

# Ethernet Switching Feature Guide on EX9200 Switches

Release  
**15.1**



Modified: 2016-07-15

Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://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.

*Ethernet Switching Feature Guide on EX9200 Switches*

15.1

Copyright © 2016, Juniper Networks, Inc.  
All rights reserved.

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

#### YEAR 2000 NOTICE

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

#### END USER LICENSE AGREEMENT

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

# Table of Contents

	About the Documentation . . . . .	xi
	Documentation and Release Notes . . . . .	xi
	Supported Platforms . . . . .	xi
	Using the Examples in This Manual . . . . .	xi
	Merging a Full Example . . . . .	xii
	Merging a Snippet . . . . .	xii
	Documentation Conventions . . . . .	xiii
	Documentation Feedback . . . . .	xv
	Requesting Technical Support . . . . .	xv
	Self-Help Online Tools and Resources . . . . .	xv
	Opening a Case with JTAC . . . . .	xvi
<b>Part 1</b>	<b>Ethernet Switching</b>	
<b>Chapter 1</b>	<b>Configuring 802.1Q VLANs . . . . .</b>	<b>3</b>
	Layer 2 VLANs Overview . . . . .	3
	Configuring a VLAN . . . . .	4
	Configuring VLAN Encapsulation . . . . .	5
	Example: Configuring VLAN Encapsulation on a Gigabit Ethernet Interface . . . . .	5
	Example: Configuring VLAN Encapsulation on an Aggregated Ethernet Interface . . . . .	6
	Configuring Inner and Outer TPIDs and VLAN IDs . . . . .	6
	Stacking a VLAN Tag . . . . .	10
	Rewriting a VLAN Tag and Adding a New Tag . . . . .	10
	Configuring VLAN Translation with a VLAN ID List . . . . .	12
	Configuring Static MAC Addresses for Logical Interfaces in a VLAN . . . . .	13
<b>Chapter 2</b>	<b>Configuring Learning and Forwarding for 802.1Q VLANs . . . . .</b>	<b>15</b>
	Layer 2 VLANs Overview . . . . .	15
	Layer 2 Learning and Forwarding Overview . . . . .	16
	Layer 2 Learning and Forwarding for VLANs Overview . . . . .	16
	Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port . . . . .	17
	Enabling MAC Accounting . . . . .	18
	Enabling MAC Accounting for a VLAN . . . . .	18
	Enabling MAC Accounting for a Set of VLANs . . . . .	18
	Verifying That MAC Accounting Is Working . . . . .	18
	Disabling Layer 2 Learning and Forwarding . . . . .	20
	Disabling MAC Learning for a VLAN or Logical Interface . . . . .	21
	Disabling MAC Learning for a Set of VLANs . . . . .	22

<b>Chapter 3</b>	<b>Configuring VLANs and VPLS Routing Instances . . . . .</b>	<b>23</b>
	Guidelines for Configuring VLAN Identifiers for VLANs and VPLS Routing Instances . . . . .	23
	Configuring VLAN Identifiers for VLANs and VPLS Routing Instances . . . . .	23
<b>Chapter 4</b>	<b>Configuring Bridging and Integrated Routing and Bridging (IRB) . . . . .</b>	<b>29</b>
	Configuring Integrated Routing and Bridging for VLANs . . . . .	29
	Configuring a Set of VLANs to Act as a Switch for a Layer 2 Trunk Port . . . . .	30
	Configuring a Logical Interface for Access Mode . . . . .	32
	Configuring the Size of the MAC Address Table . . . . .	32
	Configuring a Layer 2 Virtual Switch . . . . .	33
	Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port . . . . .	34
<b>Chapter 5</b>	<b>Configuring Private VLANs . . . . .</b>	<b>35</b>
	Understanding Private VLANs . . . . .	35
	Typical Structure and Primary Application of PVLANS . . . . .	36
	Routing Between Isolated and Community VLANs . . . . .	38
	PVLANS Use 802.1Q Tags to Identify Packets . . . . .	38
	PVLANS Use IP Addresses Efficiently . . . . .	39
	PVLAN Port Types and Forwarding Rules . . . . .	39
	Guidelines and Restrictions for PVLANS . . . . .	39
	Limitation of Private VLANs . . . . .	40
	Example: Configuring a Private VLAN on a Single Switch . . . . .	41
	Creating a Private VLAN on a Single Switch (CLI Procedure) . . . . .	45
	Creating a Private VLAN Spanning Multiple Switches (CLI Procedure) . . . . .	47
<b>Chapter 6</b>	<b>Configuring Multiple VLAN Registration Protocol (MVRP) . . . . .</b>	<b>49</b>
	Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration . . . . .	49
	How MVRP Works . . . . .	50
	Basics of MVRP . . . . .	50
	MVRP Registration Modes . . . . .	50
	MRP Timers . . . . .	51
	MRP VLAN Messages . . . . .	51
	MVRP Limitations . . . . .	51
	Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration . . . . .	52
	Enabling MVRP . . . . .	52
	Disabling MVRP . . . . .	52
	Changing the Registration Mode to Disable Dynamic VLANs . . . . .	53
	Configuring Timer Values . . . . .	53
	Configuring the Multicast MAC Address for MVRP . . . . .	54
	Configuring an MVRP Interface as a Point-to-Point Interface . . . . .	54
	Configuring MVRP Tracing Options . . . . .	54
	Verifying That MVRP Is Working Correctly . . . . .	54
<b>Chapter 7</b>	<b>Configuring Proxy ARP . . . . .</b>	<b>57</b>
	Restricted and Unrestricted Proxy ARP Overview . . . . .	57
	Restricted Proxy ARP . . . . .	57
	Unrestricted Proxy ARP . . . . .	57

	Topology Considerations for Unrestricted Proxy ARP . . . . .	58
	Configuring Restricted and Unrestricted Proxy ARP . . . . .	59
<b>Part 2</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 8</b>	<b>Configuration Statements . . . . .</b>	<b>63</b>
	[edit dynamic-profiles] Hierarchy Level . . . . .	65
	[edit interfaces] Hierarchy Level . . . . .	65
	[edit logical-systems] Hierarchy Level . . . . .	76
	[edit protocols l2-learning] Hierarchy Level . . . . .	81
	Layer 2 Routing Instances Configuration Hierarchy . . . . .	81
	[edit switch-options] Hierarchy Level . . . . .	84
	[edit vlans] Hierarchy Level . . . . .	84
	bpdu-destination-mac-address . . . . .	86
	bridge-priority . . . . .	87
	community-vlan . . . . .	88
	domain-type . . . . .	88
	encapsulation (Physical Interface) . . . . .	89
	family . . . . .	94
	fast-aps-switch . . . . .	98
	global-mac-statistics . . . . .	99
	global-mac-table-aging-time . . . . .	99
	global-no-mac-learning . . . . .	100
	inner-tag-protocol-id . . . . .	101
	inner-vlan-id . . . . .	102
	input-vlan-map . . . . .	103
	interface . . . . .	104
	interface (MVRP) . . . . .	105
	interface (Spanning Tree) . . . . .	106
	interface-mode . . . . .	107
	interfaces . . . . .	108
	inter-switch-link . . . . .	109
	isid . . . . .	109
	isid-list . . . . .	110
	isolated-vlan . . . . .	110
	join-timer (MVRP) . . . . .	111
	l2-learning . . . . .	112
	l3-interface . . . . .	112
	leaveall-timer (MVRP) . . . . .	113
	leave-timer (MVRP) . . . . .	114
	mac-statistics . . . . .	115
	mac-table-size . . . . .	117
	mvrp . . . . .	119
	native-vlan-id . . . . .	120
	no-mac-learning . . . . .	122
	output-vlan-map (Gigabit Ethernet IQ and 10-Gigabit Ethernet with SFPP) . . . . .	123
	packet-action . . . . .	124
	point-to-point (MVRP) . . . . .	127
	pop . . . . .	128

	pop-pop	129
	pop-swap	130
	private-vlan	131
	proxy-arp	132
	push	133
	push-push	134
	registration	135
	rstp	136
	service-id	137
	static-mac	138
	swap	139
	swap-push	140
	swap-swap	141
	switch-options	142
	tag-protocol-id (TPID to Rewrite)	143
	traceoptions (MVRP)	144
	unconditional-src-learn	146
	vlan-id (Bridge Domain or VLAN)	147
	vlan-id (VLAN ID to Rewrite)	149
	vlan-id-list	150
	vlan-rewrite	151
	vlan-tags	152
<b>Chapter 9</b>	<b>Operational Commands</b>	<b>153</b>
	show ethernet-switching flood	154
	show ethernet-switching interface	158
	show ethernet-switching statistics	161
	show ethernet-switching table	164
	show interfaces (10-Gigabit Ethernet)	171
	show interfaces (Gigabit Ethernet)	196
	show interfaces irb	220
	show interfaces queue	226
	show mvrp	265
	show mvrp applicant-state	267
	show mvrp dynamic-vlan-memberships	269
	show mvrp interface	270
	show mvrp registration-state	271
	show mvrp statistics	273
	show vlans	275
	traceroute ethernet	281

# List of Figures

<b>Part 1</b>	<b>Ethernet Switching</b>	
<b>Chapter 5</b>	<b>Configuring Private VLANs . . . . .</b>	<b>35</b>
	Figure 1: Private VLAN on a Single Switch . . . . .	37
	Figure 2: PVLAN Spanning Multiple Switches . . . . .	38
	Figure 3: Topology of a Private VLAN on a Single EX Series Switch . . . . .	42
<b>Chapter 7</b>	<b>Configuring Proxy ARP . . . . .</b>	<b>57</b>
	Figure 4: Edge Device Case for Unrestricted Proxy ARP . . . . .	58
	Figure 5: Core Device Case for Unrestricted Proxy ARP . . . . .	59





# List of Tables

	<b>About the Documentation</b> . . . . .	<b>xi</b>
	Table 1: Notice Icons . . . . .	xiii
	Table 2: Text and Syntax Conventions . . . . .	xiii
<b>Part 1</b>	<b>Ethernet Switching</b>	
<b>Chapter 1</b>	<b>Configuring 802.1Q VLANs</b> . . . . .	<b>3</b>
	Table 3: Rewrite Operations and Statement Usage for Input VLAN Maps . . . . .	7
	Table 4: Rewrite Operations and Statement Usage for Output VLAN Maps . . . . .	8
<b>Chapter 3</b>	<b>Configuring VLANs and VPLS Routing Instances</b> . . . . .	<b>23</b>
	Table 5: Statement Usage and Input Rewrite Operations for VLAN Identifiers for a VLAN . . . . .	27
	Table 6: Statement Usage and Output Rewrite Operations for VLAN Identifiers for a VLAN . . . . .	28
<b>Chapter 5</b>	<b>Configuring Private VLANs</b> . . . . .	<b>35</b>
	Table 7: PVLAN Ports and Layer 2 Forwarding . . . . .	39
	Table 8: Interfaces of the Topology for Configuring a PVLAN . . . . .	41
	Table 9: VLAN IDs in the Topology for Configuring a PVLAN . . . . .	42
<b>Part 2</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 9</b>	<b>Operational Commands</b> . . . . .	<b>153</b>
	Table 10: show ethernet-switching interface Output Fields . . . . .	158
	Table 11: show ethernet-switching table Output fields . . . . .	165
	Table 12: show interfaces Gigabit Ethernet Output Fields . . . . .	172
	Table 13: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type . . . . .	186
	Table 14: show interfaces (Gigabit Ethernet) Output Fields . . . . .	197
	Table 15: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type . . . . .	211
	Table 16: show interfaces irb Output Fields . . . . .	220
	Table 17: Layer 2 Overhead and Transmitted Packets or Byte Counts . . . . .	227
	Table 18: show interfaces queue Output Fields . . . . .	230
	Table 19: Byte Count by Interface Hardware . . . . .	233
	Table 20: show mvrp Output Fields . . . . .	265
	Table 21: show mvrp applicant-state Output Fields . . . . .	267
	Table 22: show mvrp dynamic-vlan-memberships Output Fields . . . . .	269
	Table 23: show mvrp interface Output Fields . . . . .	270
	Table 24: show mvrp registration-state Output Fields . . . . .	271
	Table 25: show mvrp statistics Output Fields . . . . .	273

Table 26: traceroute ethernet Output Fields . . . . .	282
---	-----

# About the Documentation

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

## 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 xiii 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 xiii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

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

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> <b>show chassis alarms</b>  No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"><li>Introduces or emphasizes important new terms.</li><li>Identifies guide names.</li><li>Identifies RFC and Internet draft titles.</li></ul>	<ul style="list-style-type: none"><li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li><li><i>Junos OS CLI User Guide</i></li><li>RFC 1997, <i>BGP Communities Attribute</i></li></ul>
<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@# <b>set system domain-name</b> <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"><li>To configure a stub area, include the <b>stub</b> statement at the [edit protocols ospf area area-id] hierarchy level.</li><li>The console port is labeled <b>CONSOLE</b>.</li></ul>
< > (angle brackets)	Encloses optional keywords or variables.	<b>stub &lt;default-metric <i>metric</i>&gt;;</b>
(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.	<b>broadcast   multicast</b>  <b>(<i>string1</i>   <i>string2</i>   <i>string3</i>)</b>
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<b>rsvp { # Required for dynamic MPLS only</b>
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	<b>community name members [</b> <i>community-ids</i> <b>]</b>
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.	
<b>GUI Conventions</b>		
<b>Bold text like this</b>	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"><li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li><li>To cancel the configuration, click <b>Cancel</b>.</li></ul>

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 <b>Protocols&gt;Ospf</b> .

## Documentation Feedback

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

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

## Requesting Technical Support

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

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

## Self-Help Online Tools and Resources

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

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

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

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

## Opening a Case with JTAC

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

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

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



## PART 1

# Ethernet Switching

- [Configuring 802.1Q VLANs on page 3](#)
- [Configuring Learning and Forwarding for 802.1Q VLANs on page 15](#)
- [Configuring VLANs and VPLS Routing Instances on page 23](#)
- [Configuring Bridging and Integrated Routing and Bridging \(IRB\) on page 29](#)
- [Configuring Private VLANs on page 35](#)
- [Configuring Multiple VLAN Registration Protocol \(MVRP\) on page 49](#)
- [Configuring Proxy ARP on page 57](#)



## CHAPTER 1

# Configuring 802.1Q VLANs

- [Layer 2 VLANs Overview on page 3](#)
- [Configuring a VLAN on page 4](#)
- [Configuring VLAN Encapsulation on page 5](#)
- [Configuring Inner and Outer TPIDs and VLAN IDs on page 6](#)
- [Stacking a VLAN Tag on page 10](#)
- [Rewriting a VLAN Tag and Adding a New Tag on page 10](#)
- [Configuring VLAN Translation with a VLAN ID List on page 12](#)
- [Configuring Static MAC Addresses for Logical Interfaces in a VLAN on page 13](#)

### Layer 2 VLANs Overview

---

You can configure one or more VLANs to perform Layer 2 bridging. The Layer 2 bridging functions include integrated routing and bridging (IRB) for support for Layer 2 bridging and Layer 3 IP routing on the same interface, and virtual switches that isolate a LAN segment with its spanning-tree protocol instance and separate its VLAN ID space.

A VLAN is a set of logical ports that share the same flooding or broadcast characteristics and span one or more ports of multiple devices.

You can configure one or more VLANs to perform Layer 2 bridging. Thus, MX Series routers or EX Series switches can function as Layer 2 switches, each with multiple bridging, or broadcast, domains that participate in the same Layer 2 network. You can also configure Layer 3 routing support for a VLAN. Integrated routing and bridging (IRB) provides support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route packets to another routed interface or to another VLAN that has a Layer 3 protocol configured.

You can also group one or more VLANs within a single instance, or virtual switch. Multiple virtual switches, each of which operates independently of other virtual switches on the device, are supported. Virtual switches isolate a LAN segment with its spanning-tree protocol instance and separate its VLAN ID space. Thus, each virtual switch can participate in a different Layer 2 network.

VLANs provide support for a Layer 2 trunk port. A Layer 2 trunk interface enables you to configure a single logical interface to represent multiple VLANs on a physical interface. You can configure a set of VLANs and VLAN identifiers that are automatically associated

with one or more Layer 2 trunk interfaces. Packets received on a trunk interface are forwarded within a VLAN that has the same VLAN identifier. A Layer 2 trunk interface also supports IRB within a VLAN. In addition, you can configure Layer 2 learning and forwarding properties that apply to the entire set of VLANs.

You can configure VPLS ports in a virtual switch instead of a dedicated routing instance of type **vpls** so that the logical interfaces of the Layer 2 VLANs in the virtual switch can handle VPLS routing instance traffic. Packets received on a Layer 2 trunk interface are forwarded within a VLAN that has the same VLAN identifier.

## Configuring a VLAN

---

A VLAN must include a set of logical interfaces that participate in Layer 2 learning and forwarding. You can optionally configure a VLAN identifier and a Layer 3 interface for the VLAN to also support Layer 3 IP routing.

To enable a VLAN, include the following statements:

```
[edit]
vpls {
  vlan-name {
    domain-type bridge;
    interface interface-name;
    l3-interface interface-name;
    vlan-id (none | all | number);
    vlan-id-list [ vlan-id-numbers ];
    vlan-tags outer number inner number;
  }
}
```

You cannot use the slash (/) character in VLAN names. If you do, the configuration does not commit and an error is generated.

For the **vlan-id** statement, you can specify either a valid VLAN identifier or the **none** or **all** options.

To include one or more logical interfaces in the VLAN, specify an **interface-name** for an Ethernet interface you configured at the **[edit interfaces]** hierarchy level.



**NOTE:** A maximum of 4096 active logical interfaces are supported for a VLAN or on each mesh group in a virtual private LAN service (VPLS) instance configured for Layer 2 bridging.

---

By default, each VLAN maintains a Layer 2 forwarding database that contains media access control (MAC) addresses learned from packets received on the ports that belong to the VLAN. You can modify Layer 2 forwarding properties, for example, disabling MAC learning for the entire system or a VLAN, adding static MAC addresses for specific logical interfaces, and limiting the number of MAC addresses learned by the entire system, the VLAN, or a logical interface.

You can also configure spanning tree protocols to prevent forwarding loops.

## Configuring VLAN Encapsulation

To configure encapsulation on an interface, enter the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
encapsulation type;
```

The following list contains important notes regarding encapsulation:

- Ethernet interfaces in VLAN mode can have multiple logical interfaces. In CCC and VPLS modes, VLAN IDs from 1 through 511 are reserved for normal VLANs, and VLAN IDs 512 through 4094 are reserved for CCC or VPLS VLANs. For 4-port Fast Ethernet interfaces, you can use VLAN IDs 512 through 1024 for CCC or VPLS VLANs.
- For encapsulation type **flexible-ethernet-services**, all VLAN IDs are valid.
- For some encapsulation types, including flexible Ethernet services, Ethernet VLAN CCC, and VLAN VPLS, you can also configure the encapsulation type that is used inside the VLAN circuit itself. To do this, include the **encapsulation** statement:

```
encapsulation (vlan-ccc | vlan-tcc | vlan-vpls);
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**
- You cannot configure a logical interface with VLAN CCC or VLAN VPLS encapsulation unless you also configure the physical device with the same encapsulation or with flexible Ethernet services encapsulation. In general, the logical interface must have a VLAN ID of 512 or higher; if the VLAN ID is 511 or lower, it will be subject to the normal destination filter lookups in addition to source address filtering. However if you configure flexible Ethernet services encapsulation, this VLAN ID restriction is removed.

In general, you configure an interface's encapsulation at the **[edit interfaces *interface-name*]** hierarchy level.

### Example: Configuring VLAN Encapsulation on a Gigabit Ethernet Interface

Configure VLAN CCC encapsulation on a Gigabit Ethernet interface:

```
interfaces ge-2/1/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 600;
  }
}
```

## Example: Configuring VLAN Encapsulation on an Aggregated Ethernet Interface

Configure VLAN CCC encapsulation on an aggregated Gigabit Ethernet interface:

```
interfaces ae0 {  
  vlan-tagging;  
  encapsulation vlan-vpls;  
  unit 0 {  
    vlan-id 100;  
  }  
}
```

### Related Documentation

- [802.1Q VLANs Overview](#)
- [Ethernet Interfaces Feature Guide for Routing Devices](#)

---

## Configuring Inner and Outer TPIDs and VLAN IDs

For some rewrite operations, you must configure the inner or outer TPID values and inner or outer VLAN ID values. These values can be applied to either the input VLAN map or the output VLAN map.

On Ethernet IQ, IQ2, and IQ2-E interfaces; on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces; and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to configure the inner TPID, include the **inner-tag-protocol-id** statement:

```
inner-tag-protocol-id tpid;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* **input-vlan-map**]
- [edit interfaces *interface-name* unit *logical-unit-number* **output-vlan-map**]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* **input-vlan-map**]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* **output-vlan-map**]

For the inner VLAN ID, include the **inner-vlan-id** statement. For the outer TPID, include the **tag-protocol-id** statement. For the outer VLAN ID, include the **vlan-id** statement:

```
input-vlan-map {  
  (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);  
  inner-tag-protocol-id tpid;  
  inner-vlan-id number;  
  tag-protocol-id tpid;  
  vlan-id number;  
}  
output-vlan-map {  
  (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);  
  inner-tag-protocol-id tpid;  
  inner-vlan-id number;
```

```

tag-protocol-id tpid;
vlan-id number;
}

```

For aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces, include the **tag-protocol-id** statement for the outer TPID. For the outer VLAN ID, include the **vlan-id** statement:

```

input-vlan-map {
  (pop | push | swap);
  tag-protocol-id tpid;
  vlan-id number;
}
output-vlan-map {
  (pop | push | swap);
  tag-protocol-id tpid;
  vlan-id number;
}

```

You can include these statements at the following hierarchy levels:

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

The VLAN IDs you define in the input VLAN maps are stacked on top of the VLAN ID bound to the logical interface. For more information about binding a VLAN ID to the logical interface, see *802.1Q VLANs Overview*.

All TPIDs you include in input and output VLAN maps must be among those you specify at the [edit interfaces *interface-name* ether-options ethernet-switch-profile tag-protocol-id [ *tpids* ]] hierarchy level.

Table 3 on page 7 and Table 4 on page 8 specify when these statements are required. Table 3 on page 7 indicates valid statement combinations for rewrite operations for the input VLAN map. “No” means the statement must not be included in the input VLAN map for the rewrite operation. “Optional” means the statement may be optionally specified for the rewrite operation in the input VLAN map. “Any” means that you must include the **vlan-id** statement, **tag-protocol-id** statement, **inner-vlan-id** statement, or **inner-tag-protocol-id** statement.

**Table 3: Rewrite Operations and Statement Usage for Input VLAN Maps**

Rewrite Operation	Input VLAN Map Statements			
	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
push	Optional	Optional	No	No
pop	No	No	No	No
swap	Any	Any	No	No
push-push	Optional	Optional	Optional	optional

Table 3: Rewrite Operations and Statement Usage for Input VLAN Maps (*continued*)

	Input VLAN Map Statements			
swap-push	Optional	Optional	Any	Any
swap-swap	Optional	Optional	Any	Any
pop-swap	No	No	Any	Any
pop-pop	No	No	No	No

Table 4 on page 8 indicates valid statement combinations for rewrite operations for the output VLAN map. “No” means the statement must not be included in the output VLAN map for the rewrite operation. “Optional” means the statement may be optionally specified for the rewrite operation in the output VLAN map.

Table 4: Rewrite Operations and Statement Usage for Output VLAN Maps

	Output VLAN Map Statements			
Rewrite Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
push	No	Optional	No	No
pop	No	No	No	No
swap	No	Optional	No	No
push-push	No	Optional	No	Optional
swap-push	No	Optional	No	Optional
swap-swap	No	Optional	No	Optional
pop-swap	No	No	No	Optional
pop-pop	No	No	No	No

The following examples use Table 3 on page 7 and Table 4 on page 8 and show how the **pop-swap** operation can be configured in an input VLAN map and an output VLAN map:

Input VLAN Map with inner-vlan-id Statement, Output VLAN Map with Optional inner-tag-protocol-id Statement	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] input-vlan-map {   pop-swap;   inner-vlan-id <i>number</i>; } output-vlan-map {   pop-swap;   inner-tag-protocol-id <i>tpid</i>;</pre>
--	---



```
}
```

Input VLAN Map with inner-tag-protocol-id Statement, Output VLAN Map with Optional inner-tag-protocol-id Statement	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] input-vlan-map {   pop-swap;   inner-tag-protocol-id <i>tpid</i>; } output-vlan-map {   pop-swap;   inner-tag-protocol-id <i>tpid</i>; }</pre>
--	---

Input VLAN Map with inner-tag-protocol-id and inner-vlan-id Statements	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] input-vlan-map {   pop-swap;   inner-vlan-id <i>number</i>;   inner-tag-protocol-id <i>tpid</i>; }</pre>
---	---

## Stacking a VLAN Tag

---

To stack a VLAN tag on all tagged frames entering or exiting the interface, include the **push**, **vlan-id**, and **tag-protocol-id** statements in the input VLAN map or the output VLAN map:

```
input-vlan-map input-vlan-map {
    push;
    vlan-id number;
    tag-protocol-id tpid;
}
output-vlan-map {
    push;
    tag-protocol-id tpid;
}
```

You can include these statements at the following hierarchy levels:

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

The VLAN IDs you define in the input VLAN maps are stacked on top of the VLAN ID bound to the logical interface.

All TPIDs you include in input and output VLAN maps must be among those you specify at the [edit interfaces *interface-name* ether-options ethernet-switch-profile tag-protocol-id [ *tpids* ]] hierarchy level.

## Rewriting a VLAN Tag and Adding a New Tag

---

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and on Gigabit Ethernet and 10-Gigabit Ethernet interfaces on EX Series switches, to replace the outer VLAN tag of the incoming frame with a user-specified VLAN tag value, include the **swap-push** statement in the input VLAN map or output VLAN map:

**swap-push**

A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* [input-vlan-map](#)]
- [edit interfaces *interface-name* unit *logical-unit-number* [output-vlan-map](#)]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* [input-vlan-map](#)]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* [output-vlan-map](#)]

See *Configuring Inner and Outer TPIDs and VLAN IDs* and *Configuring Inner and Outer TPIDs and VLAN IDs* for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

**Related  
Documentation**

- [input-vlan-map on page 103](#)
- [output-vlan-map on page 123](#)
- [swap-push on page 140](#)
- *unit*
- *Ethernet Interfaces Feature Guide for Routing Devices*

## Configuring VLAN Translation with a VLAN ID List

In many cases, the VLAN identifiers on the frames of an interface's packets are not correct. VLAN translation, or VLAN rewrite, allows you to configure bidirectional VLAN identifier translation with a list on frames arriving on and leaving from a logical interface. This lets you use unique VLAN identifiers internally and maintain legacy VLAN identifiers on logical interfaces.

To perform VLAN translation on the packets on a trunk interface, insert the **vlan-rewrite** statement at the **[edit interfaces *interface-name* unit *unit-number*]** hierarchy level. You must also include the **interface-mode trunk** statement within the **[edit interfaces *interface-name* unit *unit-number* family ethernet-switching]** hierarchy because VLAN translation is only supported on trunk interfaces. The reverse translation takes place on traffic exiting the interface. In other words, if VLAN 200 is translated to 500 on traffic entering the interface, VLAN 500 is translated to VLAN 200 on traffic leaving the interface.



**NOTE:** You can configure either flexible VLAN tagging or trunk mode on interfaces. VLAN translation does not support both. Additionally, the **inner-vlan-id-list** statement is supported only on interfaces with VLAN tagging (VLAN IDs).

The following example translates incoming trunk packets from VLAN identifier 200 to 500 and 201 to 501 (other valid VLAN identifiers are not affected):

```
[edit interfaces ge-1/0/1]
unit 0 {
  ... # Other logical interface statements
  family ethernet-switching {
    interface-mode trunk # Translation is only for trunks
    inner-vlan-id-list [ 100 500–600 ];
    vlan-rewrite {
      translate 200 500;
      translate 201 501;
    }
    ... # Other ethernet-switching statements
  }
}
```



**NOTE:** This example also translates frame VLANs from 500 to 200 and 501 to 201 on egress.

### Related Documentation

- [Rewriting a VLAN Tag and Adding a New Tag on page 10](#)

## Configuring Static MAC Addresses for Logical Interfaces in a VLAN

You can manually add static MAC entries for the logical interfaces in a VLAN. You can specify one or more static MAC addresses for each logical interface.

To add a static MAC address for a logical interface in a VLAN, include the **static-mac mac-address** statement at the **[edit vlans vlan-name switch-options interface interface-name]** hierarchy level.

```
[edit]
vlans {
  vlan-name {
    domain-type bridge;
    switch-options {
      interface interface-name {
        static-mac mac-address {
          <vlan-id number>;
        }
      }
    }
  }
}
```

You can optionally specify a VLAN identifier for the static MAC address by using the **vlan-id** statement. To specify a VLAN identifier for a static MAC address, you must use the **all** option when configuring a VLAN identifier for the VLAN.



**NOTE:** If a static MAC address you configure for a logical interface appears on a different logical interface, packets sent to that interface are dropped.

### Related Documentation

- [Disabling MAC Learning for a VLAN or Logical Interface on page 21](#)
- [Configuring the Size of the MAC Address Table on page 32](#)
- [Enabling MAC Accounting for a VLAN on page 18](#)



## CHAPTER 2

# Configuring Learning and Forwarding for 802.1Q VLANs

- [Layer 2 VLANs Overview on page 15](#)
- [Layer 2 Learning and Forwarding Overview on page 16](#)
- [Layer 2 Learning and Forwarding for VLANs Overview on page 16](#)
- [Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port on page 17](#)
- [Enabling MAC Accounting on page 18](#)
- [Enabling MAC Accounting for a VLAN on page 18](#)
- [Enabling MAC Accounting for a Set of VLANs on page 18](#)
- [Verifying That MAC Accounting Is Working on page 18](#)
- [Disabling Layer 2 Learning and Forwarding on page 20](#)
- [Disabling MAC Learning for a VLAN or Logical Interface on page 21](#)
- [Disabling MAC Learning for a Set of VLANs on page 22](#)

### Layer 2 VLANs Overview

---

You can configure one or more VLANs to perform Layer 2 bridging. The Layer 2 bridging functions include integrated routing and bridging (IRB) for support for Layer 2 bridging and Layer 3 IP routing on the same interface, and virtual switches that isolate a LAN segment with its spanning-tree protocol instance and separate its VLAN ID space.

A VLAN is a set of logical ports that share the same flooding or broadcast characteristics and span one or more ports of multiple devices.

You can configure one or more VLANs to perform Layer 2 bridging. Thus, MX Series routers or EX Series switches can function as Layer 2 switches, each with multiple bridging, or broadcast, domains that participate in the same Layer 2 network. You can also configure Layer 3 routing support for a VLAN. Integrated routing and bridging (IRB) provides support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route packets to another routed interface or to another VLAN that has a Layer 3 protocol configured.

You can also group one or more VLANs within a single instance, or virtual switch. Multiple virtual switches, each of which operates independently of other virtual switches on the device, are supported. Virtual switches isolate a LAN segment with its spanning-tree protocol instance and separate its VLAN ID space. Thus, each virtual switch can participate in a different Layer 2 network.

VLANs provide support for a Layer 2 trunk port. A Layer 2 trunk interface enables you to configure a single logical interface to represent multiple VLANs on a physical interface. You can configure a set of VLANs and VLAN identifiers that are automatically associated with one or more Layer 2 trunk interfaces. Packets received on a trunk interface are forwarded within a VLAN that has the same VLAN identifier. A Layer 2 trunk interface also supports IRB within a VLAN. In addition, you can configure Layer 2 learning and forwarding properties that apply to the entire set of VLANs.

You can configure VPLS ports in a virtual switch instead of a dedicated routing instance of type **vpls** so that the logical interfaces of the Layer 2 VLANs in the virtual switch can handle VPLS routing instance traffic. Packets received on a Layer 2 trunk interface are forwarded within a VLAN that has the same VLAN identifier.

---

## Layer 2 Learning and Forwarding Overview

You can configure Layer 2 MAC address and VLAN learning and forwarding properties in support of Layer 2 bridging. Unicast media access control (MAC) addresses are learned to avoid flooding the packets to all the ports in a VLAN. A source MAC entry is created in its source and destination MAC tables for each MAC address learned from packets received on ports that belong to the VLAN.

By default, Layer 2 address learning is enabled. You can disable MAC learning for a device or for a specific VLAN or logical interfaces. You can also configure the following Layer 2 forwarding properties:

- Timeout interval for MAC entries
- MAC accounting
- A limit to the number of MAC addresses learned from the logical interfaces

For more information about how to configure VLANs and virtual switches, see [“Configuring a VLAN” on page 4](#) and [“Configuring a Layer 2 Virtual Switch” on page 33](#).

---

## Layer 2 Learning and Forwarding for VLANs Overview

When you configure a VLAN, Layer 2 address learning is enabled by default. The VLAN learns unicast media access control (MAC) addresses to avoid flooding the packets to all the ports in the VLAN. Each VLAN creates a source MAC entry in its source and destination MAC tables for each source MAC address learned from packets received on the ports that belong to the VLAN.





**NOTE:** Traffic is not flooded back onto the interface on which it was received. However, because this “split horizon” occurs at a late stage, the packet statistics displayed by commands such as `show interfaces queue` will include flood traffic.

You can optionally disable MAC learning either for the entire device or for a specific VLAN or logical interface. You can also configure the following Layer 2 learning and forwarding properties:

- Static MAC entries for logical interfaces only
- Limit to the number of MAC addresses learned from a specific logical interface or from all the logical interfaces in a VLAN
- Size of the MAC address table for the VLAN
- MAC accounting for a VLAN

**Related  
Documentation**

- [Layer 2 Learning and Forwarding Overview on page 16](#)

## Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port

Layer 2 learning is enabled by default. A set of VLANs, configured to function as a switch with a Layer 2 trunk port, learns unicast media access control (MAC) addresses to avoid flooding packets to the trunk port.



**NOTE:** Traffic is not flooded back onto the interface on which it was received. However, because this “split horizon” occurs at a late stage, the packet statistics displayed by commands such as `show interfaces queue` will include flood traffic.

You can optionally disable Layer 2 learning for the entire set of VLANs as well as modify the following Layer 2 learning and forwarding properties:

- Limit the number of MAC addresses learned from the Layer 2 trunk port associated with the set of VLANs
- Modify the size of the MAC address table for the set of VLANs
- Enable MAC accounting for the set of VLANs

**Related  
Documentation**

- [Layer 2 Learning and Forwarding Overview on page 16](#)

## Enabling MAC Accounting

---

By default, MAC accounting is disabled. You can enable packet accounting either for a router or switch as a whole or for a specific VLAN. After you enable packet accounting, the Junos OS maintains packet counters for each MAC address learned.

To enable MAC accounting, include the **global-mac-statistics** statement at the **[edit protocols l2-learning]** hierarchy level:

```
[edit protocols l2-learning]
global-mac-statistics;
```

## Enabling MAC Accounting for a VLAN

---

By default, MAC accounting is disabled. You can enable packet counting for a VLAN. When you enable packet accounting, the Junos OS maintains packet counters for each MAC address learned on the interfaces in the VLAN.

To enable MAC accounting for a VLAN, include the **mac-statistics** statement at the **[edit vlans *vlan-name* switch-options]** hierarchy level:

```
[edit vlans vlan-name switch-options]
mac-statistics;
```

## Enabling MAC Accounting for a Set of VLANs

---

By default, MAC accounting is disabled. You can enable packet counting for a set of VLANs. After you enable packet accounting, the Junos OS maintains packet counters for each MAC address learned on the trunk port associated with the set of VLANs.

To enable MAC accounting for a set of VLANs, include the **mac-statistics** statement at the **[edit switch-options]** hierarchy level:

```
[edit switch-options on page 84]
mac-statistics;
```

## Verifying That MAC Accounting Is Working

---

**Purpose** Verify that MAC accounting is enabled and the system is counting packets and collecting statistics.

**Action** 1. Verify that MAC accounting is enabled.

```
user@switch> show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
          SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)
```

```
Routing instance : default-switch
  Vlan          MAC          MAC          Age          Logical
  name          address         flags
  VLAN101      88:e0:f3:bb:07:f0  D,SE          -          ae20.0
```

```
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
```

SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN102	88:e0:f3:bb:07:f0	D,SE	-	ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned

SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN103	88:e0:f3:bb:07:f0	D,SE	-	ae20.0

[...output truncated...]

2. Display MAC accounting statistics for all VLANs associated with an interface.

user@switch> **show ethernet-switching statistics**

Local interface: ae20.0, Index: 1039

Broadcast packets:	115
Broadcast bytes :	6900
Multicast packets:	395113
Multicast bytes :	61622869
Flooded packets :	0
Flooded bytes :	0
Unicast packets :	1419
Unicast bytes :	117924
Current MAC count:	4 (Limit 8192)

[...output truncated...]

3. Display MAC accounting statistics for each address in the MAC address table.

user@switch> **show ethernet-switching table extensive**

MAC address: 88:e0:f3:bb:07:f0

Routing instance: default-switch

VLAN ID: 101

Learning interface: ae20.0

Layer 2 flags: in\_hash,in\_ifid,in\_ifl,in\_vlan,in\_rtt,acct,kernel,in\_ifbd

Epoch: 6 Sequence number: 13

Learning mask: 0x00000020

MAC address used as destination:

Packet count:	0	Byte count:	0
---------------	---	-------------	---

MAC address used as source:

Packet count:	9	Byte count:	1116
---------------	---	-------------	------

MAC address: 88:e0:f3:bb:07:f0

Routing instance: default-switch

VLAN ID: 102

Learning interface: ae20.0

Layer 2 flags: in\_hash,in\_ifid,in\_ifl,in\_vlan,in\_rtt,acct,kernel,in\_ifbd

Epoch: 6 Sequence number: 13

Learning mask: 0x00000020

MAC address used as destination:

Packet count:	0	Byte count:	0
---------------	---	-------------	---

MAC address used as source:

Packet count:	9	Byte count:	1116
---------------	---	-------------	------

MAC address: 88:e0:f3:bb:07:f0

Routing instance: default-switch

```
VLAN ID: 103
  Learning interface: ae20/0
  Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,acct,kernel,in_ifbd
  Epoch: 6                               Sequence number: 13
  Learning mask: 0x00000020
  MAC address used as destination:
  Packet count:                          0  Byte count:                          0
  MAC address used as source:
  Packet count:                          9  Byte count:                          1116
  [...output truncated...]
```

**Meaning** In the output for **show ethernet-switching table**, the MAC flag **SE** indicates that MAC accounting is enabled for VLANs 101, 102, and 103, which are all associated with the **default-switch** routing instance.

The output for **show ethernet-switching statistics** displays packet statistics and the current number of MAC addresses learned by the VLANs associated with aggregated Ethernet interface **ae20.0**.

The output for **show ethernet-switching table extensive** shows information for each address in the MAC address table. In particular, it displays the number of packets sent to and received by an interface, which is identified by a MAC address.

The output from the three commands demonstrates that MAC accounting is working properly. That is, MAC accounting is enabled on VLANs 101, 102, and 103, and as a result, you can view statistics for each of these VLANs, aggregated Ethernet interface **ae20.0**, and each MAC address.

- Related Documentation**
- [Enabling MAC Accounting on page 18](#)
  - [Enabling MAC Accounting for a VLAN on page 18](#)
  - [Enabling MAC Accounting for a Set of VLANs on page 18](#)

---

## Disabling Layer 2 Learning and Forwarding

Disabling dynamic MAC learning on an MX Series router or an EX Series switch prevents all the logical interfaces on the router or switch from learning source and destination MAC addresses.

To disable MAC learning for an MX Series router or an EX Series switch, include the **global-no-mac-learning** statement at the **[edit protocols l2-learning]** hierarchy level:

```
[edit protocols l2-learning]
global-no-mac-learning;
```

For information about how to configure a virtual switch, see *Configuring a Layer 2 Virtual Switch*.

- Related Documentation**
- *Understanding Layer 2 Learning and Forwarding*
  - *Configuring the MAC Table Timeout Interval*
  - *Enabling MAC Accounting*

- *Limiting the Number of MAC Addresses Learned from Each Logical Interface*

## Disabling MAC Learning for a VLAN or Logical Interface

You can disable MAC learning for all logical interfaces in a specified VLAN, or for a specific logical interface in a VLAN. Disabling dynamic MAC learning prevents the specified interfaces from learning source MAC addresses.

To disable MAC learning for all logical interfaces in a VLAN in a virtual switch, include the **no-mac-learning** statement at the **[edit vlans *vlan-name* switch-options]** hierarchy level:

```
[edit]
vlans {
  vlan-name {
    domain-type bridge;
    interface interface-name;
    switch-options {
      no-mac-learning;
    }
  }
}
```

To disable MAC learning for a specific logical interface in a VLAN, include the **no-mac-learning** statement at the **[edit vlans *vlan-name* switch-options interface *interface-name*]** hierarchy level.

```
[edit]
vlans {
  vlan-name {
    domain-type bridge;
    interface interface-name;
    switch-options {
      interface interface-name {
        no-mac-learning;
      }
    }
  }
}
```



**NOTE:** When you disable MAC learning, source MAC addresses are not dynamically learned, and any packets sent to these source addresses are flooded into the VLAN.



**NOTE:** When you gather interfaces into a VLAN, the **no-mac-learn-enable** statement at the **[edit interfaces *interface-name* ether-options ethernet-switch-profile]** hierarchy level is not supported. You must use the **no-mac-learning** statement at the **[edit vlans *vlan-name* switch-options interface *interface-name*]** hierarchy level to disable MAC learning on an interface in a VLAN.



**NOTE:** When MAC learning is disabled for a VPLS routing instance, traffic is not load balanced and only one of the equal-cost next hops is used.

---

## Disabling MAC Learning for a Set of VLANs

---

You can disable MAC learning for a set of VLANs. Disabling dynamic MAC learning prevents the Layer 2 trunk port associated with the set of VLANs from learning source and destination MAC addresses. When you disable MAC learning, source MAC addresses are not dynamically learned, and any packets sent to these source addresses are flooded into the switch.

To disable MAC learning for a set of VLANs, include the **no-mac-learning** statement at the **[edit switch-options]** hierarchy level:

```
[edit switch-options]  
no-mac-learning;
```

## CHAPTER 3

# Configuring VLANs and VPLS Routing Instances

- [Guidelines for Configuring VLAN Identifiers for VLANs and VPLS Routing Instances on page 23](#)
- [Configuring VLAN Identifiers for VLANs and VPLS Routing Instances on page 23](#)

## Guidelines for Configuring VLAN Identifiers for VLANs and VPLS Routing Instances

---

For a VLAN that is performing Layer 2 switching only, you do not have to specify a VLAN identifier.

For a VLAN that is performing Layer 3 IP routing, you must specify either a VLAN identifier or dual VLAN identifier tags.

For a VPLS routing instance, you must specify either a VLAN identifier or dual VLAN identifier tags.

### Related Documentation

- [Layer 2 Learning and Forwarding Overview on page 16](#)

## Configuring VLAN Identifiers for VLANs and VPLS Routing Instances

---

You can configure VLAN identifiers for a VLAN or a VPLS routing instance in the following ways:

- By using either the **vlan-id** statement or the **vlan-tags** statement to configure a normalizing VLAN identifier. This topic describes how normalizing VLAN identifiers are processed and translated in a VLAN or a VPLS routing instance.
- By using the **input-vlan-map** and the **output-vlan-map** statements at the **[edit interfaces *interface-name* unit *logic-unit-number*]** or **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logic-unit-number*]** hierarchy level to configure VLAN mapping.

The **vlan-id** and **vlan-tags** statements are used to specify the normalizing VLAN identifier under the VLAN or VPLS routing instance. The normalizing VLAN identifier is used to perform the following functions:

- Translate, or normalize, the VLAN tags of packets received into a learn VLAN identifier.
- Create multiple learning domains that each contain a learn VLAN identifier. A learning domain is a MAC address database to which MAC addresses are added based on the learn VLAN identifier.



**NOTE:** You cannot configure VLAN mapping using the **input-vlan-map** and **output-vlan-map** statements if you configure a normalizing VLAN identifier for a VLAN or VPLS routing instance using the **vlan-id** or **vlan-tags** statements.

To configure a VLAN identifier for a VLAN, include either the **vlan-id** or the **vlan-tags** statement at the **[edit interfaces *interface-name* unit *logic-unit-number*]** or **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logic-unit-number*]** hierarchy level, and then include that logical interface in the VLAN configuration.

For a VPLS routing instance, include either the **vlan-id** or **vlan-tags** statement at the **[edit interfaces *interface-name* unit *logic-unit-number*]** or **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logic-unit-number*]** hierarchy level, and then include that logical interface in the VPLS routing instance configuration.



**NOTE:** For a single VLAN or VPLS routing instance, you can include either the **vlan-id** or the **vlan-tags** statement, but not both. If you do not configure a **vlan-id**, **vlan-tags**, or **vlan-id-list [ *vlan-id-numbers* ]** for the VLAN or the VPLS routing instance, the Layer 2 packets received are forwarded to the outbound Layer 2 interface without having the VLAN tag modified unless an **output-vlan-map** is configured on the Layer 2 interface. This results in a frame being forwarded to a Layer 2 interface with a VLAN tag that is different from what is configured for the Layer 2 interface. Note that a frame received from the Layer 2 interface is still required to match the VLAN tag(s) specified in the interface configuration. The invalid configuration may cause a Layer 2 loop to occur.

The VLAN tags associated with the inbound logical interface are compared with the normalizing VLAN identifier. If the tags are different, they are rewritten as described in [Table 5 on page 27](#). The source MAC address of a received packet is learned based on the normalizing VLAN identifier.



**NOTE:** You do not have to specify a VLAN identifier for a VLAN that is performing Layer 2 switching only. To support Layer 3 IP routing, you must specify either a VLAN identifier or a pair of VLAN tags. However, you cannot specify the same VLAN identifier for more than one VLAN within a routing instance. Each VLAN must have a unique VLAN identifier.



If the VLAN tags associated with the outbound logical interface and the normalizing VLAN identifier are different, the normalizing VLAN identifier is rewritten to match the VLAN tags of the outbound logical interface, as described in [Table 6 on page 28](#).

For the packets sent over the VPLS routing instance to be tagged by the normalizing VLAN identifier, include one of the following configuration statements:

- **vlan-id *number*** to tag all packets that are sent over the VPLS virtual tunnel (VT) interfaces with the VLAN identifier.
- **vlan-tags outer *number* inner *number*** to tag all packets sent over the VPLS VT interfaces with dual outer and inner VLAN tags.

Use the **vlan-id none** statement to have the VLAN tags removed from packets associated with an inbound logical interface when those packets are sent over VPLS VT interfaces. Note that those packets might still be sent with other customer VLAN tags.

The **vlan-id all** statement enables you to configure bridging for several VLANs with a minimum amount of configuration. Configuring this statement creates a learning domain for:

- Each inner VLAN, or learn VLAN, identifier of a logical interface configured with two VLAN tags
- Each VLAN, or learn VLAN, identifier of a logical interface configured with one VLAN tag

We recommend that you do not use customer VLAN IDs in a VPLS routing instance because customer VLAN IDs are used for learning only.

You should use the service VLAN ID in a VPLS routing instance, as in the following configuration:

```
[edit]
interface ge-1/1/1 {
  vlan-tagging;
  unit 1 {
    vlan-id s1; /* Service vlan */
    encapsulation vlan-vpls;
    input-vlan-map pop; /* Pop the service vlan on input */
    output-vlan-map push; /* Push the service vlan on output */
  }
}
interface ge-1/1/2 {
  encapsulation ethernet-vpls;
  unit 0;
}
routing-instances {
  V1 {
    instance-type vpls;
    vlan-id all;
    interface ge-1/1/1.1;
    interface ge-1/1/2.0;
  }
}
```



**NOTE:** If you configure the `vlan-id all` statement in a VPLS routing instance, we recommend using the `input-vlan-map pop` and `output-vlan-map push` statements on the logical interface to pop the service VLAN ID on input and push the service VLAN ID on output and in this way, limit the impact of double-tagged frames on scaling. You cannot use the native `vlan-id` statement when the `vlan-id all` statement is included in the configuration.

The `vlan-id-list [ vlan-id-numbers ]` statement enables you to configure bridging for multiple VLANs on a trunk interface. Configuring this statement creates a learning domain for:

- Each VLAN listed: `vlan-id-list [ 100 200 300 ]`
- Each VLAN in a range: `vlan-id-list [ 100-200 ]`
- Each VLAN in a list and range combination: `vlan-id-list [ 50, 100-200, 300 ]`

The following steps outline the process for bridging a packet received over a Layer 2 logical interface when you specify a normalizing VLAN identifier using either the `vlan-id number` or `vlan-tags` statement for a VLAN or a VPLS routing instance:

1. When a packet is received on a physical port, it is accepted only if the VLAN identifier of the packet matches the VLAN identifier of one of the logical interfaces configured on that port.
2. The VLAN tags of the received packet are then compared with the normalizing VLAN identifier. If the VLAN tags of the packet are different from the normalizing VLAN identifier, the VLAN tags are rewritten as described in [Table 5 on page 27](#).
3. If the source MAC address of the received packet is not present in the source MAC table, it is learned based on the normalizing VLAN identifier.
4. The packet is then forwarded toward one or more outbound Layer 2 logical interfaces based on the destination MAC address. A packet with a known unicast destination MAC address is forwarded only to one outbound logical interface. For each outbound Layer 2 logical interface, the normalizing VLAN identifier configured for the VLAN or VPLS routing instance is compared with the VLAN tags configured on that logical interface. If the VLAN tags associated with an outbound logical interface do not match the normalizing VLAN identifier configured for the VLAN or VPLS routing instance, the VLAN tags are rewritten as described in [Table 6 on page 28](#).

The tables below show how VLAN tags are applied for traffic sent to and from the VLAN, depending on how the `vlan-id` and `vlan-tags` statements are configured for the VLAN and on how identifiers are configured for the logical interfaces in a VLAN or VPLS routing instance. Depending on your configuration, the following rewrite operations are performed on VLAN tags:

- **pop**—Remove a VLAN tag from the top of the VLAN tag stack.
- **pop-pop**—Remove both the outer and inner VLAN tags of the frame.
- **pop-swap**—Remove the outer VLAN tag of the frame and replace the inner VLAN tag of the frame.

- **swap**—Replace the VLAN tag of the frame.
- **push**—Add a new VLAN tag to the top of the VLAN stack.
- **push-push**—Push two VLAN tags in front of the frame.
- **swap-push**—Replace the VLAN tag of the frame and add a new VLAN tag to the top of the VLAN stack.
- **swap-swap**—Replace both the outer and inner VLAN tags of the frame.

Table 5 on page 27 shows specific examples of how the VLAN tags for packets sent to the VLAN are processed and translated, depending on your configuration. “–” means that the statement is not supported for the specified logical interface VLAN identifier. “No operation” means that the VLAN tags of the received packet are not translated for the specified input logical interface.

**Table 5: Statement Usage and Input Rewrite Operations for VLAN Identifiers for a VLAN**

VLAN Identifier of Logical Interface	VLAN Configurations for a VLAN			
	vlan-id none	vlan-id 200	vlan-id all	vlan tags outer 100 inner 300
<b>none</b>	No operation	push 200	–	push 100, push 300
<b>200</b>	pop 200	No operation	No operation	swap 200 to 300, push 100
<b>1000</b>	pop 1000	swap 1000 to 200	No operation	swap 1000 to 300, push 100
<b>vlan-tags outer 2000 inner 300</b>	pop 2000, pop 300	pop 2000, swap 300 to 200	pop 2000	swap 2000 to 100
<b>vlan-tags outer 100 inner 400</b>	pop 100, pop 400	pop 100, swap 400 to 200	pop 100	swap 400 to 300
<b>vlan-id-range 10-100</b>	–	–	No operation	–
<b>vlan-tags outer 200 inner-range 10-100</b>	–	–	pop 200	–

Table 6 on page 28 shows specific examples of how the VLAN tags for packets sent from the VLAN are processed and translated, depending on your configuration. “–” means that the statement is not supported for the specified logical interface VLAN identifier. “No operation” means that the VLAN tags of the outbound packet are not translated for the specified output logical interface.

Table 6: Statement Usage and Output Rewrite Operations for VLAN Identifiers for a VLAN

VLAN Identifier of Logical Interface	VLAN Configurations for a VLAN			
	vlan-id none	vlan-id 200	vlan-id all	vlan tags outer 100 inner 300
none	no operation	pop 200	–	pop 100, pop 300
200	push 200	No operation	No operation	pop 100, swap 300 to 200
1000	push 1000	swap 200 to 1000	No operation	pop 100, swap 300 to 1000
vlan-tags outer 2000 inner 300	push 2000, push 300	swap 200 to 300, push 2000	push 2000	swap 100 to 2000
vlan-tags outer 100 inner 400	push 100, push 400	swap 200 to 400, push 100	push 100	swap 300 to 400
vlan-id-range 10-100	–	–	No operation	–
vlan-tags outer 200 inner-range 10-100	–	–	push 200	–

## CHAPTER 4

# Configuring Bridging and Integrated Routing and Bridging (IRB)

- [Configuring Integrated Routing and Bridging for VLANs on page 29](#)
- [Configuring a Set of VLANs to Act as a Switch for a Layer 2 Trunk Port on page 30](#)
- [Configuring a Logical Interface for Access Mode on page 32](#)
- [Configuring the Size of the MAC Address Table on page 32](#)
- [Configuring a Layer 2 Virtual Switch on page 33](#)
- [Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port on page 34](#)

## Configuring Integrated Routing and Bridging for VLANs

---

Integrated routing and bridging (IRB) provides simultaneous support for Layer 2 bridging and Layer 3 routing on the same interface. IRB enables you to route packets to another routed interface or to another VLAN that has an IRB interface configured. You configure a logical routing interface by specifying **irb** as an interface name at the **[edit interfaces]** hierarchy level and including that interface in the VLAN.



**NOTE:** You can include only one Layer 3 interface in a VLAN.

To configure a VLAN with IRB support, include the following statements:

```
[edit]
vlands {
  vlan-name {
    domain-type bridge;
    interface interface-name;
    l3-interface interface-name;
    vlan-id (none | number);
    vlan-tags outer number inner number;
  }
}
```

For each VLAN that you configure, specify a **vlan-name**. You must also specify the value **bridge** for the **domain-type** statement.

For the **vlan-id** statement, you can specify either a valid VLAN identifier or the **none** option.



**NOTE:** If you configure a Layer 3 interface to support IRB in a VLAN, you cannot use the `all` option for the `vlan-id` statement.

The `vlan-tags` statement enables you to specify a pair of VLAN identifiers; an **outer** tag and an **inner** tag.



**NOTE:** For a single VLAN, you can include either the `vlan-id` statement or the `vlan-tags` statement, but not both.

To include one or more logical interfaces in the VLAN, specify the *interface-name* for each Ethernet interface to include that you configured at the `[edit interfaces]` hierarchy level.



**NOTE:** A maximum of 4096 active logical interfaces are supported for a VLAN or on each mesh group in a VPLS routing instance configured for Layer 2 bridging.

To associate a Layer 3 interface with a VLAN, include the `l3-interface interface-name` statement and specify an *interface-name* you configured at the `[edit interfaces irb]` hierarchy level. You can configure only one Layer 3 interface for each VLAN.

IRB interfaces are supported for multicast snooping.

In multihomed VPLS configurations, you can configure VPLS to keep a VPLS connection up if only an IRB interface is available by configuring the `irb` option for the `connectivity-type` statement at the `[edit routing-instances routing-instance-name protocols vpls]` hierarchy level. The `connectivity-type` statement has the `ce` and `irb` options. The `ce` option is the default and specifies that a CE interface is required to maintain the VPLS connection. By default, if only an IRB interface is available, the VPLS connection is brought down.



**NOTE:** When you configure IRB interfaces in more than one logical system on a device, all of the IRB logical interfaces share the same MAC address.

## Configuring a Set of VLANs to Act as a Switch for a Layer 2 Trunk Port

You can configure a set of VLANs that are associated with a Layer 2 trunk port. The set of VLANs function as a switch. Packets received on a trunk interface are forwarded within a VLAN that has the same VLAN identifier. A trunk interface also provides support for IRB, which provides support for Layer 2 bridging and Layer 3 IP routing on the same interface.

To configure a Layer 2 trunk port and set of VLANs, include the following statements:

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

```

    unit number {
        family ethernet-switching {
            interface-mode access;
            vlan-members (vlan-name | vlan-tag);
        }
    }
}
interface-name {
    native-vlan-id number;
    unit number {
        family ethernet-switching {
            interface-mode trunk;
            vlan-members (vlan-name | vlan-tag);
        }
    }
}
[edit vlans ]
vlan-name {
    vlan-id number;
    vlan-id-list [ vlan-id-numbers ];
    ....
}

```

You must configure a VLAN and VLAN identifier for each VLAN associated with the trunk interface. You can configure one or more trunk or access interfaces at the **[edit interfaces]** hierarchy level. An access interface enables you to accept packets with no VLAN identifier.

## Configuring a Logical Interface for Access Mode

---

Enterprise network administrators can configure a single logical interface to accept untagged packets and forward the packets within a specified VLAN. A logical interface configured to accept untagged packets is called an *access interface* or *access port*.

`interface-mode access;`

You can include this statement at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number family ethernet-switching]`
- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family ethernet-switching]`

When an untagged or tagged packet is received on an access interface, the packet is accepted, the VLAN ID is added to the packet, and the packet is forwarded within the VLAN that is configured with the matching VLAN ID.

The following example configures a logical interface as an access port with a VLAN ID of 20 on routers and switches that support the enhanced Layer 2 software:

```
[edit interfaces ge-1/2/0]
unit 1 {
  family ethernet-switching {
    interface-mode access;
    vlan members 20;
  }
}
```

### Related Documentation

- [802.1Q VLANs Overview](#)
- [Ethernet Interfaces Feature Guide for Routing Devices](#)

## Configuring the Size of the MAC Address Table

---

You can modify the size of the MAC address table for each VLAN. The default table size is 5120 addresses. The minimum you can configure is 16 addresses, and the maximum is 1,048,575 addresses.

If the MAC table limit is reached, new addresses can no longer be added to the table. Unused MAC addresses are removed from the MAC address table automatically. This frees space in the table, allowing new entries to be added.

To modify the size of the MAC table, include the `mac-table-size limit` statement at the `[edit vlans vlan-name switch-options]` hierarchy level:

```
[edit]
vlans {
  vlan-name {
    domain-type bridge;
    switch-options {
      mac-table-size limit {
```



```

        packet-action drop;
    }
}
}

```

#### Related Documentation

- *Disabling MAC Learning for a Bridge Domain or Logical Interface*
- *Configuring Static MAC Addresses for Logical Interfaces in a Bridge Domain*
- *Limiting MAC Addresses Learned from an Interface in a Bridge Domain*
- *Enabling MAC Accounting for a Bridge Domain*

## Configuring a Layer 2 Virtual Switch

A Layer 2 virtual switch, which isolates a LAN segment with its spanning-tree protocol instance and separates its VLAN ID space, filters and forwards traffic only at the data link layer. Each VLAN consists of a set of logical ports that participate in Layer 2 learning and forwarding. A virtual switch represents a Layer 2 network.

Two main types of interfaces are used in virtual switch hierarchies:

- Layer 2 logical interface—This type of interface uses the VLAN-ID as a virtual circuit identifier and the scope of the VLAN-ID is local to the interface port. This type of interface is often used in service-provider-centric applications.
- Access or trunk interface—This type of interface uses a VLAN-ID with global significance. The access or trunk interface is implicitly associated with VLANs based on VLAN membership. Access or trunk interfaces are typically used in enterprise-centric applications.



**NOTE:** The difference between access interfaces and trunk interfaces is that access interfaces can be part of one VLAN only and the interface is normally attached to an end-user device (packets are implicitly associated with the configured VLAN). In contrast, trunk interfaces multiplex traffic from multiple VLANs and usually interconnect switches.

To configure a Layer 2 virtual switch, include the following statements:

```

[edit]
routing-instances {
    routing-instance-name (
        instance-type virtual-switch;
        vlans vlan-name{
            vlan-id (all | none | number);
            [...configure optional VLAN parameters]
        }
    }
}

```

To enable a virtual switch, you must specify **virtual-switch** as the **instance-type**.

The VLANs that are specified with the **vlan-id** statement are included in the virtual switch.

You can configure other optional VLAN parameters in the virtual switch.

**Related  
Documentation**

- [Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port on page 34](#)
- [Configuring a Layer 2 Virtual Switch](#)

---

## Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port

---

You can associate one or more Layer 2 trunk interfaces with a virtual switch.

A virtual switch configured with a Layer 2 trunk port also supports IRB within a VLAN. IRB provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. Only an interface configured with the **interface-mode (access | trunk)** statement can be associated with a virtual switch. An access interface enables you to accept packets with no VLAN identifier.

In addition, you can configure Layer 2 learning and forwarding properties for the virtual switch.

To configure a virtual switch with a Layer 2 trunk interface, include the following statements:

```
[edit]
routing-instances {
  routing-instance-name {
    instance-type virtual-switch;
    interface interface-name;
    vlans name {
      vlan-id (all | none | number);
      [...configure optional VLAN parameters]
    }
  }
}
```

**Related  
Documentation**

- [Configuring a Layer 2 Virtual Switch on page 33](#)

## CHAPTER 5

# Configuring Private VLANs

- [Understanding Private VLANs on page 35](#)
- [Example: Configuring a Private VLAN on a Single Switch on page 41](#)
- [Creating a Private VLAN on a Single Switch \(CLI Procedure\) on page 45](#)
- [Creating a Private VLAN Spanning Multiple Switches \(CLI Procedure\) on page 47](#)

## Understanding Private VLANs

---



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

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 is usually connected to a router, firewall, server, or provider network.

PVLANS provide Layer 2 isolation between ports within a VLAN, splitting a broadcast domain into multiple discrete broadcast subdomains by creating secondary VLANs (community VLANs and an isolated VLAN) inside a primary VLAN. Ports within the same community VLAN can communicate with each other. Ports within an isolated VLAN can communicate *only* with a single uplink port.

PVLANS are useful for restricting the flow of broadcast and unknown unicast traffic and for limiting the communication between hosts. The need to segregate a single VLAN is particularly useful in the following deployment scenarios:

- **Server farms**—A typical Internet service provider uses a server farm to provide Web hosting for numerous customers. Locating the various servers within a single server farm provides ease of management. Security concerns arise if all servers are in the same VLAN because Layer 2 broadcasts go to all servers in the VLAN.

- Metropolitan Ethernet networks—A metro service provider offers Layer 2 Ethernet access to assorted homes, rental communities, and businesses. The traditional solution of deploying one VLAN per customer is not scalable and is difficult to manage, leading to potential waste of IP addresses. PVLANS provide a more secure and more efficient solution.

This topic explains the following concepts regarding PVLANS:

- [Typical Structure and Primary Application of PVLANS on page 36](#)
- [Routing Between Isolated and Community VLANs on page 38](#)
- [PVLANS Use 802.1Q Tags to Identify Packets on page 38](#)
- [PVLANS Use IP Addresses Efficiently on page 39](#)
- [PVLAN Port Types and Forwarding Rules on page 39](#)
- [Guidelines and Restrictions for PVLANS on page 39](#)
- [Limitation of Private VLANs on page 40](#)

## Typical Structure and Primary Application of PVLANS

A PVLAN can be configured on a single switch or can be configured to span multiple switches. The types of domains and ports are:

- 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—A primary VLAN can contain only one isolated VLAN. An interface within an isolated VLAN can forward packets only to a promiscuous port or the Inter-Switch Link (ISL) port. An isolated interface cannot forward packets to another isolated interface; and an isolated interface cannot receive packets from another isolated interface. If a customer device needs to have access *only* to a gateway router, the device must be attached to an isolated trunk port.
- Community 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 ISL 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 any 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.
- Inter-Switch Link (ISL)—An ISL is a trunk port that connects multiple switches in a PVLAN and is required only when a PVLAN spans multiple switches..

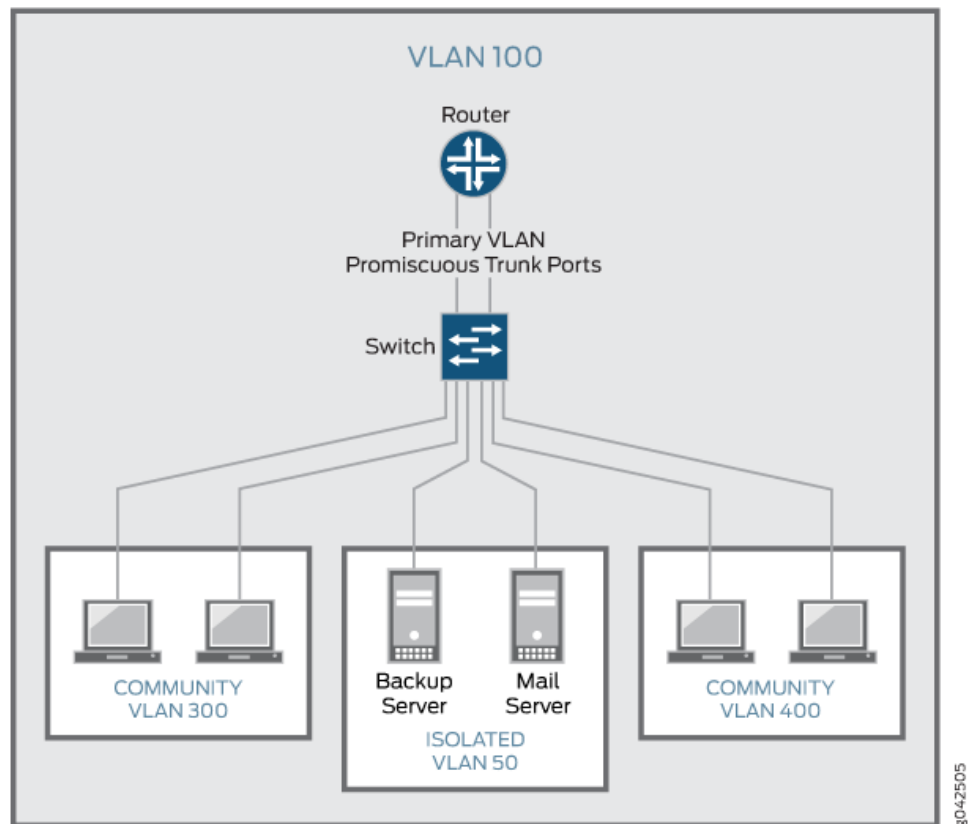


**NOTE:** The primary VLAN of the PVLAN is defined with an 802.1Q tag (VLAN ID) for the complete PVLAN. On EX9200 switches, each secondary VLAN must also be defined with its own separate VLAN ID.

For information on configuring an isolated VLAN or community VLANs, see [“Creating a Private VLAN on a Single Switch \(CLI Procedure\)”](#) on page 45 or [“Creating a Private VLAN Spanning Multiple Switches \(CLI Procedure\)”](#) on page 47.

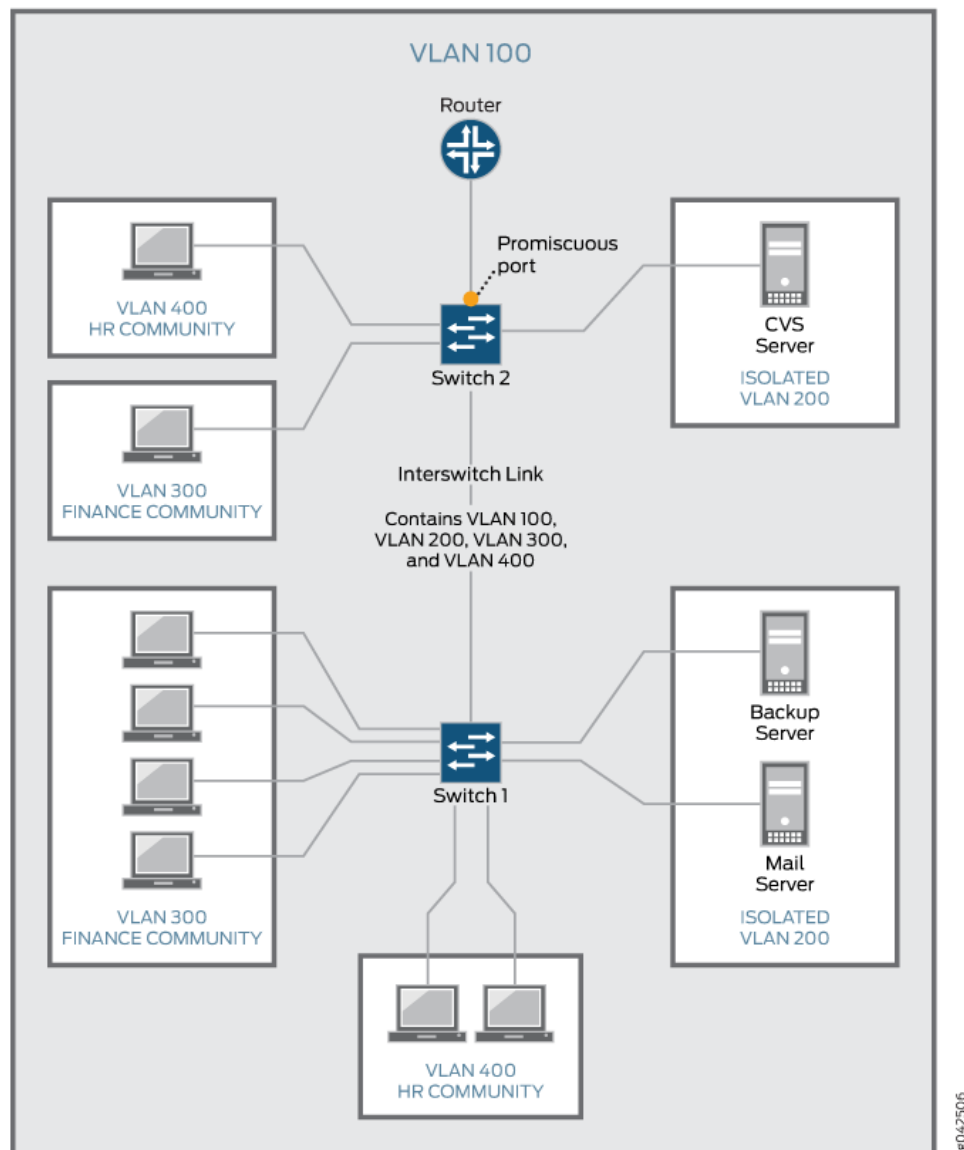
[Figure 1 on page 37](#) shows a PVLAN on a single switch in which the primary VLAN (100) contains two community VLANs and one isolated VLAN.

**Figure 1: Private VLAN on a Single Switch**



[Figure 2 on page 38](#) shows a PVLAN spanning multiple switches in which the primary VLAN (100) contains two community VLANs (300) and (400) and one isolated VLAN (200).

### Figure 2: PVLAN Spanning Multiple Switches



## Routing Between Isolated and Community VLANs

To route Layer 3 traffic between isolated and community VLANs, you must connect an external router or switch to a trunk port of the primary VLAN. The trunk port of the primary VLAN is a *promiscuous* port; therefore, it can communicate with *all* the ports in the PVLAN.

## PVLANS Use 802.1Q Tags to Identify Packets

You must specify an 802.1Q tag on the primary VLAN and on each of the secondary VLANs, regardless of whether the PVLAN is configured on a single switch or across

multiple switches. (You can also configure a VLAN name for the primary VLAN and for each of the secondary VLANs. This is optional.)

## 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.

## PVLAN Port Types and Forwarding Rules

Table 7 on page 39 summarizes Layer 2 connectivity between the different types of ports within a PVLAN.

**Table 7: PVLAN Ports and Layer 2 Forwarding**

From Port Type	To Isolated Ports?	To Promiscuous Ports?	To Community Ports?	To Inter-switch Link Port?
Isolated	Deny	Permit	Deny	Permit
Promiscuous	Permit	Permit	Permit	Permit
Community 1	Deny	Permit	Permit	Permit

## Guidelines and Restrictions for PVLANS

When configuring PVLANS, adhere to the following guidelines:

- An access interface can belong to only one PVLAN domain, that is, it cannot participate in two different primary VLANs..
- A trunk interface can be a member of two secondary VLANs as long as the secondary VLANs are in two *different* primary VLANs. A trunk interface cannot be a member of two secondary VLANs that are in the *same* primary VLAN.
- A single region of Multiple Spanning Tree Protocol (MSTP) must be configured on all VLANs that are included within the PVLAN.
- VLAN Spanning Tree Protocol (VSTP) is not supported.

Some configuration statements cannot be specified on a secondary VLAN. You can configure the following statements at the `[edit vlans vlan-name switch-options]` hierarchy level *only* on the primary PVLAN:

- `mac-table-size`
- `no-mac-learning`
- `mac-statistics`
- `interface-mac-limit`

The following features are *not* supported on PVLANS on Junos switches with support for the ELS configuration style:

- DHCP security features (DHCP snooping, dynamic ARP inspection, IP source guard)
- Egress VLAN firewall filters
- Ethernet ring protection (ERP)
- Flexible VLAN tagging
- [global-mac-statistics](#)
- Integrated routing and bridging (IRB) interface
- Multicast snooping or IGMP snooping
- Multichassis link aggregation groups (MC-LAGs)
- Port mirroring
- Q-in-Q tunneling
- VLAN Spanning Tree Protocol (VSTP)
- Voice over IP (VoIP)

You can configure the following statements at the **[edit vlans *vlan-name* switch-options]** hierarchy level only on the primary PVLAN:

- [mac-table-size](#)
- [no-mac-learning](#)
- [interface-mac-limit](#)

## Limitation of Private VLANs

If you want to change a primary VLAN to be a secondary VLAN, you must first change it to a normal VLAN and commit the change. For example, you would follow this procedure:

1. Change the primary VLAN to be a normal VLAN.
2. Commit the configuration.
3. Change the normal VLAN to be a secondary VLAN.
4. Commit the configuration.

Follow the same sequence of commits if you want to change a secondary VLAN to be a primary VLAN. That is, make the secondary VLAN a normal VLAN and commit that change and then change the normal VLAN to be a primary VLAN.

### Related Documentation

- [Example: Configuring a Private VLAN on a Single Switch on page 41](#)
- [Creating a Private VLAN on a Single Switch \(CLI Procedure\) on page 45](#)
- [Creating a Private VLAN Spanning Multiple Switches \(CLI Procedure\) on page 47](#)



## Example: Configuring a Private VLAN on a Single Switch



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



**NOTE:** Private VLANs are not supported on QFX5100 switches and QFX10002 switches running Junos OS Release 15.1X53.

For security reasons, it is often useful to restrict the flow of broadcast and unknown unicast traffic or limit the communication between known hosts. Private VLANs (PVLANS) enable you to split a broadcast domain (primary VLAN) into multiple isolated broadcast subdomains (secondary VLANs), essentially putting a VLAN inside a VLAN.

This example describes how to create a PVLAN on a single switch:

- [Requirements on page 41](#)
- [Overview and Topology on page 41](#)
- [Configuration on page 43](#)
- [Verification on page 44](#)

### Requirements

This example uses the following hardware and software components:

- One Junos OS switch
- Junos OS Release 14.1X53-D10 or later for EX Series switches  
Junos OS Release 14.1X53-D15 or later for QFX Series switches

### Overview and Topology

You can isolate groups of subscribers for improved security and efficiency. This configuration example uses a simple topology to illustrate how to create a PVLAN with one primary VLAN and three secondary VLANs (one isolated VLAN, and two community VLANs).

[Table 8 on page 41](#) lists the interfaces of the topology used in the example.

**Table 8: Interfaces of the Topology for Configuring a PVLAN**

Interface	Description
ge-0/0/0	Promiscuous member ports
ge-1/0/0	

Table 8: Interfaces of the Topology for Configuring a PVLAN (*continued*)

Interface	Description
ge-0/0/11, ge-0/0/12	HR community VLAN member ports
ge-0/0/13, ge-0/0/14	Finance community VLAN member ports
ge-0/0/15, ge-0/0/16	Isolated member ports


Table 9 on page 42 lists the VLAN IDs of the topology used in the example.

Table 9: VLAN IDs in the Topology for Configuring a PVLAN

VLAN ID	Description
100	Primary VLAN
200	HR community VLAN
300	Finance community VLAN
400	Isolated VLAN

Figure 3 on page 42 shows the topology for this example.

Figure 3: Topology of a Private VLAN on a Single EX Series Switch



## Configuration

You can use an existing VLAN as the basis for your private PVLAN and create subdomains within it. This example creates a primary VLAN—using the VLAN name **vlan-pri**—as part of the procedure.

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 vlan-pri vlan-id 100
set vlans vlan-iso private-vlan isolated vlan-id 400
set vlans vlan-hr private-vlan community vlan-id 200
set vlans vlan-finance private-vlan community vlan-id 300
set vlans vlan-pri vlan-id 100 isolated-vlan vlan-iso community-vlan vlan-hr community-vlan
vlan-finance
set interface ge-0/0/15 unit 0 family ethernet-switching interface-mode access vlan members
vlan-iso
set interface ge-0/0/16 unit 0 family ethernet-switching interface-mode access vlan members
vlan-iso
set interface ge-0/0/11 unit 0 family ethernet-switching interface-mode access vlan members
vlan-hr
set interface ge-0/0/12 unit 0 family ethernet-switching interface-mode trunk vlan members
vlan-hr
set interface ge-0/0/13 unit 0 family ethernet-switching interface-mode access vlan members
vlan-finance
set interface ge-0/0/14 unit 0 family ethernet-switching interface-mode trunk vlan members
vlan-finance
set interface ge-0/0/0 unit 0 family ethernet-switching interface-mode trunk vlan members
vlan-pri
set interface ge-1/0/0 unit 0 family ethernet-switching interface-mode trunk vlan members
vlan-pri
```

**Step-by-Step Procedure** To configure the PVLAN:

1. Create the primary VLAN (in this example, the name is **vlan-pri**) of the private VLAN:  

```
[edit vlans]
user@switch# set vlan-pri vlan-id 100
```
2. Create an isolated VLAN and assign it a VLAN ID:  

```
[edit vlans]
user@switch# set vlan-iso private-vlan isolated vlan-id 400
```
3. Create the HR community VLAN and assign it a VLAN ID:  

```
[edit vlans]
user@switch# set vlan-hr private-vlan community vlan-id 200
```
4. Create the finance community VLAN and assign it a VLAN ID:  

```
[edit vlans]
user@switch# set vlan-finance private-vlan community vlan-id 300
```
5. Associate the secondary VLANs with the primary VLAN:  

```
[edit vlans]
user@switch# set vlan-pri vlan-id 100 isolated-vlan vlan-iso community-vlan vlan-hr
community-vlan vlan-finance
```
6. Set the interfaces to the appropriate interface modes:

**[edit interfaces]**

```
user@switch# set ge-0/0/15 unit 0 family ethernet-switching interface-mode access vlan members vlan-iso
```

```
user@switch# set ge-0/0/16 unit 0 family ethernet-switching interface-mode trunk vlan members vlan-iso
```

```
user@switch# set ge-0/0/11 unit 0 family ethernet-switching interface-mode access vlan members vlan-hr
```

```
user@switch# set ge-0/0/12 unit 0 family ethernet-switching interface-mode access vlan members vlan-hr
```

```
user@switch# set ge-0/0/13 unit 0 family ethernet-switching interface-mode access vlan members vlan-finance
```

```
user@switch# set ge-0/0/14 unit 0 family ethernet-switching interface-mode trunk vlan members vlan-finance
```

7. Configure a promiscuous trunk interface of the primary VLAN. This interface is used by the primary VLAN to communicate with the secondary VLANs.

```
user@switch# set ge-0/0/0 unit 0 family ethernet-switching interface-mode trunk vlan members vlan-pri
```

8. Configure another trunk interface (it is also a promiscuous interface) of the primary VLAN, connecting the PVLAN to the router.

```
user@switch# set ge-1/0/0 unit 0 family ethernet-switching interface-mode trunk vlan members vlan-pri
```

---

## Results

Check the results of the configuration:

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

## Verification

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

- [Verifying That the Private VLAN and Secondary VLANs Were Created on page 44](#)

---

### Verifying That the Private VLAN and Secondary VLANs Were Created

<b>Purpose</b>	Verify that the primary VLAN and secondary VLANs were properly created on the switch.
<b>Action</b>	Use the <b>show vlans</b> command: <pre>user@switch&gt; show vlans extensive</pre>
<b>Meaning</b>	The output shows that the primary VLAN was created and identifies the interfaces and secondary VLANs associated with it.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Private VLANs on page 35</a></li><li>• <a href="#">Creating a Private VLAN on a Single Switch (CLI Procedure) on page 45</a></li><li>• <a href="#">Creating a Private VLAN Spanning Multiple Switches (CLI Procedure) on page 47</a></li></ul>

## Creating a Private VLAN on a Single Switch (CLI Procedure)



**NOTE:** This task uses Junos OS for switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your EX Series switch runs software that does not support ELS, see *Creating a Private VLAN on a Single EX Series Switch (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.



**NOTE:** Private VLANs are not supported on QFX5100 switches and QFX10002 switches running Junos OS Release 15.1X53.

For security reasons, it is often useful to restrict the flow of broadcast and unknown unicast traffic or limit the communication between known hosts. Private VLANs (PVLANS) enable you to split a broadcast domain (primary VLAN) into multiple isolated broadcast subdomains (secondary VLANs), essentially putting a VLAN inside a VLAN. This procedure describes how to create a PVLAN on a single switch.



**NOTE:** You must specify a VLAN ID for each secondary VLAN even if the PVLAN is configured on a single switch.

You do not need to preconfigure the primary VLAN. This topic shows the primary VLAN being configured as part of this PVLAN configuration procedure.

For a list of guidelines on configuring PVLANS, see [“Understanding Private VLANs” on page 35](#).

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. Configure at least one interface within the primary VLAN so that it communicates with all the subdomains of the PVLAN. This interface functions as a *promiscuous* port. It can be either a trunk port or an access port.

**[edit interfaces]**

```
user@switch# set interface-name unit 0 family ethernet-switching
```

```
user@switch# set interface-name unit 0 family ethernet-switching vlan members  
primary-vlan-name
```

3. Configure another promiscuous interface of the primary VLAN as a trunk port to connect the PVLAN to the external router or switch:

**[edit interfaces]**

```
user@switch# set interface-name unit 0 family ethernet-switching interface-mode trunk
```

```
user@switch# set interface-name unit 0 family ethernet-switching vlan members  
primary-vlan-name
```

4. Create an isolated VLAN by selecting the **isolated** option for **private-vlan**, and setting a VLAN ID for the isolated VLAN:

[edit vlans]

```
user@switch# set isolated-vlan-name private-vlan isolated vlan-id isolated-vlan-id
```



**NOTE:** You can create only one isolated VLAN within a private VLAN. Setting the VLAN name for the isolated VLAN is optional. Configuring the VLAN ID is required.

5. Create a community VLAN by selecting the **community** option for **private-vlan**, and setting a VLAN ID for this community VLAN:

[edit vlans]

```
user@switch# set community-vlan-name private-vlan community vlan-id community-vlan-id
```



**NOTE:** To create additional community VLANs, repeat this step and specify a different name for the community VLAN. Setting the VLAN name for the community VLAN is optional. Configuring the VLAN ID is required.

6. Associate the isolated VLAN with the primary VLAN:

[edit vlans]

```
user@switch# set primary-vlan-name vlan-id primary-vlan-id isolated-vlan isolated-vlan-name
```

7. Associate each community VLAN with the primary VLAN:

[edit vlans]

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

8. If you have not already done so, configure at least one interface of the isolated VLAN.

[edit interfaces]

```
user@switch# set interface-name unit logical-unit-number family ethernet-switching interface-mode access vlan members isolated-vlan-name
```

9. If you have not already done so, configure at least one interface of the community VLAN.

[edit interfaces]

```
user@switch# set interface-name unit logical-unit-number family ethernet-switching interface-mode access vlan members community-vlan-name
```



**NOTE:** Repeat the same step on other community VLANs that you want to include in the PVLAN.

#### Related Documentation

- [Understanding Private VLANs on page 35](#)
- [Creating a Private VLAN Spanning Multiple Switches \(CLI Procedure\) on page 47](#)

## Creating a Private VLAN Spanning Multiple Switches (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Creating a Private VLAN Spanning Multiple EX Series Switches (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.



**NOTE:** Private VLANs are not supported on QFX5100 switches and QFX10002 switches running Junos OS Release 15.1X53.

For security reasons, it is often useful to restrict the flow of broadcast and unknown unicast traffic or limit the communication between known hosts. Private VLANs (PVLANS) enable you to split a broadcast domain (primary VLAN) into multiple isolated broadcast subdomains (secondary VLANs), essentially putting a VLAN inside a VLAN. This procedure describes how to configure a PVLAN to span multiple switches.

For a list of guidelines on configuring PVLANS, see [“Understanding Private VLANs” on page 35](#).

To configure a PVLAN to span multiple switches, perform the following procedure on all the switches that will participate in the PVLAN::

1. Create the primary VLAN by setting the unique VLAN name and specify an 802.1Q tag for the VLAN:

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

2. On the switch that will connect to a router, configure a promiscuous interface as a trunk port to connect the PVLAN to the router:

```
[edit interfaces]
user@switch# set interface-name unit 0 family ethernet-switching interface-mode trunk
user@switch# set interface-name unit 0 family ethernet-switching vlan members
primary-vlan-name
```

3. On all the switches, configure a trunk interface as the Inter-Switch Link (ISL) that will be used to connect the switches to each other:

```
[edit interfaces]
user@switch# set interface-name unit 0 family ethernet-switching interface-mode trunk
interface-switch-link
user@switch# set interface-name unit 0 family ethernet-switching vlan members
name-of-private-vlan
```

4. Create an isolated VLAN within the primary VLAN by selecting the **isolated** option for **private-vlan**, and setting a VLAN ID for the isolated VLAN:

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



**NOTE:** You can create only one isolated VLAN within a private VLAN. The isolated VLAN can contain member interfaces from the multiple switches that compose the PVLAN.

Setting the VLAN name for the isolated VLAN is optional. Configuring the VLAN ID is required.

5. Create a community VLAN within the primary VLAN by selecting the **community** option for **private-vlan**, and setting a VLAN ID for this community VLAN::

[edit vlans]

```
user@switch# set community-vlan-name private-vlan community vlan-id community-vlan-id
```



**NOTE:** To create additional community VLANs, repeat this step and specify a different name for the community VLAN. Setting the VLAN name for the community VLAN is optional. Configuring the VLAN ID is required.

6. Associate the isolated VLAN with the primary VLAN:

[edit vlans *primary-vlan-name* vlan-id *primary-vlan-id*]

```
user@switch# set isolated-vlan isolated-vlan-name
```

7. Associate each community VLAN with the primary VLAN:

[edit vlans *primary-vlan-name* vlan-id *primary-vlan-id*]

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

8. If you have not already done so, configure at least one access interface to be a member of the isolated VLAN.

[edit interface]

```
user@switch# set interface-name unit logical-unit-number family ethernet-switching
interface-mode access vlan members isolated-vlan-name
```

9. If you have not already done so, configure at least one access interface to be a member of the community VLAN.

[edit interface]

```
user@switch# set interface-name unit logical-unit-number family ethernet-switching
interface-mode access vlan members community-vlan-name
```



**NOTE:** Repeat this step for the other community VLANs that you are including in the PVLAN.

#### Related Documentation

- [Understanding Private VLANs on page 35](#)
- [Example: Configuring a Private VLAN on a Single Switch on page 41](#)
- [Creating a Private VLAN on a Single Switch \(CLI Procedure\) on page 45](#)



## CHAPTER 6

# Configuring Multiple VLAN Registration Protocol (MVRP)

- [Understanding Multiple VLAN Registration Protocol \(MVRP\) for Dynamic VLAN Registration on page 49](#)
- [Configuring Multiple VLAN Registration Protocol \(MVRP\) to Manage Dynamic VLAN Registration on page 52](#)
- [Verifying That MVRP Is Working Correctly on page 54](#)

## Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration

---

You can configure Multiple VLAN Registration Protocol (MVRP) on Juniper Networks MX Series routers and EX Series switches. The primary purpose of MVRP is to manage dynamic VLAN registration in Layer 2 networks. In managing dynamic VLAN registration, MVRP also prunes VLAN information.

MVRP is an Layer 2 application protocol of the Multiple Registration Protocol (MRP) and is defined in the IEEE 802.1ak standard. MRP and MVRP were designed by IEEE to perform the same functions as Generic Attribute Registration Protocol (GARP) and GARP VLAN Registration Protocol (GVRP) while overcoming some GARP and GVRP limitations, in particular, limitations involving bandwidth usage and convergence time in large networks with large numbers of VLANs.

MVRP was created by IEEE as a replacement application for GVRP. MVRP and GVRP cannot be run concurrently to share VLAN information in a Layer 2 network.

This topic describes:

- [How MVRP Works on page 50](#)
- [Basics of MVRP on page 50](#)
- [MVRP Registration Modes on page 50](#)
- [MRP Timers on page 51](#)
- [MRP VLAN Messages on page 51](#)
- [MVRP Limitations on page 51](#)

## How MVRP Works

The VLAN registration information sent by MVRP protocol data units (PDUs) includes the current VLANs membership—that is, which routers are members of which VLANs—and which router interfaces are in which VLAN. MVRP shares all information in the PDU with all routers participating in MVRP in the Layer 2 network.

MVRP stays synchronized using these PDUs. The routers in the network participating in MVRP receive these PDUs during state changes and update their MVRP states accordingly. MVRP timers dictate when PDUs can be sent and when routers receiving MVRP PDUs can update their MVRP information.

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

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

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

## Basics of MVRP

MVRP is disabled by default. You can configure MVRP router interfaces to participate in MVRP for the Layer 2 network. MVRP can only be enabled on trunk interfaces, and dynamic VLAN configuration through MVRP is enabled by default when MVRP is enabled.

## MVRP Registration Modes

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

The following MVRP registration modes are configurable:

- **forbidden**—The interface does not register or declare VLANs (except statically configured VLANs).
- **normal**—The interface accepts MVRP messages and participates in MVRP. This is the default registration mode setting.
- **restricted**—The interface ignores all MVRP JOIN messages received for VLANs that are not statically configured on the interface.

## MRP Timers

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

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

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



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

## MRP VLAN Messages

MVRP uses MRP messages to register and declare MVRP states for a switch and to inform the Layer 2 network that a switch is leaving MVRP. These messages are communicated as part of the PDU to communicate the state of a particular switch interface on the Layer 2 network to the other switches in the network.

The following messages are communicated for MVRP:

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

## MVRP Limitations

The following limitations apply when configuring MVRP:

- MVRP works with Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP), but not with VLAN Spanning Tree Protocol (VSTP).
- MVRP is allowed only on single tagged trunk ports.

- MVRP is not allowed if a physical interface has more than one logical interface.
- MVRP is only allowed if a logical has one trunk interface (unit 0).

**Related Documentation**

- *Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers*
- [Configuring Multiple VLAN Registration Protocol \(MVRP\) to Manage Dynamic VLAN Registration on page 52](#)
- *Controlling the Management State of a VLAN in MVRP Configurations*
- [Verifying That MVRP Is Working Correctly on page 54](#)
- Deploying MVRP Learning Byte at <https://www.youtube.com/watch?v=C-JkzYbGPBk>

---

## Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration

---

Multiple VLAN Registration Protocol (MVRP) is used to manage dynamic VLAN registration in a Layer 2 network. You can use MVRP on MX Series routers or on EX Series switches.

MVRP is disabled by default on MX Series routers and EX Series switches.

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

- [Enabling MVRP on page 52](#)
- [Disabling MVRP on page 52](#)
- [Changing the Registration Mode to Disable Dynamic VLANs on page 53](#)
- [Configuring Timer Values on page 53](#)
- [Configuring the Multicast MAC Address for MVRP on page 54](#)
- [Configuring an MVRP Interface as a Point-to-Point Interface on page 54](#)
- [Configuring MVRP Tracing Options on page 54](#)

### Enabling MVRP

MVRP can only be enabled on trunk interfaces.

To enable MVRP on a specific trunk interface (here, interface **ge-3/0/5**):

```
[edit protocols mvrp]
user@host# set interface ge-3/0/5
```

### 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, use one of the following:

```
[edit]
user@host# deactivate protocols mvrp
user@host# delete protocols mvrp
```

## Changing the Registration Mode to Disable Dynamic VLANs

When the registration mode for an interface is set to **normal** (the default), dynamic VLANs are created on interfaces participating in MVRP. The dynamic VLANs created on one router or switch are then propagated by means of MVRP to other routers or switches in a topology.

However, dynamic VLAN creation through MVRP can be disabled for all trunk interfaces or for individual trunk interfaces.

For information about disabling dynamic VLAN creation on an interface so that the interface does not register and does not participate in MVRP, see *Controlling the Management State of a VLAN in MVRP Configurations*.

## 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 router or switch after receiving an MVRP PDU:

- The join timer controls the amount of time the router or switch waits to accept a registration request.
- The leave timer controls the period of time that the router or switch waits in the Leave state before changing to the unregistered state.
- 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 a specific interface:

```
[edit protocols mvrp]
user@host# set interface ge-3/0/5 join-timer 300
```

To set the leave timer for a specific interface:

```
[edit protocols mvrp]
user@host# set interface ge-3/0/5 leave-timer 1200
```

To set the leaveall timer for a specific interface:

```
[edit protocols mvrp]
user@host# set interface ge-3/0/5 leaveall-timer 12000
```

## Configuring the Multicast MAC Address for MVRP

MVRP uses the customer MVRP multicast MAC address when MVRP is enabled. However, you can configure MVRP to instead use the provider MVRP multicast MAC address.

To configure MVRP to use the provider MVRP multicast MAC address:

```
[edit protocols mvrp]
user@host# set bpd-destination-mac-address provider-bridge-group;
```

## Configuring an MVRP Interface as a Point-to-Point Interface

Specify that a configured interface is connected point-to-point. If specified, a point-to-point subset of the MRP state machine provides a simpler and more efficient method to accelerate convergence on the network.

To specify that an MVRP interface is point-to-point (here, interface **ge-3/0/5**):

```
[edit protocols mvrp]
user@host# set interface ge-3/0/5 point-to-point;
```

## Configuring MVRP Tracing Options

Set MVRP protocol-level tracing options.

To specify MVRP protocol tracing (here, the file is **/var/log/mvrp-log**, size is **2m**, number of files is **28**, the option **world-readable** indicates the log can be read by user, and MVRP is flagging **events**):

```
[edit protocols mvrp]
user@host# edit traceoptions file /var/log/mvrp-log size 2m files 28 world-readable flag
events
```

### Related Documentation

- *Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers*

---

## Verifying That MVRP Is Working Correctly

**Purpose** After configuring your MX Series router or EX Series switch to participate in Multiple VLAN Registration Protocol (MVRP), verify that the configuration is properly set and that MVRP messages are being sent and received on your switch.

**Action** 1. Confirm that the router is declaring VLANs.

Show that MVRP is enabled:

```
user@host> show mvrp
MVRP configuration for routing instance 'default-switch'
MVRP dynamic VLAN creation : Enabled
MVRP BPDU MAC address      : Customer bridge group (01-80-C2-00-00-21)
MVRP timers (ms)
  Interface   Join   Leave  LeaveAll
  ge-11/3/0   200   800    10000
```

Show the MVRP applicant state:

```
user@host> show mvrp applicant-state
```

```
MVRP applicant state for routing instance 'default-switch'
```

```
(VO) Very anxious observer, (VP) Very anxious passive, (VA) Very anxious new,
```

```
(AN) Anxious new, (AA) Anxious active, (QA) Quiet active, (LA) Leaving active,
```

```
(AO) Anxious observer, (QO) Quiet observer, (LO) Leaving observer,
```

```
(AP) Anxious passive, (QP) Quiet passive
```

VLAN Id	Interface	State
100	ge-11/3/0	Declaring (QA)
200	ge-11/3/0	Declaring (QA)
300	ge-11/3/0	Declaring (QA)

2. Confirm that VLANs are registered on interfaces.

List VLANs in the registered state:

```
user@host> show mvrp registration-state
```

```
MVRP registration state for routing instance 'default-switch'
```

VLAN Id	Interface	Registrar State	Forced State	Managed State	STP State
100	ge-11/3/0	Registered	Registered	Normal	Forwarding
200	ge-11/3/0	Registered	Registered	Normal	Forwarding
300	ge-11/3/0	Empty	Empty	Normal	Forwarding

3. Display a list of VLANs created dynamically.

List dynamic VLAN membership:

```
user@host> show mvrp dynamic-vlan-memberships
```

```
MVRP dynamic vlans for routing instance 'default-switch'
```

```
(s) static vlan, (f) fixed registration
```

VLAN Id	Interfaces
100	ge-3/3/0 ge-3/0/5
200	ge-3/3/0 ge-3/0/5

**Meaning** The output of `show mvrp applicant-state` shows that trunk interface **ge-11/3/0** is declaring (sending out) interest in the VLAN IDs **100**, **200**, and **300**, and MVRP is operating properly.

The output of `show mvrp registrant-state` shows the registrar state for VLANs **100** and **200** as **Registered**, indicating that these VLANs are receiving traffic from a customer site. VLAN **300** is in an **Empty** state and is not receiving traffic from a customer site.

The output of the `show mvrp dynamic-vlan-membership` shows that VLANs **100** and **200** are created dynamically (here, on an MX Series router operating as an aggregation switch between MX Series routers operating as edge switches). VLANs created statically are marked with an **(s)** (which is not indicated in this output).

#### Related Documentation

- *Controlling the Management State of a VLAN in MVRP Configurations*





## CHAPTER 7

# Configuring Proxy ARP

- [Restricted and Unrestricted Proxy ARP Overview on page 57](#)
- [Configuring Restricted and Unrestricted Proxy ARP on page 59](#)

## Restricted and Unrestricted Proxy ARP Overview

---

By default, the Junos OS responds to an Address Resolution Protocol (ARP) request only if the destination address of the ARP request is local to the incoming interface.

For Ethernet Interfaces, you can configure the router or switches to proxy-reply to the ARP requests using the restricted or unrestricted proxy ARP configuration.

You might want to configure restricted or unrestricted proxy ARP for routers that act as provider edge (PE) devices in Ethernet Layer 2 LAN switching domains.



**NOTE:** From Junos OS Release 10.0 onward, Junos OS does not respond to proxy ARP requests with the default route 0.0.0.0. This behavior is in compliance with RFC 1027.

## Restricted Proxy ARP

Restricted proxy ARP enables the router or switch to respond to the ARP requests in which the physical networks of the source and target are not the same and the router or switch has an active route to the target address in the ARP request. The router does not reply if the target address is on the same subnet and the same interface as the ARP requestor.

## Unrestricted Proxy ARP

Unrestricted proxy ARP enables the router or switch to respond to any ARP request, on condition that the router has an active route to the destination address of the ARP request. The route is not limited to the incoming interface of the request, nor is it required to be a direct route.



**WARNING:** If you configure unrestricted proxy ARP, the proxy router replies to ARP requests for the target IP address on the same interface as the

incoming ARP request. This behavior is appropriate for cable modem termination system (CMTS) environments, but might cause Layer 2 reachability problems if you enable unrestricted proxy ARP in other environments.

When an IP client broadcasts the ARP request across the Ethernet wire, the end node with the correct IP address responds to the ARP request and provides the correct MAC address. If the unrestricted proxy ARP feature is enabled, the router response is redundant and might fool the IP client into determining that the destination MAC address within its own subnet is the same as the address of the router.



**NOTE:** While the destination address can be remote, the source address of the ARP request must be on the same subnet as the interface upon which the ARP request is received. For security reasons, this rule applies to both unrestricted and restricted proxy ARP.

## Topology Considerations for Unrestricted Proxy ARP

In most situations, you should not configure the router or switch to perform unrestricted proxy ARP. Do so only for special situations, such as when cable modems are used.

[Figure 4 on page 58](#) and [Figure 5 on page 59](#) show examples of situations in which you might want to configure unrestricted proxy ARP.

In [Figure 4 on page 58](#), the edge device is not running any IP protocols. In this case, you configure the core router to perform unrestricted proxy ARP. The edge device is the client of the proxy.

In [Figure 5 on page 59](#), the Broadband Remote Access Server (B-RAS) routers are not running any IP protocols. In this case, you configure unrestricted proxy ARP on the B-RAS interfaces. This allows the core device to behave as though it is directly connected to the end users.

**Figure 4: Edge Device Case for Unrestricted Proxy ARP**

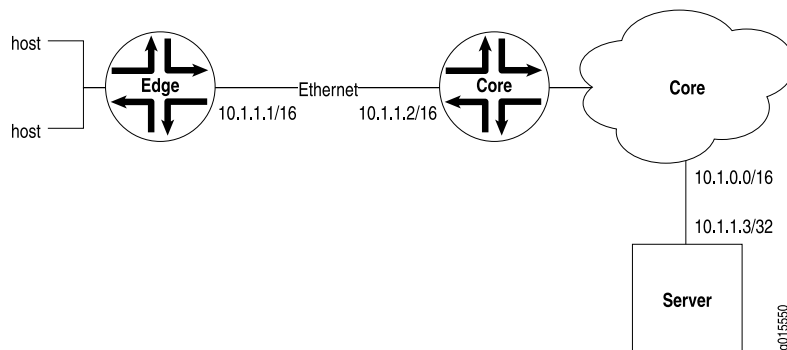



Figure 5: Core Device Case for Unrestricted Proxy ARP



- Related Documentation**
- [Configuring Restricted and Unrestricted Proxy ARP on page 59](#)
  - [Ethernet Interfaces Feature Guide for Routing Devices](#)

## Configuring Restricted and Unrestricted Proxy ARP

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

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

You can include this statement at the following hierarchy levels:

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

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

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

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



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

- Related Documentation**
- [proxy-arp on page 132](#)
  - [Restricted and Unrestricted Proxy ARP Overview on page 57](#)
  - *Configuring Gratuitous ARP*
  - *Ethernet Interfaces Feature Guide for Routing Devices*

## PART 2

# Configuration Statements and Operational Commands

- Configuration Statements on page 63
- Operational Commands on page 153



## CHAPTER 8

# Configuration Statements

- [\[edit dynamic-profiles\] Hierarchy Level on page 65](#)
- [\[edit interfaces\] Hierarchy Level on page 65](#)
- [\[edit logical-systems\] Hierarchy Level on page 76](#)
- [\[edit protocols l2-learning\] Hierarchy Level on page 81](#)
- [Layer 2 Routing Instances Configuration Hierarchy on page 81](#)
- [\[edit switch-options\] Hierarchy Level on page 84](#)
- [\[edit vlans\] Hierarchy Level on page 84](#)
- [bpdu-destination-mac-address on page 86](#)
- [bridge-priority on page 87](#)
- [community-vlan on page 88](#)
- [domain-type on page 88](#)
- [encapsulation \(Physical Interface\) on page 89](#)
- [family on page 94](#)
- [fast-aps-switch on page 98](#)
- [global-mac-statistics on page 99](#)
- [global-mac-table-aging-time on page 99](#)
- [global-no-mac-learning on page 100](#)
- [inner-tag-protocol-id on page 101](#)
- [inner-vlan-id on page 102](#)
- [input-vlan-map on page 103](#)
- [interface on page 104](#)
- [interface \(MVRP\) on page 105](#)
- [interface \(Spanning Tree\) on page 106](#)
- [interface-mode on page 107](#)
- [interfaces on page 108](#)
- [inter-switch-link on page 109](#)
- [isid on page 109](#)
- [isid-list on page 110](#)

- [isolated-vlan on page 110](#)
- [join-timer \(MVRP\) on page 111](#)
- [l2-learning on page 112](#)
- [l3-interface on page 112](#)
- [leaveall-timer \(MVRP\) on page 113](#)
- [leave-timer \(MVRP\) on page 114](#)
- [mac-statistics on page 115](#)
- [mac-table-size on page 117](#)
- [mvrp on page 119](#)
- [native-vlan-id on page 120](#)
- [no-mac-learning on page 122](#)
- [output-vlan-map \(Gigabit Ethernet IQ and 10-Gigabit Ethernet with SFPP\) on page 123](#)
- [packet-action on page 124](#)
- [point-to-point \(MVRP\) on page 127](#)
- [pop on page 128](#)
- [pop-pop on page 129](#)
- [pop-swap on page 130](#)
- [private-vlan on page 131](#)
- [proxy-arp on page 132](#)
- [push on page 133](#)
- [push-push on page 134](#)
- [registration on page 135](#)
- [rstp on page 136](#)
- [service-id on page 137](#)
- [static-mac on page 138](#)
- [swap on page 139](#)
- [swap-push on page 140](#)
- [swap-swap on page 141](#)
- [switch-options on page 142](#)
- [tag-protocol-id \(TPID to Rewrite\) on page 143](#)
- [traceoptions \(MVRP\) on page 144](#)
- [unconditional-src-learn on page 146](#)
- [vlan-id \(Bridge Domain or VLAN\) on page 147](#)
- [vlan-id \(VLAN ID to Rewrite\) on page 149](#)
- [vlan-id-list on page 150](#)
- [vlan-rewrite on page 151](#)
- [vlan-tags on page 152](#)



## [edit dynamic-profiles] Hierarchy Level

```
dynamic-profiles {
  profile-name {
    class-of-service
    ... statements from those in [edit class-of-service] Hierarchy Level.
    firewall
    ... statements from those in [edit firewall] Hierarchy Level.
    interfaces
    ... statements from those in [edit interfaces] Hierarchy Level.
    policy-options
    ... statements from those in [edit policy-options] Hierarchy Level.
    predefined-variable-defaults variable-name default-value
    profile-variable-set variable-set-name dynamic-variable-name substitute-variable-name
    protocols
    ... statements from those in [edit protocols] Hierarchy Level.
    routing-instances
    ... statements from those in [edit routing-instances] Hierarchy Level.
    routing-options
    ... statements from those in [edit routing-options] Hierarchy Level.
    services
    ... statements from those in [edit services] Hierarchy Level.
    variables {
      variable-name {
        default-value default-value;
        equals expression;
        mandatory;
      }
      uid;
      uid-reference;
    }
  }
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
  - *[edit dynamic-profiles routing-instances] Hierarchy Level*
  - *[edit dynamic-profiles routing-options] Hierarchy Level*
  - *[edit dynamic-profiles variables] Hierarchy Level*

## [edit interfaces] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```
interfaces {
  interface-name {
    ... the "interface-name" subhierarchy appears after the main [edit interfaces] hierarchy level ...
  }
  interface-set interface-set-name {
    interface interface-name {
```

```

        (unit unit-number | vlan-tags-outer vlan-tag);
    }
}
irb {
    accounting-profile name;
    arp-l2-validate;
    description text;

    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;

    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
        description text;
        disable;
        encapsulation type;
        family inet {
            accounting {
                destination-class-usage;
                source-class-usage {
                    input;
                    output;
                }
            }
        }
        address ipv4-address {
            arp ip-address (mac | multicast-mac) mac-address <publish>;
            broadcast address;
            preferred;
            primary;
            vrrp-group group-id {
                (accept-data | no-accept-data);
                advertise-interval seconds;
                advertisements-threshold number;
                authentication-key key;
                authentication-type authentication;
                fast-interval milliseconds;
                (preempt | no-preempt) {
                    hold-time seconds;
                }
            }
            priority number;
            track {
                interface interface-name {
                    bandwidth-threshold bits-per-second priority-cost priority;
                    priority-cost priority;
                }
            }
            priority-hold-time seconds;
            route prefix/prefix-length routing-instance instance-name priority-cost priority;
        }
    }
}

```

```

        virtual-address [ addresses ];
        vrrp-inherit-from vrrp-group;
    }
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bandwidth priority-cost number;
                priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
    }
    virtual-inet6-address [addresses];
}

```

```

        virtual-link-local-address ipv6-address;
        vrrp-inherit-from {
            active-group group-number;
            active-interface interface-name;
        }
    }
    (dad-disable | no-dad-disable);
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    nd6-stale-time seconds;
    no-neighbor-learn;
    no-redirects;
    policer {
        input policer-name;
        output policer-name;
    }
    rpf-check {
        fail-filter filter-name;
        mode {
            loose;
        }
    }
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}
native-inner-vlan-id vlan-id;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id-list [vlan-id's];
vlan-id-range [vlan-id-range];
}
}
traceoptions {
    file <filename> <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
    flag flag <disable>;
    no-remote-trace;
}
}

```

```

interfaces {
  interface-name {
    disable;
    accounting-profile name;
    aggregated-ether-options {
      ethernet-switch-profile {
        tag-protocol-id [ hexadecimal-identifiers ];
      }
      (flow-control | no-flow-control);
      lacp {
        (active | passive);
        admin-key key;
        fast-failover;
        link-protection {
          disable;
          (revertive | non-revertive);
        }
        periodic (fast | slow);
        system-id mac-address;
        system-priority priority;
      }
      (link-protection | no-link-protection);
      link-speed (100m | 1g | 8g | 10g | 40g | 50g | 80g | 100g | oc192);
      logical-interface-fpc-redundancy;
      (loopback | no-loopback);
      mc-ae {
        chassis-id chassis-id;
        events {
          iccp-peer-down {
            force-icl-down;
            prefer-status-control-active;
          }
        }
        mc-ae-id mc-ae-id;
        mode (active-active | active-standby);
        redundancy-group group-id;
        status-control (active | standby);
      }
      minimum-links number;
      rebalance-periodic {
        start-time time;
        interval number;
      }
      source-address-filter {
        mac-address;
      }
      (source-filtering | no-source-filtering);
    }
    auto-configure {
      remove-when-no-subscribers;
      stacked-vlan-ranges {
        access-profile profile-name;
        authentication {
          password password-string;
          username-include {

```

```

        circuit-type;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        mac-address;
        option-82 ( circuit-id | remote-id);
        radius-realm radius-realm-string;
        user-prefix user-prefix-string;
    }
}
dynamic-profile profile-name {
    accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
    ranges (any | low-tag-high-tag), (any | low-tag-high-tag);
}
}
vlan-ranges {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-name;
            mac-address;
            option-82;
            radius-realm radius-realm-string;
            user-prefix user-prefix-string;
        }
    }
    dynamic-profile profile-name {
        accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
        ranges (any | low-tag)—(any | high-tag);
    }
}
override tag vlan-tag dynamic-profile profile name;
}
encapsulation (ethernet-bridge | ethernet-vpls | extended-vlan-bridge |
    extended-vlan-vpls | flexible-ethernet-services | vlan-vpls);
ether-options {
    802.3ad {
        aex;
        (backup | primary);
        lacp {
            force-up;
            port-priority
        }
    }
}
asynchronous-notification;
(auto-negotiation | no-auto-negotiation);
ethernet-switch-profile {
    ethernet-policer-profile {
        input-priority-map {
            ieee802.1p premium [ values ];
        }
        output-priority-map {

```

```

    classifier {
        premium {
            forwarding-class class-name {
                loss-priority (high | low);
            }
        }
    }
}
policer cos-policer-name {
    aggregate {
        bandwidth-limit bps;
        burst-size-limit bytes;
    }
    premium {
        bandwidth-limit bps;
        burst-size-limit bytes;
    }
}
tag-protocol-id;
}
(mac-learn-enable | no-mac-learn-enable);
}
(flow-control | no-flow-control);
ignore-l3-incompletes;
link-mode (automatic | full-duplex | half-duplex);
(loopback | no-loopback);
keepalives <interval seconds> <down-count number> <up-count number>;
speed (1g | 10m | 100m | 10m-100m | auto-negotiation);
source-address-filter {
    mac-address;
}
source-filtering | no-source-filtering;
}
flexible-vlan-tagging;
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time (up milliseconds | down milliseconds);
interface-transmit-statistics;
(keepalives <down-count number> <interval seconds> <up-count number> |
no-keepalives);
layer2-policer {
    apply-groups [ group-names ];
    apply-groups-except [ group-names ];
}
link-mode (automatic | full-duplex);
mac mac-address;
mtu bytes;
multi-chassis-protection peer-ip-address {
    interface interface-name;
}
native-vlan-id number;
no-gratuitous-arp-request;
optics-options {
    alarm low-light-alarm {
        (link-down | syslog);
    }
}
warning low-light-warning {

```

```

        (link-down | syslog);
    }
    wavelength nm;
}
passive-monitor-mode;
per-unit-scheduler;
speed (10m | 100m | 1g | auto | oc3 | oc12 | oc48);
stacked-vlan-tagging;
traceoptions {
    flag flag;
}
transmit-bucket {
    overflow discard;
    rate percentage;
    threshold bytes;
}
(traps | no-traps);
unidirectional;
vlan-tagging;
}

interface-name {
    unit logical-unit-number {
        disable;
        accept-source-mac {
            mac-address mac-address {
                policer {
                    input policer-name;
                    output policer-name;
                }
            }
        }
    }
    accounting-profile name;
    advisory-options {
        downstream-rate rate;
        upstream-rate rate;
    }
    arp-resp (restricted|unrestricted);
    bandwidth rate;
    clear-dont-fragment-bit;
    copy-tos-to-outer-ip-header;
    demux-destination family;
    encapsulation (vlan-bridge | vlan-vpls);
    epd-threshold cells plp1 cells;
    filter filter-name;
    inner-vlan-id-range start start-id end end-id;
    input-vlan-map {
        (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
        inner-tag-protocol-id tpid;
        inner-vlan-id number;
        tag-protocol-id tpid;
        vlan-id number;
    }
    interface-shared-with psdnumerical-index;
    layer2-policer {

```



```

    input-hierarchical-policer policer-name;
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
multi-chassis-protection peer-ip-address {
    interface interface-name;
}
native-inner-vlan-id number;
output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
peer-interface interface-name;
peer-unit unit-number;
plp-to-clp;
proxy-arp <restricted | unrestricted>;
rpm {
    (client | server);
    twamp-server;
}
swap-by-poppush;
vlan-id number;
vlan-id-list [ vlan-id vlan-id-vlan-id ];
vlan-id-range number-number;
vlan-tags (inner <tpid.>vlan-id | inner-list [ vlan-id vlan-id-vlan-id ] |
    inner-range <tpid.>vlan-id-vlan-id) outer <tpid.>vlan-id;
}

unit logical-unit-number {
    family ethernet-switching {
        filter {
            group filter-group-number;
            (input filter-name | input-list [ filter-names ]);
            (output filter-name | output-list [ filter-names ]);
            (inner-vlan-id-list [ vlan-ids ] | vlan-id number | vlan-id-list [ number
                number-number ]);
            interface-mode (access | trunk);
        }
        policer {
            input policer-name;
            output policer-name;
        }
        vlan-rewrite {
            translate old-vlan-id new-vlan-id;
        }
        vlan {
            members [ all vlan-identifiers ];
        }
    }
}
family inet {
    filter {
        group filter-group-number;
    }
}

```

```

        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    input-hierarchical-policer policer-name;
    mac-validate (loose | strict);
    mtu bytes;
    no-neighbor-learn;
    no-redirects;
    policer {
        arp policer-template-name;
        input policer-name;
        output policer-name;
    }
    primary;
    receive-options-packets;
    receive-ttl-exceeded;
    rpf-check {
        fail-filter filter-name;
        mode loose;
    }
    sampling {
        (input | output | input output);
    }
    simple-filter {
        input filter-name;
    }
    targeted-broadcast {
        forward-and-send-to-re;
        forward-only;
    }
    unnumbered-address interface-name <destination address>
        <destination-profile profile-name> <preferred-source-address address>;
}

family inet6 {
    address ipv6-address {
        destination destination-address;
        eui-64;
        ndp ipv6-address <l2-interface interface-name> <(mac mac-address |
            multicast-mac multicast-mac-address) <publish>>;
        preferred;
        primary;
        vrrp-inet6-group group-number {
            (accept-data | no-accept-data);
            fast-interval milliseconds;
            inet6-advertise-interval seconds;
            (no-preempt; | ... the following preempt statement ...)
            preempt {
                hold-time seconds;
            }
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bits-per-second priority-cost priority;
                priority-cost priority;
            }
        }
    }
}

```

```

    }
    priority-hold-time seconds;
    route ip-address-prefix/prefix-length routing-instance instance-name
        priority-cost priority;
    }
    virtual-inet6-address [ addresses ];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
    }
    }
    (dad-disable | no-dad-disable);
    filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    input-hierarchical-policer policer-name;
    mtu bytes;
    nd6-stale-time seconds;
    no-neighbor-learn;
    policer {
        input policer-name;
        output policer-name;
    }
    rpf-check {
        fail-filter filter-name;
        mode loose;
    }
    sampling {
        (input | output | input output);
    }
    unnumbered-address interface-name preferred-source-address address;
}

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

family mlfr-end-to-end {
    bundle logical-interface-name;
}

family mpls {
    filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    input-hierarchical-policer policer-name;

```

```

maximum-labels maximum-labels;
mtu bytes;
policer {
    input policer-name;
    output policer-name;
}
}

family vpls {
    core-facing;
    filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    policer {
        input policer-name;
        output policer-name;
    }
}
}
}
}

```

**Related Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

## [\[edit logical-systems\] Hierarchy Level](#)

The following lists the statements that can be configured at the **[edit logical-systems]** hierarchy level that are also documented in this manual. For more information about logical systems, see the *Logical Systems Feature Guide for Routing Devices*.

```

logical-systems logical-system-name {
    interfaces interface-name {
        unit logical-unit-number {
            accept-source-mac {
                mac-address mac-address {
                    policer {
                        input cos-policer-name;
                        output cos-policer-name;
                    }
                }
            }
        }
    }
    allow-any-vci;
    atm-scheduler-map (map-name | default);
    bandwidth rate;
    backup-options {
        interface interface-name;
    }
    cell-bundle-size cells;
    clear-dont-fragment-bit;
}

```

```

compression {
  rtp {
    f-max-period number;
    port {
      minimum port-number;
      maximum port-number;
    }
    queues [ queue-numbers ];
  }
}
compression-device interface-name;
description text;
interface {
  l2tp-interface-id name;
  (dedicated | shared);
}
dialer-options {
  activation-delay seconds;
  deactivation-delay seconds;
  dial-string [ dial-string-numbers ];
  idle-timeout seconds;
  initial-route-check seconds;
  load-threshold number;
  pool pool;
  remote-name remote-callers;
  watch-list {
    [ routes ];
  }
}
disable;
dlci dlci-identifier;
drop-timeout milliseconds;
dynamic-call-admission-control {
  activation-priority priority;
  bearer-bandwidth-limit kilobits-per-second;
}
encapsulation type;
epd-threshold cells plp1 cells;
fragment-threshold bytes;
input-vlan-map {
  inner-tag-protocol-id;
  inner-vlan-id;
  (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
  tag-protocol-id tpid;
  vlan-id number;
}
interleave-fragments;
inverse-arp;
layer2-policer {
  input-policer policer-name;
  input-three-color policer-name;
  output-policer policer-name;
  output-three-color policer-name;
}
link-layer-overhead percent;
minimum-links number;

```

```
mrru bytes;
multicast-dlci dlci-identifier;
multicast-vci vpi-identifier.vci-identifier;
multilink-max-classes number;
multipoint;
oam-liveness {
    up-count cells;
    down-count cells;
}
oam-period (seconds | disable);
output-vlan-map {
    inner-tag-protocol-id;
    inner-vlan-id;
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-swap);
    tag-protocol-id tpid;
    vlan-id number;
}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
}
dynamic-profile profile-name;
pap {
    default-pap-password password;
    local-name name;
    local-password password;
    passive;
}
}
proxy-arp;
service-domain (inside | outside);
shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst length);
    queue-length number;
}
short-sequence;
transmit-weight number;
(traps | no-traps);
trunk-bandwidth rate;
trunk-id number;
tunnel {
    backup-destination address;
    destination address;
```

```

    key number;
    routing-instance {
        destination routing-instance-name;
    }
    source source-address;
    ttl number;
}
vci vpi-identifier.vci-identifier;
vlan-id number;
vlan-id-list [vlan-id vlan-id-vlan-id]
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
vlan-tags outer tpid.vlan-id inner-list [vlan-id vlan-id-vlan-id]
vpi vpi-identifier;
family family {
    accounting {
        destination-class-usage;
        source-class-usage {
            direction;
        }
    }
}
bundle interface-name;
filter {
    group filter-group-number;
    input filter-name;
    input-list {
        [filter-names];
    }
    output filter-name;
    output-list {
        [filter-names];
    }
}
ipsec-sa sa-name;
keep-address-and-control;
mtu bytes;
multicast-only;
no-redirects;
policer {
    arp policer-template-name;
    input policer-template-name;
    output policer-template-name;
}
primary;
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check <fail-filter filter-name> {
    <mode loose>;
}
sampling {
    direction;
}
service {
    input {
        service-set service-set-name <service-filter filter-name>;
    }
}

```

```

        post-service-filter filter-name;
    }
    output {
        service-set service-set-name <service-filter filter-name>;
    }
}
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
unnumbered-address interface-name destination address destination-profile
    profile-name;
address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    destination-profile name;
    eui-64;
    multipoint-destination address (dlci dlci-identifier | vci vci-identifier);
    multipoint-destination address {
        epd-threshold cells plp1 cells;
        inverse-arp;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (seconds | disable);
        shaping {
            (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
                rate burst length);
            queue-length number;
        }
        vci vpi-identifier.vci-identifier;
    }
    preferred;
    primary;
    (vrrp-group | vrrp-inet6-group) group-number {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        authentication-type authentication;
        authentication-key key;
        fast-interval milliseconds;
        (preempt | no-preempt) {
            hold-time seconds;
        }
        priority-number number;
        track {
            priority-cost seconds;
            priority-hold-time interface-name {
                interface priority;
                bandwidth-threshold bits-per-second {
                    priority;
                }
            }
        }
        route ip-address/mask routing-instance instance-name priority-cost cost;
    }
}
virtual-address [ addresses ];

```



```

    }
  }
}

```

- Related Documentation**
- *Junos OS Hierarchy and RFC Reference*
  - *Ethernet Interfaces Feature Guide for Routing Devices*
  - *Junos OS Network Interfaces Library for Routing Devices*

## [edit protocols l2-learning] Hierarchy Level

```

protocols {
  l2-learning {
    global-mac-limit {
      limit;
      packet-action drop;
    }
    global-mac-statistics;
    global-mac-table-aging-time seconds;
    global-no-mac-learning;
  }
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
  - *[edit protocols] Hierarchy Level*

## Layer 2 Routing Instances Configuration Hierarchy

Use the **vpls** routing instance type for point-to-multipoint LAN implementations between a set of sites in a VPN.

To configure routing instances for Layer 2 networks, include the following statements:

```

routing-instances {
  routing-instance-name {
    access {
      address-assignment {
        ... same statements as in the address-assignment subhierarchy in [edit access]
        Hierarchy Level ...
      }
      address-protection;
      description text;
      egress-protection {
        context-identifier context-id;
      }
      forwarding-options {
        ...forwarding-options...
      }
      instance-role role;
      instance-type type;
    }
  }
}

```

```

interface interface-name;
l2-domain-id-for-l3 id;
l2vpn-id community;
layer3-domain-identifier identifier;
multicast-snooping-options {
    ... same statements as in [edit multicast-snooping-options] Hierarchy Level EXCEPT
    FOR ...
    traceoptions {...} # NOT valid at this level
}
no-irb-layer-2-copy;
no-local-switching;
no-vrf-advertise;
no-vrf-propagate-ttl;
pbb-options {
    default-bvlan bvlan;
    peer-instance instance;
    vlan-id vlan-id isid-list [ isid-numbers ]
}
protocols {
    ... the protocols subhierarchy appears after the main [edit routing-instances
    routing-instance-name] hierarchy ...
}
provider-tunnel {
    ... the provider-tunnel subhierarchy appears after the main [edit routing-instances
    routing-instance-name] hierarchy ...
}
route-distinguisher (as-number:number | ip-address:number);
routing-interface interface;
routing-options {
    ... the routing-options subhierarchy appears after the main [edit routing-instances
    routing-instance-name] hierarchy ...
}
service-groups {
    service-group-name {
        pbb-service-options {
            default-isid isid-number;
            isid isid-number vlan-id-list [ vlan-ids ];
            mac-address mac-address;
        }
        service-type type;
    }
}
switch-options {
    ... same statements as in [edit switch-options] Hierarchy Level ...
}
vlan-id (id | all | none);
vlan-model one-to-one;
vlan-tags outer <tpid.>vlan-id inner <tpid.>vlan-id;
\[edit vlans\] Hierarchy Level on page 84 {
    ... same statements as in [edit vlans] Hierarchy Level ...
}
vrf-advertise-selective {
    family {
        inet-mvpn;
        inet6-mvpn;
    }
}

```

```

}
vrf-export [ policy-names ];
vrf-import [ policy-names ];
vrf-propagate-ttl;
vrf-table-label;
vrf-target {
    export community-name;
    import community-name;
}
protocols {
    ... protocols-configuration ...
}
routing-options {
    ... routing-options-configuration ...
}
bridge-domains {
    bridge-domain-name {
        domain-type bridge;
        interface interface-name;
        routing-interface routing-interface-name;
        vlan-id (Bridge Domain or VLAN) (none | all | number);
        vlan-tags outer number inner number;
        bridge-options {
            interface-mac-limit limit {
                packet-action drop;
            }
            interface interface-name {
                interface-mac-limit limit {
                    packet-action drop;
                }
            }
            mac-statistics;
            mac-table-size limit {
                packet-action drop;
            }
            no-mac-learning;
            static-mac mac-address;
        }
    }
}
}
}

```

With the exception of the **instance-type virtual-switch** statement (which configures a virtual-switch routing instance), you can include the statements at the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

The **instance-type virtual-switch** statement is not supported at the **[edit logical-systems *logical-system-name*]** hierarchy level.

#### Related Documentation

- [Routing Instances Overview](#)

- *Layer 2 Routing Instance Types*
- [Configuring a Layer 2 Virtual Switch on page 33](#)
- *Configuring a Layer 2 Control Protocol Routing Instance*

## [\[edit switch-options\] Hierarchy Level](#)

---

```
switch-options {  
  interface interface-name {  
    interface-mac-limit {  
      number-of-addresses;  
      packet-action drop;  
    }  
    no-mac-learning;  
  }  
  interface-mac-limit {  
    number-of-addresses;  
    packet-action drop;  
  }  
  mac-statistics;  
  mac-table-size {  
    number-of-addresses;  
    packet-action drop;  
  }  
  no-mac-learning;  
  service-id number;  
}
```

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

---

```
vlans {  
  vlan-name {  
    description text-description;  
    domain-type bridge;  
    forwarding-options {  
      dhcp-security {  
        arp-inspection;  
        group name {  
          interface interface-name {  
            static-ip ip-address {  
              mac mac-address;  
            }  
          }  
        }  
        overrides {  
          no-option82;  
          trusted;  
          untrusted;  
        }  
      }  
      ip-source-guard;  
      no-dhcp-snooping;  
      option-82 {  
        circuit-id {
```

```

    prefix {
        host-name;
        logical-system-name;
        routing-instance-name;
    }
    use-interface-description (device | logical);
    use-vlan-id;
}
remote-id {
    host-name;
    use-interface-description (device | logical);
    use-string string;
}
vendor-id {
    use-string string;
}
}
}
filter {
    input filter-name;
    output filter-name;
}
flood {
    input filter-name;
}
}
}
interface interface-name;
l3-interface interface-name;
mcae-mac-synchronize;
multicast-snooping-options {
    ... same statements as in multicast-snooping-options ...
}
no-irb-layer-2-copy;
service-id number;
switch-options {
    ... the switch-options subhierarchy appears after the main [edit vlans vlan-name]
    hierarchy ...
}
vlan-id (all | none | number);
vlan-id-list [ vlan-id-numbers ];
vlan-tags outer <tpid.>vlan-id <inner <tpid.>vlan-id>;
}

vlan-name {
    switch-options {
        interface interface-name {
            interface-mac-limit {
                limit;
                packet-action drop;
            }
            no-mac-learning;
            static-mac mac-address {
                vlan-id number;
            }
        }
    }
    interface-mac-limit {

```

```

        limit;
        packet-action drop;
    }
    mac-statistics;
    mac-table-size {
        number-of-addresses;
        packet-action drop;
    }
    no-mac-learning;
}
}
}

```


## bpdu-destination-mac-address

<b>Syntax</b>	bpdu-destination-mac-address provider-bridge-group;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <a href="#">mvrp</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp</a> ] (for virtual switch instance type), [edit protocols <a href="#">mvrp</a> ], [edit routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp</a> ] (for virtual switch instance type)
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1 for MX Series routers.
<b>Description</b>	For Multiple VLAN Registration Protocol (MVRP) configurations, specifies the multicast address for MVRP. If configured, the provider MVRP multicast MAC address is used; otherwise, the Junos OS uses the customer MVRP multicast MAC address.
<b>Default</b>	By default, the provider MVRP multicast MAC address is used (if configured). Otherwise, the customer MVRP MAC address is used.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers</a></li> <li>• <a href="#">Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration on page 52</a></li> <li>• <a href="#">Verifying That MVRP Is Working Correctly on page 54</a></li> <li>• <a href="#">Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration on page 49</a></li> </ul>

## bridge-priority

<b>Syntax</b>	<code>bridge-priority <i>priority</i>;</code>
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> protocols (mstp   <b>rstp</b>)], [edit logical-systems <i>logical-system-name</i> protocols mstp msti <i>msti-id</i>], [edit logical-systems <i>logical-system-name</i> protocols vstp vlan <i>vlan-id</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp   <b>rstp</b>)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp msti <i>msti-id</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp vlan <i>vlan-id</i>], [edit protocols (mstp   <b>rstp</b>)], [edit protocols mstp msti <i>msti-id</i>], [edit protocols vstp vlan <i>vlan-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols (mstp   <b>rstp</b>)], [edit routing-instances <i>routing-instance-name</i> protocols mstp msti <i>msti-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols vstp vlan <i>vlan-id</i>]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>
<b>Description</b>	Determine which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.
<b>Options</b>	<p><b>priority</b>—The bridge priority can be set only in increments of 4096.</p> <p><b>Range:</b> 0 through 61,440</p> <p><b>Default:</b> 32,768</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Bridge Priority for Election of Root Bridge and Designated Bridge</i></li> <li>• <i>Example: Configuring Network Regions for VLANs with MSTP</i></li> <li>• <i>Understanding MSTP for EX Series and QFX Series Switches</i></li> <li>• <i>Understanding VSTP for EX Series Switches and QFX Series Switches</i></li> </ul>

## community-vlan

<b>Syntax</b>	<code>community-vlan vlan <i>community-vlan-name</i>;</code>
<b>Hierarchy Level</b>	<code>[edit vlans <i>primary-vlan-name</i> vlan-id <i>primary-vlan-vlan-id</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.
<b>Description</b>	Configure the specified community VLAN to be a secondary VLAN of the specified primary VLAN. A <i>community</i> VLAN is used to transport frames among members of a community (a subset of users within the VLAN), and to forward frames upstream to the primary VLAN.
<div>  <p><b>NOTE:</b> Before you specify this configuration statement, you must have already configured the specified community VLAN and assigned a VLAN ID to it. See <a href="#">private-vlan</a>.</p> </div>	
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system—control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Creating a Private VLAN on a Single Switch (CLI Procedure) on page 45</a></li> <li>• <a href="#">Creating a Private VLAN Spanning Multiple Switches (CLI Procedure) on page 47</a></li> </ul>

## domain-type

<b>Syntax</b>	<code>domain-type bridge;</code>
<b>Hierarchy Level</b>	<code>[edit bridge-domains <i>bridge-domain-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i></code> <code>bridge-domains <i>bridge-domain-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6.
<b>Description</b>	(MX Series routers only) Define the type of domain for a Layer 2 bridge domain.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing—control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a Bridge Domain</a></li> <li>• <a href="#">Configuring a Layer 2 Virtual Switch</a></li> </ul>



## encapsulation (Physical Interface)

<b>Syntax</b>	encapsulation (atm-ccc-cell-relay   atm-pvc   cisco-hdlc   cisco-hdlc-ccc   cisco-hdlc-tcc   ethernet-bridge   ethernet-ccc   ethernet-over-atm   ethernet-tcc   ethernet-vpls   ethernet-vpls-fr   ether-vpls-over-atm-llc   ethernet-vpls-ppp   extended-frame-relay-ccc   extended-frame-relay-ether-type-tcc   extended-frame-relay-tcc   extended-vlan-bridge   extended-vlan-ccc   extended-vlan-tcc   extended-vlan-vpls   flexible-ethernet-services   flexible-frame-relay   frame-relay   frame-relay-ccc   frame-relay-ether-type   frame-relay-ether-type-tcc   frame-relay-port-ccc   frame-relay-tcc   generic-services   multilink-frame-relay-uni-nni   ppp   ppp-ccc   ppp-tcc   vlan-ccc   vlan-vci-ccc   vlan-vpls);
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> ], [edit interfaces rlsq <i>number:number</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers ( <b>flexible-ethernet-services</b> , <b>ethernet-ccc</b> , and <b>ethernet-tcc</b> options only).
<b>Description</b>	Specify the physical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.
<b>Default</b>	For SONET interfaces— <b>ppp</b>  For ATM interfaces— <b>atm-pvc</b>  For DS1, E1, and T1 interfaces— <b>ppp</b>  For DS3, E3, and T3 interfaces— <b>ppp</b>  For all channelized interfaces— <b>ppp</b>
<b>Options</b>	<b>atm-ccc-cell-relay</b> —Use ATM cell-relay encapsulation.  <b>atm-pvc</b> —Use ATM PVC encapsulation.  <b>cisco-hdlc</b> —Use Cisco-compatible High-Level Data Link Control (HDLC) framing.  <b>cisco-hdlc-ccc</b> —Use Cisco-compatible HDLC framing on CCC circuits.  <b>cisco-hdlc-tcc</b> —Use Cisco-compatible HDLC framing on TCC circuits for connecting different media.  <b>ethernet-bridge</b> —Use Ethernet bridge encapsulation on Ethernet interfaces that have bridging enabled and that must accept all packets.  <b>ethernet-ccc</b> —Use Ethernet CCC encapsulation on Ethernet interfaces that must accept packets carrying standard Tag Protocol ID (TPID) values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, CCC is not supported.  <b>ethernet-over-atm</b> —For interfaces that carry IPv4 traffic, use Ethernet over ATM encapsulation. When you use this encapsulation type, you cannot configure multipoint

interfaces. As defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, this encapsulation type allows ATM interfaces to connect to devices that support only bridge protocol data units (BPDUs). Junos OS does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet LLC/SNAP frames with IP or ARP in the payload, and drops the rest. For packets destined to the Ethernet LAN, a route lookup is done using the destination IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and MAC header, and the packet is forwarded to the ATM interface.

**ethernet-tcc**—For interfaces that carry IPv4 traffic, use Ethernet TCC encapsulation on interfaces that must accept packets carrying standard TPID values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.

**ethernet-vpls**—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard TPID values. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

**ethernet-vpls-fr**—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

**ethernet-vpls-ppp**—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 PPP connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use it to forward the packet into a given VPLS instance.

**ether-vpls-over-atm-llc**—For ATM intelligent queuing (IQ) interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

**extended-frame-relay-ccc**—Use Frame Relay encapsulation on CCC circuits. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC.

**extended-frame-relay-ether-type-tcc**—Use extended Frame Relay ether type TCC for Cisco-compatible Frame Relay for DLCIs 1 through 1022. This encapsulation type is used for circuits with different media on either side of the connection.

**extended-frame-relay-tcc**—Use Frame Relay encapsulation on TCC circuits to connect different media. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC.

**extended-vlan-bridge**—Use extended VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q VLAN tagging and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

**extended-vlan-ccc**—Use extended VLAN encapsulation on CCC circuits with Gigabit Ethernet and 4-port Fast Ethernet interfaces that must accept packets carrying 802.1Q values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC is not supported.

**extended-vlan-tcc**—For interfaces that carry IPv4 traffic, use extended VLAN encapsulation on TCC circuits with Gigabit Ethernet interfaces on which you want to use 802.1Q tagging. For 4-port Gigabit Ethernet PICs, extended VLAN TCC is not supported.

**extended-vlan-vpls**—Use extended VLAN VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



**NOTE:** The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

**flexible-ethernet-services**—For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. Aggregated Ethernet bundles can use this encapsulation type. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

**flexible-frame-relay**—For IQ interfaces only, use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.

**frame-relay**—Use Frame Relay encapsulation.

**frame-relay-ccc**—Use Frame Relay encapsulation on CCC circuits.

**frame-relay-ether-type**—Use Frame Relay ether type encapsulation for compatibility with the Cisco Frame Relay.

**frame-relay-ether-type-tcc**—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media.

**frame-relay-port-ccc**—Use Frame Relay port CCC encapsulation to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. When you use this encapsulation type, you can configure the **ccc** family only.

**frame-relay-tcc**—Use Frame Relay encapsulation on TCC circuits to connect different media.

**generic-services**—Use generic services encapsulation for services with a hierarchical scheduler.

**multilink-frame-relay-uni-nni**—Use MLFR UNI NNI encapsulation. This encapsulation is used on link services, voice services interfaces functioning as FRF.16 bundles, and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

**ppp**—Use serial PPP encapsulation.

**ppp-ccc**—Use serial PPP encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

**ppp-tcc**—Use serial PPP encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

**vlan-ccc**—Use Ethernet VLAN encapsulation on CCC circuits.

**vlan-vci-ccc**—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only. All logical interfaces configured on the Ethernet interface must also have the encapsulation type set to **vlan-vci-ccc**.

**vlan-vpls**—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



**NOTE:**

- Label-switched interfaces (LSIs) do not support VLAN VPLS encapsulation. Therefore, you can only use VLAN VPLS encapsulation on a PE-router-to-CE-router interface and not a core-facing interface.
  - Starting with Junos OS release 13.3, a commit error occurs when you configure **vlan-vpls** encapsulation on a physical interface and configure family **inet** on one of the logical units. Previously, it was possible to commit this invalid configuration.
- 

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

**Related  
Documentation**

- *Configuring Interface Encapsulation on Physical Interfaces*
- *Configuring CCC Encapsulation for Layer 2 VPNs*
- *Configuring Layer 2 Switching Cross-Connects Using CCC*
- *Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits*
- *Configuring ATM Interface Encapsulation*
- *Configuring ATM-to-Ethernet Interworking*
- *Configuring VLAN and Extended VLAN Encapsulation*
- *OBSOLETE: Configuring Extended VLAN Encapsulation*
- *Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces*
- *Configuring Interfaces for Layer 2 Circuits*
- *Configuring Interface Encapsulation on PTX Series Packet Transport Routers*
- *Configuring MPLS LSP Tunnel Cross-Connects Using CCC*
- *Configuring TCC*
- *Configuring VPLS Interface Encapsulation*
- *Configuring Interfaces for VPLS Routing*
- *Defining the Encapsulation for Switching Cross-Connects*

## family

