



---

Junos<sup>®</sup> OS

# Ethernet 802.1Q VLANs Feature Guide for Routing Devices

Release

14.1



---

Published: 2014-05-15

Juniper Networks, Inc.  
1194 North Mathilda Avenue  
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.

*Junos<sup>®</sup> OS Ethernet 802.1Q VLANs Feature Guide for Routing Devices*

14.1

Copyright © 2014, Juniper Networks, Inc.  
All rights reserved.

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

#### YEAR 2000 NOTICE

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

#### END USER LICENSE AGREEMENT

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

# Table of Contents

	About the Documentation . . . . .	ix
	Documentation and Release Notes . . . . .	ix
	Supported Platforms . . . . .	ix
	Using the Examples in This Manual . . . . .	ix
	Merging a Full Example . . . . .	x
	Merging a Snippet . . . . .	x
	Documentation Conventions . . . . .	xi
	Documentation Feedback . . . . .	xiii
	Requesting Technical Support . . . . .	xiii
	Self-Help Online Tools and Resources . . . . .	xiii
	Opening a Case with JTAC . . . . .	xiv
<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Ethernet 802.1Q VLANs . . . . .</b>	<b>3</b>
	802.1Q VLANs Overview . . . . .	3
<b>Part 2</b>	<b>Configuration</b>	
<b>Chapter 2</b>	<b>Ethernet 802.1Q VLANs . . . . .</b>	<b>7</b>
	Configuring Dynamic 802.1Q VLANs . . . . .	8
	802.1Q VLAN IDs and Ethernet Interface Types . . . . .	8
	Enabling VLAN Tagging . . . . .	9
	Configuring Single-Tag Framing . . . . .	10
	Configuring Dual Tagging . . . . .	10
	Configuring Mixed Tagging . . . . .	10
	Configuring Mixed Tagging Support for Untagged Packets . . . . .	11
	Example: Configuring Mixed Tagging . . . . .	11
	Example: Configuring Mixed Tagging to Support Untagged Packets . . . . .	12
	Binding VLAN IDs to Logical Interfaces . . . . .	12
	Binding VLAN IDs to Logical Interfaces Overview . . . . .	13
	Binding a VLAN ID to a Logical Interface . . . . .	13
	Binding a VLAN ID to a Single-Tag Logical Interface . . . . .	13
	Binding a VLAN ID to a Dual-Tag Logical Interface . . . . .	14
	Binding a Range of VLAN IDs to a Logical Interface . . . . .	14
	Binding a Range of VLAN IDs to a Single-Tag Logical Interface . . . . .	14
	Binding a Range of VLAN IDs to a Dual-Tag Logical Interface . . . . .	14
	Example: Binding Ranges VLAN IDs to Logical Interfaces . . . . .	15
	Binding a List of VLAN IDs to a Logical Interface . . . . .	15
	Binding a List of VLAN IDs to a Single-Tag Logical Interface . . . . .	16
	Binding a List of VLAN IDs to a Dual-Tag Logical Interface . . . . .	16

Example: Binding Lists of VLAN IDs to Logical Interfaces . . . . .	16
Associating VLAN IDs to VLAN Demux Interfaces . . . . .	17
Associating VLAN IDs to VLAN Demux Interfaces Overview . . . . .	17
Associating a VLAN ID to a VLAN Demux Interface . . . . .	18
Associating a VLAN ID to a Single-Tag VLAN Demux Interface . . . . .	18
Associating a VLAN ID to a Dual-Tag VLAN Demux Interface . . . . .	18
Configuring VLAN Encapsulation . . . . .	18
Example: Configuring VLAN Encapsulation on a Gigabit Ethernet Interface . . . . .	19
Example: Configuring VLAN Encapsulation on an Aggregated Ethernet Interface . . . . .	20
Configuring Extended VLAN Encapsulation . . . . .	20
Example: Configuring Extended VLAN Encapsulation on a Gigabit Ethernet Interface . . . . .	20
Example: Configuring Extended VLAN Encapsulation on an Aggregated Ethernet Interface . . . . .	21
Guidelines for Configuring VLAN ID List-Bundled Logical Interfaces That Connect CCCs . . . . .	21
Guidelines for Configuring Physical Link-Layer Encapsulation to Support CCCs . . . . .	22
Guidelines for Configuring Logical Link-Layer Encapsulation to Support CCCs . . . . .	22
Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface . . . . .	23
Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance . . . . .	23
Specifying the Interface Over Which VPN Traffic Travels to the CE Router . . . . .	24
Specifying the Interface to Handle Traffic for a CCC . . . . .	24
Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance . . . . .	25
Specifying the Interface Over Which VPN Traffic Travels to the CE Router . . . . .	25
Specifying the Interface to Handle Traffic for a CCC . . . . .	26
Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface . . . . .	27
Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance . . . . .	27
Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit . . . . .	27
Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance . . . . .	28
Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit . . . . .	29
Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface . . . . .	29
Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface . . . . .	31
Configuring a Logical Interface for Access Mode . . . . .	32
Example: Configuring a Logical Interface for Access Mode . . . . .	32
Configuring a Logical Interface for Trunk Mode . . . . .	33
Configuring the VLAN ID List for a Trunk Interface . . . . .	33

	Configuring a Trunk Interface on a Bridge Network . . . . .	34
	Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers . . . .	36
<b>Chapter 3</b>	<b>Network Interfaces Configuration Statements and Hierarchy . . . . .</b>	<b>37</b>
	[edit interfaces] Hierarchy Level . . . . .	37
	[edit logical-systems] Hierarchy Level . . . . .	53
<b>Chapter 4</b>	<b>Statement Summary . . