLAN to another switch or router outside the PVLAN as a member of the PVLAN, which implicitly configures it as a promiscuous port.

Configuring a PVLAN on Switch 1

CLI Quick Configuration

When configuring a PVLAN on multiple switches, these rules apply:

- The primary VLAN must be a tagged VLAN. We recommend that you configure the primary VLAN first.
- Configuring a voice over IP (VoIP) VLAN on PVLAN interfaces is not supported.
- If you are going to configure a community VLAN ID, you must first configure the primary VLAN and the PVLAN trunk port.
- If you are going to configure an isolation VLAN ID, you must first configure the primary VLAN and the PVLAN trunk port.
- Secondary VLANs and the PVLAN trunk port must be committed on a single commit if MVRP is configured on the PVLAN trunk port.

To quickly create and configure a PVLAN spanning multiple switches, copy the following commands and paste them into the terminal window of Switch 1:

```
[edit]
set vlans finance-comm vlan-id 300
set vlans finance-comm interface ge-0/0/11.0
set vlans finance-comm interface ge-0/0/12.0
set vlans finance-comm primary-vlan pvlan100
set vlans hr-comm vlan-id 400
set vlans hr-comm interface ge-0/0/13.0
set vlans hr-comm interface ge-0/0/14.0
set vlans hr-comm primary-vlan pvlan100
set vlans pvlan100 vlan-id 100
set vlans pvlan100 interface ge-0/0/15.0
set vlans pvlan100 interface ge-0/0/16.0
set vlans pvlan100 interface ge-0/0/0.0 pvlan-trunk
set vlans pvlan100 interface ge-0/0/5.0 pvlan-trunk
set vlans pvlan100 no-local-switching
set vlans pvlan100 isolation-id 50
```

**Step-by-Step
Procedure**

Complete the configuration steps below in the order shown—also, complete all steps before committing the configuration in a single commit. This is the easiest way to avoid error messages triggered by violating any of these three rules:

- If you are going to configure a community VLAN ID, you must first configure the primary VLAN and the PVLAN trunk port.
- If you are going to configure an isolation VLAN ID, you must first configure the primary VLAN and the PVLAN trunk port.
- Secondary vlans and a PVLAN trunk must be committed on a single commit.

To configure a PVLAN on Switch 1 that will span multiple switches:

1. Set the VLAN ID for the primary VLAN:

```
[edit vlans]
user@switch# set pvlan100 vlan-id 100
```

2. Set the PVLAN trunk interfaces that will connect this VLAN across neighboring switches:

```
[edit vlans]
user@switch# set pvlan100 interface ge-0/0/0.0 pvlan-trunk
user@switch# set pvlan100 interface ge-0/0/5.0 pvlan-trunk
```

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

```
[edit vlans]
user@switch# set pvlan100 no-local-switching
```

4. Set the VLAN ID for the **finance-comm** community VLAN that spans the switches:

```
[edit vlans]
user@switch# finance-comm vlan-id 300
user@switch# set pvlan100 vlan-id 100
```

5. Configure access interfaces for the **finance-comm** VLAN:

```
[edit vlans]
user@switch# set finance-comm interface ge-0/0/11.0
user@switch# set finance-comm interface ge-0/0/12.0
```

6. Set the primary VLAN of this secondary community VLAN, **finance-comm** :

```
[edit vlans]
user@switch# set vlans finance-comm primary-vlan pvlan100
```

7. Set the VLAN ID for the HR community VLAN that spans the switches.

```
[edit vlans]
user@switch# hr-comm vlan-id 400
```

8. Configure access interfaces for the **hr-comm** VLAN:

```
[edit vlans]
user@switch# set hr-comm interface ge-0/0/13.0
user@switch# set hr-comm interface ge-0/0/14.0
```

9. Set the primary VLAN of this secondary community VLAN, **hr-comm** :

```
[edit vlans]
user@switch# set vlans hr-comm primary-vlan pvlan100
```

10. Set the inter-switch isolated ID to create an inter-switch isolated domain that spans the switches:

```
[edit vlans]
user@switch# set pvlan100 isolation-id 50
```




NOTE: To configure an isolated port, include it as one of the members of the primary VLAN but do not configure it as belonging to one of the community VLANs.

Results

Check the results of the configuration:

```
[edit]
user@switch# show
vpls {
  finance-comm {
    vlan-id 300;
    interface {
      ge-0/0/11.0;
      ge-0/0/12.0;
    }
    primary-vlan pvlan100;
  }
  hr-comm {
    vlan-id 400;
    interface {
      ge-0/0/13.0;
      ge-0/0/14.0;
    }
    primary-vlan pvlan100;
  }
}
pvlan100 {
  vlan-id 100;
  interface {
    ge-0/0/15.0;
    ge-0/0/16.0;
    ge-0/0/0.0 {
      pvlan-trunk;
    }
    ge-0/0/5.0 {
      pvlan-trunk;
    }
  }
  no-local-switching;
  isolation-id 50;
}
```

Configuring a PVLAN on Switch 2

CLI Quick Configuration To quickly create and configure a private VLAN spanning multiple switches, copy the following commands and paste them into the terminal window of Switch 2:



NOTE: The configuration of Switch 2 is the same as the configuration of Switch 1 except for the interface in the inter-switch isolated domain. For Switch 2, the interface is ge-0/0/17.0.

```
[edit]
set vlans finance-comm vlan-id 300
set vlans finance-comm interface ge-0/0/11.0
set vlans finance-comm interface ge-0/0/12.0
set vlans finance-comm primary-vlan pvlan100
set vlans hr-comm vlan-id 400
set vlans hr-comm interface ge-0/0/13.0
set vlans hr-comm interface ge-0/0/14.0
set vlans hr-comm primary-vlan pvlan100
set vlans pvlan100 vlan-id 100
set vlans pvlan100 interface ge-0/0/17.0
set vlans pvlan100 interface ge-0/0/0.0 pvlan-trunk
set vlans pvlan100 interface ge-0/0/5.0 pvlan-trunk
set vlans pvlan100 no-local-switching
set vlans pvlan100 isolation-id 50
```

Step-by-Step Procedure

To configure a PVLAN on Switch 2 that will span multiple switches:

1. Set the VLAN ID for the **finance-comm** community VLAN that spans the switches:

```
[edit vlans]
user@switch# set finance-comm vlan-id (802.1Q Tagging) 300
user@switch# set pvlan100 vlan-id 100
```

2. Configure access interfaces for the **finance-comm** VLAN:

```
[edit vlans]
user@switch# set finance-comm interface (VLANs) ge-0/0/11.0
user@switch# set finance-comm interface ge-0/0/12.0
```

3. Set the primary VLAN of this secondary community VLAN, **finance-comm** :

```
[edit vlans]
user@switch# set vlans finance-comm primary-vlan pvlan100
```

4. Set the VLAN ID for the HR community VLAN that spans the switches.

```
[edit vlans]
user@switch# set hr-comm vlan-id 400
```

5. Configure access interfaces for the **hr-comm** VLAN:

```
[edit vlans]
user@switch# set hr-comm interface ge-0/0/13.0
user@switch# set hr-comm interface ge-0/0/14.0
```

6. Set the primary VLAN of this secondary community VLAN, **hr-comm** :

```
[edit vlans]
user@switch# set vlans hr-comm primary-vlan pvlan100
```

7. Set the VLAN ID for the primary VLAN:

```
[edit vlans]
user@switch# set pvlan100 vlan-id 100
```

8. Set the PVLAN trunk interfaces that will connect this VLAN across neighboring switches:

```
[edit vlans]
```

- ```

user@switch# set pvlan100 interface ge-0/0/0.0 pvlan-trunk
user@switch# set pvlan100 interface ge-0/0/5.0 pvlan-trunk

```
9. Set the primary VLAN to have no local switching:
 

```

[edit vlans]
user@switch# set pvlan100 no-local-switching

```
  10. Set the inter-switch isolated ID to create an inter-switch isolated domain that spans the switches:
 

```

[edit vlans]
user@switch# set pvlan100 isolation-id 50

```



**NOTE:** To configure an isolated port, include it as one of the members of the primary VLAN but do not configure it as belonging to one of the community VLANs.

## Results

Check the results of the configuration:

```

[edit]
user@switch# show
vlans {
 finance-comm {
 vlan-id 300;
 interface {
 ge-0/0/11.0;
 ge-0/0/12.0;
 }
 primary-vlan pvlan100;
 }
 hr-comm {
 vlan-id 400;
 interface {
 ge-0/0/13.0;
 ge-0/0/14.0;
 }
 primary-vlan pvlan100;
 }
}
pvlan100 {
 vlan-id 100;
 interface {
 ge-0/0/15.0;
 ge-0/0/16.0;
 ge-0/0/0.0 {
 pvlan-trunk;
 }
 ge-0/0/5.0 {
 pvlan-trunk;
 }
 ge-0/0/17.0;
 }
 no-local-switching;
}

```

```

 isolation-id 50;
 }
}

```

## Configuring a PVLAN on Switch 3

**CLI Quick Configuration** To quickly configure Switch 3 to function as the distribution switch of this PVLAN, copy the following commands and paste them into the terminal window of Switch 3:



**NOTE:** Interface ge-0/0/2.0 is a trunk port connecting the PVLAN to a router.

```

[edit]
set vlans finance-comm vlan-id 300
set vlans finance-comm primary-vlan pvlan100
set vlans hr-comm vlan-id 400
set vlans hr-comm primary-vlan pvlan100
set vlans pvlan100 vlan-id 100
set vlans pvlan100 interface ge-0/0/0.0 pvlan-trunk
set vlans pvlan100 interface ge-0/0/1.0 pvlan-trunk
set vlans pvlan100 no-local-switching
set vlans pvlan100 isolation-id 50

```

**Step-by-Step Procedure** To configure Switch 3 to function as the distribution switch for this PVLAN, use the following procedure:

1. Set the VLAN ID for the **finance-comm** community VLAN that spans the switches:

```

[edit vlans]
user@switch# finance-comm vlan-id (802.1Q Tagging) 300
[edit vlans]
user@switch# set pvlan100 vlan-id 100

```

2. Set the primary VLAN of this secondary community VLAN, **finance-comm**:

```

[edit vlans]
user@switch# set vlans finance-comm primary-vlan pvlan100

```

3. Set the VLAN ID for the HR community VLAN that spans the switches:

```

[edit vlans]
user@switch# hr-comm vlan-id 400

```

4. Set the primary VLAN of this secondary community VLAN, **hr-comm**:

```

[edit vlans]
user@switch# set vlans hr-comm primary-vlan pvlan100

```

5. Set the VLAN ID for the primary VLAN:

```

[edit vlans]
user@switch# set pvlan100 vlan-id 100

```

6. Set the PVLAN trunk interfaces that will connect this VLAN across neighboring switches:

```

[edit vlans]
user@switch# set pvlan100 interface (VLANs) ge-0/0/0.0 pvlan-trunk
user@switch# set pvlan100 interface ge-0/0/5.0 pvlan-trunk

```

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

```

[edit vlans]

```

- ```
user@switch# set pvlan100 no-local-switching
```
8. Set the inter-switch isolated ID to create an inter-switch isolated domain that spans the switches:

```
[edit vlans]
```

```
user@switch# set pvlan100 isolation-id 50
```



NOTE: To configure an isolated port, include it as one of the members of the primary VLAN but do not configure it as belonging to one of the community VLANs.

Results

Check the results of the configuration:

```
[edit]
user@switch# show
vlans {
  finance-comm {
    vlan-id 300;
    primary-vlan pvlan100;
  }
  hr-comm {
    vlan-id 400;
    primary-vlan pvlan100;
  }
}
pvlan100 {
  vlan-id 100;
  interface {
    ge-0/0/0.0 {
      pvlan-trunk;
    }
    ge-0/0/1.0 {
      pvlan-trunk;
    }
    ge-0/0/2.0;
  }
  no-local-switching;
  isolation-id 50;
}
}
```

Verification

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

- [Verifying That the Primary VLAN and Secondary VLANs Were Created on Switch 1 on page 130](#)
- [Verifying That the Primary VLAN and Secondary VLANs Were Created on Switch 2 on page 131](#)
- [Verifying That the Primary VLAN and Secondary VLANs Were Created on Switch 3 on page 133](#)

[Verifying That the Primary VLAN and Secondary VLANs Were Created on Switch 1](#)

Purpose Verify that the PVLAN configuration spanning multiple switches is working properly on Switch 1:

Action Use the `show vlans extensive` command:

```
user@switch> show vlans extensive
VLAN: __pvlan_pvlan100_ge-0/0/15.0__, Created at: Thu Sep 16 23:15:27 2010
Internal index: 5, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 1 (Active = 1)
    ge-0/0/0.0*, tagged, trunk, pvlan-trunk
    ge-0/0/5.0*, tagged, trunk, pvlan-trunk
    ge-0/0/15.0*, untagged, access
```

```
VLAN: __pvlan_pvlan100_ge-0/0/16.0__, Created at: Thu Sep 16 23:15:27 2010
Internal index: 6, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 1 (Active = 1)
    ge-0/0/0.0*, tagged, trunk, pvlan-trunk
    ge-0/0/5.0*, tagged, trunk, pvlan-trunk
    ge-0/0/16.0*, untagged, access
```

```
VLAN: __pvlan_pvlan100_isiv__, Created at: Thu Sep 16 23:15:27 2010
802.1Q Tag: 50, Internal index: 7, Admin State: Enabled, Origin: Static
Private VLAN Mode: Inter-switch-isolated, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 0 (Active = 0)
    ge-0/0/0.0*, tagged, trunk, pvlan-trunk
    ge-0/0/5.0*, tagged, trunk, pvlan-trunk
```

```
VLAN: default, Created at: Thu Sep 16 03:03:18 2010
Internal index: 2, Admin State: Enabled, Origin: Static
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)
```

```
VLAN: finance-comm, Created at: Thu Sep 16 23:15:27 2010
802.1Q Tag: 300, Internal index: 8, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 2 (Active = 2)
    ge-0/0/0.0*, tagged, trunk, pvlan-trunk
```

```

ge-0/0/5.0*, tagged, trunk, pvlan-trunk
ge-0/0/11.0*, untagged, access
ge-0/0/12.0*, untagged, access

VLAN: hr-comm, Created at: Thu Sep 16 23:15:27 2010
802.1Q Tag: 400, Internal index: 9, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 2 (Active = 2)
    ge-0/0/0.0*, tagged, trunk, pvlan-trunk
    ge-0/0/5.0*, tagged, trunk, pvlan-trunk
    ge-0/0/13.0*, untagged, access
    ge-0/0/14.0*, untagged, access

VLAN: pvlan100, Created at: Thu Sep 16 23:15:27 2010
802.1Q Tag: 100, Internal index: 4, Admin State: Enabled, Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 6 (Active = 6)
    ge-0/0/0.0*, tagged, trunk, pvlan-trunk
    ge-0/0/5.0*, tagged, trunk, pvlan-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
Secondary VLANs: Isolated 2, Community 2, Inter-switch-isolated 1
Isolated VLANs :
    __pvlan_pvlan100_ge-0/0/15.0__
    __pvlan_pvlan100_ge-0/0/16.0__
Community VLANs :
    finance-comm
    hr-comm
Inter-switch-isolated VLAN :
    __pvlan_pvlan100_isiv__

```

Meaning The output shows that a PVLAN was created on Switch 1 and shows that it includes two isolated VLANs, two community VLANs, and an interswitch isolated VLAN. The presence of the **pvlan-trunk** and **Inter-switch-isolated** fields indicates that this PVLAN is spanning more than one switch.

Verifying That the Primary VLAN and Secondary VLANs Were Created on Switch 2

Purpose Verify that the PVLAN configuration spanning multiple switches is working properly on Switch 2:

Action Use the **show vlans extensive** command:

```

user@switch> show vlans extensive
VLAN: __pvlan_pvlan100_ge-0/0/17.0__, Created at: Thu Sep 16 23:19:22 2010
Internal index: 5, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 1 (Active = 1)
    ge-0/0/0.0*, tagged, trunk, pvlan-trunk
    ge-0/0/5.0*, tagged, trunk, pvlan-trunk

```

```

ge-0/0/17.0*, untagged, access

VLAN: __pvlan_pvlan100_isiv__, Created at: Thu Sep 16 23:19:22 2010
802.1Q Tag: 50, Internal index: 6, Admin State: Enabled, Origin: Static
Private VLAN Mode: Inter-switch-isolated, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 0 (Active = 0)
    ge-0/0/0.0*, tagged, trunk, pvlan-trunk
    ge-0/0/5.0*, tagged, trunk, pvlan-trunk

VLAN: default, Created at: Thu Sep 16 03:03:18 2010
Internal index: 2, Admin State: Enabled, Origin: Static
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)

VLAN: finance-comm, Created at: Thu Sep 16 23:19:22 2010
802.1Q Tag: 300, Internal index: 7, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 2 (Active = 2)
    ge-0/0/0.0*, tagged, trunk, pvlan-trunk
    ge-0/0/5.0*, tagged, trunk, pvlan-trunk
    ge-0/0/11.0*, untagged, access
    ge-0/0/12.0*, untagged, access

VLAN: hr-comm, Created at: Thu Sep 16 23:19:22 2010
802.1Q Tag: 400, Internal index: 8, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 2 (Active = 2)
    ge-0/0/0.0*, tagged, trunk, pvlan-trunk
    ge-0/0/5.0*, tagged, trunk, pvlan-trunk
    ge-0/0/13.0*, untagged, access
    ge-0/0/14.0*, untagged, access

VLAN: pvlan100, Created at: Thu Sep 16 23:19:22 2010
802.1Q Tag: 100, Internal index: 4, Admin State: Enabled, Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 5 (Active = 5)
    ge-0/0/0.0*, tagged, trunk, pvlan-trunk
    ge-0/0/5.0*, tagged, trunk, pvlan-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-0/0/17.0*, untagged, access
Secondary VLANs: Isolated 1, Community 2, Inter-switch-isolated 1
Isolated VLANs :
    __pvlan_pvlan100_ge-0/0/17.0__
Community VLANs :
    finance-comm
    hr-comm
Inter-switch-isolated VLAN :
    __pvlan_pvlan100_isiv__

```

Meaning The output shows that a PVLAN was created on Switch 1 and shows that it includes two isolated VLANs, two community VLANs, and an interswitch isolated VLAN. The presence of the **pvlan-trunk** and **Inter-switch-isolated** fields indicates that this is PVLAN spanning

more than one switch. When you compare this output to the output of Switch 1, you can see that both switches belong to the same PVLAN (**pvlan100**).

Verifying That the Primary VLAN and Secondary VLANs Were Created on Switch 3

Purpose Verify that the PVLAN configuration spanning multiple switches is working properly on Switch 3:

Action Use the **show vlans extensive** command:

```
user@switch> show vlans extensive
VLAN: __pvlan_pvlan100_isiv__, Created at: Thu Sep 16 23:22:40 2010
802.1Q Tag: 50, Internal index: 5, Admin State: Enabled, Origin: Static
Private VLAN Mode: Inter-switch-isolated, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 0 (Active = 0)
    ge-0/0/0.0*, tagged, trunk, pvlan-trunk
    ge-0/0/1.0*, tagged, trunk, pvlan-trunk

VLAN: default, Created at: Thu Sep 16 03:03:18 2010
Internal index: 2, Admin State: Enabled, Origin: Static
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)

VLAN: finance-comm, Created at: Thu Sep 16 23:22:40 2010
802.1Q Tag: 300, Internal index: 6, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 0 (Active = 0)
    ge-0/0/0.0*, tagged, trunk, pvlan-trunk
    ge-0/0/1.0*, tagged, trunk, pvlan-trunk

VLAN: hr-comm, Created at: Thu Sep 16 23:22:40 2010
802.1Q Tag: 400, Internal index: 7, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan100
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 0 (Active = 0)
    ge-0/0/0.0*, tagged, trunk, pvlan-trunk
    ge-0/0/1.0*, tagged, trunk, pvlan-trunk

VLAN: pvlan100, Created at: Thu Sep 16 23:22:40 2010
802.1Q Tag: 100, Internal index: 4, Admin State: Enabled, Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 2 (Active = 2), Untagged 0 (Active = 0)
    ge-0/0/0.0*, tagged, trunk, pvlan-trunk
    ge-0/0/1.0*, tagged, trunk, pvlan-trunk
Secondary VLANs: Isolated 0, Community 2, Inter-switch-isolated 1
Community VLANs :
    finance-comm
    hr-comm
Inter-switch-isolated VLAN :
    __pvlan_pvlan100_isiv__
```

Meaning The output shows that the PVLAN (**pvlan100**) is configured on Switch 3 and that it includes two isolated VLANs, two community VLANs, and an interswitch isolated VLAN. But Switch 3 is functioning as a distribution switch, so the output does not include access

interfaces within the PVLAN. It shows only the **pvlan-trunk** interfaces that connect **pvlan100** from Switch 3 to the other switches (Switch 1 and Switch 2) in the same PVLAN.

Related Documentation

- [Example: Configuring a Private VLAN on a Single EX Series Switch on page 81](#)
- [Creating a Private VLAN on a Single EX Series Switch \(CLI Procedure\) on page 169](#)
- [Creating a Private VLAN Spanning Multiple EX Series Switches \(CLI Procedure\) on page 171](#)
- [Understanding PVLAN Traffic Flows Across Multiple Switches on page 22](#)

Example: Configuring Edge Virtual Bridging for Use with VEPA Technology

Virtual machines (VMs) can use a physical switch that is adjacent to the VMs' server to send packets both to other VMs and to the rest of the network when two conditions have been met:

- Virtual Ethernet packet aggregator (VEPA) is configured on the VM server.
- Edge virtual bridging (EVB) is configured on the switch.

This example shows how to configure EVB on the switch so that packets can flow to and from the virtual machines.

- [Requirements on page 134](#)
- [Overview and Topology on page 135](#)
- [Configuration on page 136](#)
- [Verification on page 139](#)

Requirements

This example uses the following hardware and software components:

- One EX4500 or EX8200 switch
- Junos OS Release 12.1 or later for EX Series switches

Before you configure EVB on a switch, be sure you have configured the server with virtual machines, the VLANs, and VEPA:



NOTE: The following are the numbers of components used in this example, but you can use fewer or more to configure the feature.

- On the server, configure six virtual machines, VM 1 through VM 6 as shown in [Figure 19 on page 135](#). See your server documentation.
- On the server, configure three VLANs named VLAN_Purple, VLAN_Orange, and VLAN_Blue, and add two virtual machines to each VLAN. See your server documentation.

- On the server, install and configure VEPA to aggregate the virtual machine packets.
- On the switch, configure one interface with the same three VLANs as the server (VLAN_Purple, VLAN_Orange, and VLAN_Blue). See [“Configuring VLANs for EX Series Switches \(CLI Procedure\)”](#) on page 160.

Overview and Topology


EVB is a software capability that provides multiple virtual end stations that communicate with each other and with external switches in the Ethernet network environment.

This example demonstrates the configuration that takes place on a switch when that switch is connected to a server with VEPA configured. In this example, a switch is already configured to a server hosting six virtual machines (VMs) and configured with VEPA for aggregating packets. The server's six virtual machines are VM 1 through VM 6, and each virtual machine belongs to one of the three server VLANs—VLAN_Purple, VLAN_Orange, or VLAN_Blue. Because VEPA is configured on the server, no two VMs can communicate directly—all communication between VMs must happen via the adjacent switch.

[Figure 19 on page 135](#) shows the topology for this example.

Edge Virtual Bridging Example Topology

Figure 19: Topology



g020996

The VEPA component of the server pushes all packets from any VM, regardless of whether the packets are destined to other VMs on the same server or to any external host, to the adjacent switch. The adjacent switch applies policies to incoming packets based on the interface configuration and then forwards the packets to appropriate interfaces based on the MAC learning table. If the switch has not yet learned a destination MAC, it floods the packet to all interfaces, including the source port on which the packet arrived.

[Table 19 on page 136](#) shows the components used in this example.

Table 19: Components of the Topology for Configuring EVB

Component	Description
EX Series switch	For a list of switches that support this feature, see <i>EX Series Switch Software Features Overview</i> or <i>EX Series Virtual Chassis Software Features Overview</i> .
ge-0/0/20	Switch interface to the server.
Server	Server with virtual machines and VEPA technology.
Virtual machines	Six virtual machines located on the server, named VM 1, VM 2, VM 3, VM 4, VM 5, and VM 6.
VLANs	Three VLANs, named VLAN_Purple, VLAN_Orange, and VLAN_Blue. Each VLAN has two virtual machine members.
VEPA	A virtual Ethernet port aggregator (VEPA) is a software capability on a server that collaborates with an adjacent, external switch to provide bridging support between multiple virtual machines and with external networks. The VEPA collaborates with the switch by forwarding all VM-originated frames to the adjacent bridge for frame processing and frame relay (including hairpin forwarding) and by steering and replicating frames received from the VEPA uplink to the appropriate destinations.



NOTE: Configuring EVB also enables Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP).

Configuration

CLI Quick Configuration

To quickly configure EVB, copy the following commands and paste them into the switch's CLI at the **[edit]** hierarchy level.

```
set interfaces ge-0/0/20 unit 0 family ethernet-switching port-mode tagged-access
set protocols lldp interface ge-0/0/20.0
set vlans vlan_purple interface ge-0/0/20.0
set vlans vlan_orange interface ge-0/0/20.0
set vlans vlan_blue interface ge-0/0/20.0
set protocols edge-virtual-bridging vsi-discovery interface ge-0/0/20.0
set policy-options vsi-policy P1 from vsi-manager 98 vsi-type 998 vsi-version 4 vsi-instance
09b11c53-8b5c-4eeb-8f00-c84ebb0bb998
set policy-options vsi-policy P1 then filter f2
set policy-options vsi-policy P3 from vsi-manager 97 vsi-type 997 vsi-version 3 vsi-instance
09b11c53-8b5c-4eeb-8f00-c84ebb0bb997
set policy-options vsi-policy P3 then filter f3
set firewall family ethernet-switching filter f2 term t1 then accept
set firewall family ethernet-switching filter f2 term t1 then count f2_accept
set firewall family ethernet-switching filter f3 term t1 then accept
set firewall family ethernet-switching filter f3 term t1 then count f3_accept
set protocols edge-virtual-bridging vsi-discovery vsi-policy P1
set protocols edge-virtual-bridging vsi-discovery vsi-policy P3
```

**Step-by-Step
Procedure**

To configure EVB on the switch:

1. Configure tagged-access mode for the interfaces on which you will enable EVB:

```
[edit interfaces ge-0/0/20]
user@switch# set unit 0 family ethernet-switching port-mode tagged-access
```
2. Enable the Link Layer Discovery Protocol (LLDP) on the ports interfaces on which you will enable EVB:

```
[edit protocols]
user@switch# set lldp interface ge-0/0/20.0
```
3. Configure the interface as a member of all VLANs located on the virtual machines.

```
[edit]
user@switch# set vlans vlan_purple interface ge-0/0/20.0
user@switch# set vlans vlan_orange interface ge-0/0/20.0
user@switch# set vlans vlan_blue interface ge-0/0/20.0
```
4. Enable the VSI Discovery and Control Protocol (VDP) on the interface:

```
[edit protocols]
user@switch# set edge-virtual-bridging vsi-discovery interface ge-0/0/20.0
```
5. Define policies for VSI information. VSI information is based on a VSI manager ID, VSI type, VSI version, and VSI instance ID:

```
[edit policy-options]
user@switch# set vsi-policy P1 from vsi-manager 98 vsi-type 998 vsi-version 4 vsi-instance 09b11c53-8b5c-4eeb-8f00-c84ebb0bb998
user@switch# set vsi-policy P1 then filter f2
user@switch# set vsi-policy P3 from vsi-manager 97 vsi-type 997 vsi-version 3 vsi-instance 09b11c53-8b5c-4eeb-8f00-c84ebb0bb997
user@switch# set vsi-policy P3 then filter f3
```
6. Two VSI policies were defined in the previous step, each of them mapping to different firewall filters. Define the firewall filters:

```
[edit firewall family ethernet-switching]
user@switch# set filter f2 term t1 then accept
user@switch# set filter f2 term t1 then count f2_accept
user@switch# set filter f3 term t1 then accept
user@switch# set filter f3 term t1 then count f3_accept
```
7. Associate VSI policies with VSI-discovery protocol

```
[edit]
user@switch# set protocols edge-virtual-bridging vsi-discovery vsi-policy P1
user@switch# set protocols edge-virtual-bridging vsi-discovery vsi-policy P3
```

```
Results user@switch# show protocols
edge-virtual-bridging {
  vsi-discovery {
    interface {
      ge-0/0/20.0;
    }
    vsi-policy {
      P1;
      P3;
    }
  }
}
lldp {
  interface ge-0/0/20.0;

user@switch# show policy-options
vsi-policy P1 {
  from {
    vsi-manager 98 vsi-type 998 vsi-version 4 vsi-instance 09b11c53-8b5c-4ee
b-8f00-c84ebb0bb998;
  }
  then {
    filter f2;
  }
}
vsi-policy P3 {
  from {
    vsi-manager 97 vsi-type 997 vsi-version 3 vsi-instance 09b11c53-8b5c-4ee
b-8f00-c84ebb0bb997;
  }
  then {
    filter f3;
  }
}

user@switch# show vlans
vlan_blue {
  interface {
    ge-0/0/20.0;
  }
}
vlan_orange {
  interface {
    ge-0/0/20.0;
  }
}
vlan_purple {
  interface {
    ge-0/0/20.0;
    interface;
  }
}

user@switch# show firewall
family ethernet-switching {
  filter f2 {
    term t1 {
      then {
        accept;
        count f2_accept;
      }
    }
  }
}
```

```

    }
  }
}
filter f3 {
  term t1 {
    then {
      accept;
      count f3_accept;
    }
  }
}
}

```

Verification

To confirm that EVB is enabled and working correctly, perform these tasks:

- [Verifying That EVB is Correctly Configured on page 139](#)
- [Verifying That the Virtual Machine Successfully Associated With the Switch on page 139](#)
- [Verifying That VSI Profiles Are Being Learned at the Switch on page 140](#)

Verifying That EVB is Correctly Configured

Purpose Verify that EVB is correctly configured

Action user@switch# **show edge-virtual-bridging**

Interface	Forwarding Mode	RTE	Number of VSIs	Protocols
ge-0/0/20.0	Reflective-relay	25	400	ECP, VDP, RTE

Meaning When LLDP is first enabled, an EVB LLDP exchange takes place between switch and server using LLDP. As part of this exchange the following parameters are negotiated: Number of VSIs supported, Forwarding mode, ECP support, VDP support, and Retransmission Timer Exponent (RTE). If the output has values for the negotiated parameters, EVB is correctly configured.

Verifying That the Virtual Machine Successfully Associated With the Switch

Purpose Verify that the virtual machine successfully associated with the switch. After successful association of VSI Profile with the switch interface, verify the learning of the VM's MAC address on MAC-Table or Forwarding database Table. The learn type of the VM's MAC addresses will be VDP, and upon successful shutdown of VM the corresponding MAC-VLAN entry will get flushed out from FDB table otherwise it will never shutdown.

Action user@switch# run show ethernet-switching table

Ethernet-switching table: 10 entries, 4 learned

VLAN	MAC address	Type	Age	Interfaces
v3	*	Flood	-	All-members
v3	00:02:a6:11:bb:1a	VDP		- ge-1/0/10.0
v3	00:02:a6:11:cc:1a	VDP		- ge-1/0/10.0
v3	00:23:9c:4f:70:01	Static		- Router
v4	*	Flood		- All-members
v4	00:02:a6:11:bb:bb	VDP		- ge-1/0/10.0
v4	00:23:9c:4f:70:01	Static		- Router
v5	*	Flood		- All-members
v5	00:23:9c:4f:70:01	Static		- Router
v5	52:54:00:d5:49:11	VDP		- ge-1/0/20.0

Verifying That VSI Profiles Are Being Learned at the Switch

Purpose Verify that VSI profiles are being learned at the switch.

Action user@switch# show edge-virtual-bridging vsi-profiles

Interface: ge-0/0/20.0

Manager: 97, Type: 997, Version: 3, VSI State: Associate

Instance: 09b11c53-8b5c-4eeb-8f00-c84ebb0bb997

MAC	VLAN
00:10:94:00:00:04	3

Meaning Whenever VMs configured for VEPA are started at the server, the VMs start sending VDP messages. As part of this protocol VSI profiles are learned at the switch.

If the output has values for Manager, Type, Version, VSI State, and Instance, VSI profiles are being learned at the switch.

Related Documentation

- [Configuring Edge Virtual Bridging \(CLI Procedure\) on page 188](#)
- [Understanding Edge Virtual Bridging for Use with VEPA Technology on page 33](#)

Example: Configuring Ethernet Ring Protection Switching on EX Series Switches

You can configure Ethernet ring protection switching (ERPS) on connected EX Series or QFX Series switches to prevent fatal loops from disrupting a network. ERPS is similar to the Spanning Tree Protocol, but ERPS is more efficient because it is customized for ring topologies. You must configure at least three switches to form a ring.

This example shows how to configure Ethernet ring protection switching on four switches that are connected to one another on a dedicated link in a ring topology.



NOTE: This task uses Junos OS for EX Series legacy switches without support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches* For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

- [Requirements on page 141](#)
- [Overview and Topology on page 141](#)
- [Configuration on page 142](#)
- [Verification on page 154](#)

Requirements

This example uses the following hardware and software components:

- Four connected EX Series switches without support for the Enhanced Layer 2 Software (ELS) that will function as nodes in the ring topology.
- Junos OS Release 12.1 or later for EX Series switches.

Before you begin, be sure you have:

- Configured two trunk interfaces on each of the four switches. See [Table 20 on page 142](#) for a list of the interface names used in this example.
- Configured the same VLAN (**erp-control-vlan-1**) with ID 100 on all four switches and associated two network interfaces from each of the four switches with the VLAN. See “[Configuring VLANs for EX Series Switches \(CLI Procedure\)](#)” on page 160. See [Table 20 on page 142](#) for a list of the interface names used in this example.
- Configured two VLANs (**erp-data-1** and **erp-data-2**) with IDs 101 and 102, respectively, on all four switches and associated both the east and west interfaces on each switch with **erp-data-1** and **erp-data-2**. See [Table 20 on page 142](#) for a list of the interface names used in this example.

Overview and Topology

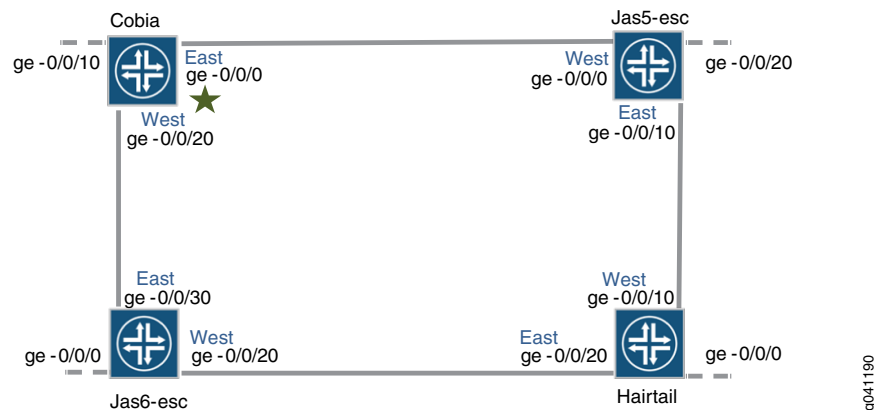
ERPS uses a dedicated physical link, including a control VLAN for trunk ports, between all of the switches to protect the active links. ERPS VLANs are all located on this link and are also blocked by default. When traffic between the switches is flowing with no problems, the active links take care of all traffic. Only if an error occurs on one of the data links would the ERPS control channel take over and start forwarding traffic.



NOTE: Trunk ports on switches use a VLAN to create individual control channels for ERPS. When multiple ERPS instances are configured for a ring, there are multiple sets of ring protection links (RPLs) and RPL owners on the ERPS link, and a different channel is blocked for each instance. Nontrunk ports use the physical link as the control channel and protocol data units (PDUs) are untagged, with no VLAN information in the packet.

This example creates one protection ring (called a node ring) named `erp1` on four switches connected in a ring by trunk ports as shown in [Figure 20 on page 142](#). Because the links are trunk ports, the VLAN named `erp-control-vlan-1` is used for `erp1` traffic. The east interface of each switch is connected with the west interface of an adjacent switch. Cobia is the RPL owner, with interface `ge-0/0/0` configured as an RPL end interface. The interface `ge-0/0/0` of `Jas5-esc` is configured as the RPL neighbor interface. In the idle state, the RPL end blocks the control VLAN and data channel VLAN for this particular ERP instance—the blocked port on Cobia is marked with a star in [Figure 20 on page 142](#).

Figure 20: Ethernet Ring Protection Switching Example



In this example, we configure the four switches with the interfaces indicated in both [Figure 20 on page 142](#) and [Table 20 on page 142](#).

Table 20: Components to Configure for This Example

Interfaces	Cobia	Jas5-esc	Jas6-esc	Hairtail
East	ge-0/0/0	ge-0/0/10	ge-0/0/30	ge-0/0/20
West	ge-0/0/20	ge-0/0/0	ge-0/0/20	ge-0/0/10
Third	ge-0/0/10	ge-0/0/20	ge-0/0/0	ge-0/0/0

Configuration

- [Configuring ERPS on Cobia, the RPL Owner Node on page 143](#)
- [Configuring ERPS on Jas5-esc on page 146](#)

- [Configuring ERPS on Hairtail on page 149](#)
- [Configuring ERPS on Jas6-esc on page 151](#)

Configuring ERPS on Cobia, the RPL Owner Node

CLI Quick Configuration

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



NOTE: RSTP and ERPS cannot both be configured on a ring port, and RSTP is configured by default. Therefore, in this example, RSTP is disabled on each ring port before configuring ERPS.

```
set protocols rstp interface ge-0/0/0 disable
set protocols rstp interface ge-0/0/20 disable
set protocols protection-group ethernet-ring erp1
set protocols protection-group ethernet-ring erp1 ring-protection-link-owner
set protocols protection-group ethernet-ring erp1 data-channel erp-data-1
set protocols protection-group ethernet-ring erp1 data-channel erp-data-2
set protocols protection-group ethernet-ring erp1 control-vlan erp-control-vlan-1
set protocols protection-group ethernet-ring erp1 east-interface control-channel
ge-0/0/0.0
set protocols protection-group ethernet-ring erp1 east-interface ring-protection-link-end
set protocols protection-group ethernet-ring erp1 west-interface control-channel
ge-0/0/20.0
set protocols protection-group ethernet-ring erp1 west-interface control-channel
ge-0/0/20.0 vlan erp-control-vlan-1
set protocols protection-group ethernet-ring erp1 east-interface control-channel
ge-0/0/20.0 vlan erp-control-vlan-1
```

Step-by-Step Procedure

To configure ERPS on Cobia:

1. Disable RSTP on the two ports that will use ERPS:


```
[edit protocols]
user@switch# set rstp interface ge-0/0/0 disable
user@switch# set rstp interface ge-0/0/20 disable
```
2. Create a node ring named erp1:


```
[edit protocols]
user@switch# set protection-group ethernet-ring erp1
```
3. Designate Cobia as the RPL owner node:


```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set ring-protection-link-owner
```
4. Configure the VLANs erp-data-1 and erp-data-2 as data channels:


```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set data-channel erp-data-1
user@switch# set data-channel erp-data-2
```

5. Configure the control VLAN `erp-control-vlan-1` for this ERP instance on the trunk interface:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set control-vlan erp-control-vlan-1
```

6. Configure the east interface of the node ring `erp1` with the control channel `ge-0/0/0.0` and indicate that this particular ring protection link ends here:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set east-interface control-channel ge-0/0/0.0
user@switch# set east-interface ring-protection-link-end
```

7. Configure the west interface of the node ring `erp1` with the control channel `ge-0/0/20.0`:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/20.0
```

8. Every ring instance on a trunk port has one control VLAN in which ERP packets traverse. The control VLAN also controls data VLANs, if any are configured. Assign `erp-control-vlan-1` as the control VLAN on both interfaces:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/20.0 vlan
erp-control-vlan-1
user@switch# set east-interface control-channel ge-0/0/20.0 vlan
erp-control-vlan-1
```

Results In configuration mode, check your ERPS configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@switch# show protocols
rstp {
  interface ge-0/0/20.0 {
    disable;
  }
  interface ge-0/0/0.0 {
    disable;
  }
}
protection-group {
  ethernet-ring erp1 {
    ring-protection-link-owner;
    east-interface {
      control-channel {
        ge-0/0/0.0;
      }
      ring-protection-link-end;
    }
    west-interface {
      control-channel {
        ge-0/0/20.0;
      }
    }
  }
}
```

```

control-vlan erp-control-vlan-1;
  data-channel {
    vlan 101-102;
  }
}

```

In configuration mode, check your VLAN configuration by entering the **show vlans** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

[edit]
user@switch# show vlans
  erp-control-vlan-1 {
    vlan-id 100;
    interface {
      ge-0/0/0.0;
      ge-0/0/20.0;
    }
  }
  erp-data-1 {
    vlan-id 101;
    interface {
      ge-0/0/10.0;
      ge-0/0/0.0;
      ge-0/0/20.0;
    }
  }
  erp-data-2 {
    vlan-id 102;
    interface {
      ge-0/0/10.0;
      ge-0/0/0.0;
      ge-0/0/20.0;
    }
  }
}

```

In configuration mode, check your interface configurations by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

[edit]
user@switch# show interfaces
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
      }
    }
  }
  ge-0/0/10 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
      }
    }
  }
}

```

```
ge-0/0/20 {  
  unit 0 {  
    family ethernet-switching {  
      port-mode trunk;  
    }  
  }  
}
```

If you are finished configuring the device, enter **commit** in configuration mode.

Configuring ERPS on Jas5-esc

CLI Quick Configuration

To quickly configure Jas5-esc, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols rstp interface ge-0/0/10 disable  
set protocols rstp interface ge-0/0/0 disable  
set protocols protection-group ethernet-ring erp1  
set protocols protection-group ethernet-ring erp1 data-channel erp-data-1  
set protocols protection-group ethernet-ring erp1 data-channel erp-data-2  
set protocols protection-group ethernet-ring erp1 control-vlan erp-control-vlan-1  
set protocols protection-group ethernet-ring erp1 east-interface control-channel  
  ge-0/0/10.0  
set protocols protection-group ethernet-ring erp1 west-interface control-channel  
  ge-0/0/0.0  
set protocols protection-group ethernet-ring erp1 west-interface control-channel  
  ge-0/0/10.0 vlan erp-control-vlan-1  
set protocols protection-group ethernet-ring erp1 east-interface control-channel  
  ge-0/0/0.0 vlan erp-control-vlan-1
```

Step-by-Step Procedure

To configure ERPS on Jas5-esc:

1. Disable RSTP on the two ports that will use ERPS:

```
[edit protocols]  
user@switch# set rstp interface ge-0/0/10 disable  
user@switch# set rstp interface ge-0/0/0 disable
```
2. Create a node ring named erp1:

```
[edit protocols]  
user@switch# set protection-group ethernet-ring erp1
```
3. Configure a control VLAN named erp-control-vlan-1 for the node ring erp1:

```
[edit protocols protection-group ethernet-ring erp1]  
user@switch# set control-vlan erp-control-vlan-1
```
4. Configure two data channels named erp-data-1 and erp-data-2 to define a set of VLAN IDs that belong to a ring instance.

```
[edit protocols protection-group ethernet-ring erp1]  
user@switch# set data-channel erp-data-1  
user@switch# set data-channel erp-data-2
```
5. Configure the east interface of the node ring erp1 with the control channel ge-0/0/10.0:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set east-interface control-channel ge-0/0/10.0
```

6. Configure the west interface of the node ring erp1 with the control channel ge-0/0/0.0:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/0.0
```

7. Every ring instance on a trunk port has one control VLAN in which ERP packets traverse. The control VLAN also controls data VLANs, if any are configured. Assign erp-control-vlan-1 as the control VLAN:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/0.0 vlan erp-control-vlan-1
user@switch# set east-interface control-channel ge-0/0/10.0 vlan erp-control-vlan-1
```

Results In configuration mode, check your ERPS configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@switch# show protocols
rstp {
  interface ge-0/0/10.0 {
    disable;
  }
  interface ge-0/0/0.0 {
    disable;
  }
}
protection-group {
  east-interface {
    control-channel {
      ge-0/0/10.0;
    }
  }
  west-interface {
    control-channel {
      ge-0/0/0.0;
    }
  }
  control-vlan erp-control-vlan-1;
  data-channel
    vlan 101-102
}
}
```

In configuration mode, check your VLAN configuration by entering the **show vlans** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@switch# show vlans
erp-control-vlan-1 {
  vlan-id 100;
}
```

```
interface {  
    ge-0/0/10.0;  
    ge-0/0/0.0;  
}  
}  
erp-data-1 {  
    vlan-id 101;  
    interface {  
        ge-0/0/20.0;  
        ge-0/0/10.0;  
        ge-0/0/0.0;  
    }  
}  
erp-data-2 {  
    vlan-id 102;  
    interface {  
        ge-0/0/20.0;  
        ge-0/0/10.0;  
        ge-0/0/0.0;  
    }  
}
```

In configuration mode, check your interface configurations by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]  
user@switch# show interfaces  
ge-0/0/0 {  
    unit 0 {  
        family ethernet-switching {  
            port-mode trunk;  
        }  
    }  
}  
ge-0/0/10 {  
    unit 0 {  
        family ethernet-switching {  
            port-mode trunk;  
        }  
    }  
}  
ge-0/0/20 {  
    unit 0 {  
        family ethernet-switching {  
            port-mode trunk;  
        }  
    }  
}
```

If you are finished configuring the device, enter **commit** in configuration mode.

Configuring ERPS on Hairtail

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

```
set protocols rstp interface ge-0/0/10 disable
set protocols rstp interface ge-0/0/20 disable
set protocols protection-group ethernet-ring erp1 data-channel erp-data-1
set protocols protection-group ethernet-ring erp1 data-channel erp-data-2
set protocols protection-group ethernet-ring erp1 control-vlan erp-control-vlan-1
set protocols protection-group ethernet-ring erp1 east-interface control-channel
ge-0/0/20.0
set protocols protection-group ethernet-ring erp1 west-interface control-channel
ge-0/0/10.0
set protocols protection-group ethernet-ring erp1 west-interface control-channel
ge-0/0/20.0 vlan erp-control-vlan-1
set protocols protection-group ethernet-ring erp1 east-interface control-channel
ge-0/0/10.0 vlan erp-control-vlan-1
```

Step-by-Step Procedure To configure ERPS on Hairtail:

1. Disable RSTP on the two ports that will use ERPS:

```
[edit protocols]
user@switch# set rstp interface ge-0/0/10 disable
user@switch# set rstp interface ge-0/0/20 disable
```
2. Create a node ring named erp1:

```
[edit protocols]
user@switch# set protection-group ethernet-ring erp1
```
3. Configure the control VLAN erp-control-vlan-1 for the node ring erp1:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set control-vlan erp-control-vlan-1
```
4. Configure two data channels named erp-data-1 and erp-data-2 to define a set of VLAN IDs that belong to a ring instance:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set data-channel erp-data-1
user@switch# set data-channel erp-data-2
```
5. Configure the east interface of the node ring erp1 with the control channel ge-0/0/20.0 and indicate that it connects to a ring protection link:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set east-interface control-channel ge-0/0/20.0
```
6. Configure the west interface of the node ring erp1 with the control channel ge-0/0/10.0 and indicate that it connects to a ring protection link:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/10.0
```

7. Every ring instance on a trunk port has one control VLAN in which ERP packets traverse. The control VLAN also controls data VLANs, if any are configured. Assign `erp-control-vlan-1` as the control VLAN:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/10.0 vlan
erp-control-vlan-1
user@switch# set east-interface control-channel ge-0/0/20.0 vlan
erp-control-vlan-1
```

Results In configuration mode, check your ERPS configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@switch# show protocols
rstp {
  interface ge-0/0/10.0 {
    disable;
  }
  interface ge-0/0/20.0 {
    disable;
  }
}
protection-group {
  ethernet-ring erp1 {
    east-interface {
      control-channel {
        ge-0/0/20.0;
      }
    }
    west-interface {
      control-channel {
        ge-0/0/10.0;
      }
    }
    control-vlan erp-control-vlan-1;
    data-channel {
      vlan 101-102;
    }
  }
}
```

In configuration mode, check your VLAN configuration by entering the **show vlans** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@switch# show vlans
erp-control-vlan-1 {
  vlan-id 100;
  interface {
    ge-0/0/20.0;
    ge-0/0/10.0;
  }
}
```

```

}
erp-data-1 {
  vlan-id 101;
  interface {
    ge-0/0/0.0;
    ge-0/0/20.0;
    ge-0/0/10.0;
  }
}
erp-data-2 {
  vlan-id 102;
  interface {
    ge-0/0/0.0;
    ge-0/0/20.0;
    ge-0/0/10.0;
  }
}

```

In configuration mode, check your interface configurations by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

[edit]
user@switch# show interfaces
ge-0/0/0 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
    }
  }
}
ge-0/0/10 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
    }
}
ge-0/0/20 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
    }
}

```

If you are finished configuring the device, enter **commit** in configuration mode.

Configuring ERPS on Jas6-esc

CLI Quick Configuration

To quickly configure Jas6-esc, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

set protocols rstp interface ge-0/0/30 disable
set protocols rstp interface ge-0/0/20 disable
set protocols protection-group ethernet-ring erp1
set protocols protection-group ethernet-ring erp1 data-channel erp-data-1

```

```
set protocols protection-group ethernet-ring erp1 data-channel erp-data-2
set protocols protection-group ethernet-ring erp1 control-vlan erp-control-vlan-1
set protocols protection-group ethernet-ring erp1 east-interface control-channel
ge-0/0/30.0
set protocols protection-group ethernet-ring erp1 west-interface control-channel
ge-0/0/20.0
set protocols protection-group ethernet-ring erp1 west-interface control-channel
ge-0/0/20.0 vlan erp-control-vlan-1
set protocols protection-group ethernet-ring erp1 east-interface control-channel
ge-0/0/30.0 vlan erp-control-vlan-1
```

**Step-by-Step
Procedure**

To configure ERPS on Jas6-esc:

1. Disable RSTP on the two ports that will use ERPS:

```
[edit protocols]
user@switch# set rstp interface ge-0/0/30 disable
user@switch# set rstp interface ge-0/0/20 disable
```
2. Create a node ring named erp1:

```
[edit protocols]
user@switch# set protection-group ethernet-ring erp1
```
3. Configure the control VLAN erp-control-vlan-1 for the node ring erp1:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set control-vlan erp-control-vlan-1
```
4. Configure two data channels named erp-data-1 and erp-data-2 to define a set of VLAN IDs that belong to a ring instance.

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set data-channel erp-data-1
user@switch# set data-channel erp-data-2
```
5. Configure the east interface of the node ring erp1 with the control channel ge-0/0/30.0 :

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set east-interface control-channel ge-0/0/30.0
```
6. Configure the west interface of the node ring erp1 with the control channel ge-0/0/20.0:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/20.0
```
7. Every ring instance on a trunk port has one control VLAN in which ERP packets traverse. The control VLAN also controls data VLANs, if any are configured. Assign erp-control-vlan-1 as the control VLAN:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/20.0 vlan
erp-control-vlan-1
user@switch# set east-interface control-channel ge-0/0/30.0 vlan
erp-control-vlan-1
```

Results In configuration mode, check your ERPS configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@switch# show protocols
  rstp {
    interface ge-0/0/20.0 {
      disable;
    }
    interface ge-0/0/30.0 {
      disable;
    }
  }
  protection-group {
    ethernet-ring erp1 {
      east-interface {
        control-channel {
          ge-0/0/30.0;
        }
      }
      west-interface {
        control-channel {
          ge-0/0/20.0;
        }
      }
      control-vlan erp-control-vlan-1;
      data-channel {
        vlan 101-102;
      }
    }
  }
}
```

In configuration mode, check your VLAN configuration by entering the **show vlans** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@switch# show vlans
  erp-control-vlan-1 {
    vlan-id 100;
    interface {
      ge-0/0/30.0;
      ge-0/0/20.0;
    }
  }
  erp-data-1 {
    vlan-id 101;
    interface {
      ge-0/0/0.0;
      ge-0/0/30.0;
      ge-0/0/20.0;
    }
  }
  erp-data-2 {
    vlan-id 102;
```

```
interface {  
  ge-0/0/0.0;  
  ge-0/0/30.0;  
  ge-0/0/20.0;  
}  
}
```

In configuration mode, check your interfaces configuration by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]  
user@switch# show interfaces  
ge-0/0/0 {  
  unit 0 {  
    family ethernet-switching {  
      port-mode trunk;  
    }  
  }  
}  
ge-0/0/20 {  
  unit 0 {  
    family ethernet-switching {  
      port-mode trunk;  
    }  
  }  
}  
ge-0/0/30 {  
  unit 0 {  
    family ethernet-switching {  
      port-mode trunk;  
    }  
  }  
}
```

Verification

Verify that ERPS is working correctly.

Verifying That ERPS Is Working Correctly

Purpose Verify that ERPS is working on the four EX switches that function as nodes in the ring topology.

Action Check the state of the ring links in the output of the **show protection-group ethernet-ring interface** command. When the ring is configured but not being used (no error exists on the data links), one ERP interface is forwarding traffic and one is discarding traffic. Discarding blocks the ring.

```
user@switch> show protection-group ethernet-ring interface  
Ethernet ring port parameters for protection group erp1  
Interface    Forward State  RPL End  Signal Failure  Admin State  
ge-0/0/2.0   discarding    yes      clear           ready  
ge-0/0/0.0   forwarding    no       clear           ready
```

To find out what has occurred since the last restart, check the RPS statistics for ring-blocked events. **NR** is a No Request ring block, which means that the switch is not blocking either of the two ERP interfaces. **NR-RB** is a No Request Ring Blocked event, which means that the switch is blocking one of its ERP interfaces and sending a packet out to notify the other switches.

```
user@switch> show protection-group ethernet-ring statistics
Ring Name Local SF Remote SF NR Event NR-RB Event
erp1      2      1      2      3
```

Meaning The **show protection-group ethernet-ring interface** command output from the RPL owner node indicates that one interface is forwarding traffic and one is discarding traffic, meaning that the ERP is ready but not active. If at least one interface in the ring is not forwarding, the ring is blocked and therefore ERP is working.

The **show protection-group ethernet-ring statistics** command output indicates that, since the last reboot, both local and remote signal failures have occurred (**Local SF** and **Remote SF**).

The **NR Event** count is 2, indicating that the NR state was entered into twice. **NR** stands for No Request. This means that the switch either originated NR PDUs or received an NR PDU from another switch and stopped blocking the interface to allow ERP to function.

The three **NR-RB** events indicate that on three occasions, this switch either sent out NR-RB PDUs or received NR-RB PDUs from another switch. This occurs when a network problem is resolved and the switch once again blocks the ERP link at one end.

- Related Documentation**
- [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 190](#)
 - [Ethernet Ring Protection Switching Overview on page 39](#)
 - [Understanding Ethernet Ring Protection Switching Functionality on page 34](#)

CHAPTER 7

Configuration Tasks

- [Configuring VLANs for EX Series Switches \(J-Web Procedure\) on page 158](#)
- [Configuring VLANs for EX Series Switches \(CLI Procedure\) on page 160](#)
- [Configuring Routed VLAN Interfaces \(CLI Procedure\) on page 164](#)
- [Configuring MAC Table Aging \(CLI Procedure\) on page 166](#)
- [Configuring the Native VLAN Identifier \(CLI Procedure\) on page 167](#)
- [Creating a Series of Tagged VLANs \(CLI Procedure\) on page 168](#)
- [Creating a Private VLAN on a Single EX Series Switch \(CLI Procedure\) on page 169](#)
- [Creating a Private VLAN Spanning Multiple EX Series Switches \(CLI Procedure\) on page 171](#)
- [Configuring a Routed VLAN Interface in a Private VLAN \(CLI Procedure\) on page 173](#)
- [Configuring Virtual Routing Instances \(CLI Procedure\) on page 174](#)
- [Configuring MAC Notification \(CLI Procedure\) on page 175](#)
- [Configuring Multiple VLAN Registration Protocol \(MVRP\) \(CLI Procedure\) on page 176](#)
- [Configuring Q-in-Q Tunneling \(CLI Procedure\) on page 180](#)
- [Configuring Layer 2 Protocol Tunneling on EX Series Switches \(CLI Procedure\) on page 181](#)
- [Configuring Redundant Trunk Groups \(J-Web Procedure\) on page 183](#)
- [Configuring Redundant Trunk Links for Faster Recovery \(CLI Procedure\) on page 185](#)
- [Configuring Proxy ARP \(CLI Procedure\) on page 186](#)
- [Adding a Static MAC Address Entry to the Ethernet Switching Table \(CLI Procedure\) on page 187](#)
- [Configuring Edge Virtual Bridging \(CLI Procedure\) on page 188](#)
- [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 190](#)
- [Enabling VLAN Pruning for Broadcast, Multicast, and Unknown Unicast Traffic in an EX Series Virtual Chassis \(CLI Procedure\) on page 192](#)

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



NOTE: This topic applies only to the J-Web Application package.

You can use the VLAN Configuration page to add a new VLAN or to edit or delete an existing VLAN on an EX Series switch.

To access the VLAN Configuration page:

1. Select **Configure > 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.



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

2. Click one of the following options:

- **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 21 on page 158](#).

Table 21: VLAN Configuration Details

Field	Function	Your Action
General tab		
VLAN Name	Specifies a unique name for the VLAN.	Enter a name.

Table 21: VLAN Configuration Details (*continued*)

Field	Function	Your Action
VLAN ID/Range/VLAN ID/List NOTE: EX4300 switches support only VLAN ID/List and not VLAN Range.	Specifies the identifier or range for the VLAN.	Select one of the following options: <ul style="list-style-type: none"> VLAN ID—Type a unique identification number from 1 through 4094. If no value is specified, the ID defaults to 0. VLAN Range/List—Type a number range to create VLANs with IDs corresponding to the numbers in the range. For example, the range 2–3 creates two VLANs with the IDs 2 and 3.
Description	Describes the VLAN.	Enter a brief description for the VLAN.
MAC-Table-Aging-Time NOTE: This option is not supported on EX4300 switches.	Specifies the maximum time that an entry can remain in the forwarding table before it <i>ages out</i> .	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 NOTE: This option is not supported on EX4300 switches.	Specifies the ports (interfaces) to be associated with this VLAN for data traffic. You can also remove the port association.	Click one of the following options: <ul style="list-style-type: none"> Add—Select the ports from the available list. For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list. Remove—Select the port that you do not want associated with the VLAN.
IP address tab		
IPv4 address	Specifies IPv4 address options for the VLAN.	Select IPv4 address to enable the IPv4 address options. To configure IPv4: <ol style="list-style-type: none"> Enter the IP address. Enter the subnet mask—for example, 255.255.255.0. You can also specify the address prefix. To apply an input firewall filter to an interface, select the firewall filter from the list. To apply an output firewall filter to an interface, select the firewall filter from the list. Click the ARP/MAC Details button. Enter the static IP address and MAC address in the window that is displayed. NOTE: In EX4300 switches, you also need to select L2 Interface in the window that is displayed.

Table 21: VLAN Configuration Details (*continued*)

Field	Function	Your Action
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	Specifies the ports to be associated with this VLAN for voice traffic. You can also remove the port association.	<p>Click one of the following options:</p> <ul style="list-style-type: none"> • Add—Select the ports from the list of available ports. For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list. • Remove—Select the port that you do not want associated with the VLAN.

Related Documentation

- [Configuring VLANs for EX Series Switches \(CLI Procedure\)](#)
- [Configuring VLANs for EX Series Switches \(CLI Procedure\) on page 160](#)
- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch](#)
- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 57](#)
- [Understanding Bridging and VLANs on EX Series Switches on page 3](#)
- [Configuring Integrated Routing and Bridging Interfaces \(CLI Procedure\)](#)
- [Configuring Routed VLAN Interfaces \(CLI Procedure\) on page 164](#)

Configuring VLANs for EX Series Switches (CLI Procedure)



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

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

- [Why Create a VLAN? on page 161](#)
- [Create a VLAN Using the Minimum Procedure on page 161](#)

- [Create a VLAN Using All of the Options on page 162](#)
- [Configuration Guidelines for VLANs on page 163](#)

Why Create a VLAN?

Some reasons to create VLANs are:

- A LAN has more than 200 devices.
- A LAN has a large amount of broadcast traffic.
- A group of clients requires that a higher-than-average level of security be applied to traffic entering or exiting the group's devices.
- A group of clients requires that the group's devices receive less broadcast traffic than they are currently receiving, so that data speed across the group is increased.

Create a VLAN Using the Minimum Procedure

Two steps are required to create a VLAN:

- Uniquely identify the VLAN. You do this by assigning either a name or an ID (or both) to the VLAN. When you assign just a VLAN name, an ID is generated by Junos OS.
- Assign at least one switch port interface to the VLAN for communication. All interfaces in a single VLAN are in a single broadcast domain, even if the interfaces are on different switches. You can assign traffic on any switch to a particular VLAN by referencing either the interface sending traffic or the MAC addresses of devices sending traffic.

The following example creates a VLAN using only the two required steps. The VLAN is created with the name `employee-vlan`. Then, three interfaces are assigned to that VLAN so that the traffic is transmitted among these interfaces.



NOTE: In this example, you could alternatively assign an ID number to the VLAN. The requirement is that the VLAN have a unique ID.

```
[edit]set vlans employee-vlan set interfaces ge-0/0/1 unit 0 family ethernet-switching
vlan members employee-vlan set interfaces ge-0/0/2 unit 0 family ethernet-switching
vlan members employee-vlan set interfaces ge-0/0/3 unit 0 family ethernet-switching
vlan members employee-vlan
```

In the example, all users connected to the interfaces `ge-0/0/1`, `ge-0/0/2`, and `ge-0/0/3` can communicate with each other, but not with users on other interfaces in this network. To configure communication between VLANs, you must configure a routed VLAN interface (RVI). See [“Configuring Routed VLAN Interfaces \(CLI Procedure\)” on page 164](#).

Create a VLAN Using All of the Options

To configure a VLAN, follow these steps:

1. In configuration mode, create the VLAN by setting the unique VLAN name:

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

2. Configure the VLAN tag ID or VLAN ID range for the VLAN. (If you assigned a VLAN name, you do not have to do this, because a VLAN ID is assigned automatically, thereby associating the name of the VLAN to an ID number. However, if you want to control the ID numbers, you can assign both a name and an ID.)

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

or

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

3. Assign at least one interface to the VLAN:

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



NOTE: You can also specify that a 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.

4. (Optional) Create a subnet for the VLAN because all computers that belong to a subnet are addressed with a common, identical, most-significant-bit group in their IP address. This makes it easy to identify VLAN members by their IP addresses. To create the subnet for the VLAN:

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

5. (Optional) Specify the description of the VLAN:

```
[edit]user@switch# set vlans vlan-name description text-description
```

6. (Optional) To avoid exceeding the maximum number of members allowed in a VLAN, specify the maximum time that an entry can remain in the forwarding table before it ages out:

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

7. (Optional) For security purposes, specify a VLAN firewall filter to be applied to incoming or outgoing packets:

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

8. (Optional) For accounting purposes, enable a counter to track the number of times this VLAN is accessed:

```
[edit]user@switch# set vlans vlan-name l3-interface ingress-counting l3-interface-name
```

9. (Optional) For Virtual Chassis bandwidth management purposes, enable VLAN Pruning to ensure all broadcast, multicast, and unknown unicast traffic entering the Virtual Chassis on the VLAN uses the shortest possible path through the Virtual Chassis:

```
[edit]
```

```
user@switch# set vlans vlan-name vlan-prune
```

Configuration Guidelines for VLANs

Two steps are required to create a VLAN. You must uniquely identify the VLAN and you must assign at least one switch port interface to the VLAN for communication.

After creating a VLAN, all users all users connected to the interfaces assigned to the VLAN can communicate with each other but not with users on other interfaces in the network. To configure communication between VLANs, you must configure a routed VLAN interface (RVI). See [“Configuring Routed VLAN Interfaces \(CLI Procedure\)” on page 164](#) to create an RVI.

The number of VLANs supported per switch varies for each switch type. Use the command **set vlans id vlan-id ?** to discover the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because each VLAN is assigned an ID number when it is created. You can, however, exceed the recommended VLAN member maximum . To determine the maximum number of VLAN members allowed on a switch, multiply the VLAN maximum obtained using **set vlans id vlan-id ?** times 8.

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

Related Documentation

- [Configuring VLANs for EX Series Switches \(J-Web Procedure\) on page 158](#)
- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 57](#)
- [Creating a Series of Tagged VLANs \(CLI Procedure\) on page 168](#)
- [Understanding Bridging and VLANs on EX Series Switches on page 3](#)
- [Understanding Integrated Routing and Bridging Interfaces and Routed VLAN Interfaces on EX Series Switches on page 12](#)

Configuring Routed VLAN Interfaces (CLI Procedure)

Routed VLAN interfaces (RVIs) allow the EX Series switch to recognize packets that are being sent to local addresses so that they are bridged (switched) whenever possible and are routed only when necessary. Whenever packets can be switched instead of routed, several layers of processing are eliminated.

An interface named `vlan` functions as a logical router on which you can configure a Layer 3 logical interface for each virtual LAN (VLAN). For redundancy, you can combine an RVI with implementations of the Virtual Router Redundancy Protocol (VRRP) in both bridging and virtual private LAN service (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 VLAN that you have associated with the RVI, as well as on the RVI itself (the interface named `vlan`).



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

To configure the RVI:

1. Create a Layer 2 VLAN by assigning it a name and a VLAN ID:

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

2. Assign an interface to the VLAN 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 interface-name unit logical-unit-number family ethernet-switching
vlan members vlan-name
```

3. Create a logical Layer 3 RVI (its name will be `vlan.logical-interface-number`, where the value for *logical-interface-number* is the value you supplied for *vlan-id* in Step 1; in the following command, it is the *logical-unit-number*) on a subnet for the VLAN's broadcast domain:

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

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

```
[edit]
user@switch# set vlans vlan-name l3-interface vlan.logical-interface-number
```




NOTE: Layer 3 interfaces on trunk ports allow the interface to transfer traffic between multiple Layer 2 VLANs. Within a VLAN, traffic is switched, while across VLANs, traffic is routed.