```

Syntax  family family {
        accounting {
            destination-class-usage;
            source-class-usage {
                (input | output | input output);
            }
        }
        access-concentrator name;
        address address {
            ... the address subhierarchy appears after the main [edit interfaces interface-name unit
                logical-unit-number family family-name] hierarchy ...
        }
        bundle interface-name;
        core-facing;
        demux-destination {
            destination-prefix;
        }
        demux-source {
            source-prefix;
        }
        direct-connect;
        duplicate-protection;
        dynamic-profile profile-name;
        filter {
            group filter-group-number;
            input filter-name;
            input-list [ filter-names ];
            output filter-name;
            output-list [ filter-names ];
        }
        interface-mode (access | trunk);
        ipsec-sa sa-name;
        keep-address-and-control;
        mac-validate (loose | strict);
        max-sessions number;
        max-sessions-vsa-ignore;
        mtu bytes;
        multicast-only;
        negotiate-address;
        no-redirects;
        policer {
            arp policer-template-name;
            input policer-template-name;
            output policer-template-name;
        }
        primary;
        protocols [inet iso mpls];
        proxy inet-address address;
        receive-options-packets;
        receive-ttl-exceeded;
        remote (inet-address address | mac-address address);
        rpf-check {

```

```

    fail-filter filter-name
    mode loose;
}
sampling {
    input;
    output;
}
service {
    input {
        post-service-filter filter-name;
        service-set service-set-name <service-filter filter-name>;
    }
    output {
        service-set service-set-name <service-filter filter-name>;
    }
}
service-name-table table-name;
short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
    maximum-seconds>;
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
translate-plp-control-word-de;
unnumbered-address interface-name destination address destination-profile profile-name;
vlan-id number;
vlan-id-list [number number-number];
address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    destination-profile name;
    eui-64;
    master-only;
    multipoint-destination address dlci dlci-identifier;
    multipoint-destination address {
        epd-threshold cells;
        inverse-arp;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (disable | seconds);
        shaping {
            (cbr rate | rtvbr burst length peak rate sustained rate | vbr burst length peak rate
                sustained rate);
            queue-length number;
        }
        vci vpi-identifier.vci-identifier;
    }
    preferred;
    primary;
    vrrp-group group-id {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
    }
}

```

```
(preempt | no-preempt) {  
    hold-time seconds;  
}  
priority number;  
track {  
    interface interface-name {  
        bandwidth-threshold bits-per-second priority-cost priority;  
        priority-cost priority;  
    }  
    priority-hold-time seconds;  
    route prefix routing-instance instance-name priority-cost priority;  
}  
}  
virtual-address [ addresses ];  
}  
virtual-link-local-address ipv6-address;  
}
```

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. Option <b>max-sessions-vs-a-ignore</b> introduced in Junos OS Release 11.4.
Description	Configure protocol family information for the logical interface.



**NOTE:** Not all subordinate stanzas are available to every protocol family.

---



**Options** *family*—Protocol family:

- **any**—Protocol-independent family used for Layer 2 packet filtering



**NOTE:** This option is not supported on T4000 Type 5 FPCs.

- **bridge**—(M Series and T Series routers only) Configure only when the physical interface is configured with **ethernet-bridge** type encapsulation or when the logical interface is configured with **vlan-bridge** type encapsulation
- **ethernet-switching**—(M Series and T Series routers only) Configure only when the physical interface is configured with **ethernet-bridge** type encapsulation or when the logical interface is configured with **vlan-bridge** type encapsulation
- **ccc**—Circuit cross-connect protocol suite
- **inet**—Internet Protocol version 4 suite
- **inet6**—Internet Protocol version 6 suite
- **iso**—International Organization for Standardization Open Systems Interconnection (ISO OSI) protocol suite
- **mlfr-end-to-end**—Multilink Frame Relay FRF.15
- **mlfr-uni-nni**—Multilink Frame Relay FRF.16
- **multilink-ppp**—Multilink Point-to-Point Protocol
- **mpls**—Multiprotocol Label Switching (MPLS)
- **pppoe**—Point-to-Point Protocol over Ethernet
- **tcc**—Translational cross-connect protocol suite
- **tnp**—Trivial Network Protocol
- **vpls**—(M Series and T Series routers only) Virtual private LAN service

The remaining statements are explained separately.

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

**Related Documentation** • *Configuring the Protocol Family*

## fast-aps-switch

---

<b>Syntax</b>	fast-aps-switch;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> sonet-options aps]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	(M320 routers with Channelized OC3/STM1 Circuit Emulation PIC with SFP only and EX Series switches) Reduce the Automatic Protection Switching (APS) switchover time in Layer 2 circuits.

**NOTE:**

- Configuring this statement reduces the APS switchover time only when the Layer 2 circuit encapsulation type for the interface receiving traffic from a Layer 2 circuit neighbor is SAToP.
  - When the fast-aps-switch statement is configured in revertive APS mode, you must configure an appropriate value for revert time to achieve reduction in APS switchover time.
  - To prevent the logical interfaces in the data path from being shut down, configure appropriate hold-time values on all the interfaces in the data path that support TDM.
  - The fast-aps-switch statement cannot be configured when the APS annex-b option is configured.
  - The interfaces that have the fast-aps-switch statement configured cannot be used in virtual private LAN service (VPLS) environments.
- 

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Reducing APS Switchover Time in Layer 2 Circuits</i></li></ul>

## global-mac-statistics

---

<b>Syntax</b>	global-mac-statistics;
<b>Hierarchy Level</b>	[edit protocols <a href="#">l2-learning</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2. Support for logical systems added in Junos OS Release 9.6.
<b>Description</b>	(MX Series routers and EX Series switches only) Enable MAC accounting for the entire router or switch.
<b>Default</b>	disabled
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Enabling MAC Accounting</i></li> </ul>

## global-mac-table-aging-time

---

<b>Syntax</b>	global-mac-table-aging-time <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit protocols <a href="#">l2-learning</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2. Support for logical systems added in Junos OS Release 9.6.
<b>Description</b>	Configure the timeout interval for entries in the MAC table.
<b>Default</b>	300 seconds
<b>Options</b>	<p><b>seconds</b>—Time elapsed before MAC table entries are timed out and entries are deleted from the table.</p> <p><b>Range:</b> For MX Series routers: 10 through 1 million; for EX Series and QFX Series switches: 60 through 1 million</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the MAC Table Timeout Interval</i></li> <li>• <i>Configuring MAC Table Aging (CLI Procedure)</i></li> <li>• <i>Configuring MAC Table Aging</i></li> </ul>

## global-no-mac-learning

---

<b>Syntax</b>	global-no-mac-learning;
<b>Hierarchy Level</b>	[edit protocols <a href="#">l2-learning</a> ], [edit protocols l2-learning]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2. Support for logical systems added in Junos OS Release 9.6.
<b>Description</b>	Disable MAC learning on the entire device.
<b>Default</b>	MAC learning is enabled.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Disabling Layer 2 Learning and Forwarding on page 20</a></li><li>• <i>Understanding Q-in-Q Tunneling on EX Series Switches</i></li></ul>

## inner-tag-protocol-id


<b>Syntax</b>	<code>inner-tag-protocol-id <i>tpid</i>;</code>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Configure the IEEE 802.1Q TPID value to rewrite for the inner tag.</p> <p>All TPIDs you include in input and output VLAN maps must be among those you specify at the [edit interfaces <i>interface-name</i> <b>gether-options ethernet-switch-profile tag-protocol-id [ <i>tpids</i> ]</b>] hierarchy level.</p> <p>On MX Series routers, you can use this statement for Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs.</p>
<b>Default</b>	If the <b>inner-tag-protocol-id</b> statement is not configured, the TPID value is 0x8100.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring Inner and Outer TPIDs and VLAN IDs</i></li> </ul>

## inner-vlan-id

---

<b>Syntax</b>	<code>inner-vlan-id <i>number</i>;</code>
<b>Hierarchy Level</b>	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code><a href="#">input-vlan-map</a>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code><a href="#">output-vlan-map</a>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	<p>For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers or 100-Gigabit Ethernet Type 5 PIC with CFP, or on Ethernet interfaces on EX Series switches, specify the VLAN ID to rewrite for the inner tag of the final packet.</p> <p>You cannot include the <b>inner-vlan-id</b> statement with the <b>swap</b> statement, <b>swap-push</b> statement, <b>push-push</b> statement, or <b>push-swap</b> statement and the <b>inner-vlan-id</b> statement at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a>]</code> hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the <b>inner-vlan-id</b> statement you include at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code> hierarchy level.</p>
<b>Options</b>	<b><i>number</i></b> —VLAN ID number. <b>Range:</b> 0 through 4094
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li><i>Configuring Inner and Outer TPIDs and VLAN IDs</i></li></ul>

## input-vlan-map

<b>Syntax</b>	<pre>input-vlan-map {   (pop   pop-pop   pop-swap   push   push-push   swap   swap-push   swap-swap);   inner-tag-protocol-id <i>tpid</i>;   inner-vlan-id <i>number</i>;   tag-protocol-id <i>tpid</i>;   vlan-id <i>number</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[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>]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>pop-pop</b>, <b>pop-swap</b>, <b>push-push</b>, <b>swap-push</b>, and <b>swap-swap</b> statements introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.</p>
<b>Description</b>	<p>For Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP interfaces, 100-Gigabit Ethernet Type 5 PIC with CFP only as well as Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces, define the rewrite profile to be applied to incoming frames on this logical interface.</p> <p>The statements are explained separately.</p>
<div style="display: flex; align-items: flex-start;"> <div style="flex: 1; text-align: center; margin-right: 10px;">  </div> <div> <p><b>NOTE:</b> Connectivity fault management (CFM) sessions for all interfaces in which <b>input-vlan-map</b> is configured are supported only if the interface also has an explicit configuration for <b>output-vlan-map</b> as <b>output-vlan-map pop</b>. See <a href="#">output-vlan-map</a>. This configuration is required for all the interfaces in the topology even when the CFM session is on that interface or on a different interface in the data path of the same topology.</p> </div> </div>	
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Stacking a VLAN Tag</i></li> <li>• <a href="#">output-vlan-map on page 123</a></li> <li>• <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i></li> </ul>

## interface

---

<b>Syntax</b>	<code>interface <i>interface-name</i>;</code>
<b>Hierarchy Level</b>	[edit bridge-domains <i>bridge-domain-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> ], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> ], [edit <a href="#">vlans on page 84</a> <i>vlan-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4. Support for top-level configuration for the <b>virtual-switch</b> type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch. Support for logical systems added in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.3R2 for EX Series switches. Statement introduced in Junos OS Release 15.1.
<b>Description</b>	(MX Series routers and EX Series switches only) Specify the logical interfaces to include in the bridge domain, VLAN, VPLS instance, or virtual switch.
<b>Options</b>	<i>interface-name</i> —Name of a logical interface. For more information about how to configure logical interfaces, see the <i>Junos OS Network Interfaces Library for Routing Devices</i> .
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring a Bridge Domain</i></li><li>• <i>Configuring a Layer 2 Virtual Switch</i></li><li>• <a href="#">Configuring a Layer 2 Virtual Switch on page 33</a></li><li>• <i>Tunnel Services Overview</i></li><li>• <i>Configuring Tunnel Interfaces on MX Series Routers</i></li></ul>



## interface (MVRP)

<b>Syntax</b>	<pre>interface (all   <i>interface-name</i>) {   join-timer <i>milliseconds</i>;   leave-timer <i>milliseconds</i>;   leaveall-timer <i>milliseconds</i>;   point-to-point;   registration (forbidden   normal   restricted); }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">mvrp</a>],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp</a>] (for virtual switch instance type),  [edit protocols <a href="#">mvrp</a>],  [edit routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp</a>] (for virtual switch instance type)</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1 for MX Series routers.
<b>Description</b>	Specify interfaces on which to configure Multiple VLAN Registration Protocol (MVRP).
<b>Default</b>	By default, MVRP is disabled.
<b>Options</b>	<p><b>all</b>—All interfaces on the router or switch.</p> <p><b><i>interface-name</i></b>—Names of interface to be configured for MVRP.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers</a></li> <li>• <a href="#">Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration on page 52</a></li> <li>• <a href="#">Verifying That MVRP Is Working Correctly on page 54</a></li> <li>• <a href="#">Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration on page 49</a></li> </ul>

## interface (Spanning Tree)

<b>Syntax</b>	<pre> interface <i>interface-name</i> {     bpdv-timeout-action {         alarm;         block;     }     cost <i>cost</i>;     edge;     mode (p2p   shared);     no-root-port;     priority <i>interface-priority</i>; } </pre>
<b>Hierarchy Level</b>	<pre> [edit logical-systems <i>logical-system-name</i> protocols (mstp   <b>rstp</b>   vstp)], [edit logical-systems <i>logical-system-name</i> protocols vstp vlan <i>vlan-id</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp   <b>rstp</b>   vstp)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp vlan <i>vlan-id</i>], [edit protocols (mstp   <b>rstp</b>   vstp)], [edit protocols vstp vlan <i>vlan-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols (mstp   <b>rstp</b>   vstp)], [edit routing-instances <i>routing-instance-name</i> protocols vstp vlan <i>vlan-id</i>] </pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>
<b>Description</b>	Configure the interface to participate in the RSTP, MSTP, or VSTP instance.
<b>Options</b>	<p><b><i>interface-name</i></b>—Name of a Gigabit Ethernet or 10-Gigabit Ethernet interface.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Spanning-Tree Instance Interface</i></li> <li>• <i>show spanning-tree interface</i></li> <li>• <i>Understanding RSTP for EX Series and QFX Series Switches</i></li> <li>• <i>Understanding MSTP for EX Series and QFX Series Switches</i></li> <li>• <i>Understanding VSTP for EX Series Switches and QFX Series Switches</i></li> </ul>

## interface-mode

<b>Syntax</b>	<code>interface-mode (access   trunk &lt;inter-switch-link&gt;);</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 15.1. <b>inter-switch-link</b> option introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

### Description



**NOTE:** This statement supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *port-mode*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

(QFX Series 3500 and 3600 standalone switches)—Determine whether the logical interface accepts or discards packets based on VLAN tags. Specify the **trunk** option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the **vlan-id** or **vlan-id-list** statement, then forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the **access** option to accept packets with no VLAN ID, then forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the **vlan-id** statement.



**NOTE:** On MX Series routers, if you want IGMP snooping to be functional for a bridge domain, then you should not configure **interface-mode** and **irb** for that bridge. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed. For more information, see *Configuring a Trunk Interface on a Bridge Network*.

<b>Options</b>	<p><b>access</b>—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the <b>vlan-id</b> statement.</p> <p><b>trunk</b>—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the <b>vlan-id</b> or <b>vlan-id-list</b> statement.</p> <p><b>trunk inter-switch-link</b>—For a private VLAN, configure the InterSwitch Link protocol (ISL) on a trunk port of the primary VLAN in order to connect the switches composing the PVLAN to each other. You do not need to configure an ISL when a PVLAN is configured</p>
----------------	---

on a single switch. This configuration specifies whether the particular interface assumes the role of interswitch link for the PVLAN domains of which it is a member. This option is supported only on MX240, MX480, and MX960 routers in enhanced LAN mode.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Access Mode on a Logical Interface</i></li><li>• <i>Configuring a Logical Interface for Trunk Mode</i></li><li>• <i>Example: Connecting Access Switches to a Distribution Switch</i></li><li>• <i>Tunnel Services Overview</i></li><li>• <i>Configuring Tunnel Interfaces on MX Series Routers</i></li></ul>

---

## interfaces

---

<b>Syntax</b>	interfaces { ... }
<b>Hierarchy Level</b>	[edit]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure interfaces on the router.
<b>Default</b>	The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Physical Interface Configuration Statements Overview</i></li><li>• <i>Configuring Aggregated Ethernet Link Protection</i></li></ul>

## inter-switch-link

<b>Syntax</b>	<code>inter-switch-link vlan members <i>primary-vlan-name</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching <b>interface-mode</b> trunk]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.
<b>Description</b>	Use this configuration statement when a private VLAN (PVLAN) spans multiple switches. The Inter-Switch Link protocol (ISL) must be configured on a trunk port of the primary VLAN in order to connect the switches composing the PVLAN to each other. You do not need to configure an ISL when a PVLAN is configured on a single switch.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Creating a Private VLAN Spanning Multiple Switches (CLI Procedure) on page 47</a></li> </ul>


## isid

<b>Syntax</b>	<code>isid <i>isid-number</i> vlan-id-list [ <i>vlan-ids</i> ] {     source-bmac &lt;<i>mac-address</i>&gt; &lt;<i>length</i>&gt; }</code>
<b>Hierarchy Level</b>	[edit routing-instances <i>routing-instance-name</i> service-groups <i>service-group-name</i> pbb-service-options]
<b>Release Information</b>	Statement introduced in JUNOS Release 10.0.
<b>Description</b>	For IEEE 802.1ah provider backbone bridge (PBB) configurations, configure the service identifier (I-SID) for the customer routing instance (I-component) service group.
<b>Options</b>	<p><i>isid</i>—Service identifier. Enter an I-SID in the range from 256 through 16777214.</p> <p><i>vlan-id-list</i> [ <i>vlan-ids</i> ]—List of service VLANs (S-VLANs).</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring E-LINE and E-LAN Services for a PBB Network on MX Series Routers</a></li> </ul>

## isid-list

<b>Syntax</b>	isid-list all-service-groups;
<b>Hierarchy Level</b>	[edit interfaces <i>pseudo-logical-interface-name</i> unit <i>logical-unit-number</i> <b>family</b> bridge]
<b>Release Information</b>	Statement introduced in JUNOS Release 10.0.
<b>Description</b>	For IEEE 802.1ah provider backbone bridge (PBB) configurations, map all service identifiers (I-SIDs) specified for the service groups.
<b>Options</b>	<b>all-service-groups</b> —Map all service identifiers (I-SIDs) for the specified service groups.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring E-LINE and E-LAN Services for a PBB Network on MX Series Routers</i></li> </ul>

## isolated-vlan

<b>Syntax</b>	isolated-vlan vlan-name <i>isolated-vlan-name</i> vlan-id <i>isolated-vlan-id</i> ;
<b>Hierarchy Level</b>	[edit vlans <i>primary-vlan-name</i> vlan-id <i>primary-vlan-vlan-id</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.
<b>Description</b>	Configure the specified isolated VLAN to be a secondary VLAN of the specified primary VLAN. An isolated VLAN receives packets only from the primary VLAN and forwards frames upstream to the primary VLAN.
<div style="display: flex; align-items: center;">  <div> <p><b>NOTE:</b> Before you specify this configuration statement, you must have already configured an isolated VLAN and assigned a VLAN ID to it. See <a href="#">private-vlan</a>.</p> </div> </div>	
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Creating a Private VLAN on a Single Switch (CLI Procedure) on page 45</a></li> <li>• <a href="#">Creating a Private VLAN Spanning Multiple Switches (CLI Procedure) on page 47</a></li> </ul>

## join-timer (MVRP)

<b>Syntax</b>	<code>join-timer <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">mvrp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp</a> interface (all   <i>interface-name</i>)] (for virtual switch instance type),</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp</a>] (for virtual switch instance type),</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">mvrp</a> interface (all   <i>interface-name</i>)],</p> <p>[edit protocols <a href="#">mvrp</a> interface (all   <i>interface-name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp</a>] (for virtual switch instance type),</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp</a> interface (all   <i>interface-name</i>)] (for virtual switch instance type)</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1 for MX Series routers.
<b>Description</b>	For Multiple VLAN Registration Protocol (MVRP), configure the maximum interval interfaces must wait before sending MVRP protocol data units (PDUs).
<b>Default</b>	200 milliseconds
<b>Options</b>	<p><i>milliseconds</i>—Interval that the interface must wait before sending MVRP 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.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers</a></li> <li>• <a href="#">Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration on page 52</a></li> <li>• <a href="#">Verifying That MVRP Is Working Correctly on page 54</a></li> <li>• <a href="#">Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration on page 49</a></li> <li>• <a href="#">leaveall-timer on page 113</a></li> <li>• <a href="#">leave-timer on page 114</a></li> <li>• <a href="#">point-to-point on page 127</a></li> <li>• <a href="#">registration on page 135</a></li> </ul>

## l2-learning

---

<b>Syntax</b>	<pre>l2-learning {     global-mac-limit <i>limit</i>;     global-mac-statistics;     global-mac-table-aging-time <i>seconds</i>;     global-no-mac-learning; }</pre>
<b>Hierarchy Level</b>	[edit protocols]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4.
<b>Description</b>	(MX Series routers, and EX Series switches, and QFX Series switches only) Configure Layer 2 address learning and forwarding properties globally.  The statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Understanding Layer 2 Learning and Forwarding</i></li></ul>

## l3-interface

---

<b>Syntax</b>	<pre>l3-interface <i>interface-name-logical-unit-number</i>;</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>ge-chassis/slot/port</i> unit <i>logical-unit-number</i> family ethernet-switching] [edit vlans <i>vlan-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. [edit vlans <i>vlan-name</i> ] hierarchy level introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers.
<b>Description</b>	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.
<b>Default</b>	No Layer 3 (routing) interface is associated with the VLAN.
<b>Options</b>	<i>interface-name-logical-unit-number</i> —Name of a logical interface.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.



## leaveall-timer (MVRP)

<b>Syntax</b>	<code>leaveall-timer <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">mvrp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp</a> interface (all   <i>interface-name</i>)] (for virtual switch instance type),</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp</a>] (for virtual switch instance type),</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">mvrp</a> interface (all   <i>interface-name</i>)],</p> <p>[edit protocols <a href="#">mvrp</a> interface (all   <i>interface-name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp</a>] (for virtual switch instance type),</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp</a> interface (all   <i>interface-name</i>)] (for virtual switch instance type)</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1 for MX Series routers.
<b>Description</b>	For Multiple VLAN Registration Protocol (MVRP), configure the interval at which the LeaveAll state operates on the interface.
<b>Default</b>	10000 milliseconds
<b>Options</b>	<i>milliseconds</i> —Interval between the sending of Leave All messages. 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.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers</a></li> <li>• <a href="#">Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration on page 52</a></li> <li>• <a href="#">Verifying That MVRP Is Working Correctly on page 54</a></li> <li>• <a href="#">Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration on page 49</a></li> <li>• <a href="#">join-timer on page 111</a></li> <li>• <a href="#">leave-timer on page 114</a></li> <li>• <a href="#">point-to-point on page 127</a></li> <li>• <a href="#">registration on page 135</a></li> </ul>

## leave-timer (MVRP)

<b>Syntax</b>	<code>leave-timer <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols <b>mvrp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>mvrp</b> interface (all   <i>interface-name</i>)] (for virtual switch instance type),</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>mvrp</b>] (for virtual switch instance type),</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>mvrp</b> interface (all   <i>interface-name</i>)],</p> <p>[edit protocols <b>mvrp</b> interface (all   <i>interface-name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>mvrp</b>] (for virtual switch instance type),</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>mvrp</b> interface (all   <i>interface-name</i>)] (for virtual switch instance type)</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1 for MX Series routers.
<b>Description</b>	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.
<b>Default</b>	1000 milliseconds
<b>Options</b>	<b><i>milliseconds</i></b> —Interval 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. 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.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers</a></li> <li>• <a href="#">Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration on page 52</a></li> <li>• <a href="#">Verifying That MVRP Is Working Correctly on page 54</a></li> <li>• <a href="#">Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration on page 49</a></li> <li>• <a href="#">join-timer on page 111</a></li> <li>• <a href="#">leaveall-timer on page 113</a></li> <li>• <a href="#">point-to-point on page 127</a></li> <li>• <a href="#">registration on page 135</a></li> </ul>

## mac-statistics

<b>Syntax</b>	mac-statistics;
<b>Hierarchy Level</b>	<p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options],</p> <p>[edit logical-systems <i>logical-system-name</i> switch-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> switch-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols evpn],</p> <p>[edit switch-options],</p> <p>[edit <a href="#">switch-options on page 84</a>],</p> <p>[edit <a href="#">vlans on page 84</a> <i>vlan-name</i> switch-options]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the <b>switch-options</b> statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the <b>virtual-switch</b> type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options] and [edit vlans <i>vlan-name</i> switch-options] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers.</p> <p>[edit switch-options] and [edit vlans <i>vlan-name</i> switch-options] hierarchy levels introduced in Junos OS Release 13.2 for the QFX Series.</p>
<b>Description</b>	(MX Series routers, EX Series switches, and QFX Series only) For bridge domains or VLANs, enable MAC accounting either for a specific bridge domain or VLAN, or for a set of bridge domains or VLANs associated with a Layer 2 trunk port.
<b>Default</b>	disabled
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Understanding Layer 2 Learning and Forwarding for Bridge Domains</i></li> <li>• <a href="#">Layer 2 Learning and Forwarding for VLANs Overview on page 16</a></li> <li>• <i>Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports</i></li> <li>• <a href="#">Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port on page 17</a></li> </ul>

- *Configuring EVPN Routing Instances*
- *Configuring EVPN Routing Instances on EX9200 Switches*

## mac-table-size

<b>Syntax</b>	<code>mac-table-size <i>limit</i> {     <code>packet-action</code> drop; }</code>
<b>Hierarchy Level</b>	<p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options],          [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options],          [edit logical-systems <i>logical-system-name</i> switch-options],          [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],          [edit routing-instances <i>routing-instance-name</i> switch-options],          [edit switch-options],  <a href="#">[edit switch-options on page 84]</a>,  <a href="#">[edit vlans on page 84 <i>vlan-name</i> switch-options]</a></p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the <b>switch-options</b> statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the <b>virtual-switch</b> type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p><b>[edit switch-options]</b> and <b>[edit vlans <i>vlan-name</i> switch-options]</b> hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the <b>[edit vlans <i>vlan-name</i> switch-options]</b> hierarchy level introduced in Junos OS Release 13.2 for the QFX Series.</p>
<b>Description</b>	Modify the size of the MAC address table for the bridge domain or VLAN, a set of bridge domains or VLANs associated with a trunk port, or a virtual switch. The default is 5120 MAC addresses.



**NOTE:** For multichassis link aggregation (MC-LAG) peers in active-active mode, configuring the **mac-table-size** statement or changing the **mac-table-size** configuration when traffic is flowing can cause the MAC entries to be out of synchronization between the two MC-LAG peers, which might result in flooding. To avoid flooding, you must either halt traffic forwarding and then configure the **mac-table-size** statement or use the **commit at** configuration statement to commit the changes at the same time in both the peer nodes.

Alternatively, if flooding does occur, you can clear the bridge MAC table on both the routers by using the **clear bridge mac-table** command. Running this command ensures that the MAC entries are re-learned and in synchronization between both the peers.

**Options** *limit*—Specify the maximum number of addresses in the MAC address table.

**Range:** 16 through 1,048,575 MAC addresses

**Default:** 5120 MAC addresses There is no default MAC address limit for the **mac-table-size** statement at the **[edit switch-options]** hierarchy level. The number of MAC addresses that can be learned is only limited by the platform, 65,535 MAC addresses for EX Series switches and 1,048,575 MAC addresses for other devices.

The remaining statement is explained separately.

**Required Privilege** routing—To view this statement in the configuration.

**Level** routing-control—To add this statement to the configuration.


**Related  
Documentation**

- *Understanding Layer 2 Learning and Forwarding for Bridge Domains*
- [Layer 2 Learning and Forwarding for VLANs Overview on page 16](#)
- *Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
- [Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port on page 17](#)

## mvrp

<b>Syntax</b>	<pre> mvrp {   bpd-destination-mac-address provider-bridge-group;   join-timer milliseconds;   leave-timer milliseconds;   leaveall-timer milliseconds;   interface (all   interface-name) {     join-timer milliseconds;     leave-timer milliseconds;     leaveall-timer milliseconds;     point-to-point;     registration (forbidden   normal   restricted);   }   no-dynamic-vlan;   traceoptions {     file filename &lt;files number &gt; &lt;size size&gt; &lt;no-stamp   world-readable         no-world-readable&gt;;     flag flag;   } } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols]  (for virtual switch instance type),  [edit protocols],  [edit routing-instances <i>routing-instance-name</i> protocols] (for virtual switch instance type),</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1 for MX Series routers.
<b>Description</b>	<p>For Layer 2 networks, configure Multiple VLAN Registration Protocol (MVRP) to dynamically share VLAN information and dynamically configure needed VLANs. Maintaining VLAN configurations based on active VLANs reduces the amount of traffic traveling in the network, saving network resources. MVRP is configured on trunk interfaces.</p> <p>The remaining statements are explained separately.</p>
<b>Default</b>	MVRP is disabled by default.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.  routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers</a></li> <li>• <a href="#">Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration on page 52</a></li> <li>• <a href="#">Verifying That MVRP Is Working Correctly on page 54</a></li> <li>• <a href="#">Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration on page 49</a></li> </ul>

## native-vlan-id

<b>Syntax</b>	<code>native-vlan-id <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>ge-fpc/pic/port</i> ], [edit interfaces <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers. Statement introduced in Junos OS Release 12.3R2 for EX Series switches. Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.
<b>Description</b>	<p>Configure mixed tagging support for untagged packets on a port for the following:</p> <ul style="list-style-type: none"> <li>• M Series routers with Gigabit Ethernet IQ PICs with SFP and Gigabit Ethernet IQ2 PICs with SFP configured for 802.1Q flexible VLAN tagging</li> <li>• MX Series routers with Gigabit Ethernet DPCs and MICs, Tri-Rate Ethernet DPCs and MICs, and 10-Gigabit Ethernet DPCs and MICs and MPCs configured for 802.1Q flexible VLAN tagging</li> <li>• T4000 routers with 100-Gigabit Ethernet Type 5 PIC with CFP</li> <li>• EX Series switches with Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces</li> </ul> <p>When the <b>native-vlan-id</b> statement is included with the <i>flexible-vlan-tagging</i> statement, untagged packets are accepted on the same mixed VLAN-tagged port.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> The logical interface on which untagged packets are received must be configured with the same VLAN ID as the native VLAN ID configured on the physical interface, otherwise the untagged packets are dropped.</p> </div> <p>To configure the logical interface, include the <b>vlan-id</b> statement (matching the <b>native-vlan-id</b> statement on the physical interface) at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.</p> <p>When the <b>native-vlan-id</b> statement is included with the <b>interface-mode</b> statement, untagged packets are accepted and forwarded within the bridge domain or VLAN that is configured with the matching VLAN ID.</p>
<b>Default</b>	By default, the untagged packets are dropped. That is, if you do not configure the <b>native-vlan-id</b> option, the untagged packets are dropped.
<b>Options</b>	<p><b>number</b>—VLAN ID number.</p> <p><b>Range:</b> (ACX Series routers and EX Series switches) 0 through 4094.</p>



<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Mixed Tagging Support for Untagged Packets</i></li><li>• <i>Configuring Access Mode on a Logical Interface</i></li><li>• <i>Configuring the Native VLAN Identifier (CLI Procedure)</i></li><li>• <i>Understanding Bridging and VLANs on EX Series Switches</i></li><li>• <i>flexible-vlan-tagging</i></li><li>• <i>Understanding Q-in-Q Tunneling on EX Series Switches</i></li></ul>

## no-mac-learning

<b>Syntax</b>	no-mac-learning;
<b>Hierarchy Level</b>	<p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options],</p> <p>[edit logical-systems <i>logical-system-name</i> switch-options],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols evpn],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols evpn interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> switch-options],</p> <p>[edit switch-options],</p> <p>[edit switch-options interface <i>interface-name</i>],</p> <p>[set vlans <i>vlan-name</i> switch-options]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the <b>switch-options</b> statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the <b>virtual-switch</b> type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options], [edit switch-options interface <i>interface-name</i>], [edit vlans <i>vlan-name</i> switch-options], and [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>] hierarchy levels introduced in Junos OS Release 12.3 R2 for EX Series switches.</p> <p>Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers.</p> <p>Hierarchy levels [edit switch-options interface <i>interface-name</i>] and [edit vlans <i>vlan-name</i> switch-options] introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p>
<b>Description</b>	<p>For MX Series routers and EX Series switches, disable MAC learning for a virtual switch, for a bridge domain or VLAN, for a specific logical interface in a bridge domain or VLAN, or for a set of bridge domains or VLANs associated with a Layer 2 trunk port. On platforms that support EVPNs, you can disable MAC learning on an EVPN.</p>



**NOTE:** When MAC learning is disabled for a VPLS routing instance, traffic is not load-balanced and only one of the equal-cost next hops is used.

<b>Default</b>	MAC learning is enabled.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring EVPN Routing Instances</i></li> <li>• <i>Configuring EVPN Routing Instances on EX9200 Switches</i></li> <li>• <i>Understanding Layer 2 Learning and Forwarding for Bridge Domains</i></li> <li>• <a href="#">Layer 2 Learning and Forwarding for VLANs Overview on page 16</a></li> <li>• <i>Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports</i></li> <li>• <i>Understanding Bridging and VLANs on EX Series Switches</i></li> <li>• <i>Understanding Q-in-Q Tunneling on EX Series Switches</i></li> </ul>

## output-vlan-map (Gigabit Ethernet IQ and 10-Gigabit Ethernet with SFPP)

<b>Syntax</b>	<pre>output-vlan-map {   (pop   pop-pop   pop-swap   push   push-push   swap   swap-push   swap-swap);   inner-tag-protocol-id <i>tpid</i>;   inner-vlan-id <i>number</i>;   tag-protocol-id <i>tpid</i>;   vlan-id <i>number</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[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>]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>pop-pop</b>, <b>pop-swap</b>, <b>push-push</b>, <b>swap-push</b>, and <b>swap-swap</b> statements added in Junos OS Release 8.1.</p>
<b>Description</b>	<p>For Gigabit Ethernet IQ and 10-Port 10-Gigabit Ethernet SFPP interfaces only, define the rewrite operation to be applied to outgoing frames on this logical interface.</p> <p>The statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Stacking and Rewriting Gigabit Ethernet VLAN Tags</i></li> <li>• <a href="#">input-vlan-map on page 103</a></li> </ul>

## packet-action

**Syntax** `packet-action action;`

**Hierarchy Level** [edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name* interface-mac-limit *limit*],  
 [edit bridge-domains *bridge-domain-name* bridge-options interface-mac-limit *limit*],  
 [edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* interface-mac-limit *limit*],  
 [edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options interface-mac-limit *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* interface-mac-limit *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface-mac-limit *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options interface *interface-name* interface-mac-limit *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options interface-mac-limit *limit*],  
 [edit logical-systems *logical-system-name* switch-options interface-mac-limit *limit*],  
 [edit protocols [l2-learning](#) global-mac-limit *limit*],  
 [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* interface-mac-limit *limit*],  
 [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface-mac-limit *limit*],  
 [edit routing-instances *routing-instance-name* protocols evpn interface-mac-limit (VPLS)],  
 [edit routing-instances *routing-instance-name* protocols evpn interface *interface-name* interface-mac-limit (VPLS)],  
 [edit routing-instances *routing-instance-name* protocols evpn mac-table-size *limit*],  
 [edit routing-instances *routing-instance-name* switch-options interface *interface-name* interface-mac-limit *limit*],  
 [edit routing-instances *routing-instance-name* switch-options interface-mac-limit *limit*],  
 [edit [switch-options on page 84](#) interface *interface-name* interface-mac-limit *limit*],  
 [edit [switch-options on page 84](#) interface-mac-limit *limit*],  
 [edit [switch-options on page 84](#) interface *interface-name* interface-mac-limit *limit*],  
 [edit [switch-options on page 84](#) interface-mac-limit *limit*],  
 [edit [switch-options on page 84 mac-table-size](#) *limit*],  
 [edit switch-options interface *interface-name* interface-mac-limit *limit*],  
 [edit [vlans on page 84](#) *vlan-name* switch-options interface *interface-name* interface-mac-limit *limit*],  
 [edit [vlans on page 84](#) *vlan-name* switch-options interface-mac-limit *limit*],  
 [edit [vlans on page 84](#) *vlan-name* switch-options [mac-table-size](#) *limit*],  
 [edit vlans *vlan-name* switch-options interface-mac-limit *limit*],  
 [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit *limit*],  
 [edit vlans *vlan-name* switch-options [mac-table-size](#) *limit*]

**Release Information** Statement introduced in Junos OS Release 8.4.  
 Support for the **switch-options** statement added in Junos OS Release 9.2.  
 Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy

supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.

Support for logical systems added in Junos OS Release 9.6.

[edit switch-options interface *interface-name* interface-mac-limit *limit*], [edit switch-options interface-mac-limit *limit*], [edit switch-options mac-table-size *limit*], [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit *limit*], [edit vlans *vlan-name* switch-options interface-mac-limit *limit*], and [edit vlans *vlan-name* switch-options mac-table-size *limit*] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

Support for EVPNs introduced in Junos OS Release 13.2 on MX Series 3D Universal Edge Routers.

Support at the [edit switch-options interface *interface-name* interface-mac-limit *limit*] hierarchy level and hierarchy levels under [edit vlans *vlan-name*] introduced in Junos OS Release 13.2X50-D10 for EX Series switches and Junos OS Release 13.2 for the QFX Series.

**Description** Specify the action taken when packets with new source MAC addresses are received after the MAC address limit is reached. If this statement is not configured, packets with new source MAC addresses are forwarded by default.

**Default**



**NOTE:** On a QFX Series Virtual Chassis, if you include the shutdown option at the [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level and issue the commit operation, the system generates a commit error. The system does not generate an error if you include the shutdown option at the [edit switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level.

Disabled. The default is for packets for new source MAC addresses to be forwarded after the MAC address limit is reached.

**Options**

- drop**—Drop packets with new source MAC addresses, and do not learn the new source MAC addresses.
- drop-and-log**—(EX Series switches and QFX Series only) Drop packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.
- log**—(EX Series switches and QFX Series only) Hold packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.
- none**—(EX Series switches and QFX Series only) Forward packets with new source MAC addresses, and learn the new source MAC address.
- shutdown**—(EX Series switches and QFX Series only) Disable the specified interface, and generate an alarm, an SNMP trap, or a system log entry.

**Required Privilege Level**

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


**Related  
Documentation**

- *Configuring EVPN Routing Instances*
- *Configuring EVPN Routing Instances on EX9200 Switches*
- *Configuring MAC Limiting (CLI Procedure)*
- *Configuring Persistent MAC Learning (CLI Procedure)*
- *Understanding Layer 2 Learning and Forwarding for Bridge Domains*
- [Layer 2 Learning and Forwarding for VLANs Overview on page 16](#)
- *Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
- [Layer 2 Learning and Forwarding for VLANs Overview on page 16](#)
- [Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port on page 17](#)

## point-to-point (MVRP)

<b>Syntax</b>	point-to-point;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <a href="#">mvrp interface</a> (all   <i>interface-name</i> )], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp interface</a> (all   <i>interface-name</i> )] (for virtual switch instance type), [edit protocols <a href="#">mvrp interface</a> (all   <i>interface-name</i> )], [edit routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp interface</a> (all   <i>interface-name</i> )] (for virtual switch instance type)
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1 for MX Series routers. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	(Optional) For Multiple VLAN Registration Protocol (MVRP) configurations, configure an interface to be recognized as a point-to-point connection. If specified, a point-to-point subset of the MRP state machine is used to provide a simpler and more efficient method to accelerate convergence on the network. Point-to-point must be enabled after enabling MVRP for the interface to be recognized as a point-to-point connection.
<b>Default</b>	MVRP is disabled by default.  point-to-point is disabled by default.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers</a></li> <li>• <a href="#">Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration on page 52</a></li> <li>• <a href="#">Verifying That MVRP Is Working Correctly on page 54</a></li> <li>• <a href="#">Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration on page 49</a></li> <li>• <a href="#">join-timer on page 111</a></li> <li>• <a href="#">leaveall-timer on page 113</a></li> <li>• <a href="#">leave-timer on page 114</a></li> <li>• <a href="#">registration on page 135</a></li> </ul>

## pop

<b>Syntax</b>	pop;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b> ], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches. Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.
<b>Description</b>	<p> <b>NOTE:</b> On EX4300 switches, pop is not supported at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b>] hierarchy level.</p> <p>For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2, and IQ2-E interfaces; 10-Gigabit Ethernet LAN/WAN PIC; aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces; 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces, specify the VLAN rewrite operation to remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Removing a VLAN Tag</i></li> <li>• <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i></li> </ul>



## pop-pop

---


<b>Syntax</b>	pop-pop;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a> ], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP, and for 10-Gigabit Ethernet SFP interfaces on EX Series switches, specify the VLAN rewrite operation to remove both the outer and inner VLAN tags of the frame.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Removing the Outer and Inner VLAN Tags</i></li> </ul>

## pop-swap


---

<b>Syntax</b>	pop-swap;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b> ], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	<p>Specify the VLAN rewrite operation to remove the outer VLAN tag of the frame, and replace the inner VLAN tag of the frame with a user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame.</p> <p>You can use this statement on Gigabit Ethernet IQ, IQ2, IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag</i></li></ul>


## private-vlan

<b>Syntax</b>	<code>private-vlan (isolated   community) vlan-id <i>number</i>;</code>
<b>Hierarchy Level</b>	<code>[edit vlans <i>vlan-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.
<b>Description</b>	Configure a secondary VLAN ( either an isolated VLAN or a community VLAN) within a private VLAN (PVLAN) and specify a VLAN ID for that secondary VLAN. This statement essentially converts a VLAN into a PVLAN, by carving out discrete subdomains (secondary VLANs) within the primary VLAN. You must specify a VLAN ID for each secondary PVLAN.
<div>  <p><b>NOTE:</b> After you have configured the secondary VLAN, you must also configure its association with a specific primary VLAN. See <a href="#">isolated-vlan</a> and <a href="#">community-vlan</a> for additional information.</p> </div>	
<b>Options</b>	<ul style="list-style-type: none"> <li>• <b>isolated</b> — The VLAN specified by <i>vlan-name</i> is defined as an <i>isolated</i> VLAN and a VLAN-ID is assigned to it. An isolated VLAN receives packets only from the primary VLAN and forwards frames upstream to the primary VLAN. The VLAN name is optional. The VLAN ID is required.</li> <li>• <b>community</b> — The VLAN specified by <i>vlan-name</i> is defined as community VLAN and a VLAN-ID is assigned to it. A <i>community</i> VLAN used to transport frames among members of a community, which is a subset of users within the VLAN, and to forward frames upstream to the primary VLAN. The VLAN name is optional. The VLAN ID is required.</li> </ul>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system—control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Creating a Private VLAN on a Single Switch (CLI Procedure) on page 45</a></li> <li>• <a href="#">Creating a Private VLAN Spanning Multiple Switches (CLI Procedure) on page 47</a></li> </ul>

## proxy-arp

<b>Syntax</b>	<code>proxy-arp (restricted   unrestricted);</code>
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.6 for EX Series switches. <b>restricted</b> added in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 12.2 for the QFX Series.
<b>Description</b>	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>  <p><b>NOTE:</b> You must configure the IP address and the inet family for the interface when you enable proxy ARP.</p> </div>	
<b>Default</b>	Proxy ARP is not enabled. The router or switch responds to an ARP request only if the destination IP address is its own.
<b>Options</b>	<ul style="list-style-type: none"> <li>• <b>none</b>—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.</li> <li>• <b>restricted</b>—(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.</li> <li>• <b>unrestricted</b>—(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.</li> </ul>
	<b>Default:</b> <code>unrestricted</code>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Restricted and Unrestricted Proxy ARP on page 59</a></li> <li>• <a href="#">Configuring Proxy ARP (CLI Procedure)</a></li> <li>• <a href="#">Example: Configuring Proxy ARP on an EX Series Switch</a></li> <li>• <a href="#">Configuring Gratuitous ARP</a></li> </ul>

## push

<b>Syntax</b>	<code>push;</code>
<b>Hierarchy Level</b>	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a>]</code>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.</p>
<b>Description</b>	<p> <b>NOTE:</b> On EX4300 switches, <code>push</code> is not supported at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a>]</code> hierarchy level.</p> <p>Specify the VLAN rewrite operation to add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.</p> <p>You can use this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces; 10-Gigabit Ethernet LAN/WAN PIC; aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces; 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces.</p> <p>If you include the <b>push</b> statement in the configuration, you must also include the <a href="#">pop</a> statement at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a>]</code> hierarchy level.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Stacking a VLAN Tag</i></li> <li>• <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i></li> </ul>

## push-push

---

<b>Syntax</b>	push-push;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a> ], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	<p>Specify the VLAN rewrite operation to push two VLAN tags in front of the frame.</p> <p>You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Stacking Two VLAN Tags</i></li></ul>

## registration

<b>Syntax</b>	registration (forbidden   normal   restricted);
<b>Hierarchy Level</b>	[edit protocols <a href="#">mvrp interface</a> (all   <i>interface-name</i> )], [edit routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp interface</a> (all   <i>interface-name</i> )] (for virtual switch instance type), [edit logical-systems <i>logical-system-name</i> protocols <a href="#">mvrp interface</a> (all   <i>interface-name</i> )], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp interface</a> (all   <i>interface-name</i> )] (for virtual switch instance type)
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1 for MX Series routers.
<b>Description</b>	For Multiple VLAN Registration Protocol (MVRP) configurations, configure the registration mode for the interface.
<b>Default</b>	normal
<b>Options</b>	<p><b>forbidden</b>—The interface or interfaces do not register and do not participate in MVRP.</p> <p><b>normal</b>—The interface or interfaces accept MVRP messages and participate in MVRP.</p> <p><b>restricted</b>—The interface or interfaces ignore all MVRP JOIN messages received for VLANs that are not statically configured for MVRP on the interface.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers</i></li> <li>• <a href="#">Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration on page 52</a></li> <li>• <a href="#">Verifying That MVRP Is Working Correctly on page 54</a></li> <li>• <a href="#">Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration on page 49</a></li> <li>• <a href="#">join-timer on page 111</a></li> <li>• <a href="#">leaveall-timer on page 113</a></li> <li>• <a href="#">leave-timer on page 114</a></li> <li>• <a href="#">point-to-point on page 127</a></li> </ul>

## rstp

<b>Syntax</b>	<pre> rstp {   bpd-block-on-edge;   bpd-destination-mac-address provider-bridge-group;   bridge-priority priority;   extended-system-id;   force-version stp;   forward-delay seconds;   hello-time seconds;   max-age seconds;   interface interface-name {     bpd-timeout-action {       alarm;       block;     }     cost cost;     edge;     mode (p2p   shared);     no-root-port;     priority interface-priority;   }   priority-hold-time seconds;   traceoptions {     file filename &lt;files number&gt; &lt;size size&gt; &lt;world-readable   no-world-readable&gt;;     flag flag &lt;flag-modifier&gt; &lt;disable&gt;;   } } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],</p> <p>[edit protocols],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.4.</p> <p><b>bpd-block-on-edge</b> statement added in Junos OS Release 9.4.</p> <p><b>bpd-timeout-action</b> statement added in Junos OS Release 9.4.</p> <p>Support for logic systems added in Junos OS Release 9.6.</p>
<b>Description</b>	Configure RSTP parameters.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Configuring Rapid Spanning Tree Protocol</li> </ul>



## service-id

---

<b>Syntax</b>	<code>service-id <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">switch-options on page 84</a> ] [edit <a href="#">vlans on page 84</a> <i>vlan-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers. Statement introduced in Junos OS Release 13.2 for the QFX Series.
<b>Description</b>	Specify a service identifier for each multichassis aggregated Ethernet interface that belongs to a link aggregation group (LAG).
<b>Options</b>	<b>number</b> —A number that identifies a particular service. <b>Range:</b> 1 through 65535
<b>Required Privilege Level</b>	<b>system</b> —To view this statement in the configuration. <b>system control</b> —To add this statement to the configuration.

## static-mac

<b>Syntax</b>	<pre>static-mac <i>mac-address</i> {     <i>vlan-id</i> <i>number</i>; }</pre>
<b>Hierarchy Level</b>	<p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols evpn interface <i>interface-name</i>]</p> <p>[edit <a href="#">vlans on page 84</a> <i>vlan-name</i> switch-options interface <i>interface-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit vlans <i>vlan-name</i> switch-options interface <i>interface name</i>] hierarchy level introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers. The <b>vlan-id</b> option is not available for EVPNs.</p> <p>[edit vlans <i>vlan-name</i> switch-options interface <i>interface name</i>] hierarchy level introduced in Junos OS Release 13.2 for the QFX Series.</p>
<b>Description</b>	<p>Configure a static MAC address for a logical interface in a bridge domain or VLAN.</p> <p>The <b>vlan-id</b> option can be specified for <b>static-macs</b> only if <b>vlan-id all</b> is configured for the bridging domain or VLAN.</p>
<b>Options</b>	<p><b><i>mac-address</i></b>—MAC address</p> <p><b><i>vlan-id number</i></b>—(Optional) VLAN identifier to associate with static MAC address.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring EVPN Routing Instances</i></li> <li>• <i>Understanding Layer 2 Learning and Forwarding for Bridge Domains</i></li> <li>• <i>Layer 2 Learning and Forwarding for VLANs Overview on page 16</i></li> </ul>

## swap

<b>Syntax</b>	swap;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b> ], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches. Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.
<b>Description</b>	Specify the VLAN rewrite operation to replace a VLAN tag. The outer VLAN tag of the frame is overwritten with the user-specified VLAN tag information.  On MX Series routers, you can enter this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, aggregated Ethernet using Gigabit Ethernet IQ interfaces, and 100-Gigabit Ethernet Type 5 PIC with CFP. On EX Series switches, you can enter this statement on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Rewriting the VLAN Tag on Tagged Frames</i></li> <li>• <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i></li> </ul>

## swap-push

---

<b>Syntax</b>	swap-push;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a> ], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	<p>Specify the VLAN rewrite operation to replace the outer VLAN tag of the frame with a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.</p> <p>You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Rewriting a VLAN Tag and Adding a New Tag on page 10</a></li></ul>

## swap-swap

<b>Syntax</b>	swap-swap;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a> ], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Specify the VLAN rewrite operation to replace both the inner and the outer VLAN tags of the frame with a user-specified VLAN tag value.  You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and for 100-Gigabit Ethernet Type 5 PIC with CFP.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Rewriting the Inner and Outer VLAN Tags</i></li> </ul>

## switch-options

<b>Syntax</b>	<pre> switch-options {   interface <i>interface-name</i> {     interface-mac-limit <i>limit</i> {       packet-action drop;     }     no-mac-learning;     static-mac <i>static-mac-address</i> {       vlan-id <i>number</i>;     }   }   interface-mac-limit <i>limit</i> {     packet-action drop;   }   mac-statistics;   mac-table-size <i>limit</i> {     packet-action drop;   }   no-mac-learning;   service-id <i>number</i>;   vtep-source-interface } </pre>
<b>Hierarchy Level</b>	<pre> [edit <i>number</i>], [edit <a href="#">vlands on page 84</a> <i>vlan--name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> <a href="#">vlands on page 84</a> <i>vlan-name</i>], [edit routing-instances <i>routing-instance-name</i> <a href="#">vlands on page 84</a> <i>vlan-name</i>] </pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p>
<b>Description</b>	<p>Configure Layer 2 learning and forwarding properties for a VLAN or a virtual switch.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

## tag-protocol-id (TPID to Rewrite)

<b>Syntax</b>	<code>tag-protocol-id <i>tpid</i>;</code>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map],  [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b>],  [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map],  [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces only, configure the outer TPID value. All TPIDs you include in input and output VLAN maps must be among those you specify at the [edit interfaces <i>interface-name</i> <b>gigether-options ethernet-switch-profile tag-protocol-id [ <i>tpids</i> ]</b>] hierarchy level.</p> <p>For 10-Gigabit Ethernet LAN/WAN PIC interfaces on T Series routers the default TPID value (<b>0x8100</b>) is supported.</p>
<b>Default</b>	If the <b>tag-protocol-id</b> statement is not configured, the TPID value is 0x8100.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Inner and Outer TPIDs and VLAN IDs</i></li> </ul>

## traceoptions (MVRP)

<b>Syntax</b>	<pre> traceoptions {     file <i>name</i> &lt;size <i>size</i>&gt; &lt;files <i>number</i>&gt; &lt;(world-readable   no-world-readable)&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">mvrp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp</a>] (for virtual switch instance type),</p> <p>[edit protocols <a href="#">mvrp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">mvrp</a>] (for virtual switch instance type)</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1 for MX Series routers.
<b>Description</b>	For Multiple VLAN Registration Protocol (MVRP), configure tracing options.
<b>Default</b>	Traceoptions is disabled.
<b>Options</b>	<p><b>disable</b> —(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. To include the <b>file</b> statement, you must specify a filename. Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place MVRP tracing output in the file <code>/var/log/mvrp-log</code>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files, in the range from 2 through 1000. The default is 1 trace file. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, 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 <b>size</b> option.</p> <p><b>flag <i>flag</i></b>—Specify which tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. You can include the following flags:</p> <ul style="list-style-type: none"> <li>• <b>all</b>—Enable all trace options flags.</li> <li>• <b>error</b>—Trace all failure conditions.</li> <li>• <b>events</b>—Trace process state change and cleanup events.</li> <li>• <b>pdu</b>—Trace RAPS PDU reception and transmission.</li> <li>• <b>socket</b>—Trace socket activity.</li> <li>• <b>state-machine</b>—Trace information about the state machine.</li> <li>• <b>timers</b>—Trace protocol timers.</li> </ul>



**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size *size***—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the ***trace-file*** again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten. The file size range is from 10240 through 4294967295. The default file size is 1 MB.

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

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

<b>Related Documentation</b>	• <i>Example: Configuring Automatic VLAN Administration Using MVRP on MX Series Routers</i>
	• <a href="#">Configuring Multiple VLAN Registration Protocol (MVRP) to Manage Dynamic VLAN Registration on page 52</a>
	• <a href="#">Verifying That MVRP Is Working Correctly on page 54</a>
	• <a href="#">Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration on page 49</a>

## unconditional-src-learn

---

<b>Syntax</b>	unconditional-src-learn;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4R16.
<b>Description</b>	Enables the router to learn IP addresses from nonvalidated sources when proxy Address Resolution Protocol (ARP) is configured.



**CAUTION:** By default, the router learns IP addresses from validated sources only. When this statement is configured and proxy ARP is enabled on an unnumbered interface, the router responds to ARP requests from any IP address, which might lead to exploitable information disclosure. An attacker can poison the ARP cache and create a fake forwarding table entry for an IP address, effectively creating a denial of service for that subscriber or interface. Therefore, exercise caution when configuring this statement.

---

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

- |                              |  |
|------------------------------|--|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Understanding Proxy ARP on EX Series Switches</i></li><li>• <i>Understanding Proxy ARP on EX Series Switches</i></li><li>• <i>Example: Configuring Proxy ARP on an EX Series Switch</i></li></ul> |
|------------------------------|--|

## vlan-id (Bridge Domain or VLAN)

<b>Syntax</b>	<code>vlan-id (all   none   <i>number</i>);</code>
<b>Hierarchy Level</b>	<code>[edit bridge-domains <i>bridge-domain-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i></code> <code>bridge-domains <i>bridge-domain-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>]</code> <code>[edit vlans <i>vlan-name</i>]</code>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for Layer 2 trunk ports added in Junos OS Release 9.2.</p> <p>Support for SRX 3400, SRX 3600, SRX 5600, and SRX 5800 devices added in Junos OS Release 9.6.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Specify a VLAN identifier (VID) to include in the packets sent to and from the bridge domain or VLAN, or a VPLS routing instance.



**NOTE:** When configuring a VLAN identifier for provider backbone bridge (PBB) routing instances, dual-tagged VIDs and the none option are not permitted.

**Options** *number*—A valid VLAN identifier. If you configure multiple bridge domains or VLANs with a valid VLAN identifier, you must specify a unique VLAN identifier for each. However, you can use the same VLAN identifier for bridge domains or VLANs that belong to different virtual switches. Use this option to send single tagged frames with the specified VLAN identifier over VPLS VT interfaces.



**NOTE:** If you specify a VLAN identifier, you cannot also use the all option. They are mutually exclusive.

**all**—Specify that the bridge domain or VLAN spans all the VLAN identifiers configured on the member logical interfaces.



**NOTE:** You cannot specify the all option if you include a routing interface in the bridge domain or VLAN.

**none**—Specify to enable shared VLAN learning or to send untagged frames over VPLS VT interfaces.



**NOTE:** Multichassis link aggregation (MC-LAG) supports the `none` option with the `vlan-id` statement with bridge domains or VLANs on Layer 2 networks and IPv4 networks. Starting with Junos OS Release 13.2R3, MC-LAG supports the `none` option with the `vlan-id` statement with bridge domains or VLANs on IPv6 networks.

---



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

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring a Bridge Domain</i></li><li>• <a href="#">Configuring a VLAN on page 4</a></li><li>• <i>Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances</i></li><li>• <a href="#">Configuring VLAN Identifiers for VLANs and VPLS Routing Instances on page 23</a></li><li>• <i>Configuring Bridge Domains as Switches for Layer 2 Trunk Ports</i></li><li>• <i>Configuring a Layer 2 Virtual Switch</i></li><li>• <a href="#">Configuring a Layer 2 Virtual Switch on page 33</a></li><li>• <i>bridge-domains</i></li><li>• <a href="#">vlans on page 84</a></li></ul>
------------------------------	---

## vlan-id (VLAN ID to Rewrite)

<b>Syntax</b>	<code>vlan-id number;</code>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces and aggregated Ethernet using Gigabit Ethernet IQ interfaces, specify the line VLAN identifiers to be rewritten at the input or output interface.</p> <p>You cannot include the <b>vlan-id</b> statement with the <b>swap</b> statement, <b>swap-push</b> statement, <b>push-push</b> statement, or <b>push-swap</b> statement at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b>] hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the <b>vlan-id</b> statement that you include at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Rewriting the VLAN Tag on Tagged Frames</i></li> <li>• <i>Binding VLAN IDs to Logical Interfaces</i></li> </ul>

## vlan-id-list

<b>Syntax</b>	<code>vlan-id-list [ <i>vlan-id-numbers</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit bridge-domains <i>bridge-domain-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>],</p> <p>[edit interfaces <i>interface-name</i> unit 0],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</p> <p>[edit vlans <i>vlan-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.4.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p>
<b>Description</b>	<p>Specify a VLAN identifier list to use for a bridge domain or VLAN in trunk mode.</p> <p>Specify the <b>trunk</b> option in the <b>interface-mode</b> statement to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the <b>vlan-id-list</b> statement to forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the <b>access</b> option to accept packets with no VLAN ID to forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the <b>vlan-id</b> statement.</p> <p>This statement also enables you to bind a logical interface to a list of VLAN IDs, thereby configuring the logical interface to receive and forward a frame with a tag that matches the specified VLAN ID list.</p>
<div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p><b>WARNING:</b> On some EX and QFX Series switches, you can apply no more than eight VLAN identifier lists to a physical interface.</p> </div> </div>	
<b>Options</b>	<p><b><i>vlan-id-numbers</i></b>—Valid VLAN identifiers. You can combine individual numbers with range lists by including a hyphen.</p> <p><b>Range:</b> 0 through 4095</p>
<div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p><b>NOTE:</b> On EX Series switches and the QFX Series, the range is 0 through 4094.</p> </div> </div>	
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Configuring a Bridge Domain](#)
  - [Configuring a VLAN on page 4](#)
  - [Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances](#)
  - [Configuring VLAN Identifiers for VLANs and VPLS Routing Instances on page 23](#)

## vlan-rewrite

<b>Syntax</b>	<code>vlan-rewrite translate (200 500   201 501)</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>number</i> family bridge interface-mode trunk] [edit interfaces <i>interface-name</i> unit <i>number</i> family ethernet-switching interface-mode trunk]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series.
<b>Description</b>	Translates an incoming VLAN to a bridge-domain VLAN, corresponding counter translation at egress. Supports translation of VLAN 200 to VLAN 500 and VLAN 201 to VLAN 501. Other valid VLANs pass through without translation.
<b>Options</b>	<b>translate 200 500</b> —Translates incoming packets with VLAN 200 to 500.  <b>translate 201 501</b> —Translates incoming packets with VLAN 201 to 501.  <b>translate 202 502</b> —Translates incoming packets with VLAN 202 to 502.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Rewriting a VLAN Tag and Adding a New Tag on page 10</a></li> </ul>

## vlan-tags

---

<b>Syntax</b>	<code>vlan-tags outer <i>number</i> inner <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit bridge-domains <i>bridge-domain-name</i> ], [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> ], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> ] [edit <a href="#">vlans on page 84</a> <i>vlan-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.3R2 for EX Series switches. Statement introduced in Junos OS Release 13.2X51-D10 for QFX Series switches.
<b>Description</b>	Specify dual VLAN identifier tags for a bridge domain, VLAN, or VPLS routing instance.
<b>Options</b>	<b>outer <i>number</i></b> —A valid VLAN identifier.  <b>inner <i>number</i></b> —A valid VLAN identifier.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring a Bridge Domain</a></li><li>• <a href="#">Configuring a VLAN on page 4</a></li><li>• <a href="#">Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances</a></li><li>• <a href="#">Configuring VLAN Identifiers for VLANs and VPLS Routing Instances on page 23</a></li><li>• <a href="#">Configuring a Layer 2 Virtual Switch .</a></li><li>• <a href="#">Configuring a Layer 2 Virtual Switch on page 33</a></li></ul>



## CHAPTER 9

# Operational Commands

- `show ethernet-switching flood`
- `show ethernet-switching interface`
- `show ethernet-switching statistics`
- `show ethernet-switching table`
- `show interfaces (10-Gigabit Ethernet)`
- `show interfaces (Gigabit Ethernet)`
- `show interfaces irb`
- `show interfaces queue`
- `show mvrp`
- `show mvrp applicant-state`
- `show mvrp dynamic-vlan-memberships`
- `show mvrp interface`
- `show mvrp registration-state`
- `show mvrp statistics`
- `show vlans`
- `traceroute ethernet`

## show ethernet-switching flood

<b>Syntax</b>	<pre>show ethernet-switching flood &lt;brief   detail   extensive&gt; &lt;event-queue&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system <i>logical-system-name</i>&gt; &lt;route (all-ce-flood   all-ve-flood   alt-root-flood   bd-flood   mlp-flood   re-flood)&gt; &lt;vlan-name <i>vlan-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 12.3R2.</p> <p>Command introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	(MX Series routers and EX Series switches only) Display Ethernet-switching flooding information.
<b>Options</b>	<p><b>none</b>—Display all Ethernet-switching flooding information for all VLANs.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>event-queue</b>—(Optional) Display the queue of pending Ethernet-switching flood events.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display Ethernet-switching flooding information for the specified routing instance.</p> <p><b>logical-system <i>logical-system-name</i></b>—(Optional) Display Ethernet-switching flooding information for the specified logical system.</p> <p><b>route (all-ce-flood   all-ve-flood   alt-root-flood   bd-flood   mlp-flood   re-flood)</b>—(Optional) Display the following:</p> <ul style="list-style-type: none"> <li><b>all-ce-flood</b>—Display the route for flooding traffic to all customer edge routers or switches if <b>no-local-switching</b> is enabled.</li> <li><b>all-ve-flood</b>—Display the route for flooding traffic to all VPLS edge routers or switches if <b>no-local-switching</b> is enabled.</li> <li><b>alt-root-flood</b>—Display the Spanning Tree Protocol (STP) alt-root flooding route used for the interface.</li> <li><b>bd-flood</b>—Display the route for flooding traffic of a VLAN if <b>no-local-switching</b> is not enabled.</li> <li><b>mlp-flood</b>—Display the route for flooding traffic to MAC learning chips.</li> <li><b>re-flood</b>—Display the route for Routing Engine flooding to all interfaces.</li> </ul> <p><b>vlan-name <i>vlan-name</i></b>—(Optional) Display Ethernet-switching flooding information for the specified VLAN.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ethernet-switching flood on page 155</a>

[show ethernet-switching flood brief on page 155](#)  
[show ethernet-switching flood detail on page 155](#)  
[show ethernet-switching flood extensive on page 156](#)

## Sample Output

### show ethernet-switching flood

```

user@host> show ethernet-switching flood
Name: __juniper_private1__
CEs: 0
VEs: 0
Name: default-switch
CEs: 9
VEs: 0
VLAN Name: VLAN101
Flood Routes:
  Prefix    Type           Owner                NhType    NhIndex
  0x3057b/51 FLOOD_GRP_COMP_NH __all_ces__         comp      12866
  0x30004/51 FLOOD_GRP_COMP_NH __re_flood__        comp      12863
VLAN Name: VLAN102
Flood Routes:
  Prefix    Type           Owner                NhType    NhIndex
  0x3057c/51 FLOOD_GRP_COMP_NH __all_ces__         comp      12875
  0x30005/51 FLOOD_GRP_COMP_NH __re_flood__        comp      12872
VLAN Name: VLAN103
Flood Routes:
  Prefix    Type           Owner                NhType    NhIndex
  0x3057d/51 FLOOD_GRP_COMP_NH __all_ces__         comp      12884
  0x30006/51 FLOOD_GRP_COMP_NH __re_flood__        comp      12881

```

### show ethernet-switching flood brief

```

user@host> show ethernet-switching flood brief
Name                Active CEs    Active VEs
__juniper_private1__ 0              0
default-switch      9              0

```

### show ethernet-switching flood detail

```

user@host> show ethernet-switching flood detail
Name: __juniper_private1__
CEs: 0
VEs: 0
Name: default-switch
CEs: 9
VEs: 0
VLAN Name: VLAN101
Flood Routes:
  Prefix    Type           Owner                NhType    NhIndex
  0x3057b/51 FLOOD_GRP_COMP_NH __all_ces__         comp      12866
  0x30004/51 FLOOD_GRP_COMP_NH __re_flood__        comp      12863
VLAN Name: VLAN102
Flood Routes:
  Prefix    Type           Owner                NhType    NhIndex
  0x3057c/51 FLOOD_GRP_COMP_NH __all_ces__         comp      12875
  0x30005/51 FLOOD_GRP_COMP_NH __re_flood__        comp      12872
VLAN Name: VLAN103
Flood Routes:
  Prefix    Type           Owner                NhType    NhIndex

```

```

0x3057d/51 FLOOD_GRP_COMP_NH __all_ces__ comp 12884
0x30006/51 FLOOD_GRP_COMP_NH __re_flood__ comp 12881

```

### show ethernet-switching flood extensive

```

user@host> show ethernet-switching flood extensive
Name: __juniper_private1__
CEs: 0
VEs: 0
Name: default-switch
CEs: 9
VEs: 0
VLAN Name: VLAN101
Flood route prefix: 0x3057b/51
Flood route type: FLOOD_GRP_COMP_NH
Flood route owner: __all_ces__
Flood group name: __all_ces__
Flood group index: 1
Nexthop type: comp
Nexthop index: 12866
Flooding to:
  Name      Type      NhType      Index
  __all_ces__ Group    comp        12860
  Composition: split-horizon
  Flooding to:
    Name      Type      NhType      Index
    ae20.0    CE        ucst        7605

Flood route prefix: 0x30004/51
Flood route type: FLOOD_GRP_COMP_NH
Flood route owner: __re_flood__
Flood group name: __re_flood__
Flood group index: 65534
Nexthop type: comp
Nexthop index: 12863
Flooding to:
  Name      Type      NhType      Index
  __all_ces__ Group    comp        12860
  Composition: split-horizon
  Flooding to:
    Name      Type      NhType      Index
    ae20.0    CE        ucst        7605

VLAN Name: VLAN102

Flood route prefix: 0x3057c/51
Flood route type: FLOOD_GRP_COMP_NH
Flood route owner: __all_ces__
Flood group name: __all_ces__
Flood group index: 1
Nexthop type: comp
Nexthop index: 12875
Flooding to:
  Name      Type      NhType      Index
  __all_ces__ Group    comp        12869
  Composition: split-horizon
  Flooding to:
    Name      Type      NhType      Index
    ae20.0    CE        ucst        7605

Flood route prefix: 0x30005/51
Flood route type: FLOOD_GRP_COMP_NH

```

```

Flood route owner: __re_flood__
Flood group name: __re_flood__
Flood group index: 65534
Nexthop type: comp
Nexthop index: 12872
Flooding to:
  Name      Type      NhType      Index
  __all_ces__  Group    comp        12869
  Composition: split-horizon
  Flooding to:
    Name      Type      NhType      Index
    ae20.0    CE        ucst        7605
VLAN Name: VLAN103

Flood route prefix: 0x3057d/51
Flood route type: FLOOD_GRP_COMP_NH
Flood route owner: __all_ces__
Flood group name: __all_ces__
Flood group index: 1
Nexthop type: comp
Nexthop index: 12884
Flooding to:
  Name      Type      NhType      Index
  __all_ces__  Group    comp        12878
  Composition: split-horizon
  Flooding to:
    Name      Type      NhType      Index
    ae20.0    CE        ucst        7605

Flood route prefix: 0x30006/51
Flood route type: FLOOD_GRP_COMP_NH
Flood route owner: __re_flood__
Flood group name: __re_flood__
Flood group index: 65534
Nexthop type: comp
Nexthop index: 12881
Flooding to:
  Name      Type      NhType      Index
  __all_ces__  Group    comp        12878
  Composition: split-horizon
  Flooding to:
    Name      Type      NhType      Index
    ae20.0    CE        ucst        7605
VLAN Name: VLAN104

```

## show ethernet-switching interface

<b>Syntax</b>	<b>show ethernet-switching interface</b> <b>&lt;brief   detail   extensive&gt;</b> <b>&lt;interface-name&gt;</b>
<b>Release Information</b>	Command introduced in Junos OS Release 12.3R2. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	(MX Series routers and EX Series switches only) Display Layer 2 learning information for all the interfaces.
<b>Options</b>	<b>none</b> —Display Ethernet-switching information for all interfaces.  <b>brief   detail   extensive</b> —(Optional) Display the specified level of output.  <b>interface-name</b> —(Optional) Display Ethernet-switching information for the specified interface.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ethernet-switching interface (Specific Interface) on page 159</a> <a href="#">show ethernet-switching interface detail on page 159</a>
<b>Output Fields</b>	<a href="#">Table 10 on page 158</a> describes the output fields for the <b>show ethernet-switching interface</b> command. Output fields are listed in the approximate order in which they appear.

**Table 10: show ethernet-switching interface Output Fields**

Field Name	Field Description
<b>Logical interface</b>	Name of the logical interface.
<b>VLAN members</b>	VLANs associated with this interface.
<b>Tag</b>	VLAN ID.
<b>MAC limit</b>	Number of MAC addresses that can be associated with the interface.
<b>STP state</b>	Spanning tree protocol (STP) state.
<b>Logical interface flags</b>	Status of Layer 2 learning properties for each interface: <ul style="list-style-type: none"> <li>• <b>DL</b>—MAC learning is disabled.</li> <li>• <b>LH</b>—MAC interface limit has been reached..</li> <li>• <b>AD</b>—Packets are dropped after the MAC interface limit is reached.</li> <li>• <b>DN</b>—The MAC interface is down.</li> </ul>
<b>Tagging</b>	Tagging state of the VLAN.

## Sample Output

### show ethernet-switching interface (Specific Interface)

```

user@host> show ethernet-switching interface ae10.0
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical  Vlan    TAG  MAC    STP    Logical    Tagging
interface members limit state  interface flags
ae10.0
      VLAN70.. 701    1024   Forwarding
      VLAN70.. 702    1024   Forwarding
      VLAN70.. 703    1024   Forwarding
      VLAN70.. 704    1024   Forwarding
      VLAN70.. 705    1024   Forwarding
      VLAN70.. 706    1024   Forwarding
      VLAN70.. 707    1024   Forwarding
      VLAN70.. 708    1024   Forwarding
      VLAN70.. 709    1024   Forwarding
      VLAN71.. 710    1024   Forwarding
      VLAN71.. 711    1024   Forwarding
      VLAN71.. 712    1024   Forwarding
      VLAN71.. 713    1024   Forwarding
      VLAN71.. 714    1024   Forwarding
      VLAN71.. 715    1024   Forwarding
[...output truncated...]

```

### show ethernet-switching interface detail

```

user@host> show ethernet-switching interface detail
Information for interface family:
Name: ge-1/0/3.0
  Type: IFF                                Handle: 0x8bba280
  Index: 331                              Generation: 159
                                           Flags: UP,
                                           Routing/Vlan index: 4
                                           Address family: 50
                                           MAC sequence number: 0
                                           MACs learned: 0
                                           Non configured static MACs learned: 0
Name: ge-1/0/3.0
  Type: IFBD (static)                     Handle: 0x8bb6e00
  Index:                                  Generation: 129
                                           Flags: UP,
                                           Routing/Vlan index: 2
                                           Address family:

```

Sequence number: 1  
MAC limit: 65535  
Static MACs learned: 0  
VSTP index: 11  
Name: ge-1/0/3.0  
Type: IFBD (static)  
Index:  
Trunk id: 0  
IFD index:  
IFL index:  
Sequence number: 1  
MAC limit: 65535  
Static MACs learned: 0  
VSTP index: 11

MAC sequence number: 1  
MACs learned: 0  
Non configured static MACs learned: 0  
Rewrite op:

Handle: 0x8bb6f00  
Generation: 130  
Flags: UP,  
Routing/Vlan index: 3  
Address family:  
MAC sequence number: 1  
MACs learned: 0  
Non configured static MACs learned: 0  
Rewrite op:



## show ethernet-switching statistics

<b>Syntax</b>	<b>show ethernet-switching statistics</b> <instance <i>instance-name</i> > <logical-system <i>logical-system-name</i> > <vlan-name <i>vlan-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 12.3R2. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	(MX Series routers and EX Series switches only) Display Ethernet-switching statistics.
<b>Options</b>	<b>none</b> —Display Ethernet-switching statistics for all VLANs in all routing instances.  <b>instance <i>instance-name</i></b> —(Optional) Display statistics for the specified routing instance.  <b>logical-system <i>logical-system-name</i></b> —(Optional) Display Ethernet-switching statistics information for the specified logical system.  <b>vlan-name <i>vlan-name</i></b> —(Optional) Display statistics for the specified VLAN.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ethernet-switching statistics on page 161</a>

## Sample Output

### show ethernet-switching statistics

```

user@host> show ethernet-switching statistics
Local interface: ae1.0, Index: 1035
Broadcast packets:          220
Broadcast bytes :          13720
Multicast packets:          130
Multicast bytes :          11700
Flooded packets :           0
Flooded bytes :            0
Unicast packets :           0
Unicast bytes :            0
Current MAC count:          0 (Limit 1024)
Local interface: vt-3/3/10.1048576, Index: 1280
Broadcast packets:          0
Broadcast bytes :           0
Multicast packets:          0
Multicast bytes :           0
Flooded packets :           2
Flooded bytes :           128
Unicast packets :          632
Unicast bytes :          39184
Current MAC count:          2
Local interface: ge-3/1/2.0, Index: 1258
Broadcast packets:          100
Broadcast bytes :          6800
Multicast packets:          200
Multicast bytes :          18000
Flooded packets :           0

```

```
Flooded bytes      :          0
Unicast packets    :          632
Unicast bytes      :        39184
Current MAC count:          2 (Limit 1024)
Local interface: ae3.0, Index: 1043
Broadcast packets:          0
Broadcast bytes    :          0
Multicast packets:          0
Multicast bytes    :          0
Flooded packets   :          0
Flooded bytes     :          0
Unicast packets   :          0
Unicast bytes     :          0
Current MAC count:          0 (Limit 1024)
Local interface: ge-3/3/8.0, Index: 1276
Broadcast packets:          0
Broadcast bytes    :          0
Multicast packets:          0
Multicast bytes    :          0
Flooded packets   :          0
Flooded bytes     :          0
Unicast packets   :          0
Unicast bytes     :          0
Current MAC count:          0 (Limit 8192)
Local interface: ae5.0, Index: 1045
Broadcast packets:          0
Broadcast bytes    :          0
Multicast packets:          0
Multicast bytes    :          0
Flooded packets   :          0
Flooded bytes     :          0
Unicast packets   :          0
Unicast bytes     :          0
Current MAC count:          0 (Limit 8192)
Local interface: ae4.0, Index: 1044
Broadcast packets:          200
Broadcast bytes    :        13600
Multicast packets:          0
Multicast bytes    :          0
Flooded packets   :          0
Flooded bytes     :          0
Unicast packets   :          0
Unicast bytes     :          0
Current MAC count:          0 (Limit 8192)
Local interface: ae26.0, Index: 1042
Broadcast packets:          0
Broadcast bytes    :          0
Multicast packets:          0
Multicast bytes    :          0
Flooded packets   :          0
Flooded bytes     :          0
Unicast packets   :          0
Unicast bytes     :          0
Current MAC count:          0 (Limit 8192)
Local interface: ae25.0, Index: 1041
Broadcast packets:          133
Broadcast bytes    :         7980
Multicast packets:        369934
Multicast bytes    :       59207572
Flooded packets   :          0
Flooded bytes     :          0
```

```
Unicast packets :          1433
Unicast bytes   :          119930
Current MAC count:           3 (Limit 8192)
Local interface: ae23.0, Index: 1040
Broadcast packets:          226
Broadcast bytes :          14464
Multicast packets:         585668
Multicast bytes :        153464476
Flooded packets :           0
Flooded bytes  :           0
Unicast packets :          26552
Unicast bytes   :        1947627
Current MAC count:           7 (Limit 8192)
Local interface: ae20.0, Index: 1039
Broadcast packets:          115
Broadcast bytes :           6900
Multicast packets:         395113
Multicast bytes :        61622869
Flooded packets :           0
Flooded bytes  :           0
Unicast packets :          1419
Unicast bytes   :        117924
Current MAC count:           4 (Limit 8192)
```

## show ethernet-switching table

**Syntax** `show ethernet-switching table`  
`<brief | count | detail | extensive | summary>`  
`<address>`  
`<instance instance-name>`  
`<interface interface-name>`  
`isid isid`  
`<logical-system logical-system-name>`  
`<persistent-learning (interface interface-name | mac mac-address)>`  
`<address>`  
`<vlan-id (all-vlan | vlan-id)>`  
`<vlan-name (all | vlan-name)>`

**Release Information** Command introduced in Junos OS Release 12.3R2.  
 Command introduced in Junos OS Release 12.3R2 for EX Series switches.  
 Options **logical-system**, **persistent-learning**, and **summary** introduced in Junos OS Release 13.2X50-D10 (ELS).

### Description



**NOTE:** If your EX Series switch CLI displays different options for the `show ethernet-switching table` command from the options shown in this document, then see *show ethernet-switching table*.

(MX Series routers and EX Series switches only) Display Layer 2 MAC address information.

**Options** **none**—Display all learned Layer 2 MAC address information.

**brief | count | detail | extensive | summary**—(Optional) Display the specified level of output.

**address**—(Optional) Display the specified learned Layer 2 MAC address information.

**instance *instance-name***—(Optional) Display learned Layer 2 MAC addresses for the specified routing instance.

**interface *interface-name***—(Optional) Display learned Layer 2 MAC addresses for the specified interface.

**isid *isid***—(Optional) Display learned Layer 2 MAC addresses for the specified ISID.

**logical-system *logical-system-name***—(Optional) Display Ethernet-switching statistics information for the specified logical system.

**persistent-learning (interface *interface-name* | mac *mac-address*)**—(Optional) Display dynamically learned MAC addresses that are retained despite device restarts and interface failures for a specified interface, or information about a specified MAC address.

**vlan-id (all-vlan | *vlan-id*)**—(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.

**vlan-name (all | *vlan-name*)**—(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.

**Additional Information** When Layer 2 protocol tunneling is enabled, the tunneling MAC address 01:00:0c:cd:cd:d0 is installed in the MAC table. When the Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) is configured for Layer 2 protocol tunneling on an interface, the corresponding protocol MAC address is installed in the MAC table.

**Required Privilege Level** view

**List of Sample Output** [show ethernet-switching table on page 166](#)  
[show ethernet-switching table brief on page 167](#)  
[show ethernet-switching table count on page 168](#)  
[show ethernet-switching table extensive on page 169](#)

**Output Fields** [Table 11 on page 165](#) describes the output fields for the **show ethernet-switching table** command. Output fields are listed in the approximate order in which they appear.

**Table 11: show ethernet-switching table Output fields**

Field Name	Field Description
<b>Routing instance</b>	Name of the routing instance.
<b>VLAN name</b>	Name of the VLAN.
<b>MAC address</b>	MAC address or addresses learned on a logical interface.
<b>MAC flags</b>	Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> <li>• <b>S</b>—Static MAC address is configured.</li> <li>• <b>D</b>—Dynamic MAC address is configured.</li> <li>• <b>L</b>—Locally learned MAC address is configured.</li> <li>• <b>SE</b>—MAC accounting is enabled.</li> <li>• <b>NM</b>—Non-configured MAC.</li> <li>• <b>R</b>—Locally learned MAC address is configured.</li> </ul>
<b>Age</b>	This field is not supported.
<b>Logical interface</b>	Name of the logical interface.
<b>MAC count</b>	Number of MAC addresses learned on the specific routing instance or interface.
<b>Learning interface</b>	Name of the logical interface on which the MAC address was learned.
<b>Learning VLAN</b>	VLAN ID of the routing instance or VLAN in which the MAC address was learned.
<b>Layer 2 flags</b>	Debugging flags signifying that the MAC address is present in various lists.

Table 11: show ethernet-switching table Output fields (*continued*)

Field Name	Field Description
<b>Epoch</b>	Spanning-tree-protocol epoch number identifying when the MAC address was learned. Used for debugging.
<b>Sequence number</b>	Sequence number assigned to this MAC address. Used for debugging.
<b>Learning mask</b>	Mask of the Packet Forwarding Engines where this MAC address was learned. Used for debugging.
<b>IPC generation</b>	Creation time of the logical interface when this MAC address was learned. Used for debugging.

## Sample Output

### show ethernet-switching table

```

user@host> show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
  Vlan      MAC      MAC      Age      Logical
  name      address   flags    interface
  VLAN101   88:e0:f3:bb:07:f0  D        -        ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
  Vlan      MAC      MAC      Age      Logical
  name      address   flags    interface
  VLAN102   88:e0:f3:bb:07:f0  D        -        ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
  Vlan      MAC      MAC      Age      Logical
  name      address   flags    interface
  VLAN103   88:e0:f3:bb:07:f0  D        -        ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
  Vlan      MAC      MAC      Age      Logical
  name      address   flags    interface
  VLAN104   88:e0:f3:bb:07:f0  D        -        ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

```

```

Routing instance : default-switch
  Vlan      MAC      MAC      Age      Logical
  name      address   flags      interface
  VLAN1101  00:1f:12:32:f5:c1 D          -      ae0.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
          SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
  Vlan      MAC      MAC      Age      Logical
  name      address   flags      interface
  VLAN1102  00:1f:12:32:f5:c1 D          -      ae0.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
          SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
  Vlan      MAC      MAC      Age      Logical
  name      address   flags      interface
  VLAN1103  00:1f:12:32:f5:c1 D          -      ae0.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
          SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
  Vlan      MAC      MAC      Age      Logical
  name      address   flags      interface
  VLAN1104  00:1f:12:32:f5:c1 D          -      ae0.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
          SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
  Vlan      MAC      MAC      Age      Logical
  name      address   flags      interface
  VLAN1105  00:1f:12:32:f5:c1 D          -      ae0.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
          SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
  Vlan      MAC      MAC      Age      Logical
  name      address   flags      interface
  VLAN1106  00:1f:12:32:f5:c1 D          -      ae0.0
[...output truncated...]

```

### show ethernet-switching table brief

```

user@host> show ethernet-switching table brief
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
          SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
  Vlan      MAC      MAC      Age      Logical
  name      address   flags      interface

```

```
VLAN101          88:e0:f3:bb:07:f0  D          -   ae20.0
```

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned  
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN102	88:e0:f3:bb:07:f0	D	-	ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned  
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN103	88:e0:f3:bb:07:f0	D	-	ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned  
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN104	88:e0:f3:bb:07:f0	D	-	ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned  
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN1101	00:1f:12:32:f5:c1	D	-	ae0.0

[...output truncated...]

### show ethernet-switching table count

```
user@host> show ethernet-switching table count
```

```
0 MAC address learned in routing instance default-switch VLAN VLAN1000
ae26.0:1000
```

```
1 MAC address learned in routing instance default-switch VLAN VLAN101
ae20.0:101
```

MAC address count per learn VLAN within routing instance:

Learn VLAN ID	MAC count	Static MAC count
101	1	0

```
1 MAC address learned in routing instance default-switch VLAN VLAN102
ae20.0:102
```

MAC address count per learn VLAN within routing instance:

Learn VLAN ID	MAC count	Static MAC count
102	1	0

```
1 MAC address learned in routing instance default-switch VLAN VLAN103
ae20.0:103
```



```

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID      MAC count      Static MAC count
          103              1              0

1 MAC address learned in routing instance default-switch VLAN VLAN104
ae20.0:104

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID      MAC count      Static MAC count
          104              1              0

0 MAC address learned in routing instance default-switch VLAN VLAN105
ae20.0:105

0 MAC address learned in routing instance default-switch VLAN VLAN106
ae20.0:106

0 MAC address learned in routing instance default-switch VLAN VLAN107
ae20.0:107

0 MAC address learned in routing instance default-switch VLAN VLAN108
ae20.0:108

0 MAC address learned in routing instance default-switch VLAN VLAN109
ae20.0:109

0 MAC address learned in routing instance default-switch VLAN VLAN110
ae20.0:110

1 MAC address learned in routing instance default-switch VLAN VLAN1101
ae0.0:1101

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID      MAC count      Static MAC count
          1101              1              0

1 MAC address learned in routing instance default-switch VLAN VLAN1102
ae0.0:1102

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID      MAC count      Static MAC count
          1102              1              0
[...output truncated...]

```

### show ethernet-switching table extensive

```

user@host> show ethernet-switching table extensive

MAC address: 88:e0:f3:bb:07:f0
Routing instance: default-switch
VLAN ID: 101
Learning interface: ae20.0
Layer 2 flags: in_hash,in_ifid,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                      Sequence number: 2
Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
Routing instance: default-switch
VLAN ID: 102
Learning interface: ae20.0

```