5. (Optional) On an EX8200 switch, enable an input counter for tracking or billing purposes:

[edit]

```
user@switch# set vlans vlan-name l3-interface vlan logical-interface-number
l3-interface-ingress-counting
```



NOTE: The input counter is maintained by a firewall filter—these counters are allocated on a first-come, first-served basis.

Related Documentation

- [Verifying Routed VLAN Interface Status and Statistics on page 291](#)
- [Understanding Integrated Routing and Bridging Interfaces and Routed VLAN Interfaces on EX Series Switches on page 12](#)

Configuring MAC Table Aging (CLI Procedure)



NOTE: This task uses Junos OS for EX Series switches that do not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Configuring MAC Table Aging (CLI Procedure)*.

The Ethernet switching table (or MAC table) aging process ensures that the EX Series switch tracks only active MAC addresses on the network and is able to flush out MAC addresses that are no longer used.

You can configure the MAC table aging time, the maximum time that an entry can remain in the Ethernet Switching table before it “ages out,” either on all VLANs on the switch or on particular VLANs. This setting can influence efficiency of network resource use by affecting the amount of traffic that is flooded to all interfaces because when traffic is received for MAC addresses no longer in the Ethernet switching table, the switch floods the traffic to all interfaces.

To configure the MAC table aging time on all VLANs on the switch:

```
[edit]
user@switch# set ethernet-switching-options mac-table-aging-time seconds
```

To configure the MAC table aging time on a VLAN:

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



NOTE: You can set the MAC table aging time to unlimited. If you specify the value as unlimited, entries are never removed from the table. Generally, use this setting only if the switch or the VLAN has a fairly static number of end devices; otherwise the table will eventually fill up. You can use this setting to minimize traffic loss and flooding that might occur when traffic arrives for MAC addresses that have been removed from the table.

Related Documentation

- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 57](#)
- [Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 65](#)
- [Example: Connecting an Access Switch to a Distribution Switch on page 72](#)
- [Controlling Authentication Session Timeouts \(CLI Procedure\)](#)
- [Understanding Bridging and VLANs on EX Series Switches on page 3](#)

Configuring the Native VLAN Identifier (CLI Procedure)



NOTE: This task uses Junos OS for EX Series switches and Junos OS for QFX3500 and QFX3600 switches that does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Configuring the Native VLAN Identifier (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

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 Documentation

- [Understanding Bridging and VLANs on EX Series Switches on page 3](#)
- [Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 65](#)
- [Example: Connecting an Access Switch to a Distribution Switch on page 72](#)
- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 57](#)

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 Documentation

- [Verifying That a Series of Tagged VLANs Has Been Created on page 287](#)
- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 57](#)
- [Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 65](#)
- [Example: Connecting an Access Switch to a Distribution Switch on page 72](#)
- [Understanding Bridging and VLANs on EX Series Switches on page 3](#)

Creating a Private VLAN on a Single EX Series Switch (CLI Procedure)

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 enables you to split a broadcast domain, also known as a primary VLAN, into multiple isolated broadcast subdomains, also known as secondary VLANs. Splitting the primary VLAN into secondary VLANs essentially nests a VLAN inside another VLAN. This topic describes how to configure a PVLAN on a single switch.

Before you begin, configure names for all secondary VLANs that will be part of the primary VLAN. (Unlike the secondary VLANs, you do not need to preconfigure the primary VLAN—this procedure provides the complete configuration of the primary VLAN.) Although tags are not needed when a secondary VLAN is configured on a single switch, configuring a secondary VLAN as tagged does not adversely affect its functionality. For instructions on configuring the secondary VLANs, see “[Configuring VLANs for EX Series Switches \(CLI Procedure\)](#)” on page 160.

Keep these rules in mind when configuring a PVLAN on a single switch:

- The primary VLAN must be a tagged VLAN.
- Configuring a VoIP VLAN on PVLAN interfaces is not supported.

To configure a private VLAN on a single switch:

1. Set the VLAN ID for the primary VLAN:

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

2. Set the interfaces and port modes:

```
[edit interfaces]
user@switch# set interface-name unit 0 family ethernet-switching port-mode mode
user@switch# set interface-name unit 0 family ethernet-switching vlan members (all | vlan-id | vlan-number)
```

3. Configure the access ports in the primary VLAN to not forward packets to one another:

```
[edit vlans]
user@switch# set vlan-id vlan-id-number no-local-switching
```

4. For each community VLAN, configure access interfaces:

```
[edit vlans]
user@switch# set community-vlan-name interface interface-name
```

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

```
[edit vlans]
user@switch# set community-vlan-name primary-vlan primary-vlan-name
```

Isolated VLANs are not configured as part of this process. Instead, they are created internally if **no-local-switching** is enabled on the primary VLAN and the isolated VLAN has access interfaces as members.

To optionally enable routing between isolated and community VLANs by using a routed VLAN interface (RVI) instead of a promiscuous port connected to a router, see [“Configuring a Routed VLAN Interface in a Private VLAN \(CLI Procedure\)”](#) on page 173.



NOTE: Only an EX8200 switch or EX8200 Virtual Chassis support the use of an RVI to route Layer 3 traffic between isolated and community VLANs in a PVLAN domain.

Related Documentation

- [Example: Configuring a Private VLAN on a Single EX Series Switch on page 81](#)
- [Creating a Private VLAN Spanning Multiple EX Series Switches \(CLI Procedure\) on page 171](#)
- [Verifying That a Private VLAN Is Working on page 292](#)
- [Understanding Private VLANs on EX Series Switches on page 16](#)

Creating a Private VLAN Spanning Multiple EX Series Switches (CLI Procedure)

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 enables an administrator to split a broadcast domain, also known as a primary VLAN, into multiple isolated broadcast subdomains, also known as secondary VLANs. Splitting the primary VLAN into secondary VLANs essentially nests a VLAN inside another VLAN. This topic describes how to configure a PVLAN to span multiple switches.

Before you begin, configure names for all secondary VLANs that will be part of the primary VLAN. (Unlike the secondary VLANs, you do not need to preconfigure the primary VLAN—this procedure provides the complete configuration of the primary VLAN.) For instructions on configuring the secondary VLANs, see [“Configuring VLANs for EX Series Switches \(CLI Procedure\)” on page 160](#).

The following rules apply to creating PVLANS:

- The primary VLAN must be a tagged VLAN.
- You must configure the primary VLAN and the PVLAN trunk port before configuring the secondary VLANs.
- Configuring a VoIP VLAN on PVLAN interfaces is not supported.
- If the Multiple VLAN Registration Protocol (MVRP) is configured on the PVLAN trunk port, the configuration of secondary VLANs and the PVLAN trunk port must be committed with the same commit operation.

To configure a private VLAN to span multiple switches:

1. Configure a name and an 802.1Q tag for the primary VLAN:

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

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

```
[edit vlans]
user@switch# set primary-vlan-name no-local-switching
```

3. Set the PVLAN trunk interface that will connect the primary VLAN to the neighboring switch:

```
[edit vlans]
user@switch# set primary-vlan-name interface interface-name pvlan-trunk
```

4. Configure a name and 802.1Q tag for a community VLAN that spans the switches:

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

5. Add access interfaces to the community VLAN:

```
[edit vlans]
user@switch# set community-vlan-name interface interface-name
```

6. Specify the primary VLAN of the specified community VLAN:

```
[edit vlans]
user@switch# set community-vlan-name primary-vlan primary-vlan-name
```

7. Add the isolated interface to the specified primary VLAN:

[edit vlans]

user@switch# **set primary-vlan-name interface interface-name**

.....



NOTE: To configure an isolated interface, include it as one of the members of the primary VLAN, but do not configure it as belonging to one of the community VLANs.

.....

8. Set the 802.1Q tag of the interswitch isolated VLAN:

[edit vlans]

user@switch# **set primary-vlan-name isolation-id number**

802.1Q tags are required for interswitch isolated VLANs because IEEE 802.1Q uses an internal tagging mechanism by which a trunking device inserts a 4-byte VLAN frame identification tag into the packet header.

To optionally enable routing between isolated and community VLANs by using a routed VLAN interface (RVI) instead of a promiscuous port connected to a router, see [“Configuring a Routed VLAN Interface in a Private VLAN \(CLI Procedure\)”](#) on page 173.

.....



NOTE: Only an EX8200 switch or EX8200 Virtual Chassis support the use of an RVI to route Layer 3 traffic between isolated and community VLANs in a PVLAN domain.

.....

Related Documentation

- [Example: Configuring a Private VLAN Spanning Multiple EX Series Switches on page 119](#)
- [Verifying That a Private VLAN Is Working on page 292](#)
- [Creating a Private VLAN on a Single EX Series Switch \(CLI Procedure\) on page 169](#)
- [Understanding Private VLANs on EX Series Switches on page 16](#)
- [Understanding PVLAN Traffic Flows Across Multiple Switches on page 22](#)

Configuring a Routed VLAN Interface in a Private VLAN (CLI Procedure)

On an EX8200 switch or EX8200 Virtual Chassis, you can configure a routed VLAN interface (RVI) for a private VLAN (PVLAN). Instead of a router connected to a promiscuous port routing Layer 3 traffic between isolated and community members, you can alternatively use an RVI.

To set up routing within a PVLAN, one RVI must be configured for the primary VLAN on one EX8200 switch or EX8200 Virtual Chassis in the PVLAN domain. This RVI serves the entire PVLAN domain regardless of whether the domain consists of one or more switches. After you configure the RVI, Layer 3 packets received by the secondary VLAN interfaces are mapped to and routed by the RVI.

When setting up the RVI, you must also enable proxy Address Resolution Protocol (ARP) so that the RVI can handle ARP requests received by the secondary VLAN interfaces.

This topic describes how to configure an RVI for a PVLAN.

Before you begin, configure the PVLAN as described in [“Creating a Private VLAN on a Single EX Series Switch \(CLI Procedure\)” on page 169](#) or [“Creating a Private VLAN Spanning Multiple EX Series Switches \(CLI Procedure\)” on page 171](#).

To configure an RVI for a PVLAN:

1. Create a logical Layer 3 RVI on a subnet for the primary VLAN's broadcast domain:

```
[edit interfaces]
```

```
user@switch# set vlan unit logical-unit-number family inet address inet-address
```

2. Enable unrestricted proxy ARP on the RVI:

```
[edit interfaces]
```

```
user@switch# set vlan unit logical-unit-number proxy-arp unrestricted
```

3. Disable sending protocol redirect messages on the RVI:

```
[edit interfaces]
```

```
user@switch# set vlan unit logical-unit-number family inet no-redirects
```

4. Link the primary VLAN to the RVI:

```
[edit vlans]
```

```
user@switch# set vlan-name l3-interface vlan.logical-unit-number
```

The value of *logical-unit-number* is the same value that you supplied for *logical-unit-number* in the previous steps.

Related Documentation

- [Understanding Private VLANs on EX Series Switches on page 16](#)

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 traffic traversing the network without using multiple devices to segment your network. VRF is supported on all Layer 3 interfaces.

Before you begin, make sure to set up your VLANs. See [“Configuring VLANs for EX Series Switches \(CLI Procedure\)” on page 160](#), [Configuring VLANs for EX Series Switches \(CLI Procedure\)](#), or [“Configuring VLANs for EX Series Switches \(J-Web Procedure\)” on page 158](#).

To configure virtual routing instances:

1. Create a routing instance:

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



NOTE: EX Series switches only support the virtual-router instance type.

2. Bind each routing instance to the corresponding physical interfaces:

```
[edit routing-instances]user@switch# set routing-instance-name interface
interface-name.logical-unit-number
```

3. Create the logical interfaces that are bound to the routing instance.

- To create a logical interface with an IPv4 address:

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

- To create a logical interface with an IPv6 address:

```
[edit interfaces]user@switch# set interface-name unit logical-unit-number family inet6
address ipv6-address
```



NOTE: Do not create a logical interface using the family ethernet-switching option in this step. Binding an interface using the family ethernet-switching option to a routing instance can cause the interface to shutdown.

4. Enable VLAN tagging on each physical interface that was bound to the routing instance:

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

Related Documentation

- [Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 88](#)
- [Verifying That Virtual Routing Instances Are Working on page 289](#)
- [Understanding Virtual Routing Instances on EX Series Switches on page 25](#)

Configuring MAC Notification (CLI Procedure)



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

When a switch learns or unlearns a MAC address, SNMP notifications can be sent to the network management system at regular intervals to record the addition or removal of the MAC address. This process is known as MAC notification.

The MAC notification interval defines how often Simple Network Management Protocol (SNMP) notifications logging the addition or removal of MAC addresses on the switch are sent to the network management system.

MAC notification is disabled by default. When MAC notification is enabled, the default MAC notification interval is 30 seconds.

To enable or disable MAC notification, or to set the MAC notification interval, perform these tasks:

- [Enabling MAC Notification on page 175](#)
- [Disabling MAC Notification on page 175](#)
- [Setting the MAC Notification Interval on page 176](#)

Enabling MAC Notification

MAC notification is disabled by default. You need to perform this procedure to enable MAC notification.

To enable MAC notification on the switch with the default MAC notification interval of 30 seconds:

```
[edit ethernet-switching-options]
user@switch# set mac-notification
```

To enable MAC notification on the switch with any other MAC notification interval (here, the MAC notification interval is set to 60 seconds):

```
[edit ethernet-switching-options]
user@switch# set mac-notification notification-interval 60
```

Disabling MAC Notification

MAC Notification is disabled by default. Perform this procedure only if MAC notification was previously enabled on your switch.

To disable MAC notification on the switch:

```
[edit ethernet-switching-options]
user@switch# delete mac-notification
```

Setting the MAC Notification Interval

The default MAC notification interval is 30 seconds. The procedure to change the MAC notification interval to a different interval is identical to the procedure to enable MAC notification on the switch with a nondefault value for the MAC notification interval.

To set the MAC notification interval on the switch (here, the MAC notification interval is set to 5 seconds):

```
[edit ethernet-switching-options]
user@switch# set mac-notification notification-interval 5
```

Related Documentation

- [Verifying That MAC Notification Is Working Properly on page 299](#)

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



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

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 176](#)
- [Disabling MVRP on page 177](#)
- [Disabling Dynamic VLANs on page 177](#)
- [Configuring Timer Values on page 177](#)
- [Configuring MVRP Registration Mode on page 178](#)
- [Using MVRP in a Mixed-Release Network on page 178](#)

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
```

Using MVRP in a Mixed-Release Network

Except for in Junos OS Releases 11.2 and earlier, MVRP has conformed with IEEE standard 802.1ak and IEEE Draft 802.1Q regarding the inclusion of an extra byte in the protocol data units (PDUs) sent and received by MVRP. As a result of the non-conformance of releases 11.2 and earlier and changes in the standards regarding the extra byte, the following mixed environments can arise:

- Mixed environment A: MVRP in Junos OS Releases 11.2 and earlier includes the extra byte, while MVRP in Junos OS Releases 11.3 and later for EX Series switches that do not support the Enhanced Layer 2 Software (ELS) configuration style does not include the extra byte.
- Mixed environment B: MVRP in Junos OS Releases 13.2 and later for EX Series switches with support for ELS includes the extra byte, while MVRP in Junos OS Releases 11.3 and later for EX Series switches that do not support ELS does not include the extra byte.

A compatibility issue arises in mixed environments A and B, wherein the versions of MVRP that include the extra byte do not recognize PDUs that do not include the extra byte. For more information about this issue, see [“Understanding Multiple VLAN Registration Protocol \(MVRP\) on EX Series Switches” on page 26](#).

If your network has a mix of MVRP versions, you can alter MVRP on the switches running Release 11.3 and later on switches that do not support ELS so they include the extra byte in the PDU and are therefore, compatible with the other MVRP versions.

You can recognize an MVRP version problem by looking at a switch running an MVRP version that includes the extra byte. Because a switch running an MVRP version that includes the extra byte cannot interpret an unmodified PDU from an MVRP version that does not include the extra byte, the switch will not add VLANs from the MVRP version that does not include the extra byte. When you execute the command **show mvrp statistics** on the MVRP version that includes the extra byte, the values for *Join Empty received* and *Join In received* will incorrectly display zero, even though the value for *MRPDU received* has been increased. Another indication that MVRP is having a version problem is that unexpected VLAN activity, such as multiple VLAN creation, takes place on the switch running the MVRP version that includes the extra byte.

To make MVRP on Release 11.3 or later compatible with MVRP in the other releases:

```
[edit protocols mvrp]
user@switch# set add-attribute-length-in-pdu
```

Related Documentation

- [Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 92](#)
- [Verifying That MVRP Is Working Correctly on page 297](#)

Configuring Q-in-Q Tunneling (CLI Procedure)



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

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.



NOTE: You cannot configure 802.1X user authentication on interfaces that have been enabled for Q-in-Q tunneling.

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 160 or “Configuring VLANs for EX Series Switches (J-Web Procedure)” on page 158.

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 (VLANs)**

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 (Q-in-Q VLANs)**

Related Documentation

- [Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 103](#)
- [Verifying That Q-in-Q Tunneling Is Working on page 290](#)
- [Understanding Q-in-Q Tunneling on EX Series Switches on page 41](#)

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



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

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: L2PT and VLAN translation configured with the **mapping** statement cannot both be configured on the same switch.



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 41 and “[Configuring Q-in-Q Tunneling \(CLI Procedure\)](#)” on page 180.

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 reenabling it using the `clear ethernet-switching layer2-protocol-tunneling error` command. Otherwise, the interface remains disabled.

Related Documentation

- [Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 107](#)
- [Understanding Layer 2 Protocol Tunneling on EX Series Switches on page 45](#)

Configuring Redundant Trunk Groups (J-Web Procedure)



NOTE: This topic applies only to the J-Web Application package.

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 Spanning Tree Protocol (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 by using the J-Web interface:

1. Select **Configure > 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.



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

2. Click one of the following options:

- **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 22 on page 184](#).

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

Table 22: RTG Configuration Fields

Field	Function	Your Action
Group Name	Specifies a unique name for the RTG.	Enter a name. NOTE: Only on EX4300 switches, you can select the name from a list.
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 select the active interface.	Select the option button.

- Related Documentation**
- *Example: Configuring Redundant Trunk Links for Faster Recovery*
 - [Example: Configuring Redundant Trunk Links for Faster Recovery on page 112](#)
 - [Understanding Redundant Trunk Links on page 49](#)

Configuring Redundant Trunk Links for Faster Recovery (CLI Procedure)

You can manage network convergence by configuring both a primary link and a secondary link on an EX Series switch; this is called a redundant trunk group (RTG). If the primary link in a redundant trunk group fails, it passes its known MAC address locations to the secondary link, which automatically takes over. You can configure a maximum of 16 redundant trunk groups on most standalone switches or on Virtual Chassis. The EX8200 switch and EX8200 Virtual Chassis, however, support up to 254 redundant trunk groups.

Generally, you configure a redundant trunk group by configuring one primary link (and its interface) and one unspecified link (and its interface) to serve as the secondary link. A second type of redundant trunk group, not shown in the procedure in this topic, consists of two unspecified links (and their interfaces); in this case, neither of the links is primary. In this second case, the software selects an active link by comparing the port numbers of the two links and activating the link with the higher port number. The procedure given here describes configuring a primary/unspecified configuration for a redundant trunk group because that configuration gives you more control and is more commonly used.

Rapid Spanning Tree Protocol (RSTP) is enabled by default on EX Series switches to create a loop-free topology, but an interface is not allowed to be in both a redundant trunk group and in a spanning-tree protocol topology at the same time.

A primary link takes over whenever it is able. You can, however, alter the number of seconds that the primary link waits before reestablishing control by configuring the primary link's preempt cutover timer.

Before you configure the redundant trunk group on the switch, be sure you have:

- Disabled RSTP on all switches that will be linked to your redundant trunk group.
- Configured at least two interfaces with their port mode set to **trunk**; be sure that these two interfaces are not part of any existing RTG. See *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*.

To configure a redundant trunk group on a switch:

1. Turn off RSTP:

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

2. Name the redundant trunk group while configuring one primary and one unspecified trunk interface:

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

3. (Optional) Change the length of time (from the default 120 seconds) that a re-enabled primary link waits to take over from an active secondary link:

```
[edit ethernet-switching-options]
set redundant-trunk-group group name preempt-cutover-timer seconds
```

- Related Documentation**
- [Example: Configuring Redundant Trunk Links for Faster Recovery on page 112](#)
 - [Understanding Redundant Trunk Links on page 49](#)

Configuring Proxy ARP (CLI Procedure)



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

You can configure proxy Address Resolution Protocol (ARP) on your EX Series switch to enable the switch to respond to ARP queries for network addresses by offering its own media access control (MAC) address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

To configure proxy ARP on a single interface:

```
[edit interfaces]
user@switch# set ge-0/0/3 unit 0 proxy-arp restricted
```



BEST PRACTICE: We recommend that you configure proxy ARP in restricted mode. In restricted mode, the switch is not a proxy if the source and target IP addresses are on the same subnet. If you use unrestricted mode, disable gratuitous ARP requests on the interface to avoid the situation of the switch's response to a gratuitous ARP request appearing to the host to be an indication of an IP conflict:

To configure proxy ARP on a routed VLAN interface (RVI):

```
[edit interfaces]
user@switch# set vlan unit 100 proxy-arp restricted
```

- Related Documentation**
- [Example: Configuring Proxy ARP on an EX Series Switch on page 117](#)
 - [Verifying That Proxy ARP Is Working Correctly on page 299](#)
 - [Configuring Routed VLAN Interfaces \(CLI Procedure\) on page 164](#)

Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)



NOTE: This task uses Junos OS for EX Series switches and Junos OS for QFX3500 and QFX3600 switches that does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

The Ethernet switching table, also known as the forwarding table, specifies the known locations of VLAN nodes. There are two ways to populate the Ethernet switching table on a switch. The easiest method is to let the switch update the table with MAC addresses.

The second way to populate the Ethernet switching table is to manually insert a VLAN node location into the table. You can do this to reduce flooding and speed up the switch's automatic learning process. To further optimize the switching process, indicate the next hop (next interface) packets will use after leaving the node.

Before configuring a static MAC address, be sure that you have:

- Set up the VLAN. See [“Configuring VLANs for EX Series Switches \(CLI Procedure\)” on page 160](#) or *Configuring VLANs*.

To add a MAC address to the Ethernet switching table:

1. Specify the MAC address to add to the table:

```
[edit ethernet-switching-options]
set static vlan vlan-name mac mac-address
```

2. Indicate the next hop MAC address for packets sent to the indicated MAC address:

```
[edit ethernet-switching-options]
set static vlan vlan-name mac mac-address next-hop interface
```

Related Documentation

- [Understanding Bridging and VLANs on EX Series Switches on page 3](#)
- *Understanding Bridging and VLANs*

Configuring Edge Virtual Bridging (CLI Procedure)

Configure edge virtual bridging (EVB) when a switch is connected to a virtual machine (VM) server using virtual Ethernet port aggregator (VEPA) technology. EVB does not convert packets; rather, it ensures that packets from one VM destined for another VM on the same VM server is switched. In other words, when the source and destination of a packet are the same port, EVB delivers the packet properly, which otherwise would not happen.



NOTE: Configuring EVB also enables Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP).

Before you begin configuring EVB, ensure that you have:

- Configured packet aggregation on the server connected to the port that you will use on the switch for EVB. See your server documentation.
- Configured the EVB interface for all VLANs located on the virtual machines. See [“Configuring VLANs for EX Series Switches \(CLI Procedure\)” on page 160](#).



NOTE: The port security features MAC move limiting and MAC limiting are supported on interfaces that are configured for EVB; however, the port security features IP source guard, dynamic ARP inspection (DAI), and DHCP snooping are not supported by EVB. For more information about these features, see *Understanding Port Security*.

To configure EVB on the switch:

1. Configure tagged-access mode for the interfaces on which you will enable EVB:

```
[edit interfaces interface-name]
user@switch# set unit 0 family ethernet-switching port-mode tagged-access
```

2. Enable the Link Layer Discovery Protocol (LLDP) on the interfaces on which you will enable EVB:

```
[edit protocols]
user@switch# set lldp interface interface-name
```

3. Configure the interfaces for EVB as members of all VLANs located on the virtual machines.

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

4. Enable VDP on the interfaces:

```
[edit protocols]
user@switch# set edge-virtual-bridging vsi-discovery interface interface-name
```

5. Define policies for VSI information, including a VSI manager ID, VSI type, VSI version, and VSI instance ID:

```
[edit policy-options]
```



```

user@switch# set vsi-policy policy-name from vsi-manager manager-number vsi-type
type-number vsi-version version-number vsi-instance instance-number
user@switch# set vsi-policy policy-name then filter filter-name

```

6. Define the firewall filters you mapped to in the previous step. When each incoming packet matches the filter, the count is incremented by 1. Other possible actions are accept and drop.

```

[edit firewall family ethernet-switching]
user@switch# set filter filter-name term term-name then action

```

7. Associate VSI policies with VDP:

```

[edit protocols]
user@switch# set edge-virtual-bridging vsi-discovery vsi-policy policy-name

```

8. Verify that the virtual machine successfully associated with the switch. After successful association of the VSI Profile with the switch interface, verify the learning of the VM's MAC address on MAC-Table or Forwarding database Table. The learn type of the VM's MAC addresses will be VDP, and upon successful shutdown of VM the corresponding MAC-VLAN entry will get flushed out from FDB table otherwise it will never shutdown.

```

admin@host# run show ethernet-switching table

```

9. Verify that VSI profiles are being learned at the switch:

```

user@switch# show edge-virtual-bridging vsi-profiles

```

10. Check the statistics of ECP packet exchanges between the switch and server:

```

user@switch# show edge-virtual-bridging ecp statistics

```

Related Documentation

- [Example: Configuring Edge Virtual Bridging for Use with VEPA Technology on page 134](#)
- [Understanding Edge Virtual Bridging for Use with VEPA Technology on page 33](#)

Configuring Ethernet Ring Protection Switching (CLI Procedure)

You can configure Ethernet ring protection switching (ERPS) on connected switches to prevent fatal loops from disrupting a network. ERPS is similar to spanning-tree protocols, but ERPS is more efficient than spanning-tree protocols because it is customized for ring topologies. You must configure at least three switches to form a ring. One of the links, called the ring protection link (RPL) end interface, is blocked until another link fails—at this time the RPL link is unblocked, ensuring connectivity.



NOTE: Ethernet OAM connectivity fault management (CFM) can be used with ERPS to detect link faults faster in some cases. See *Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)*.

The time needed for switchover to the ERPS link is affected by three settings—link failure detection time, the number of nodes in the ring, and the time it takes to unblock the RPL after a failure is detected.



NOTE: Do not configure redundant trunk groups on ERPS interfaces. You can configure VSTP on ERPS interfaces if the VSTP uses a VLAN that is not part of the ERPS control VLAN or data channel VLANs. The total number of ERPS and VSTP or MSTP instances is limited to 253.

Before you begin:

- Optionally, configure two interfaces on each switch as trunk ports. See *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*.
- Configure a VLAN to act as a control VLAN for ERPS if your interfaces are trunk ports. Configure the same VLAN on all switches and associate the two network interfaces from each of the switches with the VLAN. See “[Configuring VLANs for EX Series Switches \(CLI Procedure\)](#)” on page 160. If you have multiple ERPS instances, the control VLANs and data channel VLANs must not overlap.
- Data channels are optional on the ERPS link. If you plan to use them, configure a VLAN for each data channel.

To configure ERPS:



NOTE: You must configure at least three switches, with only one switch designated as the RPL owner node.

1. RSTP and EPRS cannot both be configured on a ring port, and RSTP is configured by default. Disable RSTP on each switch interface:

```
user@switch#setrstpinterface interface-name disable
```

2. Create a node ring on each switch:

```
[edit protocols]
user@switch# set protection-group ethernet-ring ring-name
```

3. Configure a control VLAN for the node ring if the links are trunk ports:

```
[edit protocols protection-group ethernet-ring ring name]
user@switch# set control-vlan vlan-name-or-vlan-id
```

4. Configure the east interface of the node ring with the control-channel interface. In addition, configure either the east interface or the west interface (but not both) as a link end.

```
[edit protocols protection-group ethernet-ring ring-name]
user@switch# set east-interface control-channel channel-name
user@switch# set east-interface ring-protection-link-end
```

5. Configure the west interface of the node ring with the control-channel interface. In addition, configure either the east interface or the west interface (but not both) as a link end.

```
[edit protocols protection-group ethernet-ring ring-name]
user@switch# set west-interface control-channel control-channel-interface-address
user@switch# set west-interface ring protection link end
```

6. Configure only one switch as the RPL owner node:

```
[edit protocols protection-group ethernet-ring ring-name]
user@switch# set ring-protection-link-owner
```

7. The restore interval configures the number of minutes that the node does not process any Ethernet ring protection (ERP) protocol data units (PDUs). When a link goes down, the ring protection link (RPL) activates. When the downed link comes back up, the RPL receives notification, restores the link, and waits the length of time indicated by the restore interval before issuing another block on the same link. Optionally, configure the restore interval on each switch:

```
[edit protocols protection-group ethernet-ring ring-name]
user@switch# set restore-interval restore-interval-value
```

8. The guard interval prevents ring nodes from receiving outdated messages (called RAPs). Optionally, configure the guard interval on each switch:

```
[edit protocols protection-group ethernet-ring ring name]
user@switch# set guard-interval guard-interval-value
```



NOTE: Local settings take priority over global settings.

Global settings are used when no local settings are present. Optionally, you can also configure these global settings on the switch:

- restore interval
- guard interval
- ERP traceoptions: file, page size, file size, flag name

9. Optionally, reconfigure the global guard interval on each switch:

```
[edit protocols protection-group ethernet-ring ring name]
```

```
user@switch# set guard-interval guard-interval-value
```

10. Optionally, reconfigure the global restore interval on each switch:

```
[edit protocols protection-group ethernet-ring ring name]  
user@switch# set restore-interval restore-interval-value
```

11. After detection of a link failure, switching takes place after the hold interval has expired. Optionally, reconfigure the global hold interval on each switch:

```
[edit protocols protection-group ethernet-ring ring name]  
user@switch# set hold-interval hold-interval-value
```

12. Optionally, configure VLANs for data channels on the ERPS link:

```
[edit protocols protection-group ethernet-ring ring name]  
user@switch# set data-channel vlan-name
```

**Related
Documentation**

- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches on page 140](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches](#)
- [Ethernet Ring Protection Switching Overview on page 39](#)
- [Understanding Ethernet Ring Protection Switching Functionality on page 34](#)

Enabling VLAN Pruning for Broadcast, Multicast, and Unknown Unicast Traffic in an EX Series Virtual Chassis (CLI Procedure)

You can enable VLAN pruning for VLANs assigned to interfaces in an EX Series Virtual Chassis. When you enable VLAN pruning for a VLAN in a Virtual Chassis, all broadcast, multicast, and unknown unicast traffic entering that VLAN uses the shortest possible path through the Virtual Chassis to the egress VLAN interface. Enabling VLAN pruning allows you to conserve bandwidth within the Virtual Chassis, since all broadcast, multicast, and unknown unicast traffic in a VLAN is broadcast to all Virtual Chassis member switches when VLAN pruning is disabled.



BEST PRACTICE: We recommend enabling VLAN pruning when configuring a VLAN on an EX Series Virtual Chassis.

To enable VLAN pruning when configuring a VLAN:

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

**Related
Documentation**

- [Configuring VLANs for EX Series Switches \(CLI Procedure\) on page 160](#)

CHAPTER 8

Configuration Statements

- [\[edit ethernet-switching-options\] Configuration Statement Hierarchy on EX Series Switches on page 195](#)
- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 198](#)
- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 199](#)
- [\[edit routing-instances\] Configuration Hierarchy Statement Hierarchy on EX Series Switches on page 200](#)
- [\[edit vlans\] Configuration Statement Hierarchy on EX Series Switches on page 205](#)
- [add-attribute-length-in-pdu on page 207](#)
- [arp \(System\) on page 208](#)
- [arp-on-stp on page 209](#)
- [control-channel on page 210](#)
- [control-vlan on page 211](#)
- [customer-vlans on page 212](#)
- [data-channel on page 213](#)
- [description \(VLANs\) on page 214](#)
- [disable \(MVRP\) on page 214](#)
- [dot1q-tunneling \(Ethernet Switching\) on page 215](#)
- [dot1q-tunneling \(VLANs\) on page 216](#)
- [drop-threshold on page 217](#)
- [east-interface on page 218](#)
- [edge-virtual-bridging on page 219](#)
- [ethernet-ring on page 220](#)
- [ether-type on page 221](#)
- [ethernet-switching-options on page 222](#)
- [filter \(VLANs\) on page 225](#)
- [group \(Redundant Trunk Groups\) on page 226](#)
- [guard-interval on page 227](#)
- [instance-type on page 228](#)

- [interface \(Redundant Trunk Groups\) on page 230](#)
- [interface \(Routing Instances\) on page 231](#)
- [interface \(VLANs\) on page 232](#)
- [interface \(MVRP\) on page 233](#)
- [interfaces \(Q-in-Q Tunneling\) on page 234](#)
- [isolation-id on page 234](#)
- [join-timer \(MVRP\) on page 235](#)
- [layer2-protocol-tunneling on page 236](#)
- [l3-interface \(VLANs\) on page 238](#)
- [l3-interface-ingress-counting on page 239](#)
- [leaveall-timer \(MVRP\) on page 240](#)
- [leave-timer \(MVRP\) on page 241](#)
- [mac \(Static MAC-Based VLANs\) on page 242](#)
- [mac-limit \(VLANs\) on page 243](#)
- [mac-lookup-length on page 245](#)
- [mac-notification on page 246](#)
- [mac-table-aging-time on page 247](#)
- [mapping on page 248](#)
- [members on page 249](#)
- [mvrp on page 251](#)
- [native-vlan-id on page 252](#)
- [next-hop \(Static MAC-Based VLANs\) on page 252](#)
- [no-dynamic-vlan on page 253](#)
- [no-local-switching on page 253](#)
- [no-mac-learning \(Q-in-Q VLANs\) on page 254](#)
- [no-mac-learning \(Q-in-Q Interfaces\) on page 254](#)
- [node-id on page 255](#)
- [notification-interval on page 256](#)
- [port-mode on page 257](#)
- [preempt-cutover-timer on page 258](#)
- [primary-vlan on page 259](#)
- [protection-group on page 260](#)
- [proxy-arp on page 262](#)
- [pvlan-trunk on page 263](#)
- [redundant-trunk-group on page 264](#)
- [registration on page 265](#)
- [restore-interval on page 266](#)

- [ring-protection-link-end](#) on page 267
- [ring-protection-link-owner](#) on page 267
- [routing-instances](#) on page 268
- [shutdown-threshold](#) on page 269
- [static](#) (Static MAC-Based VLANs) on page 270
- [traceoptions](#) (Ethernet Ring Protection) on page 271
- [traceoptions](#) (Edge Virtual Bridging) on page 273
- [vlan](#) (802.1Q Tagging) on page 274
- [vlan](#) (Static MAC-based VLANs) on page 275
- [vlan-id](#) (802.1Q Tagging) on page 276
- [vlan-prune](#) on page 277
- [vlan-range](#) on page 278
- [vlans](#) on page 279
- [vrf-mtu-check](#) on page 280
- [vsi-discovery](#) on page 281
- [vsi-policy](#) on page 282
- [west-interface](#) on page 283

[edit ethernet-switching-options] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit ethernet-switching-options]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit ethernet-switching-options\] Hierarchy Level](#) on page 195
- [Unsupported Statements in the \[edit ethernet-switching-options\] Hierarchy Level](#) on page 198

Supported Statements in the [edit ethernet-switching-options] Hierarchy Level

The following hierarchy shows the **[edit ethernet-switching-options]** configuration statements supported on EX Series switches:

```
ethernet-switching {
```

```
analyzer {
  name {
    input {
      egress {
        interface (all | interface-name);
      }
      ingress {
        interface (all | interface-name);
        vlan (vlan-id | vlan-name);
      }
    }
    loss-priority priority;
    output {
      interface interface-name;
      vlan (vlan-id | vlan-name);
    }
    ratio number;
  }
}
authentication-whitelist {
  interface;
  vlan-assignment;
}
bpdu-block {
  disable-timeout timeout;
  interface (all | [interface-name]) {
    (disable | drop | shutdown);
  }
}
dot1q-tunneling {
  ether-type (0x8100 | 0x88a8 | 0x9100);
}
interfaces interface-name {
  no-mac-learning;
}
mac-lookup-length number-of-entries;
}
mac-notification {
  notification-interval seconds;
}
mac-table-aging-time seconds;
port-error-disable {
  disable-timeout timeout;
}
redundant-trunk-group {
  group name {
    description;
    interface interface-name {
      primary;
    }
    preempt-cutover-timer seconds;
  }
}
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 );
    fcoe-trusted;
    mac-limit limit action action;
    no-allowed-mac-log;
    static-ip ip-address {
        mac mac-address;
        vlan vlan-name;
    }
}
uac-policy;
}
vlan (all | vlan-name) {
    (arp-inspection | no-arp-inspection );
    dhcp-option82 {
        disable;
        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);
    examine-fip {
        fc-map fc-map-value;
    }
    (ip-source-guard | no-ip-source-guard);
    mac-move-limit limit action action;
}
}
static {
    vlan vlan-id {
        mac mac-address next-hop interface-name;
    }
}
storm-control {
    action-shutdown;
    interface (all | interface-name) {
        bandwidth bandwidth;
        multicast;
        no-broadcast;
        no-multicast;
        no-registered-multicast;
        no-unknown-unicast;
    }
}

```

```
        no-unregistered-multicast;
    }
}
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]) {
        forwarding-class (assured-forwarding | best-effort | expedited-forwarding |
        network-control);
        vlan vlan-name;
    }
}
}
```

Unsupported Statements in the [edit ethernet-switching-options] Hierarchy Level

All statements in the [edit ethernet-switching-options] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

Related Documentation

- [Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 103](#)
- [Example: Configuring Redundant Trunk Links for Faster Recovery on page 112](#)
- [Configuring MAC Table Aging \(CLI Procedure\) on page 166](#)
- [Configuring MAC Notification \(CLI Procedure\) on page 175](#)
- [Configuring Q-in-Q Tunneling \(CLI Procedure\) on page 180](#)
- [Configuring Redundant Trunk Links for Faster Recovery \(CLI Procedure\) on page 185](#)
- [Configuring Nonstop Bridging on EX Series Switches \(CLI Procedure\)](#)

[edit interfaces] Configuration Statement Hierarchy on EX Series Switches

Each of the following topics lists the statements at a subhierarchy of the [edit interfaces] hierarchy:

- [\[edit interfaces ae\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit interfaces ge\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit interfaces gr\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit interfaces interface-range\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit interfaces lo\] Configuration Statement Hierarchy on EX Series Switches](#)

- [\[edit interfaces me\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit interfaces vlan\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit interfaces vme\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit interfaces xe\] Configuration Statement Hierarchy on EX Series Switches](#)

**Related
Documentation**

- [EX Series Switches Interfaces Overview](#)
- [Configuring Aggregated Ethernet Links \(CLI Procedure\)](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)
- [Configuring a Layer 3 Subinterface \(CLI Procedure\)](#)
- [Configuring Routed VLAN Interfaces \(CLI Procedure\) on page 164](#)
- [Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis \(CLI Procedure\)](#)
- [Junos OS Interfaces Fundamentals Configuration Guide](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

[\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches](#)

Each of the following topics lists the statements at a subhierarchy of the **[edit protocols]** hierarchy:

- [\[edit protocols bfd\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols bgp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols connections\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols dcbx\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols dot1x\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols igmp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols igmp-snooping\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols isis\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols lacp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols link-management\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols lldp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols lldp-med\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols mld\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols mld-snooping\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols mpls\] Configuration Statement Hierarchy on EX Series Switches](#)

- [\[edit protocols msdp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols mstp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols mvrp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols neighbor-discovery\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols oam\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols ospf\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols ospf3\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols pim\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols rip\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols ripng\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols router-advertisement\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols router-discovery\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols rstp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols rsvp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols sflow\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols stp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols uplink-failure-detection\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols vrrp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols vstp\] Configuration Statement Hierarchy on EX Series Switches](#)

**Related
Documentation**

- [EX Series Switch Software Features Overview](#)
- [EX Series Virtual Chassis Software Features Overview](#)

[\[edit routing-instances\] Configuration Hierarchy Statement Hierarchy on EX Series Switches](#)

This topic lists supported and unsupported configuration statements in the **[edit routing-instances]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.

- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit routing-instances\] Hierarchy Level on page 201](#)
- [Unsupported Statements in the \[edit routing-instances\] Hierarchy Level on page 204](#)

Supported Statements in the [edit routing-instances] Hierarchy Level

The following hierarchy shows the **[edit routing-instances]** configuration statements supported on EX Series switches:

```
routing-instances routing-instance-name {
  access {
    address-assignment {
      pool pool-name {
        family inet {
          dhcp-attributes {
            boot-file filename;
            boot-server hostname;
            domain-name domain-name;
            grace-period seconds;
            maximum-lease-time (seconds | infinite);
            name-server {
              address;
            }
            netbios-node-type (b-node | h-node | m-node | p-node);
            option option-index (array (byte | flag | integer | ip-address | short | string |
              unsigned-integer | unsigned-short) [ type-values ] | byte 8-bit-value |
              flag (false | off | on | true) | integer signed-32-bit-value | ip-address address |
              short signed-16-bit-value | string text-string | unsigned-integer 32-bit-value |
              unsigned-short 16-bit-value);
          router {
            address;
          }
          server-identifier ipv4-address;
          tftp-server hostname;
          wins-server {
            address;
          }
        }
      }
      host hostname {
        hardware-address mac-address;
        ip-address ip-address;
      }
      network ip-prefix</prefix-length>;
      range name {
        high upper-limit;
        low lower-limit;
      }
    }
    link pool-name;
  }
}
```

```

    }
  }
  access-profile profile-name;
  description text;
  forwarding-options {
    ... same statements as in [edit forwarding-options] Configuration Statement Hierarchy
    on EX Series Switches
  }
  instance-role role;
  instance-type virtual-router;
  interfaces interface-name {
    ... same statements as in [edit interfaces] Configuration Statement Hierarchy on EX
    Series Switches on page 198
  }
  l2vpn-id identifier;
  no-vrf-advertise;
  no-vrf-propagate-ttl;
  protocols {
    ... same statements as in [edit protocols] Configuration Statement Hierarchy on EX
    Series Switches on page 199
  }
  provider-tunnel {
    ingress-replication {
      create-new-ucast-tunnel;
      label-switched-path-template {
        (default-template | lsp-template-name);
      }
    }
  }
  ldp-p2mp;
  mdt {
    data-mdt-reuse;
    group-range multicast-prefix;
    threshold {
      group group-address {
        source source-address {
          rate threshold-rate;
        }
      }
    }
    tunnel-limit limit;
  }
}
pim-asm {
  group-address address;
}
pim-ssm {
  group-address address;
}
rsvp-te {
  label-switched-path-template {
    (default-template | lsp-template-name);
  }
  static-lsp point-to-multipoint-lsp-name;
}
selective {
  group multicast-prefix</prefix-length> {
    source ip-prefix</prefix-length> {

```

```

    ingress-replication {
        create-new-ucast-tunnel;
        label-switched-path {
            label-switched-path-template (default-template | template-name);
        }
    }
    ldp-p2mp;
    pim-ssm {
        group-range multicast-prefix</prefix-length>;
    }
    rsvp-te {
        label-switched-path-template {
            (default-template | lsp-template-name);
        }
        static-lsp point-to-multipoint-lsp-name;
    }
    threshold-rate kbps;
}
wildcard-source {
    ingress-replication {
        create-new-ucast-tunnel;
        label-switched-path {
            label-switched-path-template (default-template | template-name);
        }
    }
    ldp-p2mp;
    pim-ssm {
        group-range multicast-prefix</prefix-length>;
    }
    rsvp-te {
        label-switched-path-template {
            (default-template | lsp-template-name);
        }
        static-lsp point-to-multipoint-lsp-name;
    }
    threshold-rate kbps;
}
}
tunnel-limit number;
wildcard-group-inet {
    wildcard-source {
        ingress-replication {
            create-new-ucast-tunnel;
            label-switched-path {
                label-switched-path-template (default-template | template-name);
            }
        }
    }
    ldp-p2mp;
    pim-ssm {
        group-range multicast-prefix</prefix-length>;
    }
    rsvp-te {
        label-switched-path-template {
            (default-template | lsp-template-name);
        }
        static-lsp lsp-name;
    }
}

```

```

        }
        threshold-rate kbps;
    }
}
wildcard-group-inet6 {
    wildcard-source {
        pim-ssm {
            group-range multicast-prefix</prefix-length>;
        }
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
            static-lsp lsp-name;
        }
        threshold-rate kbps;
    }
}
}
route-distinguisher (as-number:number | ip-address:number);
routing-options {
    ... the routing-options subhierarchy appears after the main [edit routing-instances
        routing-instance-name] hierarchy ...
}
vlan-model one-to-one;
vrf-advertise-selective {
    family {
        inet-mvpn;
        inet6-mvpn;
    }
}
vrf-export [ policy-names ];
vrf-import [ policy-names ];
(vrf-propagate-ttl | no-vrf-propagate-ttl);
vrf-table-label;
vrf-target {
    target:community-identifier;
    export target:community-identifier;
    import target:community-identifier;
}
}
}

```

Unsupported Statements in the [edit routing-instances] Hierarchy Level

All statements in the [edit routing-instances] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 23: Unsupported [edit routing-instances] Configuration Statements on EX Series Switches

Statement	Hierarchy
interface-mac-limit	[edit routing-instances <i>routing-instance-name</i> vlans <i>vlan-name</i> switch-options] [edit routing-instances <i>routing-instance-name</i> vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>]
packet-action	[edit routing-instances <i>routing-instance-name</i> vlans <i>vlan-name</i> switch-options interface-mac-limit <i>limit</i>] [edit routing-instances <i>routing-instance-name</i> vlans <i>vlan-name</i> switch-options interface <i>interface-name</i> interface-mac-limit <i>limit</i>]

- Related Documentation**
- [Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 88](#)
 - [Configuring Virtual Routing Instances \(CLI Procedure\) on page 174](#)

[edit vlans] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit vlans]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit vlans\] Hierarchy Level on page 205](#)
- [Unsupported Statements in the \[edit vlans\] Hierarchy Level on page 206](#)

Supported Statements in the [edit vlans] Hierarchy Level

The following hierarchy shows the **[edit vlans]** configuration statements supported on one or more of the EX Series switches:

```

vlands {
  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
  output filter-name;
}
interface interface-name {
  egress;
  ingress;
  mapping (native (push | swap) | policy | tag (push | swap));
  pvlan-trunk;
}
isolation-id id-number;
l3-interface vlan.logical-interface-number;
l3-interface-ingress-counting layer-3-interface-name;
mac-limit limit action action;
mac-table-aging-time seconds;
no-local-switching;
no-mac-learning;
primary-vlan vlan-name;
vlan-id number;
vlan-prune;
vlan-range vlan-id-low-vlan-id-high;
}
}

```

Unsupported Statements in the [edit vlans] Hierarchy Level

All statements in the **[edit vlans]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 24: Unsupported [edit vlans] Configuration Statements on EX Series Switches

Statement	Hierarchy Level
NOTE: Variables, such as <i>filename</i> , are not shown in the statements or hierarchies.	
udid	[edit vlans dot1q-tunneling layer2-protocol-tunneling]

Related Documentation

- [Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 65](#)
- [Example: Connecting an Access Switch to a Distribution Switch on page 72](#)
- [Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 103](#)
- [Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 107](#)
- [Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches](#)
- [Example: Configuring a Private VLAN Spanning Multiple EX Series Switches on page 119](#)
- [Creating a Private VLAN on a Single EX Series Switch \(CLI Procedure\) on page 169](#)

add-attribute-length-in-pdu

Syntax	add-attribute-length-in-pdu;
Hierarchy Level	[edit protocols mvrp]
Release Information	Statement introduced in Junos OS Release 11.3 for EX Series switches.
Description	<p>Add an extra byte in protocol data units (PDUs) sent by the Multiple VLAN Registration Protocol (MVRP) in Junos OS Releases 11.3 and later for EX Series switches that do not support the Enhanced Layer 2 Software (ELS). By default, this MVRP version does not include the extra byte. You can add the extra byte in this MVRP version to address an incompatibility issue with the following MVRP versions:</p> <ul style="list-style-type: none"> • MVRP in Junos OS Releases 11.2 and earlier, which includes the extra byte. • MVRP in Junos OS Releases 13.2 and later for EX Series switches with support for ELS, which includes the extra byte. <p>If this incompatibility issue arises, the MVRP versions that include the extra byte do not recognize PDUs that do not include the extra byte.</p> <p>You can recognize an MVRP version problem by looking at a switch running an MVRP version that includes the extra byte. Because a switch running an MVRP version that includes the extra byte cannot interpret an unmodified PDU from an MVRP version that does not include the extra byte, the switch will not add VLANs from the MVRP version that does not include the extra byte. When you execute the command show mvrp statistics on the MVRP version that includes the extra byte, the values for <i>Join Empty received</i> and <i>Join In received</i> will incorrectly display zero, even though the value for <i>MRPDU received</i> has been increased. Another indication that MVRP is having a version problem is that unexpected VLAN activity, such as multiple VLAN creation, takes place on the switch running the MVRP version that includes the extra byte.</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 Documentation	<ul style="list-style-type: none"> • Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 176 • Understanding Multiple VLAN Registration Protocol (MVRP) on EX Series Switches on page 26

arp (System)

Syntax

```
arp {  
    aging-timer minutes;  
    gratuitous-arp-delay seconds;  
    gratuitous-arp-on-ifup;  
    interfaces {  
        interface-name {  
            aging-timer minutes;  
        }  
    }  
    passive-learning;  
    purging;  
}
```

For EX-Series switches:

```
arp {  
    aging-timer minutes;  
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Specify ARP options. You can enable backup VRRP routers to learn ARP requests for VRRP-IP to VRRP-MAC address translation. You can also set the time interval between ARP updates.

For EX-Series switches, set only the time interval between ARP updates.

Options **aging-timer**—Time interval in minutes between ARP updates. In environments where the number of ARP entries to update is high (for example, on routers only, metro Ethernet environments), increasing the time between updates can improve system performance.

passive-learning (QFX-Series only)—Configure switches to learn the ARP mappings (IP-to-MAC address) for hosts sending the requests.

Default: 20 minutes

Range: 1 to 240 minutes

The remaining statements are explained separately.

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

Related Documentation

- *Configuring Junos OS ARP Learning and Aging Options for Mapping IPv4 Network Addresses to MAC Addresses*
- *Junos OS Network Interfaces Library for Routing Devices*

- For more information about ARP updates, see the [Junos OS System Basics Configuration Guide](#).

arp-on-stp

Syntax	arp-on-stp;
Hierarchy Level	[edit protocols mstp interface (Spanning Trees) (all <i>interface-name</i>)], [edit protocols rstp interface (Spanning Trees) (all <i>interface-name</i>)], [edit protocols stp interface (Spanning Trees) (all <i>interface-name</i>)], [edit protocols vstp (all <i>vlan--id</i> <i>vlan--name</i>) interface (Spanning Trees) (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 11.2 for EX Series switches.
Description	<p>Configure the Address Resolution Protocol (ARP) in a spanning-tree network so that when a spanning-tree protocol topology change notification (TCN) is issued, the VLAN with a broken link can relearn MAC addresses from another, redundant VLAN in the network. The network must include a routed VLAN interface (RVI).</p> <p>When a link fails in a spanning-tree network (RSTP, STP, MSTP, or VSTP), a message called a TCN is issued that causes all affected Ethernet switching table entries to be flushed. The network must then relearn the MAC addresses using flooding. If you have configured an RVI on the network, you have the option of having the VLAN with the broken link relearn MAC addresses from another VLAN using ARP, thereby avoiding excessive flooding on the VLAN with the broken link.</p>
Default	ARP on STP is disabled.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring MSTP (CLI Procedure) • Configuring RSTP (CLI Procedure) • Configuring STP (CLI Procedure) • Configuring VSTP (CLI Procedure)

control-channel

Syntax	<code>control-channel <i>channel-name</i> { vlan <i>vlan-id</i>; interface name <i>interface-name</i> }</code>
Hierarchy Level	[edit protocols protection-group ethernet-ring <i>name</i> (east-interface west-interface)]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.
Description	Configure the Ethernet RPS control channel logical interface to carry the RAPS PDU. The related physical interface is the physical ring port.
Options	vlan <i>vlan-id</i> —If the control channel logical interface is a trunk port, then a dedicated vlan <i>vlan-id</i> defines the dedicated VLAN channel to carry the RAPS traffic. Only configure the vlan-id when the control channel logical interface is the trunk port. interface name <i>interface-name</i> —Interface name of the control channel.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Ethernet Ring Protection Switching Overview on page 39• Example: Configuring Ethernet Ring Protection Switching on EX Series Switches on page 140• Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches• Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 190

control-vlan

Syntax	control-vlan (<i>vlan-id</i> <i>vlan-name</i>)
Hierarchy Level	[edit protocols protection-group ethernet-ring] [edit protocols protection-group ethernet-ring name (east-interface west-interface)]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.
Description	Specify the VLAN that carries the protocol data units (PDUs) between the nodes in the protected Ethernet ring. This is a control VLAN, meaning that it carries data for one instance of an Ethernet ring protection switching (ERPS) in the control channel. Use a control VLAN on trunk port interfaces. One control channel can contain multiple control VLANs.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Ethernet Ring Protection Switching on EX Series Switches on page 140 • Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches • Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 190

customer-vlans

Syntax	<code>customer-vlans (<i>id</i> <i>native</i> <i>range</i>);</code>
Hierarchy Level	[edit vlans <i>vlan-name</i> dot1q-tunneling]
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches. Option native introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Limit the set of accepted customer VLAN tags to a range or to discrete values when mapping customer VLANs to service VLANs.
Options	<p>id—Numeric identifier for a VLAN.</p> <p>native—Accepts untagged and priority-tagged packets from access interfaces and assigns the configured S-VLAN to the packet.</p> <p>range—Range of numeric identifiers for VLANs. On the QFX series, you can include as many as eight separate customer VLAN ranges for a given service VLAN. Do not configure more than this number of ranges.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• dot1q-tunneling (Ethernet Switching) on page 215• ether-type on page 221• Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 103• Configuring Q-in-Q Tunneling (CLI Procedure) on page 180• Understanding Q-in-Q Tunneling on EX Series Switches on page 41• Configuring Q-in-Q Tunneling• Example: Setting Up Q-in-Q Tunneling• dot1q-tunneling (Ethernet Switching) on page 215• ether-type

data-channel

Syntax	data-channel { vlan <i>number</i> ; }
Hierarchy Level	[edit protocols protection-group ethernet-ring ring-name]
Release Information	Statement introduced in Junos OS Release 10.2. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.
Description	For Ethernet ring protection, configure a data channel to define a set of VLAN IDs that belong to a ring instance. VLANs specified in the data channel use the same topology used by the ERPS PDU in the control channel. Therefore, if a ring interface is blocked in the control channel, all traffic in the data channel is also blocked on that interface.
Options	vlan <i>number</i> —Specify (by VLAN ID) one or more VLANs that belong to a ring instance.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Ethernet Ring Protection Using Ring Instances for Load Balancing</i> • <i>Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers</i> • <i>Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches</i> • Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 190

description (VLANs)

Syntax	<code>description <i>text-description</i>;</code>
Hierarchy Level	[edit vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Option text-description enhanced from supporting up to 128 characters to supporting up to 256 characters in Junos OS Release 10.2 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	text-description —Text to describe the interface. It can contain letters, numbers, and hyphens (-) and can contain 256 characters. If the text includes spaces, enclose the entire text in quotation marks.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show vlans on page 362• Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 57• Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch• Understanding Bridging and VLANs on EX Series Switches on page 3

disable (MVRP)

Syntax	<code>disable;</code>
Hierarchy Level	[edit protocols mvrp], [edit protocols mvrp interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 13.1 for the QFX Series.
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 Documentation	<ul style="list-style-type: none">• Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 176

dot1q-tunneling (Ethernet Switching)

Syntax	<code>dot1q-tunneling { ether-type (0x8100 0x88a8 0x9100); }</code>
Hierarchy Level	[edit ethernet-switching-options]
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Set a global value for the EtherType for Q-in-Q tunneling. The remaining statement is explained separately.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• dot1q-tunneling on page 216• Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 103• Configuring Q-in-Q Tunneling (CLI Procedure) on page 180• Configuring Q-in-Q Tunneling• Example: Setting Up Q-in-Q Tunneling• dot1q-tunneling on page 216

dot1q-tunneling (VLANs)

Syntax	<pre>dot1q-tunneling { customer-vlans (id native range); layer2-protocol-tunneling all protocol-name { drop-threshold number; shutdown-threshold number; } }</pre>
Hierarchy Level	[edit vlans <i>vlan-name</i>]
Release Information	<p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Option native introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Options layer2-protocol-tunneling, drop-threshold, and shutdown-threshold introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
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.
- You cannot configure 802.1X user authentication on interfaces that have been enabled for Q-in-Q tunneling.

The remaining statements are explained separately.

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

Related Documentation	• Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 103
	• Configuring Q-in-Q Tunneling (CLI Procedure) on page 180
	• Configuring Q-in-Q Tunneling
	• Example: Setting Up Q-in-Q Tunneling
	• Configuring Layer 2 Protocol Tunneling
	• dot1q-tunneling (Ethernet Switching) on page 215

drop-threshold

Syntax	<code>drop-threshold <i>number</i>;</code>
Hierarchy Level	[edit vlan <i>vlan-name</i> dot1q-tunneling layer2-protocol-tunneling (all <i>protocol-name</i>)]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
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 switch begins dropping the Layer 2 PDUs. The drop threshold value must be less than or equal to the shutdown threshold value.</p> <p>L2PT processing is done by the CPU, and L2PT traffic to the CPU is rate-limited to a maximum of 1000 pps. If traffic is received at a rate faster than this limit, the rate limit causes the traffic to be dropped before it hits the threshold and the dropped packets are not reported in L2PT statistics. This can also occur if you configure a drop threshold that is less than 1000 pps but traffic is received at a faster rate. For example, if you configure a drop threshold of 900 pps and the VLAN receives traffic at rate of 1100 pps, L2PT statistics will show that 100 packets were dropped. The 100 packets dropped because of the rate limit will not be reported. Similarly, if you do not configure a drop threshold and the VLAN receives traffic at rate of 1100 pps, the 100 packets dropped because of the rate limit are not reported.</p>
	<p> NOTE: If the drop threshold value is greater than the shutdown threshold value and you try to commit the configuration, the commit operation fails.</p>
	You can specify a drop threshold value without specifying a shutdown threshold value.
Default	No drop 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 switch begins dropping the Layer 2 PDUs.</p> <p>Range: 1 through 1000</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 107 • Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) on page 181 • Configuring Layer 2 Protocol Tunneling • shutdown-threshold on page 269

east-interface

Syntax

```
east-interface {
    node-id mac-address;
    control-channel channel-name {
        vlan number;
        interface name interface-name
    }
    interface-none
    ring-protection-link-end;
}
```

Hierarchy Level [edit protocols [protection-group ethernet-ring ring-name](#)]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description Define one of the two interface ports for Ethernet ring protection, the other being defined by the **west-interface** statement at the same hierarchy level. The interface must use the control channel's logical interface name. The control channel is a dedicated VLAN channel for the ring port.

EX Series switches do not use the node-id statement--the node ID is automatically configured on the switches using the MAC address.



NOTE: Always configure this port first, before configuring the **west-interface** statement.



NOTE: The Node ID is not configurable on EX Series switches. The node ID is automatically configured using the MAC address.

The statements are explained separately.

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

Related Documentation

- [Ethernet Ring Protection Switching Overview on page 39](#)
- [Ethernet Ring Protection Using Ring Instances for Load Balancing](#)
- [west-interface on page 283](#)
- [ethernet-ring on page 220](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches on page 140](#)

- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches](#)
- [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 190](#)

edge-virtual-bridging

Syntax	<pre> edge-virtual-bridging { traceoptions { file <i>filename</i> <files <i>number</i>> <no-stamp> <replace> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> ; } vsi-discovery { interface <i>interface-name</i> vsi-policy <i>vsi-policy-name</i> } } </pre>
Hierarchy Level	[edit protocols]
Release Information	Statement introduced in Junos OS Release 12.1 for EX Series switches.
Description	<p>Configure edge virtual bridging (EVB). EVB enables a virtualized station (a physical end station, a server, connected to virtual machines (VMs)) to network with an adjacent switch so that applications residing on the virtual machines can interact with each other and external networks through a technology called virtual Ethernet packet aggregator (VEPA).</p> <p>The remaining statements are explained separately.</p>
Default	EVB is disabled by default.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Edge Virtual Bridging for Use with VEPA Technology on page 134 • Configuring Edge Virtual Bridging (CLI Procedure) on page 188

ethernet-ring

Syntax `ethernet-ring ring-name {
 control-vlan (vlan-id | vlan-name);
 data-channel {
 vlan number
 }
 east-interface {
 control-channel channel-name {
 vlan number;
 interface name interface-name
 }
 }
 guard-interval number;
 node-id mac-address;
 restore-interval number;
 ring-protection-link-owner;
 west-interface {
 control-channel channel-name {
 vlan number;
 }
 }
 }`

Hierarchy Level [edit protocols [protection-group](#)]

Release Information Statement introduced in Junos OS Release 9.4.
 Statement introduced in Junos OS Release 12.1 for EX Series switches.
 Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description For Ethernet PICs on MX Series routers or for EX Series switches, , specify the Ethernet ring in an Ethernet ring protection switching configuration.

Options *ring-name*—Name of the Ethernet protection ring.

 The remaining statements are explained separately.

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

Related Documentation

- [Ethernet Ring Protection Switching Overview on page 39](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches on page 140](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches](#)
- [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 190](#)

ether-type

Syntax	ether-type (0x8100 0x88a8 0x9100)
Hierarchy Level	[edit ethernet-switching-options dot1q-tunneling]
Release Information	Statement introduced in Junos OS 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	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• dot1q-tunneling (VLANs) on page 216• Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 103• Configuring Q-in-Q Tunneling (CLI Procedure) on page 180

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) {
                no-tag;
            }
        }
    }
    bpdu-block {
        disable-timeout timeout;
        interface (all | [interface-name]) {
            (disable | drop | shutdown);
        }
    }
    dot1q-tunneling {
        ether-type (0x8100 | 0x88a8 | 0x9100);
    }
    interfaces interface-name {
        no-mac-learning;
    }
    mac-lookup-length number-of-entries;
}
mac-notification {
    notification-interval seconds;
}
mac-table-aging-time seconds;
nonstop-bridging;
port-error-disable {
    disable-timeout timeout;
}
redundant-trunk-group {
    group name {
        interface interface-name <primary>;
        interface interface-name;
    }
}
secure-access-port {
    dhcp-snooping-file {
```

```

    location local_pathname | remote_URL;
    timeout seconds;
    write-interval seconds;
}
dhcpv6-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);
    fcoe-trusted;
    mac-limit limit action (drop | log | none | shutdown);
    no-allowed-mac-log;
    persistent-learning;
    static-ip ip-address {
        vlan vlan-name;
        mac mac-address;
    }
    static-ipv6 ip-address {
        vlan vlan-name;
        mac mac-address;
    }
}
vlan (all | vlan-name) {
    (arp-inspection | no-arp-inspection) [
        forwarding-class class-name;
    ]
    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) {
        forwarding-class class-name;
    }
    (examine-dhcpv6 | no-examine-dhcpv6) {
        forwarding-class class-name;
    }
    examine-fip {
        fc-map fc-map-value;
    }
    (ip-source-guard | no-ip-source-guard);
    (ipv6-source-guard | no-ipv6-source-guard);
    mac-move-limit limit action (drop | log | none | shutdown);
}

```

```

    }
    (neighbor-discovery-inspection | no-neighbor-discovery-inspection);
no-option-37;
static {
    vlan name {
        mac mac-address {
            next-hop interface-name;
        }
    }
}
storm-control {
    action-shutdown;
    interface (all | interface-name) {
        bandwidth bandwidth;
        level level;
        multicast;
        no-broadcast;
        no-multicast;
        no-registered-multicast;
        no-unknown-unicast;
        no-unregistered-multicast;
    }
}
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 OS Release 9.0 for EX Series switches.

Description Configure Ethernet switching options.

The remaining statements are explained separately.