```
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
Routing instance: default-switch
VLAN ID: 103
Learning interface: ae20.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
Routing instance: default-switch
VLAN ID: 104
Learning interface: ae20.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
Routing instance: default-switch
VLAN ID: 1101
Learning interface: ae0.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
Routing instance: default-switch
VLAN ID: 1102
Learning interface: ae0.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
Routing instance: default-switch
VLAN ID: 1103
Learning interface: ae0.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
Routing instance: default-switch
VLAN ID: 1104
Learning interface: ae0.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008
```

## show interfaces (10-Gigabit Ethernet)

<b>Syntax</b>	<pre>show interfaces <i>xe-fpc/pic/port</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	Command introduced in Junos OS Release 8.0.
<b>Description</b>	Display status information about the specified 10-Gigabit Ethernet interface.
<b>Options</b>	<p><i>xe-fpc/pic/port</i>—Display standard information about the specified 10-Gigabit Ethernet interface.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2) on page 186</a></p> <p><a href="#">show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode) on page 189</a></p> <p><a href="#">show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC) on page 191</a></p> <p><a href="#">show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode) on page 193</a></p> <p><a href="#">show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only) on page 193</a></p> <p><a href="#">show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only) on page 194</a></p>
<b>Output Fields</b>	See <a href="#">Table 12 on page 172</a> for the output fields for the <b>show interfaces</b> (10-Gigabit Ethernet) command.

Table 12: show interfaces Gigabit Ethernet Output Fields

Field Name	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Physical interface</b>	Name of the physical interface.	All levels
<b>Enabled</b>	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
<b>Interface index</b>	Index number of the physical interface, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Link-level type</b>	Encapsulation being used on the physical interface.	All levels
<b>MTU</b>	Maximum transmission unit size on the physical interface.	All levels
<b>Speed</b>	Speed at which the interface is running.	All levels
<b>Loopback</b>	Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .	All levels
<b>Source filtering</b>	Source filtering status: <b>Enabled</b> or <b>Disabled</b> .	All levels
<b>LAN-PHY mode</b>	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
<b>WAN-PHY mode</b>	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
<b>Unidirectional</b>	Unidirectional link mode status for 10-Gigabit Ethernet interface: <b>Enabled</b> or <b>Disabled</b> for parent interface; <b>Rx-only</b> or <b>Tx-only</b> for child interfaces.	All levels
<b>Flow control</b>	Flow control status: <b>Enabled</b> or <b>Disabled</b> .	All levels
<b>Auto-negotiation</b>	(Gigabit Ethernet interfaces) Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .	All levels
<b>Remote-fault</b>	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> <li><b>Online</b>—Autonegotiation is manually configured as online.</li> <li><b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul>	All levels
<b>Device flags</b>	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
<b>Interface flags</b>	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels

Table 12: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output	
<b>Link flags</b>	Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .	All levels	
<b>Wavelength</b>	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels	
<b>Frequency</b>	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels	
<b>CoS queues</b>	Number of CoS queues configured.	<b>detail extensive</b> none	
<b>Schedulers</b>	(Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	<b>extensive</b>	
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds.	<b>detail extensive</b>	
<b>Current address</b>	Configured MAC address.	<b>detail extensive</b> none	
<b>Hardware address</b>	Hardware MAC address.	<b>detail extensive</b> none	
<b>Last flapped</b>	Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .	<b>detail extensive</b> none	
<b>Input Rate</b>	Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified	
<b>Output Rate</b>	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified	
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>	
<b>Egress account overhead</b>	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	<b>detail extensive</b>	
<b>Ingress account overhead</b>	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	<b>detail extensive</b>	<b>detail extensive</b>

Table 12: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Traffic statistics</b>	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see <a href="#">Table 12 on page 172</a>.</p>	<b>detail extensive</b>
<b>Input errors</b>	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the <code>ignore-l3-incompletes</code> statement.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>

Table 12: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Output errors</b>	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Collisions</b>—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>
<b>Egress queues</b>	<p>Total number of egress queues supported on the specified interface.</p> <p><b>NOTE:</b> In DPCs that are not of the enhanced type, such as DPC 40x1GE R, DPCE 20x1GE + 2x10GE R, or DPCE 40x1GE R, you might notice a discrepancy in the output of the <b>show interfaces</b> command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenario, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs</p>	<b>detail extensive</b>
<b>Queue counters (Egress)</b>	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>	<b>detail extensive</b>

Table 12: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Ingress queues</b>	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	<b>extensive</b>
<b>Queue counters (Ingress)</b>	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>	<b>extensive</b>
<b>Active alarms and Active defects</b>	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value <b>None</b> or <b>Link</b> . <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul>	<b>detail extensive none</b>
<b>OTN alarms</b>	Active OTN alarms identified on the interface.	<b>detail extensive</b>
<b>OTN defects</b>	OTN defects received on the interface.	<b>detail extensive</b>
<b>OTN FEC Mode</b>	The FECmode configured on the interface. <ul style="list-style-type: none"> <li>• <b>efec</b>—Enhanced forward error correction (EFEC) is configured to detect and correct bit errors.</li> <li>• <b>gfec</b>—G.709 Forward error correction (GFEC) mode is configured to detect and correct bit errors.</li> <li>• <b>none</b>—FEC mode is not configured.</li> </ul>	<b>detail extensive</b>
<b>OTN Rate</b>	OTN mode. <ul style="list-style-type: none"> <li>• <b>fixed-stuff-bytes</b>—Fixed stuff bytes 11.0957 Gbps.</li> <li>• <b>no-fixed-stuff-bytes</b>—No fixed stuff bytes 11.0491 Gbps.</li> <li>• <b>pass-through</b>—Enable OTN passthrough mode.</li> <li>• <b>no-pass-through</b>—Do not enable OTN passthrough mode.</li> </ul>	<b>detail extensive</b>
<b>OTN Line Loopback</b>	Status of the line loopback, if configured for the DWDM OTN PIC. Its value can be: <b>enabled</b> or <b>disabled</b> .	<b>detail extensive</b>
<b>OTN FEC statistics</b>	The forward error correction (FEC) counters for the DWDM OTN PIC. <ul style="list-style-type: none"> <li>• <b>Corrected Errors</b>—The count of corrected errors in the last second.</li> <li>• <b>Corrected Error Ratio</b>—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.</li> </ul>	<b>detail extensive</b>



Table 12: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>OTN FEC alarms</b>	OTN FEC excessive or degraded error alarms triggered on the interface. <ul style="list-style-type: none"> <li><b>FEC Degrade</b>—OTU FEC Degrade defect.</li> <li><b>FEC Excessive</b>—OTU FEC Excessive Error defect.</li> </ul>	<b>detail extensive</b>
<b>OTN OC</b>	OTN OC defects triggered on the interface. <ul style="list-style-type: none"> <li><b>LOS</b>—OC Loss of Signal defect.</li> <li><b>LOF</b>—OC Loss of Frame defect.</li> <li><b>LOM</b>—OC Loss of Multiframe defect.</li> <li><b>Wavelength Lock</b>—OC Wavelength Lock defect.</li> </ul>	<b>detail extensive</b>
<b>OTN OTU</b>	OTN OTU defects detected on the interface <ul style="list-style-type: none"> <li><b>AIS</b>—OTN AIS alarm.</li> <li><b>BDI</b>—OTN OTU BDI alarm.</li> <li><b>IAE</b>—OTN OTU IAE alarm.</li> <li><b>TTIM</b>—OTN OTU TTIM alarm.</li> <li><b>SF</b>—OTN ODU bit error rate fault alarm.</li> <li><b>SD</b>—OTN ODU bit error rate defect alarm.</li> <li><b>TCA-ES</b>—OTN ODU ES threshold alarm.</li> <li><b>TCA-SES</b>—OTN ODU SES threshold alarm.</li> <li><b>TCA-UAS</b>—OTN ODU UAS threshold alarm.</li> <li><b>TCA-BBE</b>—OTN ODU BBE threshold alarm.</li> <li><b>BIP</b>—OTN ODU BIP threshold alarm.</li> <li><b>BBE</b>—OTN OTU BBE threshold alarm.</li> <li><b>ES</b>—OTN OTU ES threshold alarm.</li> <li><b>SES</b>—OTN OTU SES threshold alarm.</li> <li><b>UAS</b>—OTN OTU UAS threshold alarm.</li> </ul>	<b>detail extensive</b>
<b>Received DAPI</b>	Destination Access Port Interface (DAPI) from which the packets were received.	<b>detail extensive</b>
<b>Received SAPI</b>	Source Access Port Interface (SAPI) from which the packets were received.	<b>detail extensive</b>
<b>Transmitted DAPI</b>	Destination Access Port Interface (DAPI) to which the packets were transmitted.	<b>detail extensive</b>
<b>Transmitted SAPI</b>	Source Access Port Interface (SAPI) to which the packets were transmitted.	<b>detail extensive</b>
<b>PCS statistics</b>	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. <ul style="list-style-type: none"> <li><b>Bit errors</b>—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode.</li> <li><b>Errored blocks</b>—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.</li> </ul>	<b>detail extensive</b>

Table 12: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>MAC statistics</b>	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> <li>• <b>Total octets</b> and <b>total packets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see <a href="#">Table 13 on page 186</a></li> <li>• <b>Unicast packets</b>, <b>Broadcast packets</b>, and <b>Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—Number of frames that exceed 1518 octets.</li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>VLAN tagged frames</b>—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."</li> </ul>	<b>extensive</b>
<b>OTN Received Overhead Bytes</b>	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	<b>extensive</b>
<b>OTN Transmitted Overhead Bytes</b>	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	<b>extensive</b>

Table 12: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> <li>• <b>Input packet count</b>—Number of packets received from the MAC hardware that the filter processed.</li> <li>• <b>Input packet rejects</b>—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address.</li> <li>• <b>Input DA rejects</b>—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the routing device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local routing device (which the routing device is rejecting).</li> <li>• <b>Input SA rejects</b>—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect.</li> <li>• <b>Output packet count</b>—Number of packets that the filter has given to the MAC hardware.</li> <li>• <b>Output packet pad count</b>—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured.</li> <li>• <b>Output packet error count</b>—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment.</li> <li>• <b>CAM destination filters, CAM source filters</b>—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.</li> </ul>	extensive
PMA PHY	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. Any state other than <b>OK</b> indicates a problem.</li> </ul> <p>Subfields are:</p> <ul style="list-style-type: none"> <li>• <b>PHY Lock</b>—Phase-locked loop</li> <li>• <b>PHY Light</b>—Loss of optical signal</li> </ul>	extensive

Table 12: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>WIS section</b>	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. Any state other than <b>OK</b> indicates a problem.</li> </ul> <p>Subfields are:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B1</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>SEF</b>—Severely errored framing</li> <li>• <b>LOL</b>—Loss of light</li> <li>• <b>LOF</b>—Loss of frame</li> <li>• <b>ES-S</b>—Errored seconds (section)</li> <li>• <b>SES-S</b>—Severely errored seconds (section)</li> <li>• <b>SEFS-S</b>—Severely errored framing seconds (section)</li> </ul>	<b>extensive</b>
<b>WIS line</b>	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. State other than <b>OK</b> indicates a problem.</li> </ul> <p>Subfields are:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B2</b>—Bit interleaved parity for SONET line overhead</li> <li>• <b>REI-L</b>—Remote error indication (near-end line)</li> <li>• <b>RDI-L</b>—Remote defect indication (near-end line)</li> <li>• <b>AIS-L</b>—Alarm indication signal (near-end line)</li> <li>• <b>BERR-SF</b>—Bit error rate fault (signal failure)</li> <li>• <b>BERR-SD</b>—Bit error rate defect (signal degradation)</li> <li>• <b>ES-L</b>—Errored seconds (near-end line)</li> <li>• <b>SES-L</b>—Severely errored seconds (near-end line)</li> <li>• <b>UAS-L</b>—Unavailable seconds (near-end line)</li> <li>• <b>ES-LFE</b>—Errored seconds (far-end line)</li> <li>• <b>SES-LFE</b>—Severely errored seconds (far-end line)</li> <li>• <b>UAS-LFE</b>—Unavailable seconds (far-end line)</li> </ul>	<b>extensive</b>

Table 12: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>WIS path</b>	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. Any state other than <b>OK</b> indicates a problem.</li> </ul> <p>Subfields are:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B3</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>REI-P</b>—Remote error indication</li> <li>• <b>LOP-P</b>—Loss of pointer (path)</li> <li>• <b>AIS-P</b>—Path alarm indication signal</li> <li>• <b>RDI-P</b>—Path remote defect indication</li> <li>• <b>UNEQ-P</b>—Path unequipped</li> <li>• <b>PLM-P</b>—Path payload label mismatch</li> <li>• <b>ES-P</b>—Errored seconds (near-end STS path)</li> <li>• <b>SES-P</b>—Severely errored seconds (near-end STS path)</li> <li>• <b>UAS-P</b>—Unavailable seconds (near-end STS path)</li> <li>• <b>SES-PFE</b>—Severely errored seconds (far-end STS path)</li> <li>• <b>UAS-PFE</b>—Unavailable seconds (far-end STS path)</li> </ul>	<b>extensive</b>

Table 12: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> <li>• <b>Negotiation status:</b> <ul style="list-style-type: none"> <li>• <b>Incomplete</b>—Ethernet interface has the speed or link mode configured.</li> <li>• <b>No autonegotiation</b>—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.</li> <li>• <b>Complete</b>—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> </ul> </li> <li>• <b>Link partner status</b>—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> <li>• <b>Link partner:</b> <ul style="list-style-type: none"> <li>• <b>Link mode</b>—Depending on the capability of the attached Ethernet device, either <b>Full-duplex</b> or <b>Half-duplex</b>.</li> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is <b>None</b>. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports both <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> receive).</li> <li>• <b>Remote fault</b>—Remote fault information from the link partner—<b>Failure</b> indicates a receive link error. <b>OK</b> indicates that the link partner is receiving. <b>Negotiation error</b> indicates a negotiation error. <b>Offline</b> indicates that the link partner is going offline.</li> </ul> </li> <li>• <b>Local resolution</b>—Information from the link partner: <ul style="list-style-type: none"> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports both <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> receive).</li> <li>• <b>Remote fault</b>—Remote fault information. <b>Link OK</b> (no error detected on receive), <b>Offline</b> (local interface is offline), and <b>Link Failure</b> (link error detected on receive).</li> </ul> </li> </ul>	extensive
Received path trace, Transmitted path trace	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the routing device at the other end of the fiber. The transmitted path trace value is the message that this routing device transmits.</p>	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number.</li> </ul>	extensive

Table 12: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>CoS information</b>	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> <li>• <b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li>• <b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li>• <b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li>• <b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li>• <b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li>• <b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li>• <b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul>	<b>extensive</b>
<b>Logical Interface</b>		
<b>Logical interface</b>	Name of the logical interface.	All levels
<b>Index</b>	Index number of the logical interface, which reflects its initialization sequence.	<b>detail extensive</b> none
<b>SNMP ifIndex</b>	SNMP interface index number for the logical interface.	<b>detail extensive</b> none
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Flags</b>	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> .	All levels

Table 12: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>VLAN-Tag</b>	<p>Rewrite profile applied to incoming or outgoing frames on the outer (<b>Out</b>) VLAN tag or for both the outer and inner (<b>In</b>) VLAN tags.</p> <ul style="list-style-type: none"> <li>• <b>push</b>—An outer VLAN tag is pushed in front of the existing VLAN tag.</li> <li>• <b>pop</b>—The outer VLAN tag of the incoming frame is removed.</li> <li>• <b>swap</b>—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information.</li> <li>• <b>push</b>—An outer VLAN tag is pushed in front of the existing VLAN tag.</li> <li>• <b>push-push</b>—Two VLAN tags are pushed in from the incoming frame.</li> <li>• <b>swap-push</b>—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.</li> <li>• <b>swap-swap</b>—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value.</li> <li>• <b>pop-swap</b>—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame.</li> <li>• <b>pop-pop</b>—Both the outer and inner VLAN tags of the incoming frame are removed.</li> </ul>	<b>brief detail extensive</b> none
<b>Demux:</b>	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> <li>• Source Family Inet</li> <li>• Destination Family Inet</li> </ul>	<b>detail extensive</b> none
<b>Encapsulation</b>	Encapsulation on the logical interface.	All levels
<b>Protocol</b>	Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .	<b>detail extensive</b> none
<b>MTU</b>	Maximum transmission unit size on the logical interface.	<b>detail extensive</b> none
<b>Maximum labels</b>	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	<b>detail extensive</b> none
<b>Traffic statistics</b>	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.</li> <li>• <b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface set.</li> </ul>	<b>detail extensive</b>
<b>IPv6 transit statistics</b>	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	<b>extensive</b>
<b>Local statistics</b>	Number and rate of bytes and packets destined to the routing device.	<b>extensive</b>



Table 12: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Transit statistics</b>	Number and rate of bytes and packets transiting the switch.  <b>NOTE:</b> For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the <b>Output bytes</b> and <b>Output packets</b> interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	<b>extensive</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route Table</b>	Route table in which the logical interface address is located. For example, <b>0</b> refers to the routing table inet.0.	<b>detail extensive none</b>
<b>Flags</b>	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive</b>
<b>Donor interface</b>	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	<b>detail extensive none</b>
<b>Preferred source address</b>	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	<b>detail extensive none</b>
<b>Input Filters</b>	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	<b>detail extensive</b>
<b>Output Filters</b>	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	<b>detail extensive</b>
<b>Mac-Validate Failures</b>	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	<b>detail extensive none</b>
<b>Addresses, Flags</b>	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b><i>protocol-family</i></b>	Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.	<b>brief</b>
<b>Flags</b>	Information about address flag (possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> ).	<b>detail extensive none</b>
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive none</b>
<b>Local</b>	IP address of the logical interface.	<b>detail extensive none</b>
<b>Broadcast</b>	Broadcast address of the logical interlace.	<b>detail extensive none</b>

Table 12: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>

For Gigabit Ethernet IQ PICs, traffic and MAC statistics output varies. [Table 13 on page 186](#) describes the traffic and MAC statistics for two sample interfaces, each of which is sending traffic in packets of 500 bytes (including 478 bytes for the Layer 3 packet, 18 bytes for the Layer 2 VLAN traffic header, and 4 bytes for cyclic redundancy check [CRC] information). In [Table 13 on page 186](#), the **ge-0/3/0** interface is the inbound physical interface, and the **ge-0/0/0** interface is the outbound physical interface. On both interfaces, traffic is carried on logical unit .50 (VLAN 50).

Table 13: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound physical interface	<b>show interfaces ge-0/3/0 extensive</b>	Traffic statistics:  Input bytes: 496 bytes per packet, representing the Layer 2 packet  MAC statistics:  Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes	The additional 4 bytes are for the CRC.
Inbound logical interface	<b>show interfaces ge-0/3/0.50 extensive</b>	Traffic statistics:  Input bytes: 478 bytes per packet, representing the Layer 3 packet	
Outbound physical interface	<b>show interfaces ge-0/0/0 extensive</b>	Traffic statistics:  Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes  MAC statistics:  Received octets: 478 bytes per packet, representing the Layer 3 packet	For input bytes, the additional 12 bytes includes 6 bytes for the destination MAC address + 4 bytes for VLAN + 2 bytes for the Ethernet type.
Outbound logical interface	<b>show interfaces ge-0/0/0.50 extensive</b>	Traffic statistics:  Input bytes: 478 bytes per packet, representing the Layer 3 packet	

## Sample Output

### show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2)

```

user@host> show interfaces xe-5/0/0 extensive
Physical interface: xe-5/0/0, Enabled, Physical link is Up
  Interface index: 177, SNMP ifIndex: 99, Generation: 178
  Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Speed: 10Gbps, Loopback:

```

```

None, Source filtering: Enabled,
Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags : None
CoS queues : 8 supported, 4 maximum usable queues
Schedulers : 1024
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:14:f6:b9:f1:f6, Hardware address: 00:14:f6:b9:f1:f6
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes : 6970332384 0 bps
Output bytes : 0 0 bps
Input packets: 81050506 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Ingress traffic statistics at Packet Forwarding Engine:
Input bytes : 6970299398 0 bps
Input packets: 81049992 0 pps
Drop bytes : 0 0 bps
Drop packets: 0 0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0,
MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets

0 best-effort 81049992 81049992 0
1 expedited-fo 0 0 0
2 assured-forw 0 0 0
3 network-cont 0 0 0

Egress queues: 4 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets

0 best-effort 0 0 0
1 expedited-fo 0 0 0
2 assured-forw 0 0 0
3 network-cont 0 0 0

Active alarms : None
Active defects : None
PCS statistics Seconds
Bit errors 0
Errored blocks 0

```

```

MAC statistics:
Total octets          6970332384
Total packets        81050506
Unicast packets      81050000
Broadcast packets    506
Multicast packets    0
CRC/Align errors     0
FIFO errors          0
MAC control frames   0
MAC pause frames     0
Oversized frames     0
Jabber frames        0
Fragment frames      0
VLAN tagged frames   0
Code violations       0

Filter statistics:
Input packet count    81050506
Input packet rejects  506
Input DA rejects      0
Input SA rejects      0
Output packet count   0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0

Packet Forwarding Engine configuration:
Destination slot: 5

CoS information:
Direction : Output
CoS transmit queue   Bandwidth      Buffer Priority Limit
                        %      bps      %      usec
0 best-effort        95    950000000  95      0      low  none
3 network-control    5     50000000   5      0      low  none

Direction : Input
CoS transmit queue   Bandwidth      Buffer Priority Limit
                        %      bps      %      usec
0 best-effort        95    950000000  95      0      low  none
3 network-control    5     50000000   5      0      low  none

Logical interface xe-5/0/0.0 (Index 71) (SNMP ifIndex 95) (Generation 195)
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ] Encapsulation: ENET2
Egress account overhead: 100
Ingress account overhead: 90

Traffic statistics:
Input bytes : 0
Output bytes : 46
Input packets: 0
Output packets: 1

IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Local statistics:
Input bytes : 0
Output bytes : 46
Input packets: 0
Output packets: 1

Transit statistics:
Input bytes : 0
Output bytes : 0

```

```

Input packets:                0                0 pps
Output packets:               0                0 pps
IPv6 transit statistics:
  Input bytes :                0
  Output bytes :               0
  Input packets:              0
  Output packets:             0
Protocol inet, MTU: 1500, Generation: 253, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 192.1.1/24, Local: 192.1.1.1, Broadcast: 192.1.1.255,
Generation: 265
Protocol multiservice, MTU: Unlimited, Generation: 254, Route table: 0
  Flags: None
  Policer: Input: __default_arp_policer__

```

### show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode)

```

user@host> show interfaces xe-1/0/0 extensive
Physical interface: xe-1/0/0, Enabled, Physical link is Up
Interface index: 141, SNMP ifIndex: 34, Generation: 47
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Loopback: Disabled
WAN-PHY mode
Source filtering: Disabled, Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps 16384
Link flags : None
CoS queues : 4 supported
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:05:85:a2:10:9d, Hardware address: 00:05:85:a2:10:9d
Last flapped : 2005-07-07 11:22:34 PDT (3d 12:28 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes :                0                0 bps
  Output bytes :               0                0 bps
  Input packets:              0                0 pps
  Output packets:             0                0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  HS Link CRC errors: 0, HS Link FIFO overflows: 0,
  Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0,
  Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
  Resource errors: 0
Queue counters:
  Queued packets  Transmitted packets  Dropped packets
0 best-effort    0                0                0
1 expedited-fo   0                0                0
2 assured-forw   0                0                0
3 network-cont   0                0                0
Active alarms : LOL, LOS, LBL
Active defects: LOL, LOS, LBL, SEF, AIS-L, AIS-P
PCS statistics
  Seconds  Count
Bit errors 0        0
Errored blocks 0      0
MAC statistics:
  Receive  Transmit
Total octets 0        0
Total packets 0        0
Unicast packets 0        0
Broadcast packets 0        0
Multicast packets 0        0

```

```

CRC/Align errors                0                0
FIFO errors                     0                0
MAC control frames              0                0
MAC pause frames                0                0
Oversized frames               0
Jabber frames                  0
Fragment frames                0
VLAN tagged frames             0
Code violations                 0
Filter statistics:
  Input packet count            0
  Input packet rejects          0
  Input DA rejects              0
  Input SA rejects              0
  Output packet count           0
  Output packet pad count       0
  Output packet error count     0
CAM destination filters: 0, CAM source filters: 0
PMA PHY:
  Seconds      Count  State
  PLL lock     0      0  OK
  PHY light    63159  1  Light Missing
WIS section:
  BIP-B1        0      0
  SEF           434430  434438 Defect Active
  LOS           434430  1  Defect Active
  LOF           434430  1  Defect Active
  ES-S          434430
  SES-S         434430
  SEFS-S        434430
WIS line:
  BIP-B2        0      0
  REI-L         0      0
  RDI-L         0      0  OK
  AIS-L         434430  1  Defect Active
  BERR-SF       0      0  OK
  BERR-SD       0      0  OK
  ES-L          434430
  SES-L         434430
  UAS-L         434420
  ES-LFE        0
  SES-LFE       0
  UAS-LFE       0
WIS path:
  BIP-B3        0      0
  REI-P         0      0
  LOP-P         0      0  OK
  AIS-P         434430  1  Defect Active
  RDI-P         0      0  OK
  UNEQ-P        0      0  OK
  PLM-P         0      0  OK
  ES-P          434430
  SES-P         434430
  UAS-P         434420
  ES-PFE        0
  SES-PFE       0
  UAS-PFE       0
Received path trace:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted path trace: orissa so-1/0/0
6f 72 69 73 73 61 20 73 6f 2d 31 2f 30 2f 30 00   orissa so-1/0/0.
Packet Forwarding Engine configuration:

```

```

Destination slot: 1
CoS information:
  CoS transmit queue      Bandwidth      Buffer      Priority  Limit
                           %      bps      %      bytes
  0 best-effort           95      950000000  95        0      low      none
  3 network-control       5       50000000   5         0      low      none

```

### show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC)

```

user@host> show interfaces ge-7/0/0 extensive
Physical interface: ge-7/0/0, Enabled, Physical link is Down
Interface index: 143, SNMP ifIndex: 508, Generation: 208
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled
Device flags   : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
Link flags     : None
Wavelength     : 1550.12 nm, Frequency: 193.40 THz
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:05:85:70:2b:72, Hardware address: 00:05:85:70:2b:72
Last flapped   : 2011-04-20 15:48:54 PDT (18:39:49 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   : 0          0 bps
Output bytes  : 0          0 bps
Input packets: 0          0 pps
Output packets: 0         0 pps
IPv6 transit statistics:
Input bytes   : 0
Output bytes  : 0
Input packets: 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 2, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort           0              0              0

  1 expedited-fo         0              0              0

  2 assured-forw         0              0              0

  3 network-cont
Queue number:      Mapped forwarding classes
  0                best-effort
  1                expedited-forwarding
  2                assured-forwarding
  3                network-control
Active alarms  : LINK
Active defects : LINK
MAC statistics:
Total octets      Receive      Transmit
Total packets     0              0

```

```

Unicast packets          0          0
Broadcast packets        0          0
Multicast packets        0          0
CRC/Align errors         0          0
FIFO errors              0          0
MAC control frames       0          0
MAC pause frames         0          0
Oversized frames         0
Jabber frames            0
Fragment frames          0
VLAN tagged frames       0
Code violations           0
Total octets             0          0
Total packets            0          0
Unicast packets          0          0
Broadcast packets        0          0
Multicast packets        0          0
CRC/Align errors         0          0
FIFO errors              0          0
MAC control frames       0          0
MAC pause frames         0          0
Oversized frames         0
Jabber frames            0
Fragment frames          0
VLAN tagged frames       0
Code violations           0
OTN alarms               : None
OTN defects              : None
OTN FEC Mode             : GFEC
OTN Rate                  : Fixed Stuff Bytes 11.0957Gbps
OTN Line Loopback        : Enabled
OTN FEC statistics :
  Corrected Errors              0
  Corrected Error Ratio (      0 sec average) 0e-0
OTN FEC alarms:      Seconds      Count  State
  FEC Degrade         0           0  OK
  FEC Excessive        0           0  OK
OTN OC:              Seconds      Count  State
  LOS                  2           1  OK
  LOF                  67164       2  Defect Active
  LOM                  67164       71 Defect Active
  Wavelength Lock      0           0  OK
OTN OTU:
  AIS                  0           0  OK
  BDI                  65919       4814 Defect Active
  IAE                  67158       1  Defect Active
  TTIM                 7           1  OK
  SF                   67164       2  Defect Active
  SD                   67164       3  Defect Active
  TCA-ES               0           0  OK
  TCA-SES              0           0  OK
  TCA-UAS              80          40  OK
  TCA-BBE              0           0  OK
  BIP                  0           0  OK
  BBE                  0           0  OK
  ES                   0           0  OK
  SES                  0           0  OK
  UAS                  587         0  OK
Received DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Received SAPI:

```



```

00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
OTN Received Overhead Bytes:
  APS/PCC0: 0x02, APS/PCC1: 0x42, APS/PCC2: 0xa2, APS/PCC3: 0x48
  Payload Type: 0x03
OTN Transmitted Overhead Bytes:
  APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
  Payload Type: 0x03
Filter statistics:
  Input packet count                0
  Input packet rejects              0
  Input DA rejects                  0
  Input SA rejects                  0
  Output packet count                0
  Output packet pad count            0
  Output packet error count          0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 7
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
      0 best-effort      95      9500000000    95      0      low
none
      3 network-control   5      500000000    5      0      low
none
...

```

#### show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode)

```

user@host> show interfaces xe-7/0/0 extensive
Physical interface: xe-7/0/0, Enabled, Physical link is Up
  Interface index: 173, SNMP ifIndex: 212, Generation: 174
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
  Unidirectional: Enabled,
  Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
...

```

#### show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only)

```

user@host> show interfaces xe-7/0/0-tx extensive
Physical interface: xe-7/0/0-tx, Enabled, Physical link is Up
  Interface index: 176, SNMP ifIndex: 137, Generation: 177
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
  Unidirectional: Tx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:73:e4:83, Hardware address: 00:05:85:73:e4:83
  Last flapped   : 2007-06-01 09:08:19 PDT (3d 02:31 ago)
  Statistics last cleared: Never
Traffic statistics:
  Input bytes   :                0                0 bps

```

```

Output bytes :      322891152287160      9627472888 bps
Input packets:              0              0 pps
Output packets:    328809727380      1225492 pps

...

Filter statistics:
  Output packet count      328810554250
  Output packet pad count      0
  Output packet error count    0
...

Logical interface xe-7/0/0-tx.0 (Index 73) (SNMP ifIndex 138) (Generation 139)

Flags: SNMP-Traps Encapsulation: ENET2
Egress account overhead: 100
Ingress account overhead: 90
Traffic statistics:
  Input bytes :              0
  Output bytes :    322891152287160
  Input packets:              0
  Output packets:    328809727380
IPv6 transit statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:              0
Local statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:              0
Transit statistics:
  Input bytes :              0              0 bps
  Output bytes :    322891152287160      9627472888 bps
  Input packets:              0              0 pps
  Output packets:    328809727380      1225492 pps
IPv6 transit statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:              0
Protocol inet, MTU: 1500, Generation: 147, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.11.12/24, Local: 10.11.12.13, Broadcast: 10.11.12.255,
  Generation: 141
Protocol multiservice, MTU: Unlimited, Generation: 148, Route table: 0
  Flags: None
  Policer: Input: __default_arp_policer__

```

#### show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only)

```

user@host> show interfaces xe-7/0/0-rx extensive
Physical interface: xe-7/0/0-rx, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 118, Generation: 175
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
  Unidirectional: Rx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues

```

```

Hold-times      : Up 0 ms, Down 0 ms
Current address: 00:05:85:73:e4:83, Hardware address: 00:05:85:73:e4:83
Last flapped   : 2007-06-01 09:08:22 PDT (3d 02:31 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes :      322857456303482      9627496104 bps
Output bytes :              0              0 bps
Input packets:      328775413751      1225495 pps
Output packets:              0              0 pps

...

Filter statistics:
Input packet count      328775015056
Input packet rejects    1
Input DA rejects        0

...

Logical interface xe-7/0/0-rx.0 (Index 72) (SNMP ifIndex 120) (Generation 138)

Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
Input bytes :      322857456303482
Output bytes :              0
Input packets:      328775413751
Output packets:              0
IPv6 transit statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:              0
Local statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:              0
Transit statistics:
Input bytes :      322857456303482      9627496104 bps
Output bytes :              0              0 bps
Input packets:      328775413751      1225495 pps
Output packets:              0              0 pps
IPv6 transit statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:              0
Protocol inet, MTU: 1500, Generation: 145, Route table: 0
Addresses, Flags: Is-Preferred Is-Primary
Destination: 192.1.1/24, Local: 192.1.1.1, Broadcast: 192.1.1.255,
Generation: 139
Protocol multiservice, MTU: Unlimited, Generation: 146, Route table: 0
Flags: None
Policer: Input: __default_arp_policer__

```

## show interfaces (Gigabit Ethernet)

<b>Syntax</b>	<pre>show interfaces <i>ge-fpc/pic/port</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Display status information about the specified Gigabit Ethernet interface.
<b>Options</b>	<p><b><i>ge-fpc/pic/port</i></b>—Display standard information about the specified Gigabit Ethernet interface.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p>
<b>Additional Information</b>	In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show interfaces (Gigabit Ethernet) on page 212</a></p> <p><a href="#">show interfaces (Gigabit Ethernet on MX Series Routers) on page 213</a></p> <p><a href="#">show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration) on page 213</a></p> <p><a href="#">show interfaces brief (Gigabit Ethernet) on page 213</a></p> <p><a href="#">show interfaces detail (Gigabit Ethernet) on page 214</a></p> <p><a href="#">show interfaces extensive (Gigabit Ethernet IQ2) on page 215</a></p> <p><a href="#">show interfaces (Gigabit Ethernet Unnumbered Interface) on page 218</a></p> <p><a href="#">show interfaces (ACI Interface Set Configured) on page 219</a></p>
<b>Output Fields</b>	<p><a href="#">Table 14 on page 197</a> describes the output fields for the <b>show interfaces</b> (Gigabit Ethernet) command. Output fields are listed in the approximate order in which they appear. For Gigabit Ethernet IQ and IQE PICs, the traffic and MAC statistics vary by interface type. For more information, see <a href="#">Table 15 on page 211</a>.</p>

Table 14: show interfaces (Gigabit Ethernet) Output Fields

Field Name	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Physical interface</b>	Name of the physical interface.	All levels
<b>Enabled</b>	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
<b>Interface index</b>	Index number of the physical interface, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Link-level type</b>	Encapsulation being used on the physical interface.	All levels
<b>MTU</b>	Maximum transmission unit size on the physical interface.	All levels
<b>Speed</b>	Speed at which the interface is running.	All levels
<b>Loopback</b>	Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .	All levels
<b>Source filtering</b>	Source filtering status: <b>Enabled</b> or <b>Disabled</b> .	All levels
<b>LAN-PHY mode</b>	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
<b>WAN-PHY mode</b>	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
<b>Unidirectional</b>	Unidirectional link mode status for 10-Gigabit Ethernet interface: <b>Enabled</b> or <b>Disabled</b> for parent interface; <b>Rx-only</b> or <b>Tx-only</b> for child interfaces.	All levels
<b>Flow control</b>	Flow control status: <b>Enabled</b> or <b>Disabled</b> .	All levels
<b>Auto-negotiation</b>	(Gigabit Ethernet interfaces) Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .	All levels
<b>Remote-fault</b>	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> <li>• <b>Online</b>—Autonegotiation is manually configured as online.</li> <li>• <b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul>	All levels
<b>Device flags</b>	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
<b>Interface flags</b>	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels

Table 14: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Link flags</b>	Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .	All levels
<b>Wavelength</b>	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
<b>Frequency</b>	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
<b>CoS queues</b>	Number of CoS queues configured.	<b>detail extensive</b> none
<b>Schedulers</b>	(Gigabit Ethernet intelligent queuing 2 [IQ2] interfaces only) Number of CoS schedulers configured.	<b>extensive</b>
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds (ms).	<b>detail extensive</b>
<b>Current address</b>	Configured MAC address.	<b>detail extensive</b> none
<b>Hardware address</b>	Hardware MAC address.	<b>detail extensive</b> none
<b>Last flapped</b>	Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .	<b>detail extensive</b> none
<b>Input Rate</b>	Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None
<b>Output Rate</b>	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>
<b>Egress account overhead</b>	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	<b>detail extensive</b>
<b>Ingress account overhead</b>	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	<b>detail extensive</b>

Table 14: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Traffic statistics</b>	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the <a href="#">show interfaces (10-Gigabit Ethernet)</a> command.</p>	<b>detail extensive</b>
<b>Input errors</b>	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the <b>ignore-l3-incompletes</b> statement.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>

Table 14: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Output errors</b>	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the <b>Drops</b> field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> <ul style="list-style-type: none"> <li>• <b>Collisions</b>—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>
<b>Egress queues</b>	<p>Total number of egress queues supported on the specified interface.</p> <p><b>NOTE:</b> In DPCs that are not of the enhanced type, such as DPC 40x 1GER, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the <b>show interfaces</b> command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenario, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs</p>	<b>detail extensive</b>



Table 14: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the <b>Dropped packets</b> field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>	detail extensive
Ingress queues	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	extensive
Queue counters (Ingress)	<p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>	extensive
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul>	detail extensive none
Interface transmit statistics	<p>(On MX Series devices) Status of the <b>interface-transmit-statistics</b> configuration: Enabled or Disabled.</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—When the <b>interface-transmit-statistics</b> statement is included in the configuration. If this is configured, the interface statistics show the actual transmitted load on the interface.</li> <li>• <b>Disabled</b>—When the <b>interface-transmit-statistics</b> statement is not included in the configuration. If this is not configured, the interface statistics show the offered load on the interface.</li> </ul>	detail extensive
OTN FEC statistics	<p>The forward error correction (FEC) counters provide the following statistics:</p> <ul style="list-style-type: none"> <li>• <b>Corrected Errors</b>—Count of corrected errors in the last second.</li> <li>• <b>Corrected Error Ratio</b>—Corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.</li> </ul>	detail extensive

Table 14: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>PCS statistics</b>	<p>(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.</p> <ul style="list-style-type: none"><li>• <b>Bit errors</b>—Number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode.</li><li>• <b>Errored blocks</b>—Number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.</li></ul>	<b>detail extensive</b>

Table 14: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> <li>• <b>Total octets</b> and <b>total packets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the <a href="#">show interfaces (10-Gigabit Ethernet)</a> command.</li> <li>• <b>Unicast packets</b>, <b>Broadcast packets</b>, and <b>Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—There are two possible conditions regarding the number of oversized frames: <ul style="list-style-type: none"> <li>• Packet length exceeds 1518 octets, or</li> <li>• Packet length exceeds MRU</li> </ul> </li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>VLAN tagged frames</b>—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. <p><b>NOTE:</b> The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the <b>VLAN tagged frames</b> field displays 0 when the <b>show interfaces</b> command is executed on a 20-port Gigabit Ethernet MIC. In other words, the number of VLAN tagged frames cannot be determined for the 20-port Gigabit Ethernet MIC.</p> </li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."</li> </ul>	extensive
OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	extensive
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	extensive

Table 14: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Filter statistics</b>	<p><b>Receive</b> and <b>Transmit</b> statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> <li>• <b>Input packet count</b>—Number of packets received from the MAC hardware that the filter processed.</li> <li>• <b>Input packet rejects</b>—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address.</li> <li>• <b>Input DA rejects</b>—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the router from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local router (which the router is rejecting).</li> <li>• <b>Input SA rejects</b>—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect.</li> <li>• <b>Output packet count</b>—Number of packets that the filter has given to the MAC hardware.</li> <li>• <b>Output packet pad count</b>—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured.</li> <li>• <b>Output packet error count</b>—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment.</li> <li>• <b>CAM destination filters, CAM source filters</b>—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.</li> </ul>	<b>extensive</b>
<b>PMA PHY</b>	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. Any state other than <b>OK</b> indicates a problem.</li> </ul> <p>Subfields are:</p> <ul style="list-style-type: none"> <li>• <b>PHY Lock</b>—Phase-locked loop</li> <li>• <b>PHY Light</b>—Loss of optical signal</li> </ul>	<b>extensive</b>

Table 14: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>WIS section</b>	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. Any state other than <b>OK</b> indicates a problem.</li> </ul> <p>Subfields are:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B1</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>SEF</b>—Severely errored framing</li> <li>• <b>LOL</b>—Loss of light</li> <li>• <b>LOF</b>—Loss of frame</li> <li>• <b>ES-S</b>—Errored seconds (section)</li> <li>• <b>SES-S</b>—Severely errored seconds (section)</li> <li>• <b>SEFS-S</b>—Severely errored framing seconds (section)</li> </ul>	<b>extensive</b>
<b>WIS line</b>	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information:</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. Any state other than <b>OK</b> indicates a problem.</li> </ul> <p>Subfields are:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B2</b>—Bit interleaved parity for SONET line overhead</li> <li>• <b>REI-L</b>—Remote error indication (near-end line)</li> <li>• <b>RDI-L</b>—Remote defect indication (near-end line)</li> <li>• <b>AIS-L</b>—Alarm indication signal (near-end line)</li> <li>• <b>BERR-SF</b>—Bit error rate fault (signal failure)</li> <li>• <b>BERR-SD</b>—Bit error rate defect (signal degradation)</li> <li>• <b>ES-L</b>—Errored seconds (near-end line)</li> <li>• <b>SES-L</b>—Severely errored seconds (near-end line)</li> <li>• <b>UAS-L</b>—Unavailable seconds (near-end line)</li> <li>• <b>ES-LFE</b>—Errored seconds (far-end line)</li> <li>• <b>SES-LFE</b>—Severely errored seconds (far-end line)</li> <li>• <b>UAS-LFE</b>—Unavailable seconds (far-end line)</li> </ul>	<b>extensive</b>

Table 14: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>WIS path</b>	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information:</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. Any state other than <b>OK</b> indicates a problem.</li> </ul> <p>Subfields are:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B3</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>REI-P</b>—Remote error indication</li> <li>• <b>LOP-P</b>—Loss of pointer (path)</li> <li>• <b>AIS-P</b>—Path alarm indication signal</li> <li>• <b>RDI-P</b>—Path remote defect indication</li> <li>• <b>UNEQ-P</b>—Path unequipped</li> <li>• <b>PLM-P</b>—Path payload (signal) label mismatch</li> <li>• <b>ES-P</b>—Errored seconds (near-end STS path)</li> <li>• <b>SES-P</b>—Severely errored seconds (near-end STS path)</li> <li>• <b>UAS-P</b>—Unavailable seconds (near-end STS path)</li> <li>• <b>SES-PFE</b>—Severely errored seconds (far-end STS path)</li> <li>• <b>UAS-PFE</b>—Unavailable seconds (far-end STS path)</li> </ul>	<b>extensive</b>

Table 14: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> <li>• <b>Negotiation status:</b> <ul style="list-style-type: none"> <li>• <b>Incomplete</b>—Ethernet interface has the speed or link mode configured.</li> <li>• <b>No autonegotiation</b>—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.</li> <li>• <b>Complete</b>—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> </ul> </li> <li>• <b>Link partner status</b>—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> <li>• <b>Link partner</b>—Information from the remote Ethernet device: <ul style="list-style-type: none"> <li>• <b>Link mode</b>—Depending on the capability of the link partner, either <b>Full-duplex</b> or <b>Half-duplex</b>.</li> <li>• <b>Flow control</b>—Types of flow control supported by the link partner. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), <b>Symmetric/Asymmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> on transmit), and <b>None</b> (link partner does not support flow control).</li> <li>• <b>Remote fault</b>—Remote fault information from the link partner—<b>Failure</b> indicates a receive link error. <b>OK</b> indicates that the link partner is receiving. <b>Negotiation error</b> indicates a negotiation error. <b>Offline</b> indicates that the link partner is going offline.</li> </ul> </li> <li>• <b>Local resolution</b>—Information from the local Ethernet device: <ul style="list-style-type: none"> <li>• <b>Flow control</b>—Types of flow control supported by the local device. For Gigabit Ethernet interfaces, advertised capabilities are <b>Symmetric/Asymmetric</b> (local device supports <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> on receive) and <b>None</b> (local device does not support flow control). Depending on the result of the negotiation with the link partner, local resolution flow control type will display <b>Symmetric</b> (local device supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (local device supports <b>PAUSE</b> on receive), and <b>None</b> (local device does not support flow control).</li> <li>• <b>Remote fault</b>—Remote fault information. <b>Link OK</b> (no error detected on receive), <b>Offline</b> (local interface is offline), and <b>Link Failure</b> (link error detected on receive).</li> </ul> </li> </ul>	extensive
Received path trace, Transmitted path trace	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.</p>	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number.</li> </ul>	extensive

Table 14: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>CoS information</b>	<p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li>• <b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li>• <b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li>• <b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li>• <b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li>• <b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li>• <b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul>	<b>extensive</b>
<b>Logical Interface</b>		
<b>Logical interface</b>	Name of the logical interface.	All levels
<b>Index</b>	Index number of the logical interface, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP interface index number for the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Flags</b>	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> .	All levels



Table 14: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>VLAN-Tag</b>	<p>Rewrite profile applied to incoming or outgoing frames on the outer (<b>Out</b>) VLAN tag or for both the outer and inner (<b>In</b>) VLAN tags.</p> <ul style="list-style-type: none"> <li>• <b>push</b>—An outer VLAN tag is pushed in front of the existing VLAN tag.</li> <li>• <b>pop</b>—The outer VLAN tag of the incoming frame is removed.</li> <li>• <b>swap</b>—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information.</li> <li>• <b>push</b>—An outer VLAN tag is pushed in front of the existing VLAN tag.</li> <li>• <b>push-push</b>—Two VLAN tags are pushed in from the incoming frame.</li> <li>• <b>swap-push</b>—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.</li> <li>• <b>swap-swap</b>—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value.</li> <li>• <b>pop-swap</b>—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame.</li> <li>• <b>pop-pop</b>—Both the outer and inner VLAN tags of the incoming frame are removed.</li> </ul>	<b>brief detail extensive none</b>
<b>Demux</b>	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> <li>• Source Family Inet</li> <li>• Destination Family Inet</li> </ul>	<b>detail extensive none</b>
<b>Encapsulation</b>	Encapsulation on the logical interface.	All levels
<b>ACI VLAN: Dynamic Profile</b>	Name of the dynamic profile that defines the agent circuit identifier (ACI) interface set. If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information.	<b>brief detail extensive none</b>
<b>Protocol</b>	Protocol family. Possible values are described in the "Protocol Field" section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>MTU</b>	Maximum transmission unit size on the logical interface.	<b>detail extensive none</b>
<b>Neighbor Discovery Protocol (NDP) Queue Statistics</b>	<p>NDP statistics for protocol <b>inet6</b> under logical interface statistics.</p> <ul style="list-style-type: none"> <li>• <b>Max nh cache</b>—Maximum interface neighbor discovery nexthop cache size.</li> <li>• <b>New hold nh limit</b>—Maximum number of new unresolved nexthops.</li> <li>• <b>Curr nh cnt</b>—Current number of resolved nexthops in the NDP queue.</li> <li>• <b>Curr new hold cnt</b>—Current number of unresolved nexthops in the NDP queue.</li> <li>• <b>NH drop cnt</b>—Number of NDP requests not serviced.</li> </ul>	All levels
<b>Dynamic Profile</b>	Name of the dynamic profile that was used to create this interface configured with a Point-to-Point Protocol over Ethernet (PPPoE) family.	<b>detail extensive none</b>
<b>Service Name Table</b>	Name of the service name table for the interface configured with a PPPoE family.	<b>detail extensive none</b>

Table 14: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Max Sessions</b>	Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	<b>detail extensive none</b>
<b>Duplicate Protection</b>	State of PPPoE duplicate protection: <b>On</b> or <b>Off</b> . When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client.	<b>detail extensive none</b>
<b>Direct Connect</b>	State of the configuration to ignore DSL Forum VSAs: <b>On</b> or <b>Off</b> . When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	<b>detail extensive none</b>
<b>AC Name</b>	Name of the access concentrator.	<b>detail extensive none</b>
<b>Maximum labels</b>	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	<b>detail extensive none</b>
<b>Traffic statistics</b>	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> <li><b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.</li> <li><b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface set.</li> </ul>	<b>detail extensive</b>
<b>IPv6 transit statistics</b>	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	<b>extensive</b>
<b>Local statistics</b>	Number and rate of bytes and packets destined to the router.	<b>extensive</b>
<b>Transit statistics</b>	<p>Number and rate of bytes and packets transiting the switch.</p> <p><b>NOTE:</b> For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the <b>Output bytes</b> and <b>Output packets</b> interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.</p>	<b>extensive</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route Table</b>	Route table in which the logical interface address is located. For example, <b>0</b> refers to the routing table inet.0.	<b>detail extensive none</b>
<b>Flags</b>	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive</b>

Table 14: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Donor interface</b>	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	<b>detail extensive none</b>
<b>Preferred source address</b>	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	<b>detail extensive none</b>
<b>Input Filters</b>	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	<b>detail extensive</b>
<b>Output Filters</b>	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	<b>detail extensive</b>
<b>Mac-Validate Failures</b>	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	<b>detail extensive none</b>
<b>Addresses, Flags</b>	Information about the address flags. Possible values are described in the "Addresses Flags" section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b><i>protocol-family</i></b>	Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.	<b>brief</b>
<b>Flags</b>	Information about the address flag. Possible values are described in the "Addresses Flags" section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive none</b>
<b>Local</b>	IP address of the logical interface.	<b>detail extensive none</b>
<b>Broadcast</b>	Broadcast address of the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>

Table 15: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound physical interface	<b>show interfaces ge-0/3/0 extensive</b>	<p>Traffic statistics:</p> <p>Input bytes: 496 bytes per packet, representing the Layer 2 packet</p> <p>MAC statistics:</p> <p>Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes</p>	The additional 4 bytes are for the CRC.

Table 15: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type (*continued*)

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound logical interface	<b>show interfaces ge-0/3/0.50 extensive</b>	Traffic statistics:  Input bytes: 478 bytes per packet, representing the Layer 3 packet	
Outbound physical interface	<b>show interfaces ge-0/0/0 extensive</b>	Traffic statistics:  Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes  MAC statistics:  Received octets: 478 bytes per packet, representing the Layer 3 packet	For input bytes, the additional 12 bytes include 6 bytes for the destination MAC address plus 4 bytes for VLAN plus 2 bytes for the Ethernet type.
Outbound logical interface	<b>show interfaces ge-0/0/0.50 extensive</b>	Traffic statistics:  Input bytes: 478 bytes per packet, representing the Layer 3 packet	

## Sample Output

### show interfaces (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues    : 4 supported, 4 maximum usable queues
  Current address: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
  Last flapped  : 2006-08-10 17:25:10 PDT (00:01:08 ago)
  Input rate    : 0 bps (0 pps)
  Output rate   : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms : None
  Active defects : None

Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69)
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  Egress account overhead: 100
  Ingress account overhead: 90
  Input packets : 0
  Output packets: 0
  Protocol ccc, MTU: 1522
  Flags: Is-Primary