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

Related Documentation	<ul style="list-style-type: none"> • <i>Understanding Port Mirroring on EX Series Switches</i> • <i>Understanding Port Security</i> • <i>Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches</i> • Understanding Redundant Trunk Links on page 49 • <i>Understanding Storm Control on EX Series Switches</i> • <i>Understanding 802.1X and VoIP on EX Series Switches</i> • Understanding Q-in-Q Tunneling on EX Series Switches on page 41 • <i>Understanding Unknown Unicast Forwarding</i> • Understanding MAC Notification on EX Series Switches on page 29 • <i>Understanding FIP Snooping</i> • <i>Understanding Nonstop Bridging on EX Series Switches</i>
------------------------------	--

filter (VLANs)

Syntax	<code>filter (input output) <i>filter-name</i>;</code>
Hierarchy Level	[edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS 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>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches</i> • <i>Configuring Firewall Filters (CLI Procedure)</i> • <i>Configuring Firewall Filters (J-Web Procedure)</i> • <i>Firewall Filters for EX Series Switches Overview</i> • <i>Configuring VLANs for EX Series Switches (CLI Procedure)</i>

group (Redundant Trunk Groups)

Syntax	<pre>group name { interface interface-name <primary>; interface interface-name; preempt-cutover-timer seconds; }</pre>
Hierarchy Level	<ul style="list-style-type: none">• For platforms with ELS: [edit switch-options redundant-trunk-group]• For platforms without ELS: [edit ethernet-switching-optionsredundant-trunk-group]
Release Information	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level [edit switch-options] introduced in Junos OS Release 13.2X50-D10 (ELS). (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p>
Description	Create a redundant trunk group.
Options	<p>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.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system—control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Redundant Trunk Links for Faster Recovery on page 112• Example: Configuring Redundant Trunk Links for Faster Recovery• Understanding Redundant Trunk Links on page 49

guard-interval

Syntax	<code>guard-interval <i>number</i>;</code>
Hierarchy Level	[edit protocols protection-group ethernet-ring <i>ring-name</i>]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.
Description	When a link goes down, the ring protection link (RPL) activates. When the downed link comes back up, the RPL link receives notification, restores the link, and waits for the restore interval before issuing another block on the same link. This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.
Options	<i>number</i> —Guard timer interval, in milliseconds. Range: 10 through 2000 ms Default: 500 ms
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Ethernet Ring Protection Switching Overview on page 39 • Example: Configuring Ethernet Ring Protection Switching on EX Series Switches on page 140 • Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches • Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 190

instance-type

Syntax	<code>instance-type type;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. virtual-switch and layer2-control options introduced in Junos OS Release 8.4. Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. evpn option introduced in Junos OS Release 13.2 for MX 3D Series routers.
Description	Define the type of routing instance.

Options



NOTE: On ACX Series routers, you can configure only the forwarding, virtual router, and VRF routing instances.

type—Can be one of the following:

- **evpn**—(MX 3D Series routers only) Enable an Ethernet VPN (EVPN) on the routing instance. You cannot configure the **evpn** option under the [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* instance-type] hierarchy level.
- **forwarding**—Provide support for filter-based forwarding, where interfaces are not associated with instances. All interfaces belong to the default instance. Other instances are used for populating RPD learned routes. For this instance type, there is no one-to-one mapping between an interface and a routing instance. All interfaces belong to the default instance inet.0.
- **l2backhaul-vpn**—Provide support for Layer 2 wholesale VLAN packets with no existing corresponding logical interface. When using this instance, the router learns both the outer tag and inner tag of the incoming packets, when the **instance-role** statement is defined as **access**, or the outer VLAN tag only, when the **instance-role** statement is defined as **nni**.
- **l2vpn**—Enable a Layer 2 VPN on the routing instance. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.
- **layer2-control**—(MX Series routers only) Provide support for RSTP or MSTP in customer edge interfaces of a VPLS routing instance. This instance type cannot be used if the customer edge interface is multihomed to two provider edge interfaces. If the customer edge interface is multihomed to two provider edge interfaces, use the default BPDU tunneling.

- **no-forwarding**—This is the default routing instance. Do not create a corresponding forwarding instance. Use this routing instance type when a separation of routing table information is required. There is no corresponding forwarding table. All routes are installed into the default forwarding table. IS-IS instances are strictly nonforwarding instance types.
- **virtual-router**—Enable a virtual router routing instance. This instance type is similar to a VPN routing and forwarding instance type, but used for non-VPN-related applications. You must configure the **interface** statement for this type of routing instance. You do not need to configure the **route-distinguisher**, **vrf-import**, and **vrf-export** statements.
- **virtual-switch**—(MX Series routers only) Provide support for Layer 2 bridging. Use this routing instance type to isolate a LAN segment with its Spanning Tree Protocol (STP) instance and to separate its VLAN identifier space.
- **vpls**—Enable VPLS on the routing instance. Use this routing instance type for point-to-multipoint LAN implementations between a set of sites in a VPN. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.
- **vrf**—VPN routing and forwarding (VRF) instance. Provides support for Layer 3 VPNs, where interface routes for each instance go into the corresponding forwarding table only. Required to create a Layer 3 VPN. Create a VRF table (*instance-name.inet.0*) that contains the routes originating from and destined for a particular Layer 3 VPN. For this instance type, there is a one-to-one mapping between an interface and a routing instance. Each VRF instance corresponds with a forwarding table. Routes on an interface go into the corresponding forwarding table. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.

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

Related Documentation

- [Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 88](#)
- [Configuring Routing Instances on PE Routers in VPNs](#)
- [Configuring EVPN Routing Instances](#)
- [Configuring Virtual Routing Instances \(CLI Procedure\) on page 174](#)
- [Configuring Virtual Router Routing Instances](#)
- [Example: Configuring Filter-Based Forwarding on the Source Address](#)
- [Example: Configuring Filter-Based Forwarding on Logical Systems](#)
- [Layer 2 Routing Instance Types](#)

interface (Redundant Trunk Groups)

Syntax	<code>interface <i>interface-name</i> <primary>;</code> <code>interface <i>interface-name</i>;</code>
Hierarchy Level	For platforms with ELS: [edit switch-options redundant-trunk-group <i>group name</i>] For platforms without ELS: [edit ethernet-switching-options redundant-trunk-group <i>group name</i>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Hierarchy level [edit switch-options] introduced in Junos OS Release 13.2X50-D10 (ELS). (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.) Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
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	interface <i>interface-name</i> —A logical interface or an aggregated interface containing multiple ports. primary —(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 ge-0/1/0 and ge-0/1/1 , the software assigns ge-0/1/1 as the active link.
Required Privilege Level	system—To view this statement in the configuration. system—control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Redundant Trunk Links for Faster Recovery on page 112• Example: Configuring Redundant Trunk Links for Faster Recovery• Understanding Redundant Trunk Links on page 49

interface (Routing Instances)

Syntax	<pre>interface <i>interface-name</i> { description <i>text</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 13.2 for MX 3D Series routers.</p>
Description	Interface over which the VPN traffic travels between the PE device and CE device. You configure the interface on the PE device. If the value vrf is specified for the instance-type statement included in the routing instance configuration, this statement is required.
Options	<i>interface-name</i> —Name of the 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 Documentation	<ul style="list-style-type: none"> • <i>Configuring Routing Instances on PE Routers in VPNs</i> • <i>Configuring EVPN Routing Instances</i> • <i>Example: Configuring MPLS-Based Layer 3 VPNs on EX Series Switches</i> • <i>interface (VPLS Routing Instances)</i>

interface (VLANs)

Syntax	<pre>interface <i>interface-name</i> { egress; ingress; mapping (native (push swap) policy tag (push swap)); pvlan-trunk; }</pre>
Hierarchy Level	<pre>[edit vlan <i>vlan-name</i>], [edit vlans <i>vlan-name</i>], [edit vlan <i>vlan-name</i> vlan-id <i>number</i>], [edit vlans <i>vlan-name</i> vlan-id <i>number</i>], [edit vlans <i>vlan-name</i> vlan-id-list <i>number</i>]</pre>
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches.
Description	For a specific VLAN, configure an interface.
Options	<p><i>interface-name</i>—Name of a Gigabit Ethernet interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	system—To view this statement in the configuration.system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 57• Configuring VLANs for EX Series Switches (CLI Procedure) on page 160• Configuring Q-in-Q Tunneling (CLI Procedure) on page 180• Configuring Q-in-Q Tunneling (CLI Procedure)

interface (MVRP)

Syntax	<pre>interface (all <i>interface-name</i>) { disable; join-timer <i>milliseconds</i>; leave-timer <i>milliseconds</i>; leaveall-timer <i>milliseconds</i>; registration (forbidden normal); }</pre>
Hierarchy Level	[edit protocols mvrp]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 13.1 for the QFX Series.
Description	Specify interfaces on which to configure Multiple VLAN Registration Protocol (MVRP).



NOTE: On QFX Series switches, you must configure specific interfaces—you cannot specify interface all. You can enable MVRP on an interface range.

Default	By default, MVRP is disabled.
Options	<p>all—All interfaces on the switch.</p> <p><i>interface-name</i>—Names of interface to be configured for MVRP.</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 Documentation	<ul style="list-style-type: none"> • Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 92 • Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 176

interfaces (Q-in-Q Tunneling)

Syntax	<code>interfaces <i>interface-name</i> { no-mac-learning (Q-in-Q Interfaces); }</code>
Hierarchy Level	[edit ethernet-switching-options]
Release Information	Statement introduced in Junos OS Release 9.5 for EX Series switches.
Description	Configure settings for interfaces that have been assigned to family ethernet-switching .
Options	<i>interface-name</i> --Name of an interface that is configured for family ethernet-switching . The remaining statement is explained separately.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Q-in-Q Tunneling on EX Series Switches on page 41



isolation-id

Syntax	<code>isolation-id <i>number</i>;</code>
Hierarchy Level	[edit vlans <i>vlan-name</i> vlan-id <i>number</i>]
Release Information	Statement introduced in Junos OS Release 10.4 for EX Series switches.
Description	Configure an inter-switch isolated VLAN within a private VLAN (PVLAN) that spans multiple switches.
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 Documentation	<ul style="list-style-type: none">• Creating a Private VLAN Spanning Multiple EX Series Switches (CLI Procedure) on page 171

join-timer (MVRP)

Syntax	<code>join-timer <i>milliseconds</i>;</code>
Hierarchy Level	[edit protocols mvrp interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 13.1 for the QFX Series.
Description	Configure the maximum number of milliseconds interfaces must wait before sending Multiple VLAN Registration Protocol (MVRP) protocol data units (PDUs). 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.
Default	200 milliseconds
Options	<i>milliseconds</i> —Number of milliseconds that the interface must wait before sending MVRP PDUs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • leave-timer on page 241 • leaveall-timer on page 240 • Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 92 • Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 176

layer2-protocol-tunneling

Syntax	<pre>layer2-protocol-tunneling all <i>protocol-name</i> { drop-threshold <i>number</i>; shutdown-threshold <i>number</i>; }</pre>
Hierarchy Level	[edit vlans <i>vlan-name</i> dot1q-tunneling]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Enable Layer 2 protocol tunneling (L2PT) on the VLAN.</p> <p>The remaining statements are explained separately.</p>
Default	L2PT is not enabled.
Options	<p>all—Enable all supported Layer 2 protocols.</p> <p><i>protocol-name</i>—Name of the Layer 2 protocol. Values are:</p> <ul style="list-style-type: none"> • 802.1x—IEEE 802.1X authentication • 802.3ah—IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM)
<p> NOTE: If you enable L2PT for untagged OAM LFM packets, do not configure LFM on the corresponding access interface.</p>	
<ul style="list-style-type: none"> • cdp—Cisco Discovery Protocol • e-lmi—Ethernet local management interface • gvrp—GARP VLAN Registration Protocol • lACP—Link Aggregation Control Protocol 	
<p> NOTE: If you enable L2PT for untagged LACP packets, do not configure LACP on the corresponding access interface.</p>	
<ul style="list-style-type: none"> • lldp—Link Layer Discovery Protocol • mmp—Multiple MAC Registration Protocol • mvrp—Multiple VLAN Registration Protocol • stp—Spanning Tree Protocol, Rapid Spanning Tree Protocol, and Multiple Spanning Tree Protocol 	

- **udld**—Unidirectional Link Detection (UDLD)
- **vstp**—VLAN Spanning Tree Protocol
- **vtp**—VLAN Trunking Protocol

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

- Related Documentation**
- [show ethernet-switching layer2-protocol-tunneling interface on page 317](#)
 - [show ethernet-switching layer2-protocol-tunneling statistics on page 319](#)
 - [show ethernet-switching layer2-protocol-tunneling vlan on page 322](#)
 - [Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 107](#)
 - [Configuring Layer 2 Protocol Tunneling on EX Series Switches \(CLI Procedure\) on page 181](#)
 - *Configuring Layer 2 Protocol Tunneling*

l3-interface (VLANs)

Syntax	<code>l3-interface <i>l3-interface-name.logical-interface-number</i> { l3-interface-ingress-counting; }</code>
Hierarchy Level	[edit vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS 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	<p><i>l3-interface-name.logical-interface-number</i>—Name of the Layer 3 interface and number of the logical interface defined by using the set interfaces <i>vlan</i> unit command. The name of the Layer 3 interface is irb for an integrated routing and bridging (IRB) interface, and vlan for a routed VLAN interface (RVI). The number of the logical interface is the same number that you configure in the unit statement.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• show ethernet-switching interfaces on page 313• <i>show ethernet-switching interface</i>• show vlans on page 362• <i>show vlans</i>• Configuring Routed VLAN Interfaces (CLI Procedure) on page 164• <i>Configuring Integrated Routing and Bridging Interfaces (CLI Procedure)</i>

l3-interface-ingress-counting

Syntax	<code>l3-interface-ingress-counting <i>layer-3-interface-name</i>;</code>
Hierarchy Level	<code>[edit vlans <i>vlan-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.3 for EX Series switches.
Description	<p>(EX8200 standalone switch and EX8200 Virtual Chassis) Enable routed VLAN interface (RVI) input counters on an EX8200 switch to collect RVI source statistics for tracking or billing purposes. The input counter is maintained by a firewall filter. The switch can maintain a limited number of firewall filter counters—these counters are allocated on a first-come, first-served basis.</p> <p>Output (egress) counters for EX8200 switches are always present and cannot be removed.</p> <p>Reset ingress-counting statistics with the <code>clear interfaces statistics</code> command.</p>
Default	The input (ingress) counters (both packets and bytes) are disabled on an RVI by default.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • show vlans on page 362 • <code>clear interfaces statistics</code> • <i>Configuring Firewall Filters (CLI Procedure)</i> • <i>firewall</i> • Configuring Routed VLAN Interfaces (CLI Procedure) on page 164

leaveall-timer (MVRP)

Syntax	<code>leaveall-timer interval;</code>
Hierarchy Level	<ul style="list-style-type: none">For platforms with ELS: <code>[edit protocols mvrp],</code> <code>[edit protocols mvrp interface interface-name]</code>For platforms without ELS: <code>[edit protocols mvrp interface (all interface-name)]</code>
Release Information	<p>Statement introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Hierarchy level <code>[edit protocols mvrp]</code> introduced in Junos OS Release 13.2X50-D10 (ELS). (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.1 for the QFX Series.</p>
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. However, if you choose to change the default values, keep in mind that on an EX Series switch that uses Junos OS with support for ELS, if the timer value set on an interface level is different from the value set on a switch level, then the value on the interface level takes precedence.</p>
Options	<p>interval—Number of seconds or milliseconds between the sending of Leave All messages.</p> <p>Default: 10 seconds, or 10,000 milliseconds</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 Documentation	<ul style="list-style-type: none">join-timer on page 235leave-timer on page 241Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 92Example: Configuring Automatic VLAN Administration Using MVRP on EX Series SwitchesConfiguring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 176Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)

leave-timer (MVRP)

Syntax	<code>leave-timer <i>milliseconds</i>;</code>
Hierarchy Level	<code>[edit protocols mvrp interface (all <i>interface-name</i>)]</code>
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 13.1 for the QFX Series.
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 Documentation	<ul style="list-style-type: none"> • join-timer on page 235 • leaveall-timer on page 240 • Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 92 • Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 176

mac (Static MAC-Based VLANs)

Syntax	<code>mac <i>mac-address</i> { <i>next-hop interface-name</i>; }</code>
Hierarchy Level	[edit ethernet-switching-options static vlan <i>vlan-name</i>]
Description	Specify the MAC address to add to the Ethernet switching table. The remaining statement is explained separately.
Options	<i>mac-address</i> —MAC address
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure) on page 187

mac-limit (VLANs)

Syntax	<code>mac-limit <i>limit</i> action <i>action</i>;</code>
Hierarchy Level	[edit vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify the maximum number of MAC addresses to be associated with a VLAN—the default is unlimited , which can leave the network vulnerable to flooding. Change unlimited to any number from 2 to the switch's maximum VLAN MAC limit. The maximum number of MAC addresses allowed in a switching table per VLAN varies depending on the EX Series switch. To see the maximum number of MAC addresses per VLAN allowed on your switch, issue the set vlans <i>vlan-name</i> mac-limit ? configuration-mode command.



NOTE: Do not set the **mac-limit** value to 1. The first learned MAC address is often inserted into the forwarding database automatically—for instance, for a routed VLAN interface (RVI), the first MAC address inserted into the forwarding database is the MAC address of the RVI. For aggregated Ethernet bundles (LAGs) using LACP, the first MAC address inserted into the forwarding database in the Ethernet switching table is the source address of the protocol packet. In these cases, the switch does not learn MAC addresses other than the automatic address when **mac-limit** is set to 1, and this causes problems with MAC learning and forwarding.

When the MAC limit set by this statement is reached, no more MAC addresses are added to the Ethernet switching table. You can also, optionally, have a system log entry generated when the limit is exceeded by adding the option **action log**.



NOTE: When you reconfigure the number of MAC addresses, the Ethernet switching table is not automatically cleared. Therefore, if you reduce the number of addresses from the default (unlimited) or a previously set limit, you could already have more entries in the table than the new limit allows. Previous entries remain in the table after you reduce the number of addresses, so you should clear the Ethernet switching table for a specified interface, MAC address, or VLAN when you reduce the MAC limit. Use the command **clear ethernet-switching table** to clear existing MAC addresses from the table before using the **mac-limit** configuration statement.

Default The MAC limit is disabled, so entries are unlimited.

Options *limit*—Maximum number of MAC addresses.
Range: 1 through *switch maximum*


action—**Log** is the only action available. Configure **action log** to add a message to the system log when the mac-limit value is exceeded. A typical logged message looks like this:

```
May 5 06:18:31 bmp-199p1-dev edwd[5665]:  
ESWD_VLAN_MAC_LIMIT_EXCEEDED: vlan default mac  
00:1f:12:37:af:5b (tag 40). vlan limit exceeded
```

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

Related Documentation	<ul style="list-style-type: none">• show vlans on page 362• Understanding Bridging and VLANs on EX Series Switches on page 3
------------------------------	---

mac-lookup-length

Syntax	<code>mac-lookup-length <i>number-of-entries</i>;</code>
Hierarchy Level	[edit ethernet-switching-options]
Release Information	Statement introduced in Junos OS Release 11.4 for EX Series switches.
Description	<p>Increase the maximum number of searchable hash indexes to mitigate situations in which hash index collisions are causing problems with the learning of MAC addresses in the forwarding database (FDB).</p> <p>The FDB on EX3200, EX3300, EX4200, EX4500, EX4550, and EX6210 switches is a hash table with 8192 hash indexes (rows) of MAC addresses and four entries per hash index. When the FDB is searched, a configured hash function calculates the hash index at which to start the search. By default, after the search starts at the determined hash index, the maximum number of hash indexes that can be searched is one hash index, or four entries</p>
	<div>  <p>NOTE: Increasing the number of hash indexes increases the chances of finding an open entry in which to add a newly learned MAC address. However, searching more hash indexes requires more bandwidth and may impact the FDB performance and line-rate traffic.</p> </div>
Default	4
Options	<p><i>number-of-entries</i>—Maximum number of searchable hash indexes in the FDB.</p> <p>Range: 4, 8, 12</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding Bridging and VLANs on EX Series Switches on page 3

mac-notification

Syntax	<code>mac-notification { notification-interval seconds; }</code>
Hierarchy Level	[edit ethernet-switching-options] [edit switch-options]
Release Information	Statement introduced in Junos OS Release 9.6 for EX Series switches. Hierarchy level [edit switch-options] added in Junos OS Release 14.1X53-D10 for EX Series and QFX Series.
Description	Enable MAC notification for a switch. If you configure this statement without setting a notification interval, MAC notification is enabled with the default MAC notification interval of 30 seconds. The remaining statement is explained separately.
Default	MAC notification is disabled by default.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring MAC Notification (CLI Procedure) on page 175• <i>Configuring MAC Notification (CLI Procedure)</i>

mac-table-aging-time

Syntax	<code>mac-table-aging-time (seconds unlimited);</code>
Hierarchy Level	[edit ethernet-switching-options], [edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement updated in Junos OS Release 9.4 for EX Series switches to include [edit ethernet-switching-options] hierarchy level.
Description	<p>You configure how long MAC addresses remain in the Ethernet switching table using the mac-table-aging-time statement in either the [edit ethernet-switching-options] or the vlans hierarchy, depending on whether you want to configure it for the entire switch or only for specific VLANs.</p> <p>If you specify the time as unlimited, entries are never removed from the table. Generally, use this setting only if the switch or the VLAN has a fairly static number of end devices; otherwise the table will eventually fill up. You can use this setting to minimize traffic loss and flooding that might occur when traffic arrives for MAC addresses that have been removed from the table.</p>
Default	Entries remain in the Ethernet switching table for 300 seconds
Options	<p>seconds—Time that entries remain in the Ethernet switching table before being removed. Range: 60 through 1,000,000 seconds Default: 300 seconds</p> <p>unlimited—Entries remain in the Ethernet switching table.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • show ethernet-switching statistics aging on page 328 • Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 57 • Configuring MAC Table Aging (CLI Procedure) on page 166 • Controlling Authentication Session Timeouts (CLI Procedure) • Configuring VLANs for EX Series Switches (CLI Procedure) on page 160

mapping

Syntax	<code>mapping (native (push swap) policy tag (push swap));</code>
Hierarchy Level	<code>[edit vlan <i>vlan-name</i> interface (VLANs) <i>interface-name</i> egress],</code> <code>[edit vlan <i>vlan-name</i> interface (VLANs) <i>interface-name</i> ingress],</code> <code>[edit vlan <i>vlan-name</i> interface (VLANs) <i>interface-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.6 for EX Series switches. Option swap introduced in Junos OS Release 10.0 for EX Series switches.
Description	<p>Map a specific C-VLAN to an S-VLAN. By default, the received incoming or outgoing tag is replaced with the new tag.</p> <p>This statement is also required if you are configuring firewall filters to map traffic from an interface to a VLAN. If you are configuring firewall filters to map traffic from an interface to a VLAN, the mapping policy option must be configured using this command. The firewall filter also has to be configured using the vlan action for a match condition in the firewall filter stanza for firewall filters to map traffic from an interface for a VLAN.</p>
Options	<p>native—Maps untagged and priority-tagged packets to an S-VLAN.</p> <p>policy—Maps the interface to a firewall filter policy to an S-VLAN.</p> <p>push—Retains the incoming tag and add an additional VLAN tag instead of replacing the original tag.</p> <p>swap—Swaps the incoming VLAN tag with the VLAN ID tag of the S-VLAN. Use of this option is also referred to as VLAN ID translation.</p> <p>tag—Retains the incoming 802.1Q tag on the interface.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring VLANs for EX Series Switches (CLI Procedure) on page 160• Understanding Q-in-Q Tunneling on EX Series Switches on page 41• Understanding Bridging and VLANs on EX Series Switches on page 3

members

Syntax	<code>members [(all <i>names</i> <i>vlan-ids</i>)];</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching vlan]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement updated with enhanced ? (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches.
Description	For trunk interfaces, configure the VLANs that can carry traffic.



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



NOTE: The number of VLANs supported per switch varies for each model. Use the configuration-mode command `set vlans id vlan-id ?` to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because each VLAN is assigned an ID number when it is created. You can, however, exceed the recommended VLAN member maximum.

On an EX Series switch that runs Junos OS that does not support the Enhanced Layer 2 Software (ELS) configuration style, the maximum number of VLAN members allowed on the switch is 8 times the maximum number of VLANs the switch supports (`vmember limit = vlan max * 8`). If the switch configuration exceeds the recommended VLAN member maximum, you see a warning message when you commit the configuration. If you ignore the warning and commit such a configuration, the configuration succeeds but you run the risk of crashing the Ethernet switching process (`eswd`) due to memory allocation failure.

On an EX Series switch that runs Junos OS that supports ELS, the maximum number of VLAN members allowed on the switch is 24 times the maximum number of VLANs the switch supports (`vmember limit = vlan max * 24`). If the configuration of one of these switches exceeds the recommended VLAN member maximum, a warning message appears in the system log (`syslog`).

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



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

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



NOTE: **all** cannot be a VLAN name.

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



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

**Required Privilege
Level**

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

**Related
Documentation**

- [show ethernet-switching interfaces on page 313](#)
- *show ethernet-switching interface*
- [show vlans on page 362](#)
- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 57](#)
- *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*
- [Example: Connecting an Access Switch to a Distribution Switch on page 72](#)
- *Example: Connecting Access Switches to a Distribution Switch*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (J-Web Procedure)*
- [Configuring VLANs for EX Series Switches \(CLI Procedure\) on page 160](#)
- *Configuring VLANs for EX Series Switches (CLI Procedure)*
- [Creating a Series of Tagged VLANs \(CLI Procedure\) on page 168](#)
- [Understanding Bridging and VLANs on EX Series Switches on page 3](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

mvrp

Syntax

```
mvrp {
  add-attribute-length-in-pdu;
  disable;
  interface (all | interface-name) {
    disable;
    join-timer milliseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
    registration (forbidden | normal);
  }
  no-dynamic-vlan;
  traceoptions (Spanning Trees) {
    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 OS Release 10.0 for EX Series switches.

Description



NOTE: If your switch CLI displays different options for the mvrp statement than the options shown in this document, see *mvrp*.

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.



NOTE: At Junos OS Release 11.3, MVRP was updated to conform to the IEEE standard 802.1ak. This update might result in compatibility issues in mixed release networks. For details, see “[Configuring Multiple VLAN Registration Protocol \(MVRP\) \(CLI Procedure\)](#)” on page 176.

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 Documentation

- [Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 92](#)

- [Configuring Multiple VLAN Registration Protocol \(MVRP\) \(CLI Procedure\) on page 176](#)

native-vlan-id

Syntax	<code>native-vlan-id <i>vlan-id</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit 0 family ethernet-switching]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure the VLAN identifier to associate with untagged packets received on the interface.
Options	<i>vlan-id</i> —Numeric identifier of the VLAN. Range: 0 through 4095
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• show vlans on page 362• show ethernet-switching interfaces on page 313• Configuring Gigabit Ethernet Interfaces (CLI Procedure)• Configuring Gigabit Ethernet Interfaces (J-Web Procedure)• Understanding Bridging and VLANs on EX Series Switches on page 3• Junos OS Ethernet Interfaces Configuration Guide

next-hop (Static MAC-Based VLANs)

Syntax	<code>next-hop <i>interface-name</i>;</code>
Hierarchy Level	[edit ethernet-switching-options static vlan <i>vlan-name</i> mac <i>mac-address</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the next hop for the indicated Ethernet node.
Options	<i>interface-name</i> —Name of the next-hop interface.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure) on page 187

no-dynamic-vlan

Syntax	no-dynamic-vlan;
Hierarchy Level	[edit protocols mvrp] [edit protocols mvrp]
Release Information	Statement introduced in Junos OS 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 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	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 176 • Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)

no-local-switching

Syntax	no-local-switching
Hierarchy Level	[edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS 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	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring a Private VLAN on a Single EX Series Switch on page 81 • Creating a Private VLAN on a Single EX Series Switch (CLI Procedure) on page 169

no-mac-learning (Q-in-Q VLANs)

Syntax	no-mac-learning;
Hierarchy Level	[edit vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS 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	system—To view this statement in the configuration. system—control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Q-in-Q Tunneling (CLI Procedure) on page 180• Understanding Q-in-Q Tunneling on EX Series Switches on page 41

no-mac-learning (Q-in-Q Interfaces)

Syntax	no-mac-learning;
Hierarchy Level	[edit ethernet-switching-options interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS 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 Documentation	<ul style="list-style-type: none">• Understanding Q-in-Q Tunneling on EX Series Switches on page 41

node-id

Syntax	<code>node-id mac-address;</code>
Hierarchy Level	[edit protocols protection-group ethernet-ring ring-name]
Release Information	Statement introduced in Junos OS Release 9.4.
Description	<p>For EX Series switches and QFX Series switches, node-id is not configurable.</p> <p>For MX Series routers, optionally specify the MAC address of a node in the protection group. If this statement is not included, the router assigns the node's MAC address.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Ethernet Ring Protection Switching Overview on page 39• Example: Configuring Ethernet Ring Protection Switching on EX Series Switches on page 140• Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches

notification-interval

Syntax	notification-interval <i>seconds</i> ;
Hierarchy Level	[edit ethernet-switching-options mac-notification] [edit switch-options mac-notification]
Release Information	Statement introduced in Junos OS Release 9.6 for EX Series switches. Hierarchy level [edit switch-options] added in Junos OS Release 14.1X53-D10 for EX Series and QFX Series.
Description	<p>Configure the MAC notification interval for a switch.</p> <p>The MAC notification interval is the amount of time the switch waits before sending learned or unlearned MAC address SNMP notifications to the network management server. For instance, if the MAC notification interval is set to 10, all of the MAC address addition and removal SNMP notifications will be sent to the network management system every 10 seconds.</p>
Options	<p><i>seconds</i>—The MAC notification interval, in seconds.</p> <p>Range: 1 through 60</p> <p>Default: 30</p>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring MAC Notification (CLI Procedure) on page 175• Configuring MAC Notification (CLI Procedure)

port-mode

Syntax	<code>port-mode mode;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure whether an interface on the switch operates in access, tagged-access, or trunk mode.
Default	All switch interfaces are in access mode.
Options	<p>mode—Operating mode for an interface can be one of the following:</p> <ul style="list-style-type: none"> • access—In this mode, the interface can be in a single VLAN only. Access interfaces typically connect to single network devices such as PCs, printers, IP telephones, and IP cameras. • tagged-access—In this mode, the interface can accept tagged packets from one access device. Tagged-access interfaces typically connect to servers running Virtual machines using VEPA technology. • trunk—In this mode, the interface can be in multiple VLANs and accept tagged packets from multiple devices. Trunk interfaces typically connect to other switches and to routers on the LAN.



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

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


Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Connecting an Access Switch to a Distribution Switch on page 72 • Configuring Gigabit Ethernet Interfaces (CLI Procedure)

- [Configuring VLANs for EX Series Switches \(CLI Procedure\) on page 160](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

preempt-cutover-timer

Syntax	<code>preempt-cutover-timer seconds;</code>
Hierarchy Level	<ul style="list-style-type: none">• For platforms with ELS: <code>[edit switch-options redundant-trunk-group group name]</code>• For platforms without ELS: <code>[edit ethernet-switching-optionsredundant-trunk-group group name]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for EX Series switches. Hierarchy level <code>[edit switch-options]</code> introduced in Junos OS Release 13.2X50-D10 (ELS). (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.) Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Description	Change the length of time that a re-enabled primary link waits to take over from an active secondary link in a redundant trunk group.
Default	If you do not change the time with the preempt-cutover-timer statement, a re-enabled primary link takes over from the active secondary link after 120 seconds.
Options	seconds —Number of seconds that the primary link waits to take over from the active secondary link. Range: 1 through 600 seconds
Required Privilege Level	admin —To view this statement in the configuration. admin-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Redundant Trunk Links for Faster Recovery on page 112• Example: Configuring Redundant Trunk Links for Faster Recovery• Understanding Redundant Trunk Links on page 49

primary-vlan

Syntax	<code>primary-vlan <i>vlan-name</i>;</code>
Hierarchy Level	[edit vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement updated with enhanced ? (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches.
Description	<p>Configure the primary VLAN for this private VLAN (PVLAN). The primary VLAN is always tagged.</p> <ul style="list-style-type: none"> • If the PVLAN is configured on a single switch, do not assign a tag to the community VLANs. • If the PVLAN is configured to span multiple switches, you must assign tags to the community VLANs also.
	<div>  <p>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 name is displayed for a VLAN range.</p> </div>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system—control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring a Private VLAN on a Single EX Series Switch on page 81 • Example: Configuring a Private VLAN Spanning Multiple EX Series Switches on page 119 • Creating a Private VLAN on a Single EX Series Switch (CLI Procedure) on page 169 • Creating a Private VLAN Spanning Multiple EX Series Switches (CLI Procedure) on page 171

protection-group

```
Syntax  protection-group {
        ethernet-ring ring-name {
            control-channel channel-name {
                vlan number;
                interface name interface-name
            }
            data-channel {
                vlan number
            }
            east-interface {
                control-channel channel-name {
                    vlan number;
                    interface name interface-name
                }
            }
            guard-interval number;
            node-id mac-address;
            restore-interval number;
            ring-protection-link-owner;
            west-interface {
                control-channel channel-name {
                    vlan number;
                    interface name interface-name
                }
            }
        }
        control-vlan (vlan-id | vlan-name);
        east-interface {
            node-id mac-address;
            control-channel channel-name {
                vlan number;
                interface name interface-name
            }
            interface-none
            ring-protection-link-end;
        }
        control-channel channel-name {
            vlan number;
            interface name interface-name
        }
        data-channel {
            vlan number
        }
        guard-interval number;
        node-id mac-address;
        restore-interval number;
        ring-protection-link-owner;
        west-interface {
            node-id mac-address;
            control-channel channel-name {
```




```

        vlan number;
        interface name interface-name
    }
    interface-none
    ring-protection-link-end;
}
control-channel channel-name {
    vlan number;
    interface name interface-name
}
}
}
guard-interval number;
restore-interval number;
traceoptions {
    file filename <no-stamp> <world-readable | no-world-readable> <replace> <size size>;
    flag flag;
}
}

```

Hierarchy Level	[edit protocols]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.1 for EX Series switches.
Description	Configure Ethernet ring protection switching. The statements are explained separately. All statements apply to MX Series routers. EX Series switches do not assign node-id and use control-vlan instead of control-channel .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Ethernet Ring Protection Switching Overview on page 39 • Ethernet Ring Protection Using Ring Instances for Load Balancing • Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers • Example: Configuring Ethernet Ring Protection Switching on EX Series Switches on page 140 • Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches • Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 190

proxy-arp

Syntax	<code>proxy-arp (restricted unrestricted);</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.6 for EX Series switches. restricted added in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 12.2 for the QFX Series.
Description	For Ethernet interfaces only, configure the router or switch to respond to any ARP request, as long as the router or switch has an active route to the ARP request's target address.
<div>  NOTE: You must configure the IP address and the inet family for the interface when you enable proxy ARP. </div>	
Default	Proxy ARP is not enabled. The router or switch responds to an ARP request only if the destination IP address is its own.
Options	<ul style="list-style-type: none"> • none—The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address. • restricted—(Optional) The router or switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are in the same subnet. The router or switch must also have a route to the target IP address. • unrestricted—(Optional) The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.
	Default: unrestricted
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Restricted and Unrestricted Proxy ARP</i> • Configuring Proxy ARP (CLI Procedure) on page 186 • <i>Configuring Proxy ARP (CLI Procedure)</i> • Example: Configuring Proxy ARP on an EX Series Switch on page 117 • <i>Configuring Gratuitous ARP</i>

pvlan-trunk

Syntax	pvlan-trunk;
Hierarchy Level	[edit vlan <i>vlan-name</i> vlan-id <i>number</i> interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 10.4 for EX Series switches.
Description	Configure an interface to be the trunk port, connecting switches that are configured with a private VLAN (PVLAN) across these switches.
Required Privilege Level	routing—To view this statement in the configuration. routing—control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Creating a Private VLAN Spanning Multiple EX Series Switches (CLI Procedure) on page 171

redundant-trunk-group

Syntax	<pre>redundant-trunk-group { group name { interface interface-name <primary>; interface interface-name; preempt-cutover-timer seconds; } }</pre>
Hierarchy Level	<ul style="list-style-type: none">• For platforms with ELS: [edit switch-options]• For platforms without ELS: [edit ethernet-switching-options]
Release Information	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level [edit switch-options] introduced in Junos OS Release 13.2X50-D10 (ELS). (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p>
Description	<p>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 spanning-tree protocol convergence.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system—control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Redundant Trunk Links for Faster Recovery on page 112• Example: Configuring Redundant Trunk Links for Faster Recovery• Understanding Redundant Trunk Links on page 49

registration

Syntax	registration (forbidden normal);
Hierarchy Level	[edit protocols mvrp interface (all <i>interface-name</i>)], [edit protocols mvrp interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS 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	forbidden —The interface or interfaces do not register and do not participate in MVRP. normal —The interface or interfaces accept MVRP messages and participate in MVRP.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 176 • <i>Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)</i>

restore-interval

Syntax	<code>restore-interval <i>number</i>;</code>
Hierarchy Level	[edit protocols protection-group ethernet-ring <i>ring-name</i>]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.
Description	Configures the number of minutes that the node does not process any Ethernet ring protection (ERP) protocol data units (PDUs).. This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.
Options	<i>number</i> —Specify the restore interval. Range: 5 through 12 minutes
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Ethernet Ring Protection Switching Overview on page 39• Example: Configuring Ethernet Ring Protection Switching on EX Series Switches on page 140• Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches• Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 190

ring-protection-link-end

Syntax	ring-protection-link-end;
Hierarchy Level	[edit protocols protection-group ethernet-ring <i>ring-name</i> (east-interface west-interface)]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.1 for EX Series switches.
Description	Specify that the port is one side of a ring protection link (RPL) by setting the RPL end flag.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Ethernet Ring Protection Switching Overview on page 39 • Example: Configuring Ethernet Ring Protection Switching on EX Series Switches on page 140 • Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches • Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 190

ring-protection-link-owner

Syntax	ring-protection-link-owner;
Hierarchy Level	[edit protocols protection-group ethernet-ring <i>ring-name</i>]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.
Description	Specify the ring protection link (RPL) owner flag in the Ethernet protection ring. Include this statement only once for each ring (only one node can function as the RPL owner).
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Ethernet Ring Protection Switching Overview on page 39 • Example: Configuring Ethernet Ring Protection Switching on EX Series Switches on page 140 • Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches

routing-instances

Syntax	<code>routing-instances <i>routing-instance-name</i> { <i>instance-type</i> virtual-router; interface <i>interface-name</i>; }</code>
Hierarchy Level	[edit]
Release Information	Statement introduced in Junos OS Release 9.2 for EX Series switches.
Description	Configure a virtual routing entity.
Options	<i>routing-instance-name</i> —Name for this routing instance. 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 Documentation	<ul style="list-style-type: none">• Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 88• Configuring Virtual Routing Instances (CLI Procedure) on page 174

shutdown-threshold

Syntax	<code>shutdown-threshold <i>number</i>;</code>
Hierarchy Level	<code>[edit vlan <i>vlan-name</i> dot1q-tunneling layer2-protocol-tunneling (all <i>protocol-name</i>)]</code>
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
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 clear ethernet-switching layer2-protocol-tunneling error 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>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • drop-threshold on page 217 • Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 107 • Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) on page 181 • Configuring Layer 2 Protocol Tunneling

static (Static MAC-Based VLANs)

Syntax static {
 vlan *vlan-name* {
 mac *mac-address* {
 next-hop *interface-name*;
 }
 }
 }

Hierarchy Level [edit ethernet-switching-options]

Release Information Statement introduced in Junos OS Release 11.1 for EX Series switches.

Description Specify VLAN and MAC addresses to add to the Ethernet switching table.

 The remaining statements are explained separately.

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

Related Documentation • [Adding a Static MAC Address Entry to the Ethernet Switching Table \(CLI Procedure\) on page 187](#)

traceoptions (Ethernet Ring Protection)

Syntax	<pre>traceoptions { file <i>filename</i> <no-stamp> <world-readable no-world-readable> <replace> <size <i>size</i>>; flag <i>flag</i>; }</pre>
Hierarchy Level	[edit protocols protection-group]
Release Information	<p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.</p>
Description	Configure trace options for the protection group.
Options	<p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. All files are placed in the directory <code>/var/log</code>. You can include the following file options:</p> <ul style="list-style-type: none"> • no-stamp—(Optional) Do not timestamp trace file. • no-world-readable—(Optional) Do not allow any user to read the log file. • replace—(Optional) Replace the trace file rather than appending to it. • size—(Optional) Maximum trace file size (10240..4294967295). • world-readable—(Optional) Allow any user to read the log file. <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"> • all—Trace all SBC process operations. • configuration—Trace configuration events. • debug—Trace device monitor events. • events—Trace events to the protocol state machine • normal—Trace normal messages. • pdu—Trace RAPS PDU reception and transmission. • periodic-packet-management—Trace periodic packet management state and events. • state-machine—Trace RAPS state machine. • timers—Trace protocol timers.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 190 • Example: Configuring Ethernet Ring Protection Switching on EX Series Switches on page 140

- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches*

traceoptions (Edge Virtual Bridging)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <no-stamp> <replace> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; } </pre>
Hierarchy Level	[edit protocols edge-virtual-bridging]
Release Information	Statement introduced in Junos OS Release 12.1 for EX Series switches.
Description	Define global tracing operations for edge virtual bridging (EVB) features on Ethernet switches.
Default	Tracing operations are disabled by default.
Options	<p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <code>/var/log</code>.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached (xk to specify KB, xm to specify MB, or xg to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000</p> <p>Default: 3 files</p> <p>no-world-readable—(Optional) Restrict file access to the user who created the file.</p> <p>replace—(Optional) Replace an existing trace file if there is one rather than appending output to it.</p> <p>Default: If you do not include this option, tracing output is appended to an existing trace file.</p> <p>size <i>size</i>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the files option.</p> <p>Syntax: xk to specify KB, xm to specify MB, or xg to specify gigabytes</p> <p>Range: 10 KB through 1 gigabyte</p> <p>Default: 128 KB</p> <p>world-readable—(Optional) Enable unrestricted file access.</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p>

- **all**—Trace everything.
- **ecp**—Trace Edge Control Protocol (ECP) events.
- **evb-tlv**— Trace EVB type, length, and value (TLV) events.
- **parse**—Trace configuration parsing.
- **policy**—Trace policy events.
- **vdp**—Trace Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP) events.

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

Related Documentation • [Configuring Edge Virtual Bridging \(CLI Procedure\) on page 188](#)

vlan (802.1Q Tagging)

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

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

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

Description Bind an 802.1Q VLAN tag ID to a logical interface.

The remaining statement is explained separately.

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

Related Documentation • [show ethernet-switching interfaces on page 313](#)
 • *show ethernet-switching interface*
 • [Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 65](#)
 • [Configuring Routed VLAN Interfaces \(CLI Procedure\) on page 164](#)
 • [Configuring Integrated Routing and Bridging Interfaces \(CLI Procedure\)](#)
 • [Understanding Bridging and VLANs on EX Series Switches on page 3](#)
 • [Junos OS Ethernet Interfaces Configuration Guide](#)

vlan (Static MAC-based VLANs)

Syntax	<pre>vlan <i>vlan-name</i> { <i>mac mac-address</i> { <i>next-hop interface-name</i>; } }</pre>
Hierarchy Level	[edit ethernet-switching-options static]
Release Information	Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the name of a VLAN to add to the Ethernet switching table.
Options	<i>vlan-name</i> —Name of the VLAN to add to the Ethernet switching table. The remaining statements are explained separately.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure) on page 187

vlan-id (802.1Q Tagging)

Syntax	<code>vlan-id <i>number</i>;</code>
Hierarchy Level	[edit vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Configure an 802.1Q tag to apply to all traffic that originates on the VLAN.</p> <p>The number zero is reserved for priority tagging and the number 4095 is also reserved.</p>
Default	If you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of 1.
Options	<p><i>number</i> —VLAN tag identifier</p> <p>Range:</p> <ul style="list-style-type: none">• 1 through 4094 (all switches except EX8200 Virtual Chassis)• 1 through 4092 (EX8200 Virtual Chassis only) <p>Default: 1</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 65• Example: Connecting Access Switches to a Distribution Switch• Example: Configuring a Private VLAN on a Single EX Series Switch on page 81• Example: Configuring a Private VLAN Spanning Multiple EX Series Switches on page 119• Creating a Private VLAN on a Single EX Series Switch (CLI Procedure) on page 169• Creating a Private VLAN Spanning Multiple EX Series Switches (CLI Procedure) on page 171

vlan-prune

Syntax	vlan-prune;
Hierarchy Level	[edit vllans <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3 for EX Series switches.
Description	<p>Prune the Virtual Chassis port (VCP) paths in a Virtual Chassis to ensure received broadcast, multicast, and unknown unicast traffic in a VLAN uses the shortest possible path through the Virtual Chassis to the egress VLAN interface.</p> <p>By default, all broadcast, multicast, and unknown unicast traffic in a VLAN on an EX Series Virtual Chassis is broadcast to all member switches in the Virtual Chassis. This behavior unnecessarily consumes bandwidth within the Virtual Chassis because unneeded traffic is sent to all Virtual Chassis member switches.</p> <p>Enabling this option allows you to conserve bandwidth within the Virtual Chassis. Broadcast, multicast, and unknown unicast traffic still enters and exits the Virtual Chassis within the same VLAN, without the added bandwidth consumption that results from broadcasting this traffic to all member switches.</p>
Default	Disabled
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring VLANs for EX Series Switches (CLI Procedure) on page 160

vlan-range

Syntax	<code>vlan-range <i>vlan-id-low-vlan-id-high</i>;</code>
Hierarchy Level	[edit vlan <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS 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	system—To view this statement in the configuration. system—control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VLANs for EX Series Switches (CLI Procedure) on page 160• Configuring VLANs for EX Series Switches (J-Web Procedure) on page 158• Configuring Routed VLAN Interfaces (CLI Procedure) on page 164• 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 {
                egress;
                ingress;
                mapping (native (push | swap) | policy | tag (push | swap));
                pvlan-trunk;
            }
            isolation-id id-number;
            l3-interface l3-interface-name.logical-interface-number;
            l3-interface-ingress-counting layer-3-interface-name;
            mac-limit limit action action;
            mac-table-aging-time seconds;
            no-local-switching;
            no-mac-learning;
            primary-vlan vlan-name;
            vlan-id number;
            vlan-prune;
            vlan-range vlan-id-low-vlan-id-high;
        }
    }
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 9.0 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 Q-in-Q-tunneled and the packet is arriving from an access interface.
- You cannot use a firewall filter to assign an integrated routing and bridging (IRB) interface or a routed VLAN interface (RVI) to a VLAN.
- VLAN assignments performed using a firewall filter override all other VLAN assignments.

Options *vlan-name*—Name of the VLAN. The name can include letters, numbers, hyphens (-), and periods (.) and can contain up to 255 characters long.

The remaining statements are explained separately.

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VLANs for EX Series Switches (CLI Procedure) on page 160• Configuring VLANs for EX Series Switches (CLI Procedure)• Configuring Q-in-Q Tunneling (CLI Procedure) on page 180• Creating a Series of Tagged VLANs (CLI Procedure) on page 168• Configuring Routed VLAN Interfaces (CLI Procedure) on page 164• Configuring Integrated Routing and Bridging Interfaces (CLI Procedure)• Understanding Bridging and VLANs on EX Series Switches on page 3

vrf-mtu-check

Syntax	vrf-mtu-check;
Hierarchy Level	[edit chassis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches.
Description	On M Series routers (except the M120 and M320 router), T Series routers, and on EX Series 8200 switches, configure path maximum transmission unit (MTU) checks on the outgoing interface for unicast traffic routed on a virtual private network (VPN) routing and forwarding (VRF) instance.
Default	Disabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Path MTU Checks for VPNs• Configuring the Junos OS to Enable MTU Path Check for a Routing Instance on M Series Routers


vsi-discovery

Syntax	<pre>vsi-discovery { interface <i>interface-name</i> vsi-policy <i>vsi-policy-name</i> }</pre>
Hierarchy Level	[edit protocols edge-virtual-bridging]
Description	Configure Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP). VDP is used to program policies for each individual station interface (VSI).
Default	VDP is disabled by default.
Options	interface-name —Name of the interface on which VDP is configured. The remaining statement is explained separately.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Edge Virtual Bridging for Use with VEPA Technology on page 134• Configuring Edge Virtual Bridging (CLI Procedure) on page 188

vsi-policy

Syntax	<code>vsi-policy <i>vsi-policy-name</i> from vsi-manager <i>vsi-manager-id</i> vsi-type <i>vsi-type</i> vsi-version <i>vsi-version</i> vsi-instance <i>instance-number</i>;</code>
Hierarchy Level	[edit policy-options]
Description	<p>Define and apply the named VSI policy to the edge virtual bridging (EVB) configuration. For use with edge virtual bridging, each virtual machine (VM) on the server is uniquely identified by following four parameters, which are contained in a VSI policy:</p> <ul style="list-style-type: none">• vsi-manager-id• vsi-type• vsi-version• vsi-instance-id <p>The vsi-policy command manually configures these four parameters on the EX switch for the successful association of VM-VSI. VDP protocol helps determine the parameters defined for the virtual machines on the server and configure them on the switch. Use policy options to define the VM-VSI parameters. Configure a firewall filter for each of the VM profiles and use it in this statement.</p>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Edge Virtual Bridging for Use with VEPA Technology on page 134• Configuring Edge Virtual Bridging (CLI Procedure) on page 188

west-interface

Syntax	<pre> west-interface { node-id mac-address; control-channel channel-name { vlan number; interface name interface-name } interface-name ring-protection-link-end; } </pre>
Hierarchy Level	[edit protocols protection-group ethernet-ring ring-name]
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.</p>
Description	<p>Define one of the two interface ports for Ethernet ring protection, the other being defined by the east-interface statement at the same hierarchy level. The interface must use the control channel's logical interface name. The control channel is a dedicated VLAN channel for the ring port.</p>
<div>  NOTE: Always configure this port second, after configuring the east-interface statement. </div>	
<p>The statements are explained separately.</p>	
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Ethernet Ring Protection Switching Overview on page 39 • Ethernet Ring Protection Using Ring Instances for Load Balancing • east-interface on page 218 • ethernet-ring on page 220 • Example: Configuring Ethernet Ring Protection Switching on EX Series Switches on page 140 • Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches • Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 190

PART 3

Administration

- [Routine Monitoring on page 287](#)
- [Operational Commands on page 303](#)

CHAPTER 9

Routine Monitoring

- [Verifying That a Series of Tagged VLANs Has Been Created on page 287](#)
- [Verifying That Virtual Routing Instances Are Working on page 289](#)
- [Verifying That Q-in-Q Tunneling Is Working on page 290](#)
- [Verifying Routed VLAN Interface Status and Statistics on page 291](#)
- [Verifying That a Private VLAN Is Working on page 292](#)
- [Verifying That MVRP Is Working Correctly on page 297](#)
- [Verifying That MAC Notification Is Working Properly on page 299](#)
- [Verifying That Proxy ARP Is Working Correctly on page 299](#)
- [Monitoring Ethernet Switching on page 300](#)

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 Documentation

- [Creating a Series of Tagged VLANs \(CLI Procedure\) on page 168](#)

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	Primary RIB	Type	Active/holddown/hidden
master	inet.0	forwarding	3/0/0
__juniper_private1__	__juniper_private1__.inet.0	forwarding	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          Type Index NhRef Netif
0.0.0.0/32       perm  0          dscd  537   1
103.1.1.0/24     ifdn  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          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          mdsc  538   1
224.0.0.1/32     perm  0 224.0.0.1  mcst  534   1
255.255.255.255/32 perm  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 Documentation

- [Configuring Virtual Routing Instances \(CLI Procedure\) on page 174](#)
- [Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 88](#)

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)
xe-0/0/1, tagged, trunk
xe-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 Documentation

- [Configuring Q-in-Q Tunneling \(CLI Procedure\) on page 180](#)
- [Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 103](#)

Verifying Routed VLAN Interface Status and Statistics

Purpose Determine status information and traffic statistics for routed VLAN interfaces (RVIs) by using the following commands:

Action Display RVI interfaces and their current states:

```

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

```

Display Layer 2 VLANs, including any tags assigned to the VLANs and the interfaces associated with the VLANs:

```

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*

```

Display Ethernet switching table entries for the VLAN that is attached to the RVI:

```

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

```

Display an RVI's ingress-counting statistics with either the **show interfaces vlan detail** command or the **show interfaces vlan extensive** command. Ingress counting is displayed as **Input bytes** and **Input packets** under **Transit Statistics**.

```

user@switch> show interfaces vlan.100 detail

Logical interface vlan.100 (Index 65) (SNMP ifIndex 503) (HW Token 100) (Generation 131)
Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
Traffic statistics:
Input bytes:    17516756
Output bytes:   411764

```

```
Input packets: 271745
Output packets: 8256
Local statistics:
Input bytes: 3240
Output bytes: 411764
Input packets: 54
Output packets: 8256
Transit statistics:
Input bytes: 17513516 0 bps
Output bytes: 0 0 bps
Input packets: 271745 0 pps
Output packets: 0 0 pps
Protocol inet, Generation: 148, Route table: 0
Flags: None
Addresses, Flags: is-Preferred is-Primary
Destination: 50.1.1/24, Local: 50.1.1.1, Broadcast: 50.1.1.255, Generation: 136
```

- Meaning**
- **show interfaces vlan** displays a list of interfaces, including RVI interfaces, and their current states (up, down).
 - **show vlans** displays a list of VLANs, including any tags assigned to the VLANs and the interfaces associated with the VLANs.
 - **show ethernet-switching table** displays the Ethernet switching table entries, including VLANs attached to the RVI.
 - **show interfaces vlan detail** displays RVI ingress counting as Input Bytes and Input Packets under Transit Statistics.

- Related Documentation**
- [Configuring Routed VLAN Interfaces \(CLI Procedure\) on page 164](#)

Verifying That a Private VLAN Is Working

Purpose After creating and configuring private VLANs (PVLANS), verify that they are set up properly.

- Action**
1. To determine whether you successfully created the primary and secondary VLAN configurations:

- For a PVLAN on a single switch, use the **show configuration vlans** command:

```
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;
}

```

- For a PVLAN spanning multiple switches, use the [show vlans extensive](#) command:

```

user@switch> show vlans extensive
VLAN: COM1, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 100, Internal index: 3, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 1)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/7.0*, untagged, access

```

```

VLAN: __pvlan_primary_ge-0/0/0.0__, Created at: Tue May 11 18:16:05 2010
Internal index: 5, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 1)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/0.0*, untagged, access

```

```

VLAN: __pvlan_primary_ge-0/0/2.0__, Created at: Tue May 11 18:16:05 2010
Internal index: 6, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 0)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/2.0, untagged, access

```

```

VLAN: __pvlan_primary_isiv__, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 50, Internal index: 7, Admin State: Enabled, Origin: Static
Private VLAN Mode: Inter-switch-isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 0 (Active = 0)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk

```

```

VLAN: community2, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 20, Internal index: 8, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 2 (Active = 2)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk

```

```

ge-0/0/1.0*, untagged, access
ge-1/0/6.0*, untagged, access

```

```

VLAN: primary, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 10, Internal index: 2, Admin State: Enabled, Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 5 (Active = 4)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/0.0*, untagged, access
    ge-0/0/1.0*, untagged, access
    ge-0/0/2.0, untagged, access
    ge-0/0/7.0*, untagged, access
    ge-1/0/6.0*, untagged, access

```

```

Secondary VLANs: Isolated 2, Community 2, Inter-switch-isolated 1
Isolated VLANs :
    __pvlan_primary_ge-0/0/0.0__
    __pvlan_primary_ge-0/0/2.0__
Community VLANs :
    COM1
    community2
Inter-switch-isolated VLAN :
    __pvlan_primary_isiv__

```

2. Use the **show vlans extensive** command to view VLAN information and link status for a PVLAN on a single switch or for a PVLAN spanning multiple switches.

- For a PVLAN on a single switch:

```

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

```

- For a PVLAN spanning multiple switches:

```

user@switch> show vlans extensive
VLAN: COM1, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 100, Internal index: 3, Admin State: Enabled, Origin: Static

```

Private VLAN Mode: Community, Primary VLAN: primary
 Protocol: Port Mode, Mac aging time: 300 seconds
 Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 1)
 ge-0/0/20.0*, tagged, trunk
 ge-0/0/22.0*, tagged, trunk, pvlan-trunk
 ge-0/0/23.0*, tagged, trunk, pvlan-trunk
 ge-0/0/7.0*, untagged, access

VLAN: __pvlan_primary_ge-0/0/0.0__, Created at: Tue May 11 18:16:05 2010
 Internal index: 5, Admin State: Enabled, Origin: Static
 Private VLAN Mode: Isolated, Primary VLAN: primary
 Protocol: Port Mode, Mac aging time: 300 seconds
 Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 1)
 ge-0/0/20.0*, tagged, trunk
 ge-0/0/22.0*, tagged, trunk, pvlan-trunk
 ge-0/0/23.0*, tagged, trunk, pvlan-trunk
 ge-0/0/0.0*, untagged, access

VLAN: __pvlan_primary_ge-0/0/2.0__, Created at: Tue May 11 18:16:05 2010
 Internal index: 6, Admin State: Enabled, Origin: Static
 Private VLAN Mode: Isolated, Primary VLAN: primary
 Protocol: Port Mode, Mac aging time: 300 seconds
 Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 0)
 ge-0/0/20.0*, tagged, trunk
 ge-0/0/22.0*, tagged, trunk, pvlan-trunk
 ge-0/0/23.0*, tagged, trunk, pvlan-trunk
 ge-0/0/2.0, untagged, access

VLAN: __pvlan_primary_isiv__, Created at: Tue May 11 18:16:05 2010
 802.1Q Tag: 50, Internal index: 7, Admin State: Enabled, Origin: Static
 Private VLAN Mode: Inter-switch-isolated, Primary VLAN: primary
 Protocol: Port Mode, Mac aging time: 300 seconds
 Number of interfaces: Tagged 3 (Active = 3), Untagged 0 (Active = 0)
 ge-0/0/20.0*, tagged, trunk
 ge-0/0/22.0*, tagged, trunk, pvlan-trunk
 ge-0/0/23.0*, tagged, trunk, pvlan-trunk

VLAN: community2, Created at: Tue May 11 18:16:05 2010
 802.1Q Tag: 20, Internal index: 8, Admin State: Enabled, Origin: Static
 Private VLAN Mode: Community, Primary VLAN: primary
 Protocol: Port Mode, Mac aging time: 300 seconds
 Number of interfaces: Tagged 3 (Active = 3), Untagged 2 (Active = 2)
 ge-0/0/20.0*, tagged, trunk
 ge-0/0/22.0*, tagged, trunk, pvlan-trunk
 ge-0/0/23.0*, tagged, trunk, pvlan-trunk
 ge-0/0/1.0*, untagged, access
 ge-1/0/6.0*, untagged, access

VLAN: primary, Created at: Tue May 11 18:16:05 2010
 802.1Q Tag: 10, Internal index: 2, Admin State: Enabled, Origin: Static
 Private VLAN Mode: Primary
 Protocol: Port Mode, Mac aging time: 300 seconds
 Number of interfaces: Tagged 3 (Active = 3), Untagged 5 (Active = 4)
 ge-0/0/20.0*, tagged, trunk
 ge-0/0/22.0*, tagged, trunk, pvlan-trunk
 ge-0/0/23.0*, tagged, trunk, pvlan-trunk
 ge-0/0/0.0*, untagged, access

```

ge-0/0/1.0*, untagged, access
ge-0/0/2.0, untagged, access
ge-0/0/7.0*, untagged, access
ge-1/0/6.0*, untagged, access

```

Secondary VLANs: Isolated 2, Community 2, Inter-switch-isolated 1

Isolated VLANs :

```
__pvlan_primary_ge-0/0/0.0__
```

```
__pvlan_primary_ge-0/0/2.0__
```

Community VLANs :

```
COM1
```

```
community2
```

Inter-switch-isolated VLAN :

```
__pvlan_primary_isiv__
```

3. Use the **show ethernet-switching table** command to view logs for MAC learning on the VLANs:

```

user@switch> show ethernet-switching table
Ethernet-switching table: 8 entries, 1 learned

```

VLAN	MAC address	Type	Age	Interfaces
default	*	Flood	-	All-members
pvlan	*	Flood	-	All-members
pvlan	MAC1	Replicated	-	interface a
pvlan	MAC2	Replicated	-	interface c
pvlan	MAC3	Replicated	-	isolated2
pvlan	MAC4	Learn	0	trunk1
__pvlan_pvlan_isolated1__	*	Flood	-	All-members
__pvlan_pvlan_isolated1__	MAC4	Replicated	-	trunk1
__pvlan_pvlan_isolated2__	*	Flood	-	All-members
__pvlan_pvlan_isolated2__	MAC3	Learn	0	isolated2
__pvlan_pvlan_isolated2__	MAC4	Replicated	-	trunk1
community1	*	Flood	-	All-members
community1	MAC1	Learn	0	interface a
community1	MAC4	Replicated	-	trunk1
community2	*	Flood	-	All-members
community2	MAC2	Learn	0	interface c
community2	MAC4	Replicated	-	trunk1



NOTE: If you have configured a PVLAN spanning multiple switches, you can use the same command on all the switches to check the logs for MAC learning on those switches.

Meaning In the output displays for a PVLAN on a single switch, you can see that the primary VLAN contains two community domains (**community1** and **community2**), two isolated ports, and two trunk ports. The PVLAN on a single switch has only one tag (**1000**), which is for the primary VLAN.

The PVLAN that spans multiple switches contains multiple tags:

- The community domain **COM1** is identified with tag **100**.
- The community domain **community2** is identified with tag **20**.
- The interswitch isolated domain is identified with tag **50**.
- The primary VLAN **primary** is identified with tag **10**.

Also, for the PVLAN that spans multiple switches, the trunk interfaces are identified as **pvlan-trunk**.

- Related Documentation**
- [Creating a Private VLAN on a Single EX Series Switch \(CLI Procedure\) on page 169](#)
 - [Creating a Private VLAN Spanning Multiple EX Series Switches \(CLI Procedure\) on page 171](#)
 - *Creating a Private VLAN on a Single Switch*
 - *Creating a Private VLAN Spanning Multiple Switches*

Verifying That MVRP Is Working Correctly

Purpose After configuring your 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
-----
a11            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.



NOTE: You can identify an MVRP compatibility issue on EX Series switches by looking at the output from this command. If *Join Empty received* and *Join In received* incorrectly display zero, even though the value for *MRPDU received* has been increased, you are probably running different versions of Junos OS, including Release 11.3, on the switches in this network. Another indication that MVRP is having a version problem is that unexpected VLAN activity, such as multiple VLAN creation, takes place on the switch running the earlier release version. To remedy these problems, see [“Configuring Multiple VLAN Registration Protocol \(MVRP\) \(CLI Procedure\)” on page 176](#).

- Related Documentation**
- [Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 92](#)
 - [Example: Configuring Automatic VLAN Administration Using MVRP](#)
 - [Configuring Multiple VLAN Registration Protocol \(MVRP\) \(CLI Procedure\) on page 176](#)

Verifying That MAC Notification Is Working Properly

Purpose	Verify that MAC notification is enabled or disabled, and that the MAC notification interval is set to the specified value.
Action	<p>Verify that MAC notification is enabled while also verifying the MAC notification interval setting.</p> <pre>user@switch> show ethernet-switching mac-notification Notification Status: Enabled Notification Interval: 30</pre>
Meaning	<p>The output in the Notification Status field shows that MAC notification is enabled. The output in the Notification Status field would display Disabled if MAC notification was disabled.</p> <p>The Notification Interval field output shows that the MAC notification interval is set to 30 seconds.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring MAC Notification (CLI Procedure) on page 175 • Configuring MAC Notification (CLI Procedure)

Verifying That Proxy ARP Is Working Correctly

Purpose	Verify that the switch is sending proxy ARP messages.
Action	<p>List the system statistics for ARP:</p> <pre>user@switch> show system statistics arp arp: 90060 datagrams received 34 ARP requests received 610 ARP replies received 2 resolution request received 0 unrestricted proxy requests 0 restricted proxy requests 0 received proxy requests 0 unrestricted proxy requests not proxied 0 restricted proxy requests not proxied 0 datagrams with bogus interface 0 datagrams with incorrect length 0 datagrams for non-IP protocol 0 datagrams with unsupported op code 0 datagrams with bad protocol address length 0 datagrams with bad hardware address length 0 datagrams with multicast source address 0 datagrams with multicast target address 0 datagrams with my own hardware address 0 datagrams for an address not on the interface 0 datagrams with a broadcast source address 294 datagrams with source address duplicate to mine 89113 datagrams which were not for me 0 packets discarded waiting for resolution 0 packets sent after waiting for resolution</pre>