```

**show interfaces (Gigabit Ethernet on MX Series Routers)**

```

user@host> show interfaces ge-2/2/2
Physical interface: ge-2/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 188
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, MAC-REWRITE Error: None,
  Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 0
  Current address: 00:1f:12:b7:d7:c0, Hardware address: 00:1f:12:b7:d6:76
  Last flapped   : 2008-09-05 16:44:30 PDT (3d 01:04 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  Logical interface ge-2/2/2.0 (Index 82) (SNMP ifIndex 219)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Input packets : 10232
    Output packets: 10294
    Protocol inet, MTU: 1500
      Flags: Sendbroadcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
    Protocol inet6, MTU: 1500
      Max nh cache: 4, New hold nh limit: 100000, Curr nh cnt: 4, Curr new hold
      cnt: 4, NH drop cnt: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 2001:db8::/32, Local: 2001:db8::5
      Addresses, Flags: Is-Preferred
        Destination: 2001:db8::/32, Local: 2001:db8:223:9cff:fe9f:3e78
    Protocol multiservice, MTU: Unlimited
      Flags: Is-Primary

```

**show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration)**

```

user@host> show interfaces ge-2/1/2 extensive | match "output|interface"
Physical interface: ge-2/1/2, Enabled, Physical link is Up
  Interface index: 151, SNMP ifIndex: 530, Generation: 154
  Interface flags: SNMP-Traps Internal: 0x4000
  Output bytes   :          240614363944          772721536 bps
  Output packets:          3538446506          1420444 pps
  Direction     : Output
  Interface transmit statistics: Enabled

  Logical interface ge-2/1/2.0 (Index 331) (SNMP ifIndex 955) (Generation 146)
    Output bytes   :          195560312716          522726272 bps
    Output packets:          4251311146          1420451 pps

```

**show interfaces brief (Gigabit Ethernet)**

```

user@host> show interfaces ge-3/0/2 brief
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,

```

Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,  
 Remote fault: Online  
 Device flags : Present Running  
 Interface flags: SNMP-Traps Internal: 0x4000  
 Link flags : None

Logical interface ge-3/0/2.0  
 Flags: SNMP-Traps 0x4000  
 VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push  
 0x8100.512 0x8100.513)  
 Encapsulation: VLAN-CCC  
 ccc

Logical interface ge-3/0/2.32767  
 Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2

### show interfaces detail (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2 detail
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35, Generation: 177
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags : None
  CoS queues : 4 supported, 4 maximum usable queues
  Hold-times : Up 0 ms, Down 0 ms
  Current address: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
  Last flapped : 2006-08-09 17:17:00 PDT (01:31:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes : 0 0 bps
    Output bytes : 0 0 bps
    Input packets: 0 0 pps
    Output packets: 0 0 pps
  Ingress traffic statistics at Packet Forwarding Engine:
    Input bytes : 0 0 bps
    Input packets: 0 0 pps
    Drop bytes : 0 0 bps
    Drop packets: 0 0 pps
  Ingress queues: 4 supported, 4 in use
  Queue counters:
    Queued packets Transmitted packets Dropped packets

    0 best-effort 0 0 0
    1 expedited-fo 0 0 0
    2 assured-forw 0 0 0
    3 network-cont 0 0 0

  Egress queues: 4 supported, 4 in use
  Queue counters:
    Queued packets Transmitted packets Dropped packets

    0 best-effort 0 0 0
    1 expedited-fo 0 0 0
    2 assured-forw 0 0 0
```

```

3 network-cont                                0                                0                                0

Active alarms : None
Active defects : None

Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69) (Generation 140)
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530)
Out(swap-push 0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  Egress account overhead: 100
  Ingress account overhead: 90
  Traffic statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Local statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Transit statistics:
    Input bytes : 0 0 bps
    Output bytes : 0 0 bps
    Input packets: 0 0 pps
    Output packets: 0 0 pps
  Protocol ccc, MTU: 1522, Generation: 149, Route table: 0
  Flags: Is-Primary

Logical interface ge-3/0/2.32767 (Index 71) (SNMP ifIndex 70)
(Generation 139)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2
  Traffic statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Local statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Transit statistics:
    Input bytes : 0 0 bps
    Output bytes : 0 0 bps
    Input packets: 0 0 pps
    Output packets: 0 0 pps

```

### show interfaces extensive (Gigabit Ethernet IQ2)

```

user@host> show interfaces ge-7/1/3 extensive
Physical interface: ge-7/1/3, Enabled, Physical link is Up
Interface index: 170, SNMP ifIndex: 70, Generation: 171
Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4004000
Link flags : None

```

```

CoS queues      : 8 supported, 4 maximum usable queues
Schedulers      : 256
Hold-times      : Up 0 ms, Down 0 ms
Current address: 00:14:f6:30:5e:74, Hardware address: 00:14:f6:30:5e:74
Last flapped    : 2007-11-07 21:31:41 PST (02:03:33 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes :          38910844056          7952 bps
Output bytes :           7174605          8464 bps
Input packets:         418398473           11 pps
Output packets:         78903           12 pps
IPv6 transit statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:              0
Ingress traffic statistics at Packet Forwarding Engine:
Input bytes :         38910799145          7952 bps
Input packets:         418397956           11 pps
Drop bytes :              0           0 bps
Drop packets:              0           0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

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

0 best-effort          418390823          418390823              0

1 expedited-fo              0              0              0

2 assured-forw              0              0              0

3 network-cont           7133           7133              0

Egress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

0 best-effort          1031           1031              0

1 expedited-fo              0              0              0

2 assured-forw              0              0              0

3 network-cont          77872          77872              0

Active alarms : None
Active defects : None
MAC statistics:
Total octets          38910844056          7174605
Total packets         418398473          78903
Unicast packets       408021893366          1026
Broadcast packets           10           12
Multicast packets     418398217          77865
CRC/Align errors              0              0
FIFO errors              0              0

```



```

MAC control frames          0          0
MAC pause frames            0          0
Oversized frames            0
Jabber frames                0
Fragment frames              0
VLAN tagged frames          0
Code violations              0 OTN Received Overhead Bytes:
APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58
Payload Type: 0x08
OTN Transmitted Overhead Bytes:
APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
Payload Type: 0x08
Filter statistics:
  Input packet count        418398473
  Input packet rejects      479
  Input DA rejects          479
  Input SA rejects          0
  Output packet count              78903
  Output packet pad count         0
  Output packet error count       0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
    Link mode: Full-duplex, Flow control: Symmetric/Asymmetric,
    Remote fault: OK
  Local resolution:
    Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 7
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer      Priority      Limit
                           %      bps      %      usec
  0 best-effort           95      950000000  95      0
low  none
  3 network-control       5      500000000  5      0
low  none
  Direction : Input
  CoS transmit queue      Bandwidth      Buffer      Priority      Limit
                           %      bps      %      usec
  0 best-effort           95      950000000  95      0
low  none
  3 network-control       5      500000000  5      0
low  none

Logical interface ge-7/1/3.0 (Index 70) (SNMP ifIndex 85) (Generation 150)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
  Input bytes :      812400
  Output bytes :    1349206
  Input packets:      9429
  Output packets:     9449
IPv6 transit statistics:
  Input bytes :      0
  Output bytes :      0
  Input packets:      0
  Output packets:      0
Local statistics:
  Input bytes :      812400
  Output bytes :    1349206

```

```

Input packets:          9429
Output packets:         9449
Transit statistics:
Input bytes :           0          7440 bps
Output bytes :          0          7888 bps
Input packets:          0          10 pps
Output packets:         0          11 pps
IPv6 transit statistics:
Input bytes :           0
Output bytes :          0
Input packets:          0
Output packets:         0
Protocol inet, MTU: 1500, Generation: 169, Route table: 0
Flags: Is-Primary, Mac-Validate-Strict
Mac-Validate Failures: Packets: 0, Bytes: 0
Addresses, Flags: Is-Preferred Is-Primary
Input Filters: F1-ge-3/0/1.0-in, F3-ge-3/0/1.0-in
Output Filters: F2-ge-3/0/1.0-out (53)
Destination: 10.74.2/24, Local: 10.74.2.2, Broadcast: 10.74.2.255,
Generation: 196
Protocol multiservice, MTU: Unlimited, Generation: 170, Route table: 0
Flags: Is-Primary
Policer: Input: __default_arp_policer__

```

**NOTE:** For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics displayed in the **show interfaces** command output might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the interface counters. For detailed information, see the description of the logical interface **Transit statistics** fields in [Table 14 on page 197](#).

### show interfaces (Gigabit Ethernet Unnumbered Interface)

```

user@host> show interfaces ge-3/2/0
Physical interface: ge-3/2/0, Enabled, Physical link is Up
Interface index: 148, SNMP ifIndex: 50
Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags     : None
CoS queues    : 8 supported, 4 maximum usable queues
Current address: 00:14:f6:11:26:f8, Hardware address: 00:14:f6:11:26:f8
Last flapped   : 2006-10-27 04:42:23 PDT (08:01:52 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 624 bps (1 pps)
Active alarms  : None
Active defects : None

Logical interface ge-3/2/0.0 (Index 67) (SNMP ifIndex 85)
Flags: SNMP-Traps Encapsulation: ENET2
Input packets : 0
Output packets: 6
Protocol inet, MTU: 1500
Flags: Unnumbered
Donor interface: lo0.0 (Index 64)
Preferred source address: 192.02.5

```

### show interfaces (ACI Interface Set Configured)

```
user@host> show interfaces ge-1/0/0.4001
Logical interface ge-1/0/0.4001 (Index 340) (SNMP ifIndex 548)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4001 ] Encapsulation: PPP-over-
    Ethernet
  ACI VLAN:
    Dynamic Profile: aci-vlan-set-profile
  PPPoE:
    Dynamic Profile: aci-vlan-pppoe-profile,
    Service Name Table: None,
    Max Sessions: 32000, Max Sessions VSA Ignore: Off,
    Duplicate Protection: On, Short Cycle Protection: Off,
    Direct Connect: Off,
    AC Name: nbc
  Input packets : 9
  Output packets: 8
  Protocol multiservice, MTU: Unlimited
```

## show interfaces irb

<b>Syntax</b>	<pre>show interfaces irb &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;routing-instance <i>instance-name</i>&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 12.3R2.</p> <p>Command introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2 for the QFX Series</p>
<b>Description</b>	Display integrated routing and bridging interfaces information.
<b>Options</b>	<p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>routing-instance <i>instance-name</i></b>—(Optional) Display information for the interface with the specified SNMP index.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the interface with the specified SNMP index.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p>
<b>Additional Information</b>	Integrated routing and bridging (IRB) provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route local packets to another routed interface or to another VLAN that has a Layer 3 protocol configured.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show interfaces irb extensive on page 224</a></p> <p><a href="#">show interfaces irb snmp-index on page 225</a></p>
<b>Output Fields</b>	<a href="#">Table 16 on page 220</a> lists the output fields for the <b>show interfaces irb</b> command. Output fields are listed in the approximate order in which they appear.

Table 16: show interfaces irb Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the physical interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels

Table 16: show interfaces irb Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Proto</b>	Protocol configured on the interface.	<b>terse</b>
<b>Interface index</b>	Physical interface index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive none</b>
<b>Type</b>	Physical interface type.	<b>detail extensive none</b>
<b>Link-level type</b>	Encapsulation being used on the physical interface.	<b>detail extensive brief none</b>
<b>MTU</b>	MTU size on the physical interface.	<b>detail extensive brief none</b>
<b>Clocking</b>	Reference clock source: <b>Internal</b> or <b>External</b> . Always unspecified on IRB interfaces.	<b>detail extensive brief</b>
<b>Speed</b>	Speed at which the interface is running. Always unspecified on IRB interfaces.	<b>detail extensive brief</b>
<b>Device flags</b>	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive brief none</b>
<b>Interface flags</b>	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive brief none</b>
<b>Link type</b>	Physical interface link type: <b>full duplex</b> or <b>half duplex</b> .	<b>detail extensive none</b>
<b>Link flags</b>	Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>Physical Info</b>	Physical interface information.	All levels
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds.	<b>detail extensive</b>
<b>Current address</b>	Configured MAC address.	<b>detail extensive none</b>
<b>Hardware address</b>	MAC address of the hardware.	<b>detail extensive none</b>
<b>Alternate link address</b>	Backup address of the link.	<b>detail extensive</b>
<b>Last flapped</b>	Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .	<b>detail extensive none</b>
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>

Table 16: show interfaces irb Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Traffic statistics</b>	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>IPv6 transit statistics</b>	<p>Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Input errors</b>	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>detail extensive</b>
<b>Output errors</b>	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the DPC is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>detail extensive</b>

---

#### Logical Interface

---

Table 16: show interfaces irb Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Logical interface</b>	Name of the logical interface.	All levels
<b>Index</b>	Index number of the logical interface (which reflects its initialization sequence).	<b>detail extensive</b> none
<b>SNMP ifIndex</b>	SNMP interface index number of the logical interface.	<b>detail extensive</b> none
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Flags</b>	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive</b>
<b>Encapsulation</b>	Encapsulation on the logical interface.	<b>detail extensive</b>
<b>Bandwidth</b>	Speed at which the interface is running.	<b>detail extensive</b>
<b>Routing Instance</b>	Routing instance IRB is configured under.	<b>detail extensive</b>
<b>Bridging Domain</b>	Bridging domain IRB is participating in.	<b>detail extensive</b>
<b>Traffic statistics</b>	Number and rate of bytes and packets received and transmitted on the logical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>IPv6 transit statistics</b>	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Local statistics</b>	Statistics for traffic received from and transmitted to the Routing Engine.	<b>detail extensive</b>
<b>Transit statistics</b>	Statistics for traffic transiting the router.	<b>detail extensive</b>
<b>Protocol</b>	Protocol family configured on the local interface. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .	<b>detail extensive</b>
<b>MTU</b>	Maximum transmission unit size on the logical interface.	<b>detail extensive</b>
<b>Maximum labels</b>	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	<b>detail extensive</b> none

Table 16: show interfaces irb Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route table</b>	Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	<b>detail extensive</b>
<b>Addresses, Flags</b>	Information about address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive</b>
<b>Policer</b>	The policer that is to be evaluated when packets are received or transmitted on the interface.	<b>detail extensive</b>
<b>Flags</b>	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive</b>

## Sample Output

### show interfaces irb extensive

```

user@host> show interfaces irb extensive
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 23, Generation: 130
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: Unspecified
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   : 0
    Output bytes  : 0
    Input packets : 0
    Output packets: 0
  IPv6 transit statistics:
    Input bytes   : 0
    Output bytes  : 0
    Input packets : 0
    Output packets: 0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
  0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
  0

  Logical interface irb.0 (Index 68) (SNMP ifIndex 70) (Generation 143)
    Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
    Bandwidth: 1000mbps
    Routing Instance: customer_0 Bridging Domain: bd0

```



```

Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Protocol inet, MTU: 1500, Generation: 154, Route table: 0
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255,
    Generation: 155
Protocol multiservice, MTU: 1500, Generation: 155, Route table: 0
  Flags: Is-Primary
  Policer: Input: __default_arp_policer

```

#### show interfaces irb snmp-index

```

user@host> show interfaces irb snmp-index 25
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 25
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514
  Device flags : Present Running
  Interface flags: SNMP-Traps
  Link type : Full-Duplex
  Link flags : None
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Last flapped : Never
    Input packets : 0
    Output packets: 0

Logical interface irb.0 (Index 68) (SNMP ifIndex 70)
  Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
  Bandwidth: 1000mbps
  Routing Instance: customer_0 Bridging Domain: bd0
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 1500
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255
  Protocol multiservice, MTU: 1500
    Flags: Is-Primary

```

## show interfaces queue

<b>Syntax</b>	<pre>show interfaces queue &lt;aggregate   remaining-traffic&gt; &lt;both-ingress-egress&gt; &lt;egress&gt; &lt;forwarding-class forwarding-class&gt; &lt;ingress&gt; &lt;interface-name interface-name&gt; &lt;l2-statistics&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p><b>both-ingress-egress</b>, <b>egress</b>, and <b>ingress</b> options introduced in Junos OS Release 7.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p><b>l2-statistics</b> option introduced in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	Display class-of-service (CoS) queue information for physical interfaces.
<b>Options</b>	<p><b>none</b>—Show detailed CoS queue statistics for all physical interfaces.</p> <p><b>aggregate</b>—(Optional) Display the aggregated queuing statistics of all logical interfaces that have traffic-control profiles configured. (Not on the QFX Series.)</p> <p><b>both-ingress-egress</b>—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics. (Not on the QFX Series.)</p> <p><b>egress</b>—(Optional) Display egress queue statistics.</p> <p><b>forwarding-class forwarding-class</b>—(Optional) Forwarding class name for this queue. Shows detailed CoS statistics for the queue associated with the specified forwarding class.</p> <p><b>ingress</b>—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics. (Not on the QFX Series.)</p> <p><b>interface-name interface-name</b>—(Optional) Show detailed CoS queue statistics for the specified interface.</p> <p><b>l2-statistics</b>—(Optional) Display Layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles</p> <p><b>remaining-traffic</b>—(Optional) Display the remaining-traffic queue statistics of all logical interfaces that have traffic-control profiles configured.</p>

### Overhead for Layer 2 Statistics

Transmitted packets and transmitted byte counts are displayed for the Layer 2 level with the addition of encapsulation overheads applied for fragmentation, as shown in [Table 17 on page 227](#). Others counters, such as packets and bytes queued (input) and drop counters, are displayed at the Layer 3 level. In the case of link fragmentation and interleaving (LFI) for which fragmentation is not applied, corresponding Layer 2 overheads are added, as shown in [Table 17 on page 227](#).

Table 17: Layer 2 Overhead and Transmitted Packets or Byte Counts

Protocol	Fragmentation		LFI
	First fragmentation	Second to <i>n</i> fragmentations	
	Bytes	Bytes	
MLPPP (Long)	13	12	8
MLPPP (short)	11	10	8
MLFR (FRF15)	12	10	8
MFR (FRF16)	10	8	-
MCMLPPP(Long)	13	12	-
MCMLPPP(Short)	11	10	-

#### Layer 2 Statistics—Fragmentation Overhead Calculation

##### MLPPP/MC-MLPPP Overhead details:

=====

##### Fragment 1:

```

Outer PPP header           : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
Inner PPP header           : 1 byte
HDLC flag and FCS bytes    : 4 bytes

```

##### Fragments 2 .. n :

```

Outer PPP header           : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
HDLC flag and FCS bytes    : 4 bytes

```

##### MLFR (FRF15) Overhead details:

=====

##### Fragment 1:

```

Framereley header         : 2 bytes
Control,NLPID             : 2 bytes
Fragmentaion header       : 2 bytes
Inner proto               : 2 bytes
HDLC flag and FCS         : 4 bytes

```

##### Fragments 2 ...n :

```

Framereley header         : 2 bytes
Control,NLPID             : 2 bytes
Fragmentaion header       : 2 bytes
HDLC flag and FCS         : 4 bytes

```

##### MFR (FRF16) Overhead details:

=====

```

Fragment 1:
  Fragmentation header : 2 bytes
  Framereelay header   : 2 bytes
  Inner proto          : 2 bytes
  HDLC flag and FCS    : 4 bytes

Fragments 2 ...n :
  Fragmentation header : 2 bytes
  Framereelay header   : 2 bytes
  HDLC flag and FCS    : 4 bytes

```

## Overhead with LFI

```

MLPPP(Long & short sequence):
=====
  Outer PPP header : 4 bytes
  HDLC flag and FCS : 4 bytes

MLFR (FRF15):
=====
  Framereelay header : 2 bytes
  Control,NLPID      : 2 bytes
  HDLC flag and FCS  : 4 bytes

```

The following examples show overhead for different cases:

- A 1000-byte packet is sent to a mlppp bundle without any fragmentation. At the Layer 2 level, bytes transmitted is 1013 in 1 packet. This overhead is for MLPPP long sequence encap.
- A 1000-byte packet is sent to a mlppp bundle with a fragment threshold of 250byte. At the Layer 2 level, bytes transmitted is 1061 bytes in 5 packets.
- A 1000-byte LFI packet is sent to an mlppp bundle. At the Layer 2 level, bytes transmitted is 1008 in 1 packet.

**remaining-traffic**—(Optional) Display the queuing statistics of all logical interfaces that do not have traffic-control profiles configured. (Not on the QFX Series.)

## Additional Information

For rate-limited interfaces hosted on Modular Interface Cards (MICs), Modular Port Concentrators (MPCs), or Enhanced Queuing DPCs, rate-limit packet-drop operations occur *before* packets are queued for transmission scheduling. For such interfaces, the statistics for queued traffic do not include the packets that have already been dropped due to rate limiting, and consequently the displayed statistics for queued traffic are the same as the displayed statistics for transmitted traffic.



**NOTE:** For rate-limited interfaces hosted on other types of hardware, rate-limit packet-drop operations occur *after* packets are queued for transmission scheduling. For these other interface types, the statistics for queued traffic include the packets that are later dropped due to rate limiting, and consequently the displayed statistics for queued traffic equals the sum of the statistics for transmitted and rate-limited traffic.

On M Series routers (except for the M320 and M120 routers), this command is valid only for a PIC installed on an enhanced Flexible PIC Concentrator (FPC).

Queue statistics for aggregated interfaces are supported on the M Series and T Series routers only. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can view the statistics for a child interface by using the **show interfaces statistics** command for that child interface.

When you configure tricolor marking on a 10-port 1-Gigabit Ethernet PIC, for queues 6 and 7 only, the output does not display the number of queued bytes and packets, or the number of bytes and packets dropped because of RED. If you do not configure tricolor marking on the interface, these statistics are available for all queues.

For the 4-port Channelized OC12 IQE PIC and 1-port Channelized OC48 IQE PIC, the **Packet Forwarding Engine Chassis Queues** field represents traffic bound for a particular physical interface on the PIC. For all other PICs, the **Packet Forwarding Engine Chassis Queues** field represents the total traffic bound for the PIC.

For Gigabit Ethernet IQ2 PICs, the **show interfaces queue** command output does not display the number of tail-dropped packets. This limitation does not apply to Packet Forwarding Engine chassis queues.

When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (under the **Packet Forwarding Engine Chassis Queues** field) shows the prefragmentation values.

The behavior of the **egress** queues for the **Routing Engine-Generated Traffic** is not same as the configured queue for MLPPP and MFR configurations.

For information about how to configure CoS, see the *Junos OS Network Interfaces Library for Routing Devices*. For related CoS operational mode commands, see the [CLI Explorer](#).

<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show interfaces queue (Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC) on page 234</a> <a href="#">show interfaces queue (Aggregated Ethernet on a T320 Router) on page 235</a> <a href="#">show interfaces queue (Gigabit Ethernet on a T640 Router) on page 237</a> <a href="#">show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC) on page 237</a> <a href="#">show interfaces queue (Gigabit Ethernet IQ2 PIC) on page 241</a> <a href="#">show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC) on page 244</a> <a href="#">show interfaces queue ingress (Gigabit Ethernet IQ2 PIC) on page 246</a> <a href="#">show interfaces queue egress (Gigabit Ethernet IQ2 PIC) on page 247</a> <a href="#">show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC) on page 249</a> <a href="#">show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode) on page 251</a> <a href="#">show interfaces queue (QFX Series) on page 261</a> <a href="#">show interfaces queue l2-statistics (lsq interface) on page 262</a> <a href="#">show interfaces queue lsq (lsq-ifd) on page 263</a>
<b>Output Fields</b>	Table 18 on page 230 lists the output fields for the <b>show interfaces queue</b> command. Output fields are listed in the approximate order in which they appear.

Table 18: show interfaces queue Output Fields

Field Name	Field Description
<b>Physical interface</b>	Name of the physical interface.
<b>Enabled</b>	State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> .
<b>Interface index</b>	Physical interface's index number, which reflects its initialization sequence.
<b>SNMP ifIndex</b>	SNMP index number for the interface.
<b>Forwarding classes supported</b>	Total number of forwarding classes supported on the specified interface.
<b>Forwarding classes in use</b>	Total number of forwarding classes in use on the specified interface.
<b>Ingress queues supported</b>	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues supported on the specified interface.
<b>Ingress queues in use</b>	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues in use on the specified interface.
<b>Output queues supported</b>	Total number of output queues supported on the specified interface.
<b>Output queues in use</b>	Total number of output queues in use on the specified interface.
<b>Egress queues supported</b>	Total number of egress queues supported on the specified interface.
<b>Egress queues in use</b>	Total number of egress queues in use on the specified interface.
<b>Queue counters (Ingress)</b>	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>
<b>Burst size</b>	(Logical interfaces on IQ PICs only) Maximum number of bytes up to which the logical interface can burst. The burst size is based on the shaping rate applied to the interface.
The following output fields are applicable to both interface component and Packet Forwarding component in the <b>show interfaces queue</b> command:	
<b>Queue</b>	Queue number.
<b>Forwarding classes</b>	Forwarding class name.

Table 18: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
<b>Queued Packets</b>	<p>Number of packets queued to this queue.</p> <p><b>NOTE:</b> For Gigabit Ethernet IQ2 interfaces, the Queued Packets count is calculated by the Junos OS interpreting one frame buffer as one packet. If the queued packets are very large or very small, the calculation might not be completely accurate for transit traffic. The count is completely accurate for traffic terminated on the router.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see <a href="#">“Additional Information” on page 228</a>.</p>
<b>Queued Bytes</b>	<p>Number of bytes queued to this queue. The byte counts vary by interface hardware. For more information, see <a href="#">Table 19 on page 233</a>.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see <a href="#">“Additional Information” on page 228</a>.</p>
<b>Transmitted Packets</b>	<p>Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the <b>Packet Forwarding Engine Chassis Queues</b> field) shows the prefragmentation values.</p> <p><b>NOTE:</b> For Layer 2 statistics, see <a href="#">“Overhead for Layer 2 Statistics” on page 226</a></p>
<b>Transmitted Bytes</b>	<p>Number of bytes transmitted by this queue. The byte counts vary by interface hardware. For more information, see <a href="#">Table 19 on page 233</a>.</p> <p><b>NOTE:</b> On MX Series routers, this number can be inaccurate when you issue the command for a physical interface repeatedly and in quick succession, because the statistics for the child nodes are collected infrequently. Wait ten seconds between successive iterations to avoid this situation.</p> <p><b>NOTE:</b> For Layer 2 statistics, see <a href="#">“Overhead for Layer 2 Statistics” on page 226</a></p>
<b>Tail-dropped packets</b>	<p>Number of packets dropped because of tail drop.</p> <p><b>NOTE:</b> The <b>Tail-dropped packets</b> counter is not supported on the PTX Series Packet Transport Routers.</p>
<b>RL-dropped packets</b>	<p>Number of packets dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see <a href="#">“Additional Information” on page 228</a>.</p> <p><b>NOTE:</b> The <b>RL-dropped packets</b> counter is not supported on the PTX Series Packet Transport Routers, and is omitted from the output.</p>
<b>RL-dropped bytes</b>	<p>Number of bytes dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see <a href="#">“Additional Information” on page 228</a>.</p>

Table 18: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
RED-dropped packets	<p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP packets dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP packets dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP packets dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP packets dropped because of RED.</li> </ul> </li> <li>(MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li><b>Low</b>—Number of low-loss priority packets dropped because of RED.</li> <li><b>Medium-low</b>—Number of medium-low loss priority packets dropped because of RED.</li> <li><b>Medium-high</b>—Number of medium-high loss priority packets dropped because of RED.</li> <li><b>High</b>—Number of high-loss priority packets dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>
RED-dropped bytes	<p>Number of bytes dropped because of RED. The byte counts vary by interface hardware. For more information, see <a href="#">Table 19 on page 233</a>.</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP bytes dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP bytes dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP bytes dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP bytes dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>

Byte counts vary by interface hardware. [Table 19 on page 233](#) shows how the byte counts on the outbound interfaces vary depending on the interface hardware. [Table 19 on page 233](#) is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.



Table 19: Byte Count by Interface Hardware

Interface Hardware	Output Level	Byte Count Includes	Comments
Gigabit Ethernet IQ and IQE PICs	Interface	<p>Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes</p>	<p>The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type.</p> <p>For RED dropped, 6 bytes are added for the source MAC address.</p>
	Packet forwarding component	<p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p>	—
Non-IQ PIC	Interface	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 478 bytes of Layer 3 packet.</li> </ul> <p>T4000 routers with Type 5 FPCs :</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Inter frame Gap.</li> <li>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Interframe Gap.</li> </ul> <p>M Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead.</li> </ul> <p>PTX Series Packet Transport Routers:</p> <ul style="list-style-type: none"> <li>Queued: The sum of the transmitted bytes and the RED dropped bytes.</li> <li>Transmitted: Full Layer 2 overhead (including all L2 encapsulation and CRC) + 12 inter-packet gap + 8 for the preamble.</li> <li>RED dropped: Full Layer 2 overhead (including all L2 encapsulation and CRC) + 12 inter-packet gap + 8 for the preamble (does not include the VLAN header or MPLS pushed bytes).</li> </ul>	<p>The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.</p>

Table 19: Byte Count by Interface Hardware (*continued*)

Interface Hardware	Output Level	Byte Count Includes	Comments
IQ and IQE PICs with a SONET/SDH interface	Interface	Queued: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes  Transmitted: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes  RED dropped: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes	The additional 4 bytes are for the Layer 2 Point-to-Point Protocol (PPP) header.
	Packet forwarding component	Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet  Transmitted: 486 bytes per packet, representing 478 bytes of Layer 3 packet + 8 bytes	For transmitted packets, the additional 8 bytes includes 4 bytes for the PPP header and 4 bytes for a cookie.
Non-IQ PIC with a SONET/SDH interface	Interface	T Series, TX Series, T1600, and MX Series routers: <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 478 bytes of Layer 3 packet.</li> </ul> M Series routers: <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 483 bytes per packet, representing 478 bytes of Layer 3 packet + 5 bytes</li> <li>RED dropped: 478 bytes per packet, representing 478 bytes of Layer 3 packet</li> </ul>	For transmitted packets, the additional 5 bytes includes 4 bytes for the PPP header and 1 byte for the packet loss priority (PLP).
Interfaces configured with Frame Relay Encapsulation	Interface	The default Frame Relay overhead is 7 bytes. If you configure the Frame Check Sequence (FCS) to 4 bytes, then the overhead increases to 10 bytes.	
1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs	Interface	Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.  Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.	The Layer 2 overhead is 18 bytes for non-VLAN traffic and 22 bytes for VLAN traffic.
4-port 1G IQ2 and IQ2-E PICs	Packet forwarding component	Queued: 478 bytes of Layer 3 packet.	—
8-port 1G IQ2 and IQ2-E PICs		Transmitted: 478 bytes of Layer 3 packet.	

## Sample Output

### show interfaces queue (Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC)

The following example shows queue information for the rate-limited interface ge-4/2/0 on a Gigabit Ethernet MIC in an MPC. For rate-limited queues for interfaces hosted on MICs or MPCs, rate-limit packet drops occur prior to packet output queuing. In the

command output, the nonzero statistics displayed in the **RL-dropped packets** and **RL-dropped bytes** fields quantify the traffic dropped to rate-limit queue 0 output to 10 percent of 1 gigabyte (100 megabits) per second. Because the RL-dropped traffic is not included in the **Queued** statistics, the statistics displayed for queued traffic are the same as the statistics for transmitted traffic.

```
user@host> show interfaces queue ge-4/2/0
Physical interface: ge-4/2/0, Enabled, Physical link is Up
  Interface index: 203, SNMP ifIndex: 1054
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets          :          131300649          141751 pps
    Bytes            :          11287964840        99793248 bps
  Transmitted:
    Packets          :          131300649          141751 pps
    Bytes            :          11287964840        99793248 bps
    Tail-dropped packets :              0              0 pps
    RL-dropped packets  :          205050862        602295 pps
    RL-dropped bytes    :          13595326612      327648832 bps
    RED-dropped packets :              0              0 pps
      Low              :              0              0 pps
      Medium-low       :              0              0 pps
      Medium-high      :              0              0 pps
      High              :              0              0 pps
    RED-dropped bytes   :              0              0 bps
      Low              :              0              0 bps
      Medium-low       :              0              0 bps
      Medium-high      :              0              0 bps
      High              :              0              0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets          :              0              0 pps
    Bytes            :              0              0 bps
```

### show interfaces queue (Aggregated Ethernet on a T320 Router)

The following example shows that the aggregated Ethernet interface, **ae1**, has traffic on queues **af1** and **af12**:

```
user@host> show interfaces queue ae1
Physical interface: ae1, Enabled, Physical link is Up
  Interface index: 158, SNMP ifIndex: 33 Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets          :              5              0 pps
    Bytes            :              242              0 bps
  Transmitted:
    Packets          :              5              0 pps
    Bytes            :              242              0 bps
    Tail-dropped packets :              0              0 pps
    RED-dropped packets :              0              0 pps
    RED-dropped bytes   :              0              0 bps
Queue: 1, Forwarding classes: af1
  Queued:
    Packets          :          42603765          595484 pps
```

```

Bytes                :          5453281920          609776496 bps
Transmitted:
Packets              :          42603765           595484 pps
Bytes                :          5453281920          609776496 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 2, Forwarding classes: ef1
Queued:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Transmitted:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 3, Forwarding classes: nc
Queued:
Packets              :              45              0 pps
Bytes                :             3930              0 bps
Transmitted:
Packets              :              45              0 pps
Bytes                :             3930              0 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 4, Forwarding classes: af11
Queued:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Transmitted:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 5, Forwarding classes: ef11
Queued:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Transmitted:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 6, Forwarding classes: af12
Queued:
Packets              :          31296413          437436 pps
Bytes                :          4005940864          447935200 bps
Transmitted:
Packets              :          31296413          437436 pps
Bytes                :          4005940864          447935200 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 7, Forwarding classes: nc2
Queued:
Packets              :              0              0 pps
Bytes                :              0              0 bps

```

```

Transmitted:
Packets      :                0                0 pps
Bytes        :                0                0 bps
Tail-dropped packets :                0                0 pps
RED-dropped packets :                0                0 pps
RED-dropped bytes  :                0                0 bps

```

### show interfaces queue (Gigabit Ethernet on a T640 Router)

```

user@host> show interfaces queue
Physical interface: ge-7/0/1, Enabled, Physical link is Up
  Interface index: 150, SNMP ifIndex: 42
  Forwarding classes: 8 supported, 8 in use
  Output queues: 8 supported, 8 in use
  Queue: 0, Forwarding classes: be
    Queued:
      Packets      :                13                0 pps
      Bytes        :               622                0 bps
    Transmitted:
      Packets      :                13                0 pps
      Bytes        :               622                0 bps
      Tail-dropped packets :                0                0 pps
      RED-dropped packets :                0                0 pps
      RED-dropped bytes  :                0                0 bps
  Queue: 1, Forwarding classes: af1
    Queued:
      Packets      :            1725947945            372178 pps
      Bytes        :          220921336960          381110432 bps
    Transmitted:
      Packets      :            1725947945            372178 pps
      Bytes        :          220921336960          381110432 bps
      Tail-dropped packets :                0                0 pps
      RED-dropped packets :                0                0 pps
      RED-dropped bytes  :                0                0 bps
  Queue: 2, Forwarding classes: ef1
    Queued:
      Packets      :                0                0 pps
      Bytes        :                0                0 bps
    Transmitted:
      Packets      :                0                0 pps
      Bytes        :                0                0 bps
      Tail-dropped packets :                0                0 pps
      RED-dropped packets :                0                0 pps
      RED-dropped bytes  :                0                0 bps
  Queue: 3, Forwarding classes: nc
    Queued:
      Packets      :                571                0 pps
      Bytes        :             49318             336 bps
    Transmitted:
      Packets      :                571                0 pps
      Bytes        :             49318             336 bps
      Tail-dropped packets :                0                0 pps
      RED-dropped packets :                0                0 pps
      RED-dropped bytes  :                0                0 bps

```

### show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 aggregate

```

```

Physical interface: ge-2/2/9, Enabled, Physical link is Up
  Interface index: 238, SNMP ifIndex: 71
  Forwarding classes: 16 supported, 4 in use
  Ingress queues: 4 supported, 4 in use
  Queue: 0, Forwarding classes: best-effort
    Queued:
      Packets      :      148450735      947295 pps
      Bytes        :      8016344944    409228848 bps
    Transmitted:
      Packets      :      76397439      487512 pps
      Bytes        :    4125461868    210602376 bps
      Tail-dropped packets : Not Available
      RED-dropped packets :      72053285      459783 pps
        Low        :      72053285      459783 pps
        Medium-low  :           0          0 pps
        Medium-high :           0          0 pps
        High        :           0          0 pps
      RED-dropped bytes  :    3890877444    198626472 bps
        Low        :    3890877444    198626472 bps
        Medium-low  :           0          0 bps
        Medium-high :           0          0 bps
        High        :           0          0 bps
  Queue: 1, Forwarding classes: expedited-forwarding
    Queued:
      Packets      :           0          0 pps
      Bytes        :           0          0 bps
    Transmitted:
      Packets      :           0          0 pps
      Bytes        :           0          0 bps
      Tail-dropped packets : Not Available
      RED-dropped packets :           0          0 pps
        Low        :           0          0 pps
        Medium-low  :           0          0 pps
        Medium-high :           0          0 pps
        High        :           0          0 pps
      RED-dropped bytes  :           0          0 bps
        Low        :           0          0 bps
        Medium-low  :           0          0 bps
        Medium-high :           0          0 bps
        High        :           0          0 bps
  Queue: 2, Forwarding classes: assured-forwarding
    Queued:
      Packets      :      410278257      473940 pps
      Bytes        :    22156199518    204742296 bps
    Transmitted:
      Packets      :      4850003      4033 pps
      Bytes        :    261900162    1742256 bps
      Tail-dropped packets : Not Available
      RED-dropped packets :      405425693      469907 pps
        Low        :      405425693      469907 pps
        Medium-low  :           0          0 pps
        Medium-high :           0          0 pps
        High        :           0          0 pps
      RED-dropped bytes  :    21892988124    203000040 bps
        Low        :    21892988124    203000040 bps
        Medium-low  :           0          0 bps
        Medium-high :           0          0 bps
        High        :           0          0 bps
  Queue: 3, Forwarding classes: network-control
    Queued:
      Packets      :           0          0 pps

```

```

Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : 76605230 485376 pps
Bytes : 5209211400 264044560 bps
Transmitted:
Packets : 76444631 484336 pps
Bytes : 5198235612 263478800 bps
Tail-dropped packets : Not Available
RED-dropped packets : 160475 1040 pps
Low : 160475 1040 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 10912300 565760 bps
Low : 10912300 565760 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 4836136 3912 pps
Bytes : 333402032 2139056 bps
Transmitted:
Packets : 3600866 1459 pps
Bytes : 244858888 793696 bps
Tail-dropped packets : Not Available

```

```

RED-dropped packets :          1225034          2450 pps
  Low                :          1225034          2450 pps
  Medium-low         :              0              0 pps
  Medium-high        :              0              0 pps
  High               :              0              0 pps
RED-dropped bytes   :          83302312        1333072 bps
  Low                :          83302312        1333072 bps
  Medium-low         :              0              0 bps
  Medium-high        :              0              0 bps
  High               :              0              0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets            :              0              0 pps
  Bytes              :              0              0 bps
Transmitted:
  Packets            :              0              0 pps
  Bytes              :              0              0 bps
Tail-dropped packets : Not Available
RED-dropped packets :              0              0 pps
  Low                :              0              0 pps
  Medium-low         :              0              0 pps
  Medium-high        :              0              0 pps
  High               :              0              0 pps
RED-dropped bytes   :              0              0 bps
  Low                :              0              0 bps
  Medium-low         :              0              0 bps
  Medium-high        :              0              0 bps
  High               :              0              0 bps

```

Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

```

Queued:
  Packets            :          77059796        486384 pps
  Bytes              :          3544750624      178989576 bps
Transmitted:
  Packets            :          77059797        486381 pps
  Bytes              :          3544750670      178988248 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets :              0              0 pps
  Low                :              0              0 pps
  Medium-low         :              0              0 pps
  Medium-high        :              0              0 pps
  High               :              0              0 pps
RED-dropped bytes   :              0              0 bps
  Low                :              0              0 bps
  Medium-low         :              0              0 bps
  Medium-high        :              0              0 bps
  High               :              0              0 bps

```

Queue: 1, Forwarding classes: expedited-forwarding

```

Queued:
  Packets            :              0              0 pps
  Bytes              :              0              0 bps
Transmitted:
  Packets            :              0              0 pps
  Bytes              :              0              0 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets :              0              0 pps
  Low                :              0              0 pps
  Medium-low         :              0              0 pps
  Medium-high        :              0              0 pps

```



```

      High : 0 0 pps
    RED-dropped bytes : 0 0 bps
      Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
      High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets : 4846580 3934 pps
    Bytes : 222942680 1447768 bps
  Transmitted:
    Packets : 4846580 3934 pps
    Bytes : 222942680 1447768 bps
    Tail-dropped packets : 0 0 pps
    RED-dropped packets : 0 0 pps
      Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
      High : 0 0 pps
    RED-dropped bytes : 0 0 bps
      Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
      High : 0 0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : 0 0 pps
    RED-dropped packets : 0 0 pps
      Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
      High : 0 0 pps
    RED-dropped bytes : 0 0 bps
      Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
      High : 0 0 bps

```

#### show interfaces queue (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-7/1/3
Physical interface: ge-7/1/3, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 70 Forwarding classes: 16 supported, 4 in
  use Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets : 418390039 10 pps
    Bytes : 38910269752 7440 bps
  Transmitted:
    Packets : 418390039 10 pps
    Bytes : 38910269752 7440 bps
    Tail-dropped packets : Not Available
    RED-dropped packets : 0 0 pps
    RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding

```