```

309 ARP requests sent
35 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
0 replies for from unnumbered interfaces
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor

```

Meaning The statistics show that two proxy ARP requests were received. The **unrestricted proxy requests not proxied** and **restricted proxy requests not proxied** fields indicate that all the unproxied ARP requests received have been proxied by the switch.

Related Documentation

- *Configuring Proxy ARP*
- *Configuring Proxy ARP (CLI Procedure)*

Monitoring Ethernet Switching

Purpose



NOTE: This topic applies only to the J-Web Application package.

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 25 on page 300](#) summarizes the Ethernet switching output fields.

Table 25: Ethernet Switching Output Fields

Field	Value
Ethernet Switching Table Information or MAC Table Summary	
MAC Table Count	The number of entries added to the Ethernet switching table.

Table 25: Ethernet Switching Output Fields (*continued*)

Field	Value
MAC Table Learned	The number of dynamically learned MAC addresses in the Ethernet switching table.
Ethernet Switching Table Information or MAC Table Information	
VLAN	The VLAN name.
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 NOTE: This option is not supported on EX4300 switches.	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.
MAC Flag NOTE: This option is supported only on EX4300 switches.	Status of MAC address learning properties for each interface. Values are: <ul style="list-style-type: none"> • S - Static MAC address is configured • D - Dynamic MAC address is configured • L - Locally learned MAC address is configured • P - Persistent static • SE - MAC accounting is enabled • NM - Nonconfigured MAC • R - Remotely learned MAC address is configured
Age	The time remaining before the entry ages out and is removed from the Ethernet switching table.
Interfaces	The associated interfaces.
MAC Learning Log or Interface Information	
Interface Name NOTE: This option is supported only on EX4300 switches.	The name of the Interface.
VLAN-Name	The VLAN name.
VLAN ID NOTE: This option is supported only on EX4300 switches.	The VLAN ID.

Table 25: Ethernet Switching Output Fields (*continued*)

Field	Value
MAC Address NOTE: This option is not supported on EX4300 switches.	The learned MAC address associated with the VLAN ID.
MAC Limit NOTE: This option is supported only on EX4300 switches.	Maximum number of MAC addresses.
STP State NOTE: This option is supported only on EX4300 switches.	The state of the STP interface. Values are Discarding or Receiving.
Flags NOTE: This option is supported only on EX4300 switches.	Specifies the logical interface flags. The flag options are: <ul style="list-style-type: none"> • DL - disable learning • AD - packet action drop • LH - MAC limit hit • DN - interface down
Tagging NOTE: This option is supported only on EX4300 switches.	States whether the VLAN associated with interface is tagged or untagged.
Time NOTE: This option is not supported on EX4300 switches.	Time at which the MAC address was added or deleted from the MAC learning log.
State NOTE: This option is not supported on EX4300 switches.	Operating state of the interface. Values are Up or Down.

- Related Documentation**
- [Configuring MAC Table Aging \(CLI Procedure\) on page 166](#)
 - [Understanding Bridging and VLANs on EX Series Switches on page 3](#)

CHAPTER 10

Operational Commands

- clear edge-virtual-bridging
- clear ethernet-switching layer2-protocol-tunneling error
- clear ethernet-switching layer2-protocol-tunneling statistics
- clear ethernet-switching table
- clear mvrp statistics
- show edge-virtual-bridging
- show ethernet-switching interfaces
- show ethernet-switching layer2-protocol-tunneling interface
- show ethernet-switching layer2-protocol-tunneling statistics
- show ethernet-switching layer2-protocol-tunneling vlan
- show ethernet-switching mac-learning-log
- show ethernet-switching mac-notification
- show ethernet-switching statistics aging
- show ethernet-switching statistics mac-learning
- show ethernet-switching table
- show mvrp
- show mvrp dynamic-vlan-memberships
- show mvrp statistics
- show protection-group ethernet-ring aps
- show protection-group ethernet-ring configuration
- show protection-group ethernet-ring interface
- show protection-group ethernet-ring node-state
- show protection-group ethernet-ring statistics
- show redundant-trunk-group
- show system statistics arp
- show vlans

clear edge-virtual-bridging

Syntax	<code>clear edge-virtual-bridging</code> <code><edge-control-protocol-statistics></code> <code><firewall <interface <i>interface-name</i>></code> <code><vsi-profiles <interface <i>interface-name</i>></code>
Release Information	Statement introduced in Junos OS Release 12.1 for EX Series switches.
Description	Clear edge-virtual-bridging (EVB).
Options	none —Clear EVB. edge-control-protocol-statistics —(Optional) Clear Edge Control Protocol (ECP) statistics. firewall <interface <i>interface-name</i>> —(Optional) Clear EVB implicit filter counters on all interfaces or on a specific interface.
Required Privilege Level	system —To view this statement in the configuration. system-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Edge Virtual Bridging for Use with VEPA Technology on page 134• Configuring Edge Virtual Bridging (CLI Procedure) on page 188

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 OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
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 shutdown threshold 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 Documentation	<ul style="list-style-type: none"> • Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 107 • Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) on page 181 • Configuring Layer 2 Protocol Tunneling
List of Sample Output	clear ethernet-switching layer2-protocol-tunneling error on page 305 clear ethernet-switching layer2-protocol-tunneling error interface xe-0/0/1.0 on page 305

Sample Output

clear ethernet-switching layer2-protocol-tunneling error

```
user@switch> clear ethernet-switching layer2-protocol-tunneling error
```

clear ethernet-switching layer2-protocol-tunneling error interface xe-0/0/1.0

```
user@switch> clear ethernet-switching layer2-protocol-tunneling error interface xe-0/0/1.0
```

clear ethernet-switching layer2-protocol-tunneling statistics

Syntax	<code>clear ethernet-switching layer2-protocol-tunneling statistics</code> <code><interface <i>interface-name</i>></code> <code><vlan <i>vlan-name</i>></code>
Release Information	Command introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
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 Documentation	<ul style="list-style-type: none">• show ethernet-switching layer2-protocol-tunneling statistics on page 319• Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 107• Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) on page 181• Configuring Layer 2 Protocol Tunneling
List of Sample Output	clear ethernet-switching layer2-protocol-tunneling statistics on page 306 clear ethernet-switching layer2-protocol-tunneling error interface ge-0/1/1.0 on page 306 clear ethernet-switching layer2-protocol-tunneling error vlan v2 on page 306

Sample Output

`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 xe-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 ethernet-switching table

Syntax	clear ethernet-switching table <interface <i>interface-name</i> > <mac <i>mac-address</i> > <management-vlan> <persistent-mac < <i>interface</i> <i>mac-address</i> >> <vlan <i>vlan-name</i> >
Syntax (QFX Series)	clear ethernet-switching table <interface <i>interface-name</i> > <mac <i>mac-address</i> > <persistent-mac < <i>interface</i> <i>mac-address</i> >> <vlan <i>vlan-name</i> >
Release Information	Command introduced in Junos OS Release 9.3 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	<div>  <p>NOTE: On a QFabric system, using this command on an FCoE-enabled VLAN when FCoE sessions are active can cause traffic flooding and FCoE traffic drop. The FCoE sessions are not terminated and the traffic reconverges after a short period of time.</p> </div> <p>Clear learned entries, which are media access control (MAC) addresses, in the Ethernet switching table (also called the forwarding database table).</p>
Options	<p>none—Clear learned entries in the Ethernet switching table, except for persistent MAC addresses.</p> <p>interface <i>interface-name</i>—(Optional) Clear all learned MAC addresses for the specified interface from the Ethernet switching table.</p> <p>mac <i>mac-address</i>—(Optional) Clear the specified learned MAC address from the Ethernet switching table.</p> <p>management-vlan—(Optional) Clear all MAC addresses learned for the management VLAN from the Ethernet switching table. Note that you do not specify a VLAN name because only one management VLAN exists.</p> <p>persistent-mac <<i>interface</i> <i>mac-address</i>>—(Optional) Clear all MAC addresses, including persistent MAC addresses. Use the interface option to clear all MAC addresses on an interface, or use the mac-address option to clear all entries for a specific MAC address.</p> <p>Use this command whenever you move a device in your network that has a persistent MAC address on the switch. If you move the device to another port on the switch and do not clear the persistent MAC address from the original port it was learned on, then the new port will not learn the MAC address and the device will not be able to connect. If the original port is down when you move the device, then the new port</p>

will learn the MAC address and the device can connect—however, unless you cleared the MAC address on the original port, when the port comes back up, the system reinstalls the persistent MAC address in the forwarding table for that port. If this occurs, the address is removed from the new port and the device loses connectivity.

vlan *vlan-name*—(Optional) Clear all MAC addresses learned for the specified VLAN from the Ethernet switching table.

Required Privilege Level

view

Related Documentation

- [show ethernet-switching table on page 334](#)
- *show ethernet-switching table*
- *Verifying That Persistent MAC Learning Is Working Correctly*

List of Sample Output

[clear ethernet-switching table on page 308](#)

Output Fields

This command produces no output.

Sample Output

[clear ethernet-switching table](#)

```
user@switch> clear ethernet-switching table
```


clear mvrp statistics

Syntax	<code>clear mvrp statistics <interface <i>interface-name</i>></code>
Release Information	Command introduced in Junos OS 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 Documentation	<ul style="list-style-type: none"> • show mvrp statistics on page 342 • Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 92
List of Sample Output	<p>clear mvrp statistics on page 309</p> <p>clear mvrp statistics interface ge-0/0/1.0 on page 309</p>
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

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 edge-virtual-bridging

Syntax	<pre>show edge-virtual-bridging <detail> <edge-control-protocol statistics <interface interface-name>> <firewall> <interface interface-name> vsi-profiles <interface interface-name></pre>
Release Information	Command introduced in Junos OS Release 12.1 for EX Series switches.
Description	Display information about edge virtual bridging (EVB).
Options	<p>none—Display EVB parameters for all interfaces configured with EVB.</p> <p>detail—(Optional) Display EVB parameters and virtual station interface (VSI) profiles associated with each interface.</p> <p>edge-control-protocol statistics <interface <interface-name>>—(Optional) Display Edge Control Protocol (ECP) statistics for all configured EVB interfaces or for the specified interface.</p> <p>firewall—Display the firewall filters created by EVB.</p> <p>interface <interface-name>—(Optional) Display EVB parameters for the specified interface.</p> <p>vsi-profiles <interface interface-name>—(Optional) Display VSI profiles associated on each interface or for the specified interface.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Edge Virtual Bridging for Use with VEPA Technology on page 134 • Example: Configuring Edge Virtual Bridging for Use with VEPA Technology on page 134
List of Sample Output	show edge-virtual-bridging on page 311 show edge-virtual-bridging interface on page 311 show edge-virtual-bridging edge-control-protocol statistics on page 311 show edge-virtual-bridging vsi-profiles on page 312 show edge-virtual-bridging vsi-profiles interface on page 312 show edge-virtual-bridging firewall on page 312
Output Fields	Table 26 on page 310 lists the output fields for the show edge-virtual-bridging command. Output fields are listed in the approximate order in which they appear.

Table 26: show edge-virtual-bridging Output Field Descriptions

Field Name	Field Description
Interface	Switch interface configured for EVB.

Table 26: show edge-virtual-bridging Output Field Descriptions (*continued*)

Field Name	Field Description
Interface input ECP Packets	Number of ECP packets received by the switch. ECP is a Layer 2 protocol that is used to carry VSI Discovery and Configuration Protocol (VDP) messages.
Interface output ECP Packets	Number of ECP packets sent by the switch. ECP is a Layer 2 protocol that is used to carry VDP messages.
Forwarding Mode	Mode by which packets are forwarded to their destination. The value for forwarding mode is either Standard (meaning the forwarding is done through 802.1Q) or Reflective-relay , meaning that both the source and destination addresses are located on the same VM server.
RTE	Retransmission timer exponent (RTE) is an EVB interface attribute used to calculate the minimum VDP protocol data unit (PDU) retransmission time.
Number of VSIs	Number of virtual station interfaces on the switch connected to the VEPA.
Protocols	EVB protocols currently enabled. The values can be VDP , ECP or RTE . Protocols are configured during the capabilities exchange via an EVB type, length, and value (TLV) carried by the Link Layer Discovery Protocol (LLDP) between the switch and the server.
VSI profile	EVB profile including parameters that uniquely identify each VSI entry (VSI manager, VSI type, VSI version, VSI instance, VSI state).
Filter Name	Name of the filter defined in the firewall stanza.
Counters	Number of packets and bytes that have satisfied the match conditions defined by the filter.

Sample Output

show edge-virtual-bridging

```

user@switch#show edge-virtual-bridging
Interface      Forwarding Mode  RTE  Number of VSIs  Protocols
ge-0/0/20.0    Reflective-relay  25   400              ECP, VDP, RTE

```

show edge-virtual-bridging interface

```

user@switch#show edge-virtual-bridging interface ge-0/0/20.0
Interface: ge-0/0/20.0, Forwarding mode: Reflective-relay RTE: 25, Number of VSIs:
400, Protocols: ECP, VDP, RTE
VSI profiles:
Manager: 97, Type: 997, Version: 3, VSI State: Associate
Instance: 09b11c53-8b5c-4eeb-8f00-c84ebb0bb997
          MAC                      VLAN
          00:10:94:00:00:04

```

show edge-virtual-bridging edge-control-protocol statistics

```

user@switch#show edge-virtual-bridging edge-control-protocol-statistics
Interface: ge-0/0/20.0
Input ECP packets: 302

```

Output ECP packets: 303

show edge-virtual-bridging vsi-profiles

```
user@switch#show edge-virtual-bridging vsi-profiles
Interface: ge-0/0/20.0
Manager: 97, Type: 997, Version: 3, VSI State: Associate
Instance: 09b11c53-8b5c-4eeb-8f00-c84ebb0bb997
      MAC                               VLAN
      00:10:94:00:00:04                 3
```

show edge-virtual-bridging vsi-profiles interface

```
user@switch#show edge-virtual-bridging vsi-profiles interface ge-0/0/20.0
Interface: ge-0/0/20.0
Manager: 97, Type: 997, Version: 3, VSI State: Associate
Instance: 09b11c53-8b5c-4eeb-8f00-c84ebb0bb997
      MAC
      VLAN    00:10:94:00:00:04        3
```

show edge-virtual-bridging firewall

```
user@switch#show edge-virtual-bridging firewall
Filter name: evb_filter_ge-0/0/20
Counters:
  Name: evb_filter_term_3_00:10:94:00:00:04_default
    Bytes: 0, Packets: 0
  Name: f3_accept__evb_filter_term_3_00:10:94:00:00:04-f3-t1
    Bytes: 1028, Packets: 14
```

show ethernet-switching interfaces

Syntax	<pre>show ethernet-switching interfaces <brief detail summary> <interface <i>interface-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>In Junos OS Release 9.6 for EX Series switches, the following updates were made:</p> <ul style="list-style-type: none"> • Blocking field output was updated. • The default view was updated to include information about 802.1Q tags. • The detail view was updated to include information on VLAN mapping. <p>In Junos OS Release 11.1 for EX Series switches, the detail view was updated to include reflective relay information.</p>
Description	Display information about Ethernet switching interfaces.
Options	<p>none—Display brief information for Ethernet switching interfaces.</p> <p>brief detail summary—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display Ethernet switching information for a specific interface.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ethernet-switching mac-learning-log on page 324 • show ethernet-switching table on page 334 • <i>Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)</i>
List of Sample Output	<p>show ethernet-switching interfaces on page 315</p> <p>show ethernet-switching interfaces ge-0/0/15 brief on page 315</p> <p>show ethernet-switching interfaces ge-0/0/2 detail (Blocked by RTG rtggroup) on page 315</p> <p>show ethernet-switching interfaces ge-0/0/15 detail (Blocked by STP) on page 316</p> <p>show ethernet-switching interfaces ge-0/0/17 detail (Disabled by bpdu-control) on page 316</p> <p>show ethernet-switching interfaces detail (C-VLAN to S-VLAN Mapping) on page 316</p> <p>show ethernet-switching interfaces detail (Reflective Relay Is Configured) on page 316</p>
Output Fields	Table 27 on page 314 lists the output fields for the show ethernet-switching interfaces command. Output fields are listed in the approximate order in which they appear.

Table 27: show ethernet-switching interfaces Output Fields

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

Table 27: show ethernet-switching interfaces Output Fields (*continued*)

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

Sample Output

show ethernet-switching interfaces

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

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

show ethernet-switching interfaces ge-0/0/15 brief

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

Interface	State	VLAN members	Tag	Tagging	Blocking
ge-0/0/15.0	up	vlan100	100	tagged	blocked by STP
		vlan200	200	tagged	blocked by STP

show ethernet-switching interfaces ge-0/0/2 detail (Blocked by RTG rtggroup)

```
user@switch> show ethernet-switching interfaces ge-0/0/2 detail
```

```
Interface: ge-0/0/2.0, Index: 65, State: up, Port mode: Access
Ether type for the interface: 0X8100
VLAN membership:
    vlan300, 802.1Q Tag: 300, untagged, msti-id: 0, blocked by RTG(rtggroup)
Number of MACs learned on IFL: 0
```

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

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

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

Number of MACs learned on IFL: 0
```

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

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

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

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

```
user@switch> show ethernet-switching interfaces ge-0/0/6.0 detail
Interface: ge-0/0/6.0, Index: 73, State: up, Port mode: Access
Ether type for the interface: 0X8100
VLAN membership:
    map, 802.1Q Tag: 134, Mapped Tag: native, push, dot1q-tunneled, unblocked
    map, 802.1Q Tag: 134, Mapped Tag: 20, push, dot1q-tunneled, unblocked
```

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

```
user@switch1> show ethernet-switching interfaces ge-7/0/2 detail
Interface: ge-7/0/2, Index: 66, State: down, Port mode: Tagged-access
Ether type for the interface: 0X8100
Reflective Relay Status: Enabled
Ether type for the interface: 0x8100
VLAN membership:
    VLAN_Purple VLAN_Orange VLAN_Blue, 802.1Q Tag: 450, tagged, unblocked
Number of MACs learned on IFL: 0
```


show ethernet-switching layer2-protocol-tunneling interface

Syntax	<code>show ethernet-switching-layer2-protocol-tunneling interface</code> <code><interface-name></code>
Release Information	Command introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
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. interface-name —(Optional) Display L2PT information for the specified interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ethernet-switching layer2-protocol-tunneling statistics on page 319 • show ethernet-switching layer2-protocol-tunneling vlan on page 322 • Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) on page 181 • show ethernet-switching layer2-protocol-tunneling statistics on page 319 • show ethernet-switching layer2-protocol-tunneling vlan on page 322 • Configuring Layer 2 Protocol Tunneling
List of Sample Output	show ethernet-switching layer2-protocol-tunneling interface on page 318 show ethernet-switching layer2-protocol-tunneling interface xe-0/0/0.0 on page 318
Output Fields	Table 28 on page 317 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 28: 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.

Sample Output

show ethernet-switching layer2-protocol-tunneling interface

```
user@switch> show ethernet-switching layer2-protocol-tunneling interface
```

```
Layer2 Protocol Tunneling information:
```

Interface	Operation	State	Description
xe-0/0/0.0	Encapsulation	Shutdown	Shutdown threshold exceeded
xe-0/0/1.0	Decapsulation	Shutdown	Loop detected
xe-0/0/2.0	Decapsulation	Active	

show ethernet-switching layer2-protocol-tunneling interface xe-0/0/0.0

```
user@switch> show ethernet-switching layer2-protocol-tunneling interface xe-0/0/0.0
```

```
Layer2 Protocol Tunneling information:
```

Interface	Operation	State	Description
xe-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 <i>interface-name</i> > <vlan <i>vlan-name</i> >
Release Information	Command introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display Layer 2 protocol tunneling (L2PT) statistics for Layer 2 PDU packets received by the switch.
<div>  <p>NOTE: The <code>show ethernet-switching-layer2-protocol-tunneling statistics</code> command does not display L2PT statistics for Layer 2 PDU packets transmitted from the switch.</p> </div>	
Options	<p>none—Display L2PT statistics for all interfaces on which you enabled L2PT.</p> <p>interface <i>interface-name</i>—(Optional) Display L2PT statistics for the specified interface.</p> <p>vlan <i>vlan-name</i>—(Optional) Display L2PT statistics for the specified VLAN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear ethernet-switching layer2-protocol-tunneling statistics on page 306 • show ethernet-switching layer2-protocol-tunneling interface on page 317 • show ethernet-switching layer2-protocol-tunneling vlan on page 322 • show vlans on page 362 • Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 107 • Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) on page 181 • show vlans • Configuring Layer 2 Protocol Tunneling
List of Sample Output	show ethernet-switching layer2-protocol-tunneling statistics on page 320 show ethernet-switching layer2-protocol-tunneling statistics interface xe-0/0/0.0 on page 320 show ethernet-switching layer2-protocol-tunneling statistics vlan v2 on page 320
Output Fields	Table 29 on page 320 lists the output fields for the <code>show ethernet-switching layer2-protocol-tunneling statistics</code> command. Output fields are listed in the approximate order in which they appear.

Table 29: show ethernet-switching layer2-protocol-tunneling statistics Output Fields

VLAN	Field Description
VLAN	Name of a VLAN on which L2PT has been configured.
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 , 802.1x , 802.3ah , cdp , e-lmi , gvrp , lACP , lldp , mmrp , mvrp , stp , udld , 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 de-encapsulated.
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.

Sample Output

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    xe-0/0/0.0  mvrp     Encapsulation  0        0      0
v1    xe-0/0/1.0  mvrp     Decapsulation  0        0      0
v1    xe-0/0/2.0  mvrp     Decapsulation  60634    0      0
v2    xe-0/0/0.0  cdp      Encapsulation  0        0      0
v2    xe-0/0/0.0  gvrp     Encapsulation  0        0      0
v2    xe-0/0/0.0  lldp     Encapsulation  0        0      0
```

show ethernet-switching layer2-protocol-tunneling statistics interface xe-0/0/0.0

```
user@switch> show ethernet-switching layer2-protocol-tunneling statistics interface xe-0/0/0.0
```

```
Layer2 Protocol Tunneling Statistics:
VLAN  Interface  Protocol  Operation  Packets  Drops  Shutdowns
v1    xe-0/0/0.0  mvrp     Encapsulation  0        0      0
v2    xe-0/0/0.0  cdp      Encapsulation  0        0      0
v2    xe-0/0/0.0  gvrp     Encapsulation  0        0      0
v2    xe-0/0/0.0  lldp     Encapsulation  0        0      0
v2    xe-0/0/0.0  mvrp     Encapsulation  0        0      0
v2    xe-0/0/0.0  stp      Encapsulation  0        0      0
v2    xe-0/0/0.0  vtp      Encapsulation  0        0      0
v2    xe-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    xe-0/0/0.0  cdp      Encapsulation  0        0      0
v2    xe-0/0/0.0  gvrp     Encapsulation  0        0      0
```

v2	xe-0/0/0.0	lldp	Encapsulation	0	0	0
v2	xe-0/0/0.0	mvrp	Encapsulation	0	0	0
v2	xe-0/0/0.0	stp	Encapsulation	0	0	0
v2	xe-0/0/0.0	vtp	Encapsulation	0	0	0
v2	xe-0/0/0.0	vstp	Encapsulation	0	0	0
v2	xe-0/0/1.0	cdp	Decapsulation	0	0	0
v2	xe-0/0/1.0	gvrp	Decapsulation	0	0	0
v2	xe-0/0/1.0	lldp	Decapsulation	0	0	0
v2	xe-0/0/1.0	mvrp	Decapsulation	0	0	0
v2	xe-0/0/1.0	stp	Decapsulation	0	0	0
v2	xe-0/0/1.0	vtp	Decapsulation	0	0	0

show ethernet-switching layer2-protocol-tunneling vlan

Syntax	<code>show ethernet-switching-layer2-protocol-tunneling vlan <vlan-name></code>
Release Information	Command introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display information about Layer 2 protocol tunneling (L2PT) on VLANs that have been configured for L2PT.
Options	none —Display information about L2PT for the VLANs on which you have configured L2PT. vlan-name —(Optional) Display information about L2PT for the specified VLAN.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ethernet-switching layer2-protocol-tunneling interface on page 317 • show ethernet-switching layer2-protocol-tunneling statistics on page 319 • show vlans on page 362 • Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 107 • Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) on page 181 • show vlans • Configuring Layer 2 Protocol Tunneling
List of Sample Output	show ethernet-switching layer2-protocol-tunneling vlan on page 323 show ethernet-switching layer2-protocol-tunneling vlan v2 on page 323
Output Fields	Table 30 on page 322 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 30: 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 , 802.1x , 802.3ah , cdp , e-lmi , gvrp , lacp , lldp , mmrp , 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.

Sample Output

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 OS Release 9.0 for EX Series switches.
Description	Displays the event log of learned MAC addresses.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ethernet-switching table on page 334 • show ethernet-switching interfaces on page 313 • Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 57 • Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 65 • Example: Connecting an Access Switch to a Distribution Switch on page 72
List of Sample Output	show ethernet-switching mac-learning-log on page 324
Output Fields	Table 31 on page 324 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 31: 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	<p>The forwarding state of the interface:</p> <ul style="list-style-type: none"> • blocked—Traffic is not being forwarded on the interface. • unblocked—Traffic is forwarded on the interface.

Sample Output

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 mac-notification

Syntax	show ethernet-switching mac-notification
Release Information	Command introduced in Junos OS Release 9.6 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display information about MAC notification.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Verifying That MAC Notification Is Working Properly on page 299
List of Sample Output	show ethernet-switching mac-notification (MAC Notification Enabled) on page 326 show ethernet-switching mac-notification (MAC Notification Disabled) on page 326
Output Fields	Table 32 on page 326 lists the output fields for the show ethernet-switching mac-notification command. Output fields are listed in the order in which they appear.

Table 32: show ethernet-switching mac-notification Output Fields

Field Name	Field Description
Notification Status	MAC notification status: <ul style="list-style-type: none"> • Enabled—MAC notification is enabled. • Disabled—MAC notification is disabled.
Notification Interval	MAC notification interval in seconds.
Notifications Sent	Number of notifications sent to SNMP when MACs are learned or when MACs age out.
Notifications Table Maxsize	Maximum size of the notification table, which is populated when notifications are sent to the SNMP server.

Sample Output

show ethernet-switching mac-notification (MAC Notification Enabled)

```

user@switch> show ethernet-switching mac-notification
Notification Status           : Enabled
Notification Interval         : 30
Notifications Sent            : 0
Notifications Table Maxsize   : 256

```

Sample Output

show ethernet-switching mac-notification (MAC Notification Disabled)

```

user@switch> show ethernet-switching mac-notification
Notification Status           : Disabled
Notification Interval         : 0

```

Notifications Sent : 0
Notifications Table Maxsize : 256

show ethernet-switching statistics aging

Syntax	show ethernet-switching statistics aging
Release Information	Command introduced in Junos OS 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 Documentation	<ul style="list-style-type: none"> • show ethernet-switching statistics mac-learning on page 330 • Configuring MAC Table Aging (CLI Procedure) on page 166
List of Sample Output	show ethernet-switching statistics aging on page 328
Output Fields	Table 33 on page 328 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 33: 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

Sample Output

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	<code>show ethernet-switching statistics mac-learning</code> <code><brief detail></code> <code><interface <i>interface-name</i>></code>
Release Information	Command introduced in Junos OS Release 9.4 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display media access control (MAC) learning statistics.
Options	none —(Optional) Display MAC learning statistics for all interfaces. brief detail —(Optional) Display the specified level of output. The default is brief . interface <i>interface-name</i> —(Optional) Display MAC learning statistics for the specified interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show ethernet-switching statistics aging on page 328• show ethernet-switching mac-learning-log on page 324• show ethernet-switching table on page 334• show ethernet-switching interfaces on page 313• Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 57• Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 65• 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 on the QFX Series• Example: Setting Up Bridging with Multiple VLANs
List of Sample Output	show ethernet-switching statistics mac-learning on page 331 show ethernet-switching statistics mac-learning detail on page 332 show ethernet-switching statistics mac-learning interface ge-0/0/28 detail on page 332 show ethernet-switching statistics mac-learning interface on page 332 show ethernet-switching statistics mac-learning detail (QFX Series) on page 332
Output Fields	Table 34 on page 331 lists the output fields for the show ethernet-switching statistics mac-learning command. Output fields are listed in the approximate order in which they appear.

Table 34: 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. (Displayed in the output under the heading Interface .)	All levels
Learning message from local packets	MAC learning message generated due to packets coming in on the management interface. (Displayed in the output under the heading Local pkts .)	All levels
Learning message from transit packets	MAC learning message generated due to packets coming in on network interfaces. (Displayed in the output under the heading Transit pkts .)	All levels
Learning message with error	<p>MAC learning messages received with errors (Displayed under the heading Error):</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. • VLAN membership limit—The number of MAC addresses learned on the interface as a member of the specified VLAN (VLAN membership MAC limit) has exceeded the limit. • Invalid VLAN index—The VLAN of the packet, although 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

Sample Output

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 statistics mac-learning interface ge-0/0/28 detail

```
user@switch> show ethernet-switching statistics mac-learning interface ge-0/0/28 detail
```

```
Interface: ge-0/0/28.0
Learning message from local packets: 0
Learning message from transit packets: 5
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
                          VLAN membership limit: 20
  Invalid VLAN index: 0   Interface not learning: 0
  No nexthop: 0          MAC learning disabled: 0
  Others: 0
```

show ethernet-switching statistics mac-learning interface

```
user@switch> show ethernet-switching statistics mac-learning interface ge-0/0/1
Interface      Local pkts  Transit pkts  Error
ge-0/0/1.0      0           1             1
```

show ethernet-switching statistics mac-learning detail (QFX Series)

```
user@switch> show ethernet-switching statistics mac-learning detail
Learning stats: 0 learn msg rcvd, 0 error
```

```
Interface: xe-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: xe-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 table

Syntax	<pre>show ethernet-switching table <brief detail extensive summary> <interface <i>interface-name</i>> <management-vlan> <persistent-mac <interface <i>interface-name</i>>> <sort-by (<i>name</i> <i>tag</i>)> <vlan <i>vlan-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Options summary, management-vlan, and vlan <i>vlan-name</i> introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Option sort-by and field name tag introduced in Junos OS Release 10.1 for EX Series switches.</p> <p>Option persistent-mac introduced in Junos OS Release 11.4 for EX Series switches.</p>
Description	<p> NOTE: If your EX Series switch CLI displays different options for the show ethernet-switching table command than the options shown in this document, see <i>show ethernet-switching table</i>.</p> <p>Display the Ethernet switching table.</p>
Options	<p>none—(Optional) Display brief information about the Ethernet switching table.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display the Ethernet switching table for a specific interface.</p> <p>management-vlan—(Optional) Display the Ethernet switching table for a management VLAN.</p> <p>persistent-mac <interface <i>interface-name</i>>—(Optional) Display the persistent MAC addresses learned for all interfaces or a specified interface. You can use this command to view entries that you want to clear for an interface that you intentionally disabled.</p> <p>sort-by (<i>name</i> <i>tag</i>)—(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.</p> <p>vlan <i>vlan-name</i>—(Optional) Display the Ethernet switching table for a specific VLAN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear ethernet-switching table on page 307 • Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 57

- [Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 65](#)
- [Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 103](#)

List of Sample Output [show ethernet-switching table on page 336](#)
[show ethernet-switching table brief on page 336](#)
[show ethernet-switching table detail on page 337](#)
[show ethernet-switching table extensive on page 337](#)
[show ethernet-switching table persistent-mac on page 338](#)
[show ethernet-switching table persistent-mac interface ge-0/0/16.0 on page 338](#)

Output Fields [Table 35 on page 335](#) lists the output fields for the **show ethernet-switching table** command. Output fields are listed in the approximate order in which they appear.

Table 35: show ethernet-switching table Output Fields

Field Name	Field Description	Level of Output
VLAN	The name of a VLAN.	All levels
Tag	The VLAN ID tag name or number.	extensive
MAC or MAC address	The MAC address associated with the VLAN.	All levels
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. • persistent—The learned MAC addresses that will persist across restarts of the switch or interface-down events. 	All levels except persistent-mac
Type	The type of MAC address. Values are: <ul style="list-style-type: none"> • installed—addresses that are in the Ethernet switching table. • uninstalled—addresses that could not be installed in the table or were uninstalled in an interface-down event and will be reinstalled in the table when the interface comes back up. 	persistent-mac
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
Nexthop index	The next-hop index number.	detail, extensive

Table 35: show ethernet-switching table Output Fields (*continued*)

Field Name	Field Description	Level of Output
persistent-mac	installed indicates MAC addresses that are in the Ethernet switching table and uninstalled indicates MAC addresses that could not be installed in the table or were uninstalled in an interface-down event (and will be reinstalled in the table when the interface comes back up).	

Sample Output

show ethernet-switching table

```

user@switch> show ethernet-switching table
Ethernet-switching table: 57 entries, 15 learned, 2 persistent
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 Persistent 0 ge-0/0/46.0
T1         00:00:5e:00:01:00 Static    - Router
T1         00:19:e2:50:63:e0 Persistent 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, 15 learned, 2 persistent entries
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 Persistent 0 ge-0/0/46.0
T1          00:00:5e:00:01:00 Static      - Router
T1          00:19:e2:50:63:e0 Persistent 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: 5 entries, 2 learned entries
VLAN: default, Tag: 0, MAC: *, Interface: All-members
  Interfaces:
    ge-0/0/11.0, ge-0/0/20.0, ge-0/0/30.0, ge-0/0/36.0, ge-0/0/3.0
  Type: Flood
  Nexthop index: 1307

VLAN: default, Tag: 0, MAC: 00:1f:12:30:b8:83, Interface: ge-0/0/3.0
  Type: Learn, Age: 0, Learned: 20:09:26
  Nexthop index: 1315

VLAN: v1, Tag: 101, MAC: *, Interface: All-members
  Interfaces:
    ge-0/0/31.0
  Type: Flood
  Nexthop index: 1313

VLAN: v1, Tag: 101, MAC: 00:1f:12:30:b8:89, Interface: ge-0/0/31.0
  Type: Learn, Age: 0, Learned: 20:09:25
  Nexthop index: 1312

VLAN: v2, Tag: 102, MAC: *, Interface: All-members
  Interfaces:
    ae0.0
  Type: Flood
  Nexthop index: 1317

```

show ethernet-switching table extensive

```

user@switch> show ethernet-switching table extensive
Ethernet-switching table: 3 entries, 1 learned, 5 persistent entries

VLAN: v1, Tag: 10, MAC: *, Interface: All-members

```

Interfaces:
ge-0/0/14.0, 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-0/0/10.0,
ge-0/0/0.0
Type: Flood
Nexthop index: 567

VLAN: v1, Tag: 10, MAC: 00:21:59:c6:93:22, Interface: Router
Type: Static
Nexthop index: 0

VLAN: v1, Tag: 10, MAC: 00:21:59:c9:9a:4e, Interface: ge-0/0/14.0
Type: Learn, Age: 0, Learned: 18:40:50
Nexthop index: 564

show ethernet-switching table persistent-mac

```
user@switch> show ethernet-switching table persistent-mac
```

VLAN	MAC address	Type	Interface
default	00:10:94:00:00:02	installed	ge-0/0/42.0
default	00:10:94:00:00:03	installed	ge-0/0/42.0
default	00:10:94:00:00:04	installed	ge-0/0/42.0
default	00:10:94:00:00:05	installed	ge-0/0/42.0
default	00:10:94:00:00:06	installed	ge-0/0/42.0
default	00:10:94:00:05:02	uninstalled	ge-0/0/16.0
default	00:10:94:00:06:03	uninstalled	ge-0/0/16.0
default	00:10:94:00:07:04	uninstalled	ge-0/0/16.0

show ethernet-switching table persistent-mac interface ge-0/0/16.0

VLAN	MAC address	Type	Interface
default	00:10:94:00:05:02	uninstalled	ge-0/0/16.0
default	00:10:94:00:06:03	uninstalled	ge-0/0/16.0
default	00:10:94:00:07:04	uninstalled	ge-0/0/16.0

show mvrp

Syntax	show mvrp
Release Information	Command introduced in Junos OS Release 10.0 for EX Series switches.
Description	Display Multiple VLAN Registration Protocol (MVRP) configuration information.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show mvrp statistics on page 342 • Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 92 • Verifying That MVRP Is Working Correctly on page 297
List of Sample Output	show mvrp on page 340
Output Fields	Table 36 on page 339 lists the output fields for the show mvrp command. Output fields are listed in the approximate order in which they appear.

Table 36: show mvrp Output Fields

Field Name	Field Description
Global MVRP configuration	<p>Displays global MVRP information:</p> <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 Enabled or Disabled.
MVRP Timers (ms)	<p>Displays MVRP timer information:</p> <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	<p>Displays interface-specific MVRP information:</p> <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.

Sample Output

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 OS 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 Documentation	<ul style="list-style-type: none"> • show mvrp on page 339 • show mvrp statistics on page 342 • Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 92 • Verifying That MVRP Is Working Correctly on page 297
List of Sample Output	show mvrp dynamic-vlan-memberships on page 341
Output Fields	Table 37 on page 341 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 37: 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.

Sample Output

show mvrp dynamic-vlan-memberships

```

user@switch> 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

```

show mvrp statistics

Syntax	show mvrp statistics <interface <i>interface-name</i> >
Release Information	Command introduced in Junos OS 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>—(Optional) Show MVRP statistics for the specified interface.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show mvrp on page 339 • clear mvrp statistics on page 309 • Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 92 • Verifying That MVRP Is Working Correctly on page 297
List of Sample Output	show mvrp statistics interface xe-0/1/1.0 on page 343
Output Fields	Table 38 on page 342 lists the output fields for the show mvrp statistics command. Output fields are listed in the approximate order in which they appear.

Table 38: 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 JoinEmpty messages received on the switch. Either this value or the value for <i>JoinIn received</i> should increase when the value for <i>MRPDU received</i> increases. If this value is not incrementing when it should, you might have a Junos OS release version compatibility issue. To fix a version compatibility issue, see “Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)” on page 176.
Join In received	Number of MRP JoinIn messages received on the switch. Either this value or the value for <i>JoinEmpty received</i> should increase when the value for <i>MRPDU received</i> increases. If this value is not incrementing when it should, you might have a Junos OS release version compatibility issue. To fix a version compatibility issue, see “Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)” on page 176.
Empty received	Number of MRP Empty messages received on the switch.

Table 38: show mvrp statistics Output Fields (*continued*)

Field Name	Field Description
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.
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 JoinEmpty messages sent from the switch.
Join In transmitted	Number of MRP JoinIn 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.

Sample Output

show mvrp statistics interface xe-0/1/1.0

```

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

```


show protection-group ethernet-ring aps

Syntax	show protection-group ethernet-ring aps
Release Information	Command introduced in Junos OS Release 9.4. Command introduced in Junos OS Release 12.1 for EX Series switches.
Description	Display the status of the Automatic Protection Switching (APS) and Ring APS (RAPS) messages on an Ethernet ring.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show protection-group ethernet-ring data-channel • show protection-group ethernet-ring interface on page 350 • show protection-group ethernet-ring node-state on page 353 • show protection-group ethernet-ring statistics on page 356 • show protection-group ethernet-ring vlan
List of Sample Output	show protection-group ethernet-ring aps (EX Switches) on page 346 show protection-group ethernet-ring aps (Owner Node, Normal Operation on MX Routers) on page 346 show protection-group ethernet-ring aps (Ring Node, Normal Operation on MX Routers) on page 346 show protection-group ethernet-ring aps (Owner Node, Failure Condition on MX Routers) on page 346 show protection-group ethernet-ring aps (Ring Node, Failure Condition on MX Routers) on page 346
Output Fields	Table 39 on page 345 lists the output fields for the show protection-group ethernet-ring aps command. Output fields are listed in the approximate order in which they appear.

Table 39: show protection-group ethernet-ring aps Output Fields

Field Name	Field Description
Ethernet Ring Name	Name configured for the Ethernet ring.
Request/State	Status of the Ethernet ring RAPS messages. <ul style="list-style-type: none"> • NR—Indicates there is no request for APS on the ring. • SF—Indicates there is a signal failure on the ring.
No Flush	State of the ring flushing: No (normal) or Yes (failure).
Ring Protection Link Blocked	Blocking on the ring protection link: Yes or No .

Table 39: show protection-group ethernet-ring aps Output Fields (*continued*)

Field Name	Field Description
Originator	Whether this node is the ring originator: Yes or No .
Remote Node ID	Identifier (in MAC address format) of the remote node.

Sample Output

show protection-group ethernet-ring aps (EX Switches)

```

user@switch>> show protection-group ethernet-ring aps
Ring Name    Request/state  No Flush  RPL Blocked  Originator  Remote Node ID
erp1         NR             no        yes          no          00:1F:12:30:B8:81

```

Sample Output

show protection-group ethernet-ring aps (Owner Node, Normal Operation on MX Routers)

```

user@host> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg101              NR             No        Yes

Originator  Remote Node ID
Yes

```

show protection-group ethernet-ring aps (Ring Node, Normal Operation on MX Routers)

```

user@host> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg102              NR             No        Yes

Originator  Remote Node ID
No          00:01:01:00:00:01

```

show protection-group ethernet-ring aps (Owner Node, Failure Condition on MX Routers)

```

user@host> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg101              SF             No        No

Originator  Remote Node ID
No          00:01:02:00:00:01

```

show protection-group ethernet-ring aps (Ring Node, Failure Condition on MX Routers)

```

user@host> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg102              SF             No        Yes

Originator  Remote Node ID
Yes         00:00:00:00:00:00

```

show protection-group ethernet-ring configuration

Syntax	show protection-group ethernet-ring configuration
Release Information	Command introduced in Junos OS Release 12.1 for EX Series switches. Command introduced in Junos OS Release 14.1 for MX Series routers.
Description	Display the configuration of Ethernet ring protection group on EX Switches and MX Series routers.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show protection-group ethernet-ring aps on page 345 • show protection-group ethernet-ring data-channel • show protection-group ethernet-ring interface on page 350 • show protection-group ethernet-ring node-state on page 353 • show protection-group ethernet-ring statistics on page 356 • show protection-group ethernet-ring vlan
List of Sample Output	show protection-group ethernet-ring configuration (EX Switch) on page 349 show protection-group ethernet-ring configuration (MX Series Router) on page 349
Output Fields	Table 40 on page 347 lists the output fields for the show protection-group ethernet-ring configuration command. Output fields are listed in the approximate order in which they appear.

Table 40: show protection-group ethernet-ring configuration Output Fields

Output Fields	Field Description
G8032 Compatibility Version	This is the compatibility version mode of ERP. This parameter always takes the value 1 in the case of G8032v1. This parameter is valid only for MX Series routers.
East Interface	One of the two switch interfaces that participates in a ring link. When Junos supports G8032v2, this interface is treated as interface 0.
West Interface	One of the two interfaces in a switch that participates in a ring link. When Junos supports G8032v2, this interface is treated as interface 1.
Restore Interval	Configured interval of wait time after a link is restored. When the down link becomes active again, the RPL owner receives a notification. The RPL owner waits for the restore interval before issuing a block on the RPL link. The configured restore interval can be 5 through 12 minutes for ER Pv1 and 1 through 12 minutes for ER Pv2. This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.

Table 40: show protection-group ethernet-ring configuration Output Fields (*continued*)

Output Fields	Field Description
Guard Interval	Configured number of milliseconds (in 10 millisecond intervals, 10 milliseconds through 2000 milliseconds) that the node does not process any Ethernet ring protection protocol data units (PDUs). This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.
Hold off interval	This is the interval at which the link is held down even before declaring that the link is down. Because the parameter is not supported at present, its value is always considered 0. This parameter is valid only for MX Series routers.
Node ID	Node ID for the switch or router. If the node ID is not configured, it is assigned by default. For EX Series switches, the Node ID value cannot be configured, whereas for MX Series routers, it can be configured.
Ring ID	In G8032v2, the ring ID can be within the range 1–239. All the nodes in a ring should have the same ring ID. In the case of G8032v1, the value of the ring ID is always 1. This parameter is valid only for MX Series routers.
Node Role	This parameter indicates whether the ring node is operating in normal or RPL-owner node. This parameter is valid only for MX Series routers.
Node RPL End	The ring link which is configured as Ring Protection Link (RPL) end.
Revertive mode of operation	This parameter indicates whether the ring is operating in revertive mode or nonrevertive mode. In nonrevertive mode of operation, when all links in the ring and Ethernet Ring Nodes have recovered and no external requests are active, the Ethernet Ring does not automatically revert. G8032v1 supports only revertive mode of operation. This parameter is valid only for MX Series routers.
RAPS Tx Dot1p priority	The RAPS Tx Dot1p priority is a parameter with which the RAPS is transmitted from the ring node. For G8032v1, the value of this parameter is always 0. For G8032v2, the value of this parameter can be within the range 0–7. This parameter is valid only for MX Series routers.
Node type	This parameter indicates the type of the node whether it is a normal ring node having two links in the ring or an open-ring node having just a single link in the ring node. This parameter is valid only for MX Series routers.
Control Vlan	The VLAN that transfers ERP PDUs from one node to another.
Physical Ring	Physical ring if the east and west interfaces are nontrunk ports. For MX Series routers, the ring is termed a physical ring if no data channels are defined for the ring and the entire physical port forwarding is controlled by ERP.
Data Channel VLAN(s)	Data VLANs for which forwarding behavior is controlled by the ring instance.

Sample Output

show protection-group ethernet-ring configuration (EX Switch)

```
user@switch> show protection-group ethernet-ring configuration
Ethernet ring configuration parameters for protection group erp1
East Interface   : ge-0/0/3.0
West Interface   : ge-0/0/9.0
Restore Interval : 5 minutes
Guard Interval   : 500 ms
Node Id          : 00:1F:12:30:B8:81
Control Vlan     : 101
Physical Ring     : yes
```

show protection-group ethernet-ring configuration (MX Series Router)

```
user@switch> show protection-group ethernet-ring configuration
Ethernet Ring configuration information for protection group pg_major
G8032 Compatibility Version      : 1
East interface (interface 0)    : xe-2/3/0.1
West interface (interface 1)    : xe-2/2/1.1
Restore interval                 : 5 minutes
Guard interval                   : 500 ms
Hold off interval                : 0 ms
Node ID                         : 64:87:88:65:37:D0
Ring ID (1 ... 239)             : 1
Node role (normal/rpl-owner)    : rpl-owner
Node RPL end                     : east-port
Revertive mode of operation     : 1
RAPS Tx Dot1p priority (0 .. 7) : 0
Node type (normal/open)         : Normal
Control Vlan                     : 100
Physical Ring                   : No
Data Channel Vlan(s)            : 200,300
```

show protection-group ethernet-ring interface

Syntax	show protection-group ethernet-ring interface
Release Information	Command introduced in Junos OS Release 9.4.
Description	Displays the status of the Automatic Protection Switching (APS) interfaces on an Ethernet ring.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show protection-group ethernet-ring data-channel • show protection-group ethernet-ring aps on page 345 • show protection-group ethernet-ring node-state on page 353 • show protection-group ethernet-ring statistics on page 356 • show protection-group ethernet-ring vlan
List of Sample Output	show protection-group ethernet-ring interface (EX Series Switch Owner Node) on page 351 show protection-group ethernet-ring interface (Owner Node MX Series Router) on page 351 show protection-group ethernet-ring interface (EX Series Switch Ring Node) on page 351 show protection-group ethernet-ring interface (MX Series Router Ring Node) on page 351
Output Fields	Table 41 on page 350 lists the output fields for both the EX Series switch and the MX Series router show protection-group ethernet-ring interface commands. Output fields are listed in the approximate order in which they appear.

Table 41: MX Series Routers show protection-group ethernet-ring interface Output Fields

Field Name	Field Description
Ethernet ring port parameters for protection group <i>group-name</i>	Output is organized by configured protection group.
Interface	Physical interfaces configured for the Ethernet ring.
Control Channel	(MX Series router only) Logical unit configured on the physical interface. <ul style="list-style-type: none"> • NR—Indicates there is no request for APS on the ring. • SF—Indicates there is a signal failure on the ring.
Forward State	State of the ring forwarding on the interface: discarding or forwarding .

Table 41: MX Series Routers show protection-group ethernet-ring interface Output Fields (*continued*)

Field Name	Field Description
Ring Protection Link End	Whether this interface is the end of the ring: Yes or No .
Signal Failure	Whether there a signal failure exists on the link: Clear or Set .
Admin State	State of the interface: For EX switches, ready , ifl ready , or waiting . For MX routers, IFF ready or IFF disabled .

Sample Output

show protection-group ethernet-ring interface (EX Series Switch Owner Node)

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101
```

Interface	Forward State	RPL End	Signal Failure	Admin State
ge-0/0/3.0	discarding	Yes	Clear	ready
ge-0/0/9.0	forwarding	No	Clear	ready

show protection-group ethernet-ring interface (Owner Node MX Series Router)

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101
```

Interface	Control Channel	Forward State	Ring Protection Link End
ge-1/0/1	ge-1/0/1.1	discarding	Yes
ge-1/2/4	ge-1/2/4.1	forwarding	No

Signal Failure	Admin State
Clear	IFF ready
Clear	IFF ready

show protection-group ethernet-ring interface (EX Series Switch Ring Node)

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102
```

Ethernet ring port parameters for protection group pg101

Interface	Forward State	RPL End	Signal Failure	Admin State
ge-0/0/3.0	discarding	Yes	Clear	ready
ge-0/0/9.0	forwarding	No	Clear	ready

show protection-group ethernet-ring interface (MX Series Router Ring Node)

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102
```

Interface	Control Channel	Forward State	Ring Protection Link End
ge-1/2/1	ge-1/2/1.1	forwarding	No
ge-1/0/2	ge-1/0/2.1	forwarding	No

Signal	Failure	Admin	State
Clear		IFF	ready
Clear		IFF	ready

show protection-group ethernet-ring node-state

Syntax	show protection-group ethernet-ring node-state
Release Information	Command introduced in Junos OS Release 9.4. Command introduced in Junos OS Release 12.1 for EX Series switches.
Description	Display the status of the Automatic Protection Switching (APS) nodes on an Ethernet ring.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show protection-group ethernet-ring data-channel • show protection-group ethernet-ring aps on page 345 • show protection-group ethernet-ring interface on page 350 • show protection-group ethernet-ring statistics on page 356 • show protection-group ethernet-ring vlan
List of Sample Output	show protection-group ethernet-ring node-state (EX Series Switch) on page 354 show protection-group ethernet-ring node-state (MX Series Router - Owner Node, Normal Operation) on page 354 show protection-group ethernet-ring node-state (MX Series Router - Ring Node, Normal Operation) on page 354 show protection-group ethernet-ring node-state (MX Series Router - Owner Node, Remote Signal Failure Condition) on page 355 show protection-group ethernet-ring node-state (MX Series Router - Ring Node, Local Signal Failure Condition) on page 355 show protection-group ethernet-ring node-state detail (MX Series Router - Node state at RPL-owner after signal failure condition is cleared in the ring and before reversion) on page 355
Output Fields	Table 42 on page 353 lists the output fields for the show protection-group ethernet-ring node-state command. Output fields are listed in the approximate order in which they appear.

Table 42: show protection-group ethernet-ring node-state Output Fields

Field Name	Field Description
Ring Name/Ethernet Ring	Name configured for the Ethernet ring.

Table 42: show protection-group ethernet-ring node-state Output Fields (*continued*)

Field Name	Field Description
APS State	State of the Ethernet ring APS. <ul style="list-style-type: none"> idle—Indicates that the ring is working in normal condition and no protection-switching request active or pending in the ring. When the ring is in idle state, it is blocked at RPL link. protected—Indicates that there is a protection switch on the ring due to signal failure condition on the ring link.
Event	Events on the ring. <ul style="list-style-type: none"> NR-RB—Indicates there is no APS request and the ring link is blocked on the ring owner node. NR—Indicates there is no APS request pending in the ring. Local SF—Indicates there is signal failure on one or both the ring links of the node. Remote SF—Indicates there is signal failure on ring links of any other node of the ring.
RPL Owner / Ring Protection Link Owner	Whether this node is the ring owner: Yes or No .
WTR Timer / Restore Timer	Restoration timer: running or disabled .
Guard Timer	Guard timer: running or disabled .
Op state / Operational State	State of the node: Operational or any internal wait state.

Sample Output

show protection-group ethernet-ring node-state (EX Series Switch)

```

user@switch> show protection-group ethernet-ring node-state
Ring Name APS State Event RPL Owner WTR Timer Guard Timer Op State
erp1      idle      NR-RB  yes    disabled  disabled  operational

```

show protection-group ethernet-ring node-state (MX Series Router - Owner Node, Normal Operation)

```

user@host> show protection-group ethernet-ring node-state
Ethernet ring  APS State  Event      RPL Owner
pg101         idle      NR-RB      Yes

Restore Timer  Guard Timer  Operation state
disabled       disabled    operational

```

show protection-group ethernet-ring node-state (MX Series Router - Ring Node, Normal Operation)

```

user@host> show protection-group ethernet-ring node-state
Ethernet ring  APS State  Event      RPL Owner
pg102         idle      NR-RB      No

```

Restore Timer	Guard Timer	Operation state
disabled	disabled	operational

show protection-group ethernet-ring node-state (MX Series Router - Owner Node, Remote Signal Failure Condition)

```
user@host> show protection-group ethernet-ring node-state
Ethernet ring    APS State    Event    RPL Owner
pg101           protected    remote SF    Yes

Restore Timer    Guard Timer    Operation state
disabled         disabled      operational
```

show protection-group ethernet-ring node-state (MX Series Router - Ring Node, Local Signal Failure Condition)

```
user@host> show protection-group ethernet-ring node-state
Ethernet ring    APS State    Event    RPL Owner
pg102           protected    local SF    No

Restore Timer    Guard Timer    Operation state
disabled         disabled      operational
```

show protection-group ethernet-ring node-state detail (MX Series Router - Node state at RPL-owner after signal failure condition is cleared in the ring and before reversion)

```
user@host> show protection-group ethernet-ring node-state detail
Ethernet-Ring name      : pg_major
APS State               : protected
Event                  : WTR running (time to expire: 269 sec)
Ring Protection Link Owner : Yes
Restore Timer          : running
Guard Timer            : disabled
Operation state         : operational
```

show protection-group ethernet-ring statistics

Syntax	show protection-group ethernet-ring statistics <group-name <i>group-name</i> >
Release Information	Command introduced in Junos OS Release 9.4. Command introduced in Junos OS Release 12.1 for EX Series switches.
Description	Display statistics regarding Automatic Protection Switching (APS) protection groups on an Ethernet ring.
Options	group-name —Protection group for which to display statistics. In you omit this optional field, all protection group statistics for configured groups will be displayed.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show protection-group ethernet-ring data-channel • show protection-group ethernet-ring aps on page 345 • show protection-group ethernet-ring node-state on page 353 • show protection-group ethernet-ring interface on page 350 • show protection-group ethernet-ring vlan
List of Sample Output	show protection-group ethernet-ring statistics (EX Switch) on page 357 show protection-group ethernet-ring statistics (Owner Node, Normal Operation on MX Router) on page 357 show protection-group ethernet-ring statistics (Ring Node, Normal Operation on MX Router) on page 357 show protection-group ethernet-ring statistics (Owner Node, Failure Condition on MX Router) on page 357 show protection-group ethernet-ring statistics (Ring Node, Failure Condition on MX Router) on page 358
Output Fields	Table 43 on page 356 lists the output fields for the show protection-group ethernet-ring statistics command. Output fields are listed in the approximate order in which they appear.

Table 43: show protection-group ethernet-ring statistics Output Fields

Field Name	Field Description
Ethernet Ring Statistics for PG	Name of the protection group for which statistics are displayed.
RAPS sent	Number of Ring Automatic Protection Switching (RAPS) messages sent. (On MX Series switches only)
RAPS received	Number of RAPS messages received. (On MX Series switches only)

Table 43: show protection-group ethernet-ring statistics Output Fields (*continued*)

Field Name	Field Description
Local SF	Number of times a signal failure (SF) has occurred locally.
Remote SF	Number of times a signal failure (SF) has occurred anywhere else on the ring.
NR event	Number of times a No Request (NR) event has occurred on the ring.
NR-RB event	Number of times a No Request, Ring Blocked (NR-RB) event has occurred on the ring.

Sample Output

show protection-group ethernet-ring statistics (EX Switch)

```
user@switch> show protection-group ethernet-ring statistics
Ring Name Local SF Remote SF NR Event NR-RB Event
erp1      2      1      2      3
```

show protection-group ethernet-ring statistics (Owner Node, Normal Operation on MX Router)

```
user@host> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent           : 1
RAPS received       : 0
Local SF happened   : 0
Remote SF happened   : 0
NR event happened    : 0
NR-RB event happened : 1
```

show protection-group ethernet-ring statistics (Ring Node, Normal Operation on MX Router)

```
user@host> show protection-group ethernet-ring statistics group-name pg102
Ethernet Ring statistics for PG pg102
RAPS sent           : 0
RAPS received       : 1
Local SF happened   : 0
Remote SF happened   : 0
NR event happened    : 0
NR-RB event happened : 1
```

show protection-group ethernet-ring statistics (Owner Node, Failure Condition on MX Router)

```
user@host> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent           : 1
RAPS received       : 1
Local SF happened   : 0
Remote SF happened   : 1
NR event happened    : 0
NR-RB event happened : 1
```

show protection-group ethernet-ring statistics (Ring Node, Failure Condition on MX Router)

```
user@host> show protection-group ethernet-ring statistics group-name pg102
Ethernet Ring statistics for PG pg102
RAPS sent                               : 1
RAPS received                           : 1
Local SF happened                        : 1
Remote SF happened                       : 0
NR event happened                        : 0
NR-RB event happened                     : 1
```

show redundant-trunk-group

Syntax	<code>show redundant-trunk-group <group-name group-name></code>
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Description	Display information about redundant trunk groups.
Options	<code>group-name group-name</code> —Display information about the specified redundant trunk group.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Redundant Trunk Links for Faster Recovery on page 112 • Example: Configuring Redundant Trunk Links for Faster Recovery • Understanding Redundant Trunk Links on page 49
List of Sample Output	show redundant-trunk-group group-name Group1 on page 359
Output Fields	Table 44 on page 359 lists the output fields for the <code>show redundant-trunk-group</code> command. Output fields are listed in the approximate order in which they appear.

Table 44: 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.
State	Operating state of the interface. <ul style="list-style-type: none"> • Up denotes the interface is up. • Down denotes the interface is down. • Pri denotes a primary interface. • Act denotes an active interface.
Time of last flap	Date and time at which the advertised link became unavailable, and then, available again.
Flap count	Total number of flaps since the last switch reboot.

Sample Output

show redundant-trunk-group group-name Group1

```
user@switch> show redundant-trunk-group group-name Group1
```

Group name	Interface	State	Time of last flap	Flap Count
------------	-----------	-------	-------------------	------------

Group1	ge-0/0/45.0	UP/Pri/Act	Never	0
	ge-0/0/47.0	UP	Never	0

show system statistics arp

Syntax	show system statistics arp
Release Information	Command introduced in Junos OS Release 9.6 for EX Series switches.
Description	Display system-wide Address Resolution Protocol (ARP) statistics.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Proxy ARP on an EX Series Switch on page 117 • Verifying That Proxy ARP Is Working Correctly on page 299

show system statistics arp

```

user@switch> show system statistics arp
arp:
    90060 datagrams received
    34 ARP requests received
    610 ARP replies received
    0 resolution request received
    0 unrestricted proxy requests
    0 restricted proxy requests
    0 received proxy requests
    0 unrestricted proxy requests not proxied
    0 restricted proxy requests not proxied
    0 datagrams with bogus interface
    0 datagrams with incorrect length
    0 datagrams for non-IP protocol
    0 datagrams with unsupported op code
    0 datagrams with bad protocol address length
    0 datagrams with bad hardware address length
    0 datagrams with multicast source address
    0 datagrams with multicast target address
    0 datagrams with my own hardware address
    0 datagrams for an address not on the interface
    0 datagrams with a broadcast source address
    294 datagrams with source address duplicate to mine
    89113 datagrams which were not for me
    0 packets discarded waiting for resolution
    0 packets sent after waiting for resolution
    309 ARP requests sent
    35 ARP replies sent
    0 requests for memory denied
    0 requests dropped on entry
    0 requests dropped during retry
    0 requests dropped due to interface deletion
    0 requests on unnumbered interfaces
    0 new requests on unnumbered interfaces
    0 replies for from unnumbered interfaces
    0 requests on unnumbered interface with non-subnetted donor
    0 replies from unnumbered interface with non-subnetted donor

```

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 OS Release 9.0 for EX Series switches.

Description



NOTE: If your switch CLI displays different options for the `show vlans` command than the options shown in this document, see *show vlans*.

Display information about VLANs configured on bridged Ethernet interfaces. For interfaces configured to support a voice over IP (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 name of 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-tagged, 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 display 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 Documentation**
- [show ethernet-switching interfaces on page 313](#)
 - [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 57](#)
 - [Example: Setting Up Bridging with Multiple VLANs for EX Series Switches on page 65](#)
 - [Example: Configuring a Private VLAN on a Single EX Series Switch on page 81](#)
 - [Example: Configuring a Private VLAN Spanning Multiple EX Series Switches on page 119](#)
 - [Example: Setting Up Q-in-Q Tunneling on EX Series Switches on page 103](#)
 - [Understanding Bridging and VLANs on EX Series Switches on page 3](#)

List of Sample Output

[show vlans on page 366](#)
[show vlans brief on page 366](#)
[show vlans detail on page 367](#)
[show vlans extensive \(for a PVLAN spanning multiple switches\) on page 367](#)
[show vlans extensive \(MAC-based\) on page 369](#)
[show vlans extensive \(Port-based\) on page 369](#)
[show vlans sort-by tag on page 370](#)
[show vlans sort-by name on page 371](#)
[show vlans employee \(vlan-range-name\) on page 371](#)
[show vlans summary on page 372](#)

Output Fields [Table 45 on page 363](#) lists the output fields for the **show vlans** command. Output fields are listed in the approximate order in which they appear.

Table 45: 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

Table 45: show vlans Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
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. Also lists the following attributes of the interfaces: <ul style="list-style-type: none"> • tagged or untagged • trunk or access port mode • pvlan-trunk 	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 (type of broadcast domain) for this VLAN. Values are Primary , Isolated , Inter-switch-isolated , and Community .	detail, extensive
Primary VLAN	The primary VLAN tag for this secondary VLAN.	extensive
Internal Index	VLAN index internal to Junos OS.	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

Table 45: show vlans Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 VLAN	The isolated VLANs associated with a primary VLAN.	extensive
Inter-switch isolated VLAN	The inter-switch isolated VLAN associated with a primary VLAN.	extensive
Community VLANs	The community VLANs associated with a primary VLAN.	extensive
VLANs summary	VLAN counts: <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	802.1Q VLAN counts: <ul style="list-style-type: none"> • Total—Total number of 802.1Q-tagged and untagged VLANs on the switch. • Tagged VLANs—Number of 802.1Q-tagged 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 community transporting and forwarding private VLANs. • Isolated VLANs—Number of isolated receiving and forwarding private VLANs. • Inter-switch-isolated VLANs—Number of inter-switch isolated receiving and forwarding private VLANs. 	All levels
Dot1q Tunneled VLANs summary	Q-in-Q-tunneled VLAN counts: <ul style="list-style-type: none"> • Total—Total number of Q-in-Q-tunneled VLANs on the switch. • Private VLAN—Counts of primary, community, and isolated Q-in-Q-tunneled private VLANs (PVLANS). 	All levels

Table 45: show vlans Output Fields (*continued*)

Field Name	Field Description	Level of Output
Dynamic VLANs	Counts of VLANs assigned or created dynamically by a protocol: <ul style="list-style-type: none"> • Total—Total number of dynamic VLANs on the switch. • Dot1x—Number of 802.1Q-tagged 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

Sample Output

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	ge-0/0/24.0, ge-0/0/23.0, ge-0/0/22.0, ge-0/0/21.0
v0002	2	None
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 (for a PVLAN spanning multiple switches)

```

user@switch> show vlans extensive
VLAN: COM1, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 100, Internal index: 3, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 1)
  ge-0/0/20.0*, tagged, trunk
  ge-0/0/22.0*, tagged, trunk, pvlan-trunk
  ge-0/0/23.0*, tagged, trunk, pvlan-trunk
  ge-0/0/7.0*, untagged, access

VLAN: __pvlan_primary_ge-0/0/0.0__, Created at: Tue May 11 18:16:05 2010
Internal index: 5, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 1)

```

```
ge-0/0/20.0*, tagged, trunk
ge-0/0/22.0*, tagged, trunk, pvlan-trunk
ge-0/0/23.0*, tagged, trunk, pvlan-trunk
ge-0/0/0.0*, untagged, access
```

```
VLAN: __pvlan_primary_ge-0/0/2.0__, Created at: Tue May 11 18:16:05 2010
Internal index: 6, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 0)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/2.0, untagged, access
```

```
VLAN: __pvlan_primary_isiv__, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 50, Internal index: 7, Admin State: Enabled, Origin: Static
Private VLAN Mode: Inter-switch-isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 0 (Active = 0)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
```

```
VLAN: community2, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 20, Internal index: 8, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 2 (Active = 2)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/1.0*, untagged, access
    ge-1/0/6.0*, untagged, access
```

```
VLAN: primary, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 10, Internal index: 2, Admin State: Enabled, Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 5 (Active = 4)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/0.0*, untagged, access
    ge-0/0/1.0*, untagged, access
    ge-0/0/2.0, untagged, access
    ge-0/0/7.0*, untagged, access
    ge-1/0/6.0*, untagged, access
```

```
Secondary VLANs: Isolated 2, Community 2, Inter-switch-isolated 1
Isolated VLANs :
    __pvlan_primary_ge-0/0/0.0__
    __pvlan_primary_ge-0/0/2.0__
Community VLANs :
    COM1
    community2
Inter-switch-isolated VLAN :
```

__pvlan_primary_isiv__

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 (vlan-range-name)

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

Name	Tag	Interfaces
__employee_120__	120	

__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 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 Tunnelled VLANs summary:
  Total: 0
  Private VLAN:
    Primary VLANs: 0,   Community VLANs: 0, Isolated VLANs: 0

Dynamic VLANs:
  Total: 2,   Dot1x: 2, MVRP: 0
```


PART 4

Troubleshooting

- [Troubleshooting Procedure on page 375](#)

Troubleshooting Procedure

- [Troubleshooting Ethernet Switching on page 375](#)

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 375](#)

MAC Address in the Switch's Ethernet Switching Table Is Not Updated After a MAC Address Move

Problem **Description:** 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 OS 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 Documentation**
- [arp \(System\) on page 208](#)
 - [mac-table-aging-time on page 247](#)