```

Queued:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
Transmitted:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :          0          0 pps
  RED-dropped bytes  :          0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
Transmitted:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :          0          0 pps
  RED-dropped bytes  :          0          0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets      :        7055          1 pps
  Bytes       :    451552        512 bps
Transmitted:
  Packets      :        7055          1 pps
  Bytes       :    451552        512 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :          0          0 pps
  RED-dropped bytes  :          0          0 bps
Forwarding classes: 16 supported, 4 in use Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets      :        1031          0 pps
  Bytes       :    143292          0 bps
Transmitted:
  Packets      :        1031          0 pps
  Bytes       :    143292          0 bps
  Tail-dropped packets : Not Available
  RL-dropped packets  :          0          0 pps
  RL-dropped bytes   :          0          0 bps
  RED-dropped packets :          0          0 pps
  RED-dropped bytes  :          0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
Transmitted:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
  Tail-dropped packets : Not Available
  RL-dropped packets  :          0          0 pps
  RL-dropped bytes   :          0          0 bps
  RED-dropped packets :          0          0 pps
  RED-dropped bytes  :          0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
Transmitted:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps

```

```

Tail-dropped packets : Not Available
RL-dropped packets   :                0                0 pps
RL-dropped bytes     :                0                0 bps
RED-dropped packets   :                0                0 pps
RED-dropped bytes     :                0                0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets               :                77009            11 pps
Bytes                 :            6894286            7888 bps
Transmitted:
Packets               :                77009            11 pps
Bytes                 :            6894286            7888 bps
Tail-dropped packets : Not Available
RL-dropped packets   :                0                0 pps
RL-dropped bytes     :                0                0 bps
RED-dropped packets   :                0                0 pps
RED-dropped bytes     :                0                0 bps

```

#### Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

```

Queued:
Packets               :                1031            0 pps
Bytes                 :            147328            0 bps
Transmitted:
Packets               :                1031            0 pps
Bytes                 :            147328            0 bps
Tail-dropped packets :                0                0 pps
RED-dropped packets   :                0                0 pps
Low, non-TCP          :                0                0 pps
Low, TCP              :                0                0 pps
High, non-TCP         :                0                0 pps
High, TCP             :                0                0 pps
RED-dropped bytes     :                0                0 bps
Low, non-TCP          :                0                0 bps
Low, TCP              :                0                0 bps
High, non-TCP         :                0                0 bps
High, TCP             :                0                0 bps

```

Queue: 1, Forwarding classes: expedited-forwarding

```

Queued:
Packets               :                0                0 pps
Bytes                 :                0                0 bps
Transmitted:
Packets               :                0                0 pps
Bytes                 :                0                0 bps
Tail-dropped packets :                0                0 pps
RED-dropped packets   :                0                0 pps
Low, non-TCP          :                0                0 pps
Low, TCP              :                0                0 pps
High, non-TCP         :                0                0 pps
High, TCP             :                0                0 pps
RED-dropped bytes     :                0                0 bps
Low, non-TCP          :                0                0 bps
Low, TCP              :                0                0 bps
High, non-TCP         :                0                0 bps
High, TCP             :                0                0 bps

```

Queue: 2, Forwarding classes: assured-forwarding

```

Queued:
Packets               :                0                0 pps
Bytes                 :                0                0 bps
Transmitted:

```

```

Packets          : 0 0 pps
Bytes            : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
  Low, non-TCP    : 0 0 pps
  Low, TCP        : 0 0 pps
  High, non-TCP   : 0 0 pps
  High, TCP       : 0 0 pps
RED-dropped bytes : 0 0 bps
  Low, non-TCP    : 0 0 bps
  Low, TCP        : 0 0 bps
  High, non-TCP   : 0 0 bps
  High, TCP       : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets          : 94386 12 pps
  Bytes            : 13756799 9568 bps
Transmitted:
  Packets          : 94386 12 pps
  Bytes            : 13756799 9568 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
  Low, non-TCP    : 0 0 pps
  Low, TCP        : 0 0 pps
  High, non-TCP   : 0 0 pps
  High, TCP       : 0 0 pps
RED-dropped bytes : 0 0 bps
  Low, non-TCP    : 0 0 bps
  Low, TCP        : 0 0 bps
  High, non-TCP   : 0 0 bps
  High, TCP       : 0 0 bps

```

#### show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 both-ingress-egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets          : Not Available
  Bytes            : 0 0 bps
Transmitted:
  Packets          : 254 0 pps
  Bytes            : 16274 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets          : Not Available
  Bytes            : 0 0 bps
Transmitted:
  Packets          : 0 0 pps
  Bytes            : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding

```

```

Queued:
  Packets          : Not Available
  Bytes           :                0                0 bps
Transmitted:
  Packets          :                0                0 pps
  Bytes           :                0                0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                0                0 pps
RED-dropped bytes  :                0                0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets          : Not Available
  Bytes           :                0                0 bps
Transmitted:
  Packets          :                0                0 pps
  Bytes           :                0                0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                0                0 pps
RED-dropped bytes  :                0                0 bps
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets          : Not Available
  Bytes           :                0                0 bps
Transmitted:
  Packets          :                3                0 pps
  Bytes           :               126                0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                0                0 pps
RED-dropped bytes  :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets          : Not Available
  Bytes           :                0                0 bps
Transmitted:
  Packets          :                0                0 pps
  Bytes           :                0                0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                0                0 pps
RED-dropped bytes  :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets          : Not Available
  Bytes           :                0                0 bps
Transmitted:
  Packets          :                0                0 pps
  Bytes           :                0                0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                0                0 pps
RED-dropped bytes  :                0                0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets          : Not Available
  Bytes           :                0                0 bps
Transmitted:
  Packets          :                0                0 pps
  Bytes           :                0                0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                0                0 pps
RED-dropped bytes  :                0                0 bps

```

```

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :      80564692      0 pps
    Bytes        :      3383717100    0 bps
  Transmitted:
    Packets      :      80564692      0 pps
    Bytes        :      3383717100    0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :      80564685      0 pps
    Bytes        :      3383716770    0 bps
  Transmitted:
    Packets      :      80564685      0 pps
    Bytes        :      3383716770    0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      :      0      0 pps
    Bytes        :      0      0 bps
  Transmitted:
    Packets      :      0      0 pps
    Bytes        :      0      0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      :      9397      0 pps
    Bytes        :      3809052      232 bps
  Transmitted:
    Packets      :      9397      0 pps
    Bytes        :      3809052      232 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps

```

### show interfaces queue ingress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 ingress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      : Not Available
    Bytes        :      0      0 bps
  Transmitted:
    Packets      :      288      0 pps
    Bytes        :      18450      0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps

```

```

Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      : Not Available
    Bytes        :                0                0 bps
  Transmitted:
    Packets      :                0                0 pps
    Bytes        :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      : Not Available
    Bytes        :                0                0 bps
  Transmitted:
    Packets      :                0                0 pps
    Bytes        :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      : Not Available
    Bytes        :                0                0 bps
  Transmitted:
    Packets      :                0                0 pps
    Bytes        :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps

```

#### show interfaces queue egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      : Not Available
    Bytes        :                0                0 bps
  Transmitted:
    Packets      :                3                0 pps
    Bytes        :               126                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      : Not Available
    Bytes        :                0                0 bps
  Transmitted:
    Packets      :                0                0 pps
    Bytes        :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:

```

```

Packets          : Not Available
Bytes            :                      0          0 bps
Transmitted:
Packets          :                      0          0 pps
Bytes            :                      0          0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets          : Not Available
Bytes            :                      0          0 bps
Transmitted:
Packets          :                      0          0 pps
Bytes            :                      0          0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets          :                      80564692      0 pps
Bytes            :                      3383717100     0 bps
Transmitted:
Packets          :                      80564692      0 pps
Bytes            :                      3383717100     0 bps
Tail-dropped packets :                      0          0 pps
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets          :                      80564685      0 pps
Bytes            :                      3383716770     0 bps
Transmitted:
Packets          :                      80564685      0 pps
Bytes            :                      3383716770     0 bps
Tail-dropped packets :                      0          0 pps
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets          :                      0          0 pps
Bytes            :                      0          0 bps
Transmitted:
Packets          :                      0          0 pps
Bytes            :                      0          0 bps
Tail-dropped packets :                      0          0 pps
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets          :                      9538          0 pps
Bytes            :                      3819840         0 bps
Transmitted:
Packets          :                      9538          0 pps
Bytes            :                      3819840         0 bps
Tail-dropped packets :                      0          0 pps
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps

```



## show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 remaining-traffic
Physical interface: ge-2/2/9, Enabled, Physical link is Up
  Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :          110208969          472875 pps
    Bytes        :          5951284434        204282000 bps
  Transmitted:
    Packets      :          110208969          472875 pps
    Bytes        :          5951284434        204282000 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :
    Low          :              0              0 pps
    Medium-low   :              0              0 pps
    Medium-high  :              0              0 pps
    High         :              0              0 pps
  RED-dropped bytes :
    Low          :              0              0 bps
    Medium-low   :              0              0 bps
    Medium-high  :              0              0 bps
    High         :              0              0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps
  Transmitted:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :
    Low          :              0              0 pps
    Medium-low   :              0              0 pps
    Medium-high  :              0              0 pps
    High         :              0              0 pps
  RED-dropped bytes :
    Low          :              0              0 bps
    Medium-low   :              0              0 bps
    Medium-high  :              0              0 bps
    High         :              0              0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps
  Transmitted:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :
    Low          :              0              0 pps
    Medium-low   :              0              0 pps
    Medium-high  :              0              0 pps
    High         :              0              0 pps
  RED-dropped bytes :
    Low          :              0              0 bps
    Medium-low   :              0              0 bps
    Medium-high  :              0              0 bps

```

```

      High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets : 0 0 pps
  Bytes : 0 0 bps
Transmitted:
  Packets : 0 0 pps
  Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
  Low : 0 0 pps
  Medium-low : 0 0 pps
  Medium-high : 0 0 pps
  High : 0 0 pps
RED-dropped bytes : 0 0 bps
  Low : 0 0 bps
  Medium-low : 0 0 bps
  Medium-high : 0 0 bps
  High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets : 109355853 471736 pps
  Bytes : 7436199152 256627968 bps
Transmitted:
  Packets : 109355852 471736 pps
  Bytes : 7436198640 256627968 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
  Low : 0 0 pps
  Medium-low : 0 0 pps
  Medium-high : 0 0 pps
  High : 0 0 pps
RED-dropped bytes : 0 0 bps
  Low : 0 0 bps
  Medium-low : 0 0 bps
  Medium-high : 0 0 bps
  High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets : 0 0 pps
  Bytes : 0 0 bps
Transmitted:
  Packets : 0 0 pps
  Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
  Low : 0 0 pps
  Medium-low : 0 0 pps
  Medium-high : 0 0 pps
  High : 0 0 pps
RED-dropped bytes : 0 0 bps
  Low : 0 0 bps
  Medium-low : 0 0 bps
  Medium-high : 0 0 bps
  High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets : 0 0 pps
  Bytes : 0 0 bps

```

```

Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                0                0 pps
  Low            :                0                0 pps
  Medium-low     :                0                0 pps
  Medium-high    :                0                0 pps
  High           :                0                0 pps
RED-dropped bytes  :                0                0 bps
  Low            :                0                0 bps
  Medium-low     :                0                0 bps
  Medium-high    :                0                0 bps
  High           :                0                0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                0                0 pps
  Low            :                0                0 pps
  Medium-low     :                0                0 pps
  Medium-high    :                0                0 pps
  High           :                0                0 pps
RED-dropped bytes  :                0                0 bps
  Low            :                0                0 bps
  Medium-low     :                0                0 bps
  Medium-high    :                0                0 bps
  High           :                0                0 bps

```

#### show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode)

```

user@host> show interfaces queue t3-1/1/0:7
Physical interface: t3-1/1/0:7, Enabled, Physical link is Up

  Interface index: 192, SNMP ifIndex: 1948

  Description: full T3 interface connect to 6ce13 t3-3/1/0:7 for FR testing -
Lam

  Forwarding classes: 16 supported, 9 in use

  Egress queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

Queued:

  Packets          :                214886            13449 pps

  Bytes            :                9884756           5164536 bps

Transmitted:

  Packets          :                214886            13449 pps

  Bytes            :                9884756           5164536 bps

```

Tail-dropped packets :	0	0 pps
RED-dropped packets :	0	0 pps
Low :	0	0 pps
Medium-low :	0	0 pps
Medium-high :	0	0 pps
High :	0	0 pps
RED-dropped bytes :	0	0 bps
Low :	0	0 bps
Medium-low :	0	0 bps
Medium-high :	0	0 bps
High :	0	0 bps

Queue: 1, Forwarding classes: REALTIME

Queued:

Packets :	0	0 pps
Bytes :	0	0 bps

Transmitted:

Packets :	0	0 pps
Bytes :	0	0 bps
Tail-dropped packets :	0	0 pps
RED-dropped packets :	0	0 pps
Low :	0	0 pps
Medium-low :	0	0 pps
Medium-high :	0	0 pps
High :	0	0 pps
RED-dropped bytes :	0	0 bps
Low :	0	0 bps
Medium-low :	0	0 bps
Medium-high :	0	0 bps
High :	0	0 bps

Queue: 2, Forwarding classes: PRIVATE

## Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

## Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

## Queue: 3, Forwarding classes: CONTROL

## Queued:

Packets	:	60	0 pps
Bytes	:	4560	0 bps

## Transmitted:

Packets	:	60	0 pps
Bytes	:	4560	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps

RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 6, Forwarding classes: CLASS\_V\_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 7, Forwarding classes: CLASS\_S\_OUTPUT, GETS

Queued:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Packet Forwarding Engine Chassis Queues:

Queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

Queued:

Packets	:	371365	23620 pps
Bytes	:	15597330	7936368 bps

Transmitted:

Packets	:	371365	23620 pps
Bytes	:	15597330	7936368 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps



High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 1, Forwarding classes: REALTIME

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 2, Forwarding classes: PRIVATE

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Tail-dropped packets :	0	0 pps
RED-dropped packets :	0	0 pps
Low :	0	0 pps
Medium-low :	0	0 pps
Medium-high :	0	0 pps
High :	0	0 pps
RED-dropped bytes :	0	0 bps
Low :	0	0 bps
Medium-low :	0	0 bps
Medium-high :	0	0 bps
High :	0	0 bps

Queue: 3, Forwarding classes: CONTROL

Queued:

Packets :	32843	0 pps
Bytes :	2641754	56 bps

Transmitted:

Packets :	32843	0 pps
Bytes :	2641754	56 bps
Tail-dropped packets :	0	0 pps
RED-dropped packets :	0	0 pps
Low :	0	0 pps
Medium-low :	0	0 pps
Medium-high :	0	0 pps
High :	0	0 pps
RED-dropped bytes :	0	0 bps
Low :	0	0 bps
Medium-low :	0	0 bps
Medium-high :	0	0 bps
High :	0	0 bps

Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

## Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

## Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

## Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

## Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

## Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 6, Forwarding classes: CLASS\_V\_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 7, Forwarding classes: CLASS\_S\_OUTPUT, GETS

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps

RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

### show interfaces queue (QFX Series)

```

user@switch> show interfaces queue xe-0/0/15
Physical interface: xe-0/0/15, Enabled, Physical link is Up
Interface index: 49165, SNMP ifIndex: 539
Forwarding classes: 12 supported, 8 in use
Egress queues: 12 supported, 8 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
Queue: 3, Forwarding classes: fcoe
  Queued:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
0 bps
Queue: 4, Forwarding classes: no-loss
  Queued:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
Queue: 7, Forwarding classes: network-control

```

```

Queued:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
Transmitted:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
  Tail-dropped packets : Not Available
  Total-dropped packets:          0          0 pps
  Total-dropped bytes  :          0          0 bps
Queue: 8, Forwarding classes: mcast
Queued:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
Transmitted:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
  Tail-dropped packets : Not Available
  Total-dropped packets:          0          0 pps
  Total-dropped bytes  :          0          0 bps

```

#### show interfaces queue l2-statistics (lsq interface)

```

user@switch> show interfaces queue lsq-2/2/0.2 l2-statistics
Logical interface lsq-2/2/0.2 (Index 69) (SNMP ifIndex 1598)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
Queued:
  Packets      :          1          0 pps
  Bytes       :        1001          0 bps
Transmitted:
  Packets      :          5          0 pps
  Bytes       :        1062          0 bps
  Tail-dropped packets :          0          0 pps
  RED-dropped packets :          0          0 pps
  RED-dropped bytes  :          0          0 bps
Queue: 1, Forwarding classes: ef
Queued:
  Packets      :          1          0 pps
  Bytes       :        1500          0 bps
Transmitted:
  Packets      :          6          0 pps
  Bytes       :       1573          0 bps
  Tail-dropped packets :          0          0 pps
  RED-dropped packets :          0          0 pps
  RED-dropped bytes  :          0          0 bps
Queue: 2, Forwarding classes: af
Queued:
  Packets      :          1          0 pps
  Bytes       :         512          0 bps
Transmitted:
  Packets      :          3          0 pps
  Bytes       :         549          0 bps
  Tail-dropped packets :          0          0 pps
  RED-dropped packets :          0          0 pps
  RED-dropped bytes  :          0          0 bps
Queue: 3, Forwarding classes: nc
Queued:
  Packets      :          0          0 pps

```

```

Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
=====

```

### show interfaces queue lsq (lsq-ifd)

```

user@switch> show interfaces queue lsq-1/0/0
Logical interface lsq-1/0/0 (Index 348) (SNMP ifIndex 660)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
  Queued:
    Packets : 55576 1206 pps
    Bytes : 29622008 5145472 bps
  Transmitted:
    Packets : 55576 1206 pps
    Bytes : 29622008 5145472 bps
    Tail-dropped packets : 0 0 pps
    RL-dropped packets : 0 0 pps
    RL-dropped bytes : 0 0 bps
    RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
    RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps
Queue: 1, Forwarding classes: ef
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : 0 0 pps
    RL-dropped packets : 0 0 pps
    RL-dropped bytes : 0 0 bps
    RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
    RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps
Queue: 2, Forwarding classes: af
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps

```

```
Transmitted:
Packets          :          0          0 pps
Bytes            :          0          0 bps
Tail-dropped packets :          0          0 pps
RL-dropped packets :          0          0 pps
RL-dropped bytes  :          0          0 bps
RED-dropped packets :          0          0 pps
  Low            :          0          0 pps
  Medium-low     :          0          0 pps
  Medium-high    :          0          0 pps
  High           :          0          0 pps
RED-dropped bytes :          0          0 bps
  Low            :          0          0 bps
  Medium-low     :          0          0 bps
  Medium-high    :          0          0 bps
  High           :          0          0 bps
Queue: 3, Forwarding classes: nc
Queued:
Packets          :        22231        482 pps
Bytes            :       11849123     2057600 bps
Transmitted:
Packets          :        22231        482 pps
Bytes            :       11849123     2057600 bps
Tail-dropped packets :          0          0 pps
RL-dropped packets :          0          0 pps
RL-dropped bytes  :          0          0 bps
RED-dropped packets :          0          0 pps
  Low            :          0          0 pps
  Medium-low     :          0          0 pps
  Medium-high    :          0          0 pps
  High           :          0          0 pps
RED-dropped bytes :          0          0 bps
  Low            :          0          0 bps
  Medium-low     :          0          0 bps
  Medium-high    :          0          0 bps
  High           :          0          0 bps
```



## show mvrp

<b>Syntax</b>	show mvrp
<b>Release Information</b>	Command introduced in Junos OS Release 10.1. Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Display Multiple VLAN Registration Protocol (MVRP) configuration information.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show mvrp applicant-state on page 267</a></li> <li>• <a href="#">show mvrp dynamic-vlan-memberships on page 269</a></li> <li>• <a href="#">show mvrp interface on page 270</a></li> <li>• <a href="#">show mvrp registration-state on page 271</a></li> <li>• <a href="#">show mvrp statistics on page 273</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show mvrp on page 265</a>
<b>Output Fields</b>	Table 20 on page 265 lists the output fields for the <b>show mvrp</b> command. Output fields are listed in the approximate order in which they appear.

Table 20: show mvrp Output Fields

Field Name	Field Description
MVRP dynamic VLAN creation	Displays whether global MVRP dynamic VLAN creation is <b>Enabled</b> or <b>Disabled</b> .
MVRP BPDU MAC address	Displays the multicast media access control (MAC) address for MVRP. If configured, the provider MVRP multicast MAC address is used; otherwise, the customer MVRP multicast MAC address is used.
MVRP timers (ms)	Displays MVRP timer information: <ul style="list-style-type: none"> <li>• <b>Interface</b>—The interface on which MVRP is configured.</li> <li>• <b>Join</b>—The maximum number of milliseconds the interfaces must wait before sending VLAN advertisements.</li> <li>• <b>Leave</b>—The number of milliseconds an interface must wait after receiving a Leave message to remove the interface from the VLAN specified in the message.</li> <li>• <b>LeaveAll</b>—The interval at which LeaveAll messages are sent on interfaces. LeaveAll messages maintain current MVRP VLAN membership information in the network.</li> </ul>

## Sample Output

### show mvrp

```
user@host> show mvrp
```

```
MVRP configuration for routing instance 'default-switch'
MVRP dynamic VLAN creation : Enabled
MVRP BPDU MAC address      : Customer bridge group (01-80-C2-00-00-21)
MVRP timers (ms)
  Interface      Join   Leave  LeaveAll
  ge-11/2/8      200    800    10000
  ge-11/0/9      200    800    10000
  ge-11/3/0      200    800    10000
```

## show mvrp applicant-state

<b>Syntax</b>	show mvrp applicant-state
<b>Release Information</b>	Command introduced in Junos OS Release 10.1.
<b>Description</b>	For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) applicant state information.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show mvrp on page 265</a></li> <li>• <a href="#">show mvrp interface on page 270</a></li> <li>• <a href="#">show mvrp registration-state on page 271</a></li> <li>• <a href="#">show mvrp statistics on page 273</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show mvrp applicant-state on page 268</a>
<b>Output Fields</b>	Table 21 on page 267 lists the output fields for the <b>show mvrp applicant-state</b> command. Output fields are listed in the approximate order in which they appear.

**Table 21: show mvrp applicant-state Output Fields**

Field Name	Field Description
<b>VLAN Id</b>	Displays the VLAN ID number.
<b>Interface</b>	Displays the interface number associated with the VLAN ID.
<b>State</b>	Displays one of the following MVRP registrar states: <ul style="list-style-type: none"> <li>• <b>VO</b>—Very anxious observer.</li> <li>• <b>VP</b>—Very anxious passive.</li> <li>• <b>VA</b>—Very anxious new.</li> <li>• <b>AN</b>—Anxious new.</li> <li>• <b>AA</b>—Anxious active.</li> <li>• <b>QA</b>—Quiet active.</li> <li>• <b>LA</b>—Leaving active.</li> <li>• <b>AO</b>—Anxious observer.</li> <li>• <b>QO</b>—Quiet observer.</li> <li>• <b>LO</b>—Leaving observer.</li> <li>• <b>AP</b>—Anxious passive.</li> <li>• <b>QA</b>—Quiet passive.</li> </ul>

## Sample Output

### show mvrp applicant-state

```
user@host> show mvrp applicant-state
MVRP applicant state for routing instance 'default-switch'
(V0) Very anxious observer, (VP) Very anxious passive, (VA) Very anxious new,
(AN) Anxious new, (AA) Anxious active, (QA) Quiet active, (LA) Leaving active,
(A0) Anxious observer, (Q0) Quiet observer, (L0) Leaving observer,
(AP) Anxious passive, (QP) Quiet passive

VLAN Id      Interface      State
-----
100          ge-11/3/0      Declaring (QA)
200          ge-11/3/0      Declaring (QA)
300          ge-11/3/0      Declaring (QA)
```

## show mvrp dynamic-vlan-memberships

<b>Syntax</b>	show mvrp dynamic-vlan-memberships
<b>Release Information</b>	Command introduced in Junos OS Release 10.1. Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Display all VLANs that have been created dynamically using Multiple VLAN Registration Protocol (MVRP) on the router or switch.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show mvrp on page 265</a></li> <li>• <a href="#">show mvrp applicant-state on page 267</a></li> <li>• <a href="#">show mvrp interface on page 270</a></li> <li>• <a href="#">show mvrp registration-state on page 271</a></li> <li>• <a href="#">show mvrp statistics on page 273</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show mvrp dynamic-vlan-memberships on page 269</a>
<b>Output Fields</b>	<a href="#">Table 22 on page 269</a> lists the output fields for the <b>show mvrp dynamic-vlan-memberships</b> command. Output fields are listed in the approximate order in which they appear.

**Table 22: show mvrp dynamic-vlan-memberships Output Fields**

Field Name	Field Description
VLAN Id	The VLAN ID 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@host> show mvrp dynamic-vlan-memberships
MVRP dynamic vlans for routing instance 'default-switch'
(s) static vlan, (f) fixed registration

VLAN Id      Interfaces
  100 (s)    ge-11/3/0
  200 (s)    ge-11/3/0
  300 (s)

```

## show mvrp interface

<b>Syntax</b>	show mvrp interface
<b>Release Information</b>	Command introduced in Junos OS Release 10.1.
<b>Description</b>	For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) interface-specific information.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show mvrp on page 265</a></li> <li>• <a href="#">show mvrp applicant-state on page 267</a></li> <li>• <a href="#">show mvrp dynamic-vlan-memberships on page 269</a></li> <li>• <a href="#">show mvrp registration-state on page 271</a></li> <li>• <a href="#">show mvrp statistics on page 273</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show mvrp interface on page 270</a>
<b>Output Fields</b>	<a href="#">Table 23 on page 270</a> lists the output fields for the <b>show mvrp interface</b> command. Output fields are listed in the approximate order in which they appear.

**Table 23: show mvrp interface Output Fields**

Field Name	Field Description
Interface	Interface on which MVRP is configured.
Status	Status of the MVRP: <b>Enabled</b> or <b>Disabled</b> .
Registration Mode	Registration for the interface: <b>Fixed</b> , <b>Forbidden</b> , or <b>Normal</b> .
Applicant Mode	Applicant mode.

## Sample Output

### show mvrp interface

```

user@host> show mvrp interface
MVRP interface information for routing instance 'default-switch'

Interface      Status      Registration Mode      Applicant
Mode
ge-11/2/8      Enabled     Normal                 Normal
ge-11/0/9      Enabled     Normal                 Normal
ge-11/3/0      Enabled     Normal                 Normal

```

## show mvrp registration-state

<b>Syntax</b>	show mvrp registration-state
<b>Release Information</b>	Command introduced in Junos OS Release 10.1.
<b>Description</b>	For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) registration state information.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show mvrp on page 265</a></li> <li>• <a href="#">show mvrp dynamic-vlan-memberships on page 269</a></li> <li>• <a href="#">show mvrp interface on page 270</a></li> <li>• <a href="#">show mvrp statistics on page 273</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show mvrp registration-state on page 271</a>
<b>Output Fields</b>	Table 24 on page 271 lists the output fields for the <b>show mvrp registration-state</b> command. Output fields are listed in the approximate order in which they appear.

Table 24: show mvrp registration-state Output Fields

Field Name	Field Description
<b>VLAN Id</b>	Displays the VLAN ID number.
<b>Interface</b>	Displays the interface number associated with the VLAN ID.
<b>Registrar State</b>	Displays whether the registrar state is <b>Registered</b> or <b>Empty</b> .
<b>Forced State</b>	Displays whether the forced state is <b>Registered</b> or <b>Empty</b> .
<b>Managed State</b>	Displays one of the following states: <ul style="list-style-type: none"> <li>• <b>Fixed</b>—VLANs always stay in a registered state and are declared as such on all other forwarding ports.</li> <li>• <b>Normal</b>—VLANs participate in the MVRP protocol and honor incoming join requests normally.</li> <li>• <b>Forbidden</b>—VLANs ignore the incoming join requests and always stay in an unregistered state.</li> </ul>
<b>STP State</b>	Displays whether the Spanning Tree Protocol (STP) is <b>Blocking</b> or <b>Forwarding</b> .

## Sample Output

### show mvrp registration-state

```
user@host> show mvrp registration-state
MVRP registration state for routing instance 'default-switch'
```

VLAN Id	Interface	Registrar State	Forced State	Managed State	STP State
100	ge-11/2/8	Empty	Registered	Fixed	Forwarding
	ge-11/0/9	Empty	Empty	Normal	Forwarding
	ge-11/3/0	Registered	Registered	Normal	Forwarding
101	ge-11/2/8	Empty	Registered	Fixed	Forwarding
	ge-11/0/9	Empty	Empty	Normal	Forwarding
	ge-11/3/0	Registered	Registered	Normal	Forwarding



## show mvrp statistics

<b>Syntax</b>	show mvrp statistics
<b>Release Information</b>	Command introduced in Junos OS Release 10.1.
<b>Description</b>	For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) statistics in the form of Multiple Registration Protocol data unit (MRPDU) messages.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show mvrp on page 265</a></li> <li>• <a href="#">show mvrp applicant-state on page 267</a></li> <li>• <a href="#">show mvrp dynamic-vlan-memberships on page 269</a></li> <li>• <a href="#">show mvrp interface on page 270</a></li> <li>• <a href="#">show mvrp registration-state on page 271</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show mvrp statistics on page 273</a>
<b>Output Fields</b>	<a href="#">Table 25 on page 273</a> lists the output fields for the <b>show mvrp statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 25: show mvrp statistics Output Fields**

Field Name	Field Description
interface name	Interface for which MVRP statistics are displayed.
VLAN IDs registered	Number of VLAN IDs registered.
Sent MVRP PDUs	Number of MRPDU messages transmitted from the router.
Received MVRP PDUs without error	Number of MRPDU messages received on the router.
Received MVRP PDUs with error	Number of invalid MRPDU messages received on the router.

## Sample Output

### show mvrp statistics

```

user@host> show mvrp statistics
MVRP statistics for routing instance 'default-switch'

Interface name           : ge-11/2/8
VLAN IDs registered      : 0
Sent MVRP PDUs           : 1467

```

```
Received MVRP PDUs without error: 0
Received MVRP PDUs with error   : 0


Interface name                   : ge-11/0/9
VLAN IDs registered              : 0
Sent MVRP PDUs                  : 1418
Received MVRP PDUs without error: 702
Received MVRP PDUs with error   : 0

Interface name                   : ge-11/3/0
VLAN IDs registered              : 2
Sent MVRP PDUs                  : 1524
Received MVRP PDUs without error: 1366
Received MVRP PDUs with error   : 0
```

## show vlans

**Syntax** `show vlans`  
`<brief | detail | extensive>`  
`<instance instance-name>`  
`<logical-system logical-system-name>`  
`<operational>`  
`<vlan-name>`  
`<interface interface-name>`

**Release Information** Command introduced in Junos OS Release 12.3R2.  
 Command introduced in Junos OS Release 12.3R2 for EX Series switches.  
 Option **interface** introduced in Junos OS Release 13.2X50-D10 (ELS).

**Description**  **NOTE:** If your EX Series switch CLI displays different options for the `show vlans` command from the options shown in this document, see *show vlans*.

(MX Series routers and EX Series switches only) Display VLAN information.

**Options** **none**—Display information for all VLANs.

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

**instance *instance-name***—(Optional) Display information for the specified routing instance.

**logical-system *logical-system-name***—(Optional) Display Ethernet-switching statistics information for the specified logical system.

**operational**—(Optional) Display information for the operational routing instances.

***vlan-name***— (Optional) Display information about the specified VLAN.

**interface *interface-name***—(Optional) Display information about the specified interface.

**Required Privilege Level** view

**List of Sample Output** [show vlans brief \(EX Series Switch\) on page 275](#)  
[show vlans brief on page 276](#)  
[show vlans detail \(EX Series Switch\) on page 276](#)  
[show vlans detail on page 278](#)  
[show vlans extensive \(EX Series Switch\) on page 278](#)  
[show vlans extensive on page 280](#)

## Sample Output

`show vlans brief (EX Series Switch)`

```
user@switch> show vlans brief
Routing instance   VLAN name   Tag   Interfaces
default-switch    c1         20
```

			ge-0/0/0.0*
			ge-1/0/0.0*
			ge-2/0/0.0*
default-switch	c2	30	
			ge-0/0/0.0*
			ge-2/0/0.0*
default-switch	default	1	
default-switch	iso	10	
			ge-0/0/1.0*
default-switch	iso1	50	
			ge-0/0/0.0*
			ge-2/0/0.0*
default-switch	pri	100	
			ge-0/0/0.0*
			ge-1/0/0.0*
			ge-2/0/0.0*

### show vlans brief

```

user@host> show vlans brief
Routing instance  VLAN name  Tag  Interfaces
VPLS-1           __VPLS-1__  all  ae1.0
VPLS-2           __VPLS-2__  all  ae3.0
                                   ge-3/1/2.0
                                   vt-3/3/10.1048576
default-switch   VLAN1000    1000 ae26.0
default-switch   VLAN101     101  ae20.0
default-switch   VLAN102     102  ae20.0
default-switch   VLAN103     103  ae20.0
default-switch   VLAN104     104  ae20.0
default-switch   VLAN105     105  ae20.0
default-switch   VLAN106     106  ae20.0
default-switch   VLAN107     107  ae20.0
default-switch   VLAN108     108  ae20.0
[...output truncated...]

```

### show vlans detail (EX Series Switch)

```

user@switch> show vlans detail
Routing instance: default-switch
  VLAN Name: c1                      State: Active
  Tag: 20
  PVLAN type : Community
  Internal index: 16, Generation Index: 21, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-1/0/0.0*,tagged,trunk

```

```

    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 3      , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: c2                                State: Active
  Tag: 30
  PVLAN type : Community
  Internal index: 17, Generation Index: 22, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 2      , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: default                          State: Active
  Tag: 1
  Internal index: 5, Generation Index: 5, Origin: Static
  MAC aging time: 300 seconds
  Number of interfaces: Tagged 0      , Untagged 0
  Total MAC count: 0

Routing instance: default-switch
  VLAN Name: iso                              State: Active
  Tag: 10
  Internal index: 14, Generation Index: 19, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-0/0/1.0*,untagged,access
Number of interfaces: Tagged 0      , Untagged 1
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: iso1                             State: Active
  Tag: 50
  PVLAN type : Isolated
  Internal index: 15, Generation Index: 20, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 2      , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: pri                              State: Active
  Tag: 100
  PVLAN type : Primary
  Isolated VLAN :
  vlan-id : 50 vlan name : iso1
  Community VLAN :
  vlan-id : 20 vlan name : c1
  vlan-id : 30 vlan name : c2
  Internal index: 9, Generation Index: 14, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-1/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk

```

```
Number of interfaces: Tagged 3      , Untagged 0
Total MAC count: 0
```

### show vlans detail

```
user@host> show vlans detail
Routing instance: VPLS-1
  VLAN Name: __VPLS-1__                State: Active
  Tag: all
  Internal index: 2, Generation Index:  , Origin: Dynamic
  Interfaces:
    ae1.0,tagged
  Number of interfaces: Tagged 1      , Untagged 0
  Total MAC count: 0

Routing instance: VPLS-2
  VLAN Name: __VPLS-2__                State: Active
  Tag: all
  Internal index: 3, Generation Index:  , Origin: Dynamic
  Interfaces:
    ae3.0,tagged
    ge-3/1/2.0,tagged
    vt-3/3/10.1048576,tagged
  Number of interfaces: Tagged 3      , Untagged 0
  Total MAC count: 4

Routing instance: default-switch
  VLAN Name: VLAN1000                  State: Active
  Tag: 1000
  Internal index: 4, Generation Index: 1, Origin: Static
  Layer 3 interface: irb.1000
  Interfaces:
    ae26.0,tagged,trunk
  Number of interfaces: Tagged 1      , Untagged 0
  Total MAC count: 0

Routing instance: default-switch
  VLAN Name: VLAN101                   State: Active
  Tag: 101
  Internal index: 5, Generation Index: 2, Origin: Static
  Layer 3 interface: irb.101
  Interfaces:
    ae20.0,tagged,trunk
  Number of interfaces: Tagged 1      , Untagged 0
  Total MAC count: 1

Routing instance: default-switch
  VLAN Name: VLAN102                   State: Active
  Tag: 102
  Internal index: 6, Generation Index: 3, Origin: Static
  Layer 3 interface: irb.102
  Interfaces:
    ae20.0,tagged,trunk
  Number of interfaces: Tagged 1      , Untagged 0
  Total MAC count: 1
[...output truncated...]
```

### show vlans extensive (EX Series Switch)

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

```

Routing instance: default-switch
  VLAN Name: c1                               State: Active
  Tag: 20
  PVLAN type : Community
  Internal index: 16, Generation Index: 21, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-1/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
  Number of interfaces: Tagged 3      , Untagged 0
  Total MAC count: 0

```

```

Routing instance: default-switch
  VLAN Name: c2                               State: Active
  Tag: 30
  PVLAN type : Community
  Internal index: 17, Generation Index: 22, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
  Number of interfaces: Tagged 2      , Untagged 0
  Total MAC count: 0

```

```

Routing instance: default-switch
  VLAN Name: default                          State: Active
  Tag: 1
  Internal index: 5, Generation Index: 5, Origin: Static
  MAC aging time: 300 seconds
  Number of interfaces: Tagged 0      , Untagged 0
  Total MAC count: 0

```

```

Routing instance: default-switch
  VLAN Name: iso                             State: Active
  Tag: 10
  Internal index: 14, Generation Index: 19, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-0/0/1.0*,untagged,access
  Number of interfaces: Tagged 0      , Untagged 1
  Total MAC count: 0

```

```

Routing instance: default-switch
  VLAN Name: iso1                            State: Active
  Tag: 50
  PVLAN type : Isolated
  Internal index: 15, Generation Index: 20, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
  Number of interfaces: Tagged 2      , Untagged 0
  Total MAC count: 0

```

```

Routing instance: default-switch
  VLAN Name: pri                             State: Active
  Tag: 100
  PVLAN type : Primary
  Isolated VLAN :
  vlan-id : 50 vlan name : iso1

```

```
Community VLAN :
vlan-id : 20 vlan name : c1
vlan-id : 30 vlan name : c2
Internal index: 9, Generation Index: 14, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-1/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 3      , Untagged 0
Total MAC count: 0
```

### show vlans extensive

```
user@host> show vlans extensive
Routing instance: default-switch
  VLAN Name: VLAN_10                                State: Active
  Tag: 10
  Internal index: 2, Generation Index: 1, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-1/0/3.0*,tagged,trunk
  Number of interfaces: Tagged 1      , Untagged 0
  Total MAC count: 0

Routing instance: default-switch
  VLAN Name: VLAN_20                                State: Active
  Tag: 20
  Internal index: 3, Generation Index: 2, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-1/0/3.0*,tagged,trunk
  Number of interfaces: Tagged 1      , Untagged 0
  Total MAC count: 0
```



## traceroute ethernet

<b>Syntax</b>	<pre> traceroute ethernet local-mep <i>mep-id</i> maintenance-association <i>ma-name</i> maintenance-domain <i>md-name</i> &lt;ttl <i>value</i>&gt; &lt;wait <i>seconds</i>&gt; <i>mac-address</i>   <i>mep-id</i> &lt;detail&gt; </pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.0.</p> <p><b>mep-id</b> option introduced in Junos OS Release 9.1.</p> <p><b>local-mep</b> option introduced in Junos OS Release 15.1</p>
<b>Description</b>	<p>Triggers the linktrace protocol to trace the route between two maintenance points. The result of the traceroute protocol is stored in the path database. To display the path database, use the <b>show oam ethernet connectivity-fault-management path-database</b> command.</p> <p>Before using the traceroute command, you can verify the remote MEP's MAC address using the <b>show oam ethernet connectivity-fault-management path-database</b> command.</p>
<b>Options</b>	<p><b>local-mep <i>mep-id</i></b>—(Required when multiple MEPs are configured) Identifier for the local maintenance endpoint.</p> <p><b>detail</b>—(Optional) Provide detailed information of the responder hostname, ingress port name, egress port name, TTL, and relay action.</p> <p><b>mac-address</b>—Destination unicast MAC address of the remote maintenance point.</p> <p><b>mep-id</b>—MEP identifier of the remote maintenance point. The range of values is 1 through 8191.</p> <p><b>maintenance-association <i>ma-name</i></b>—Specifies an existing maintenance association from the set of configured maintenance associations.</p> <p><b>maintenance-domain <i>md-name</i></b>—Specifies an existing maintenance domain from the set of configured maintenance domains.</p> <p><b>ttl <i>value</i></b>—Number of hops to use in the linktrace request. The range is 1 to 255 hops. The default is 4.</p> <p><b>wait <i>seconds</i></b>—(Optional) Maximum time to wait for a response to the traceroute request. The range is 1 to 255 seconds. The default is 5.</p>
<b>Required Privilege Level</b>	network
<b>List of Sample Output</b>	<p><a href="#">traceroute ethernet on page 283</a></p> <p><a href="#">traceroute ethernet detail on page 283</a></p>

**Output Fields** Table 26 on page 282 lists the output fields for the **traceroute ethernet** command. Output fields are listed in the approximate order in which they appear.

**Table 26: traceroute ethernet Output Fields**

Field Name	Field Description
<b>Linktrace to</b>	MAC address of the destination maintenance point.
<b>Interface</b>	Local interface used to send the linktrace message (LTM).
<b>Maintenance Domain</b>	Maintenance domain specified in the traceroute command.
<b>Level</b>	Maintenance domain level configured.
<b>Maintenance Association</b>	Maintenance association specified in the traceroute command.
<b>Local Mep</b>	The local maintenance end point identifier.
<b>Transaction Identifier</b>	4-byte identifier maintained by the MEP. Each LTM uses a transaction identifier. The transaction identifier is maintained globally across all Maintenance Domains. Use the transaction identifier to match an incoming linktrace response (LTR), with a previously sent LTM.
<b>Hop</b>	Sequential hop count of the linktrace path.
<b>TTL</b>	Number of hops remaining in the linktrace message. The time to live (TTL) is decremented at each hop.
<b>Source MAC address</b>	MAC address of the 802.1ag node responding to the LTM or the source MAC address of the LTR.
<b>Next-hop MAC address</b>	MAC address of the egress interface of the node to which the LTM is forwarded or  the next-hop MAC address derived from the next egress identifier in the Egress-ID TLV of the LTR PDU.
<b>Responder Hostname</b>	The hostname of the responding router. A valid hostname is received only when the responding system is a Juniper Networks router.
<b>Ingress port name</b>	The port name for ingress connections.
<b>Egress port name</b>	The port name for egress connections.
<b>Flags</b>	The configurable flags can include: <ul style="list-style-type: none"> <li>• H— Hardware only, incoming LT frame has hardware bit set.</li> <li>• T— Terminal MEP, responder is a terminating MEP.</li> <li>• F— FWD yes, LTM frame is relayed further.</li> </ul>

Table 26: traceroute ethernet Output Fields (*continued*)

Field Name	Field Description
<b>Relay Action</b>	The associated relay action. Relay action can be one of the following: <ul style="list-style-type: none"> <li>RlyHit— Relay hit; target MAC address matches the MP mac address.</li> <li>RlyFDB— Relay FDB; output port decided by consulting forwarding database.</li> <li>RlyMPDB— Relay MIP; output port decided by consulting MIP database.</li> </ul>

## Sample Output

### traceroute ethernet

```
user@host> traceroute ethernet maintenance-domain md1 maintenance-association ma1
00:01:02:03:04:05
```

```
Linktrace to 00:01:02:03:04:05, Interface : ge-5/0/0.0
Maintenance Domain: MD1, Level: 7
Maintenance Association: MA1, Local Mep: 1
```

Hop	TTL	Source MAC address	Next hop MAC address
Transaction Identifier:100001			
1	63	00:00:aa:aa:aa:aa	00:00:ab:ab:ab:ab
2	62	00:00:bb:bb:bb:bb	00:00:bc:bc:bc:bc
3	61	00:00:cc:cc:cc:cc	00:00:cd:cd:cd:cd
4	60	00:01:02:03:04:05	00:00:00:00:00:00

### traceroute ethernet detail

```
user@host> run traceroute ethernet maintenance-domain md6 maintenance-association ma6
mep 101 detail
```

```
Linktrace to 00:00:5E:00:53:CC, Interface : ge-1/0/0.1
Maintenance Domain: md6, Level: 6
Maintenance Association: ma6, Local Mep: 201
Transaction Identifier: 2077547465
```

#### Legend for RelayAction:

```
RlyHit -- Relay hit, Target MAC address matches the MP mac address
RlyFDB -- Relay FDB, output port decided by consulting FDB database
RlyMPDB -- Relay MIP, output port decided by consulting MIP database
```

#### Legend for Flags:

```
H -- Hardware only, incoming LT frame has hardware bit set
T -- Terminal MEP, responder is a terminating MEP
F -- FWD yes, LTM frame is relayed further
```

TTL	Responder	Hostname	Ingress port name	Egress port name	RelayAction
Responder Service					
62	host1		ge-1/0/0.1	ge-2/3/0.1	RlyFDB
br1			00:00:5E:00:53:00	00:00:5E:00:53:A0	HF-
63	host2		ge-2/3/0.1	ge-1/0/0.1	RlyFDB
br1			00:00:5E:00:53:AA	00:00:5E:00:53:A2	HF-
61	host3		ge-1/0/0.1	--:--	RlyHit
br1			00:00:5E:00:53:B0	--:--	H-T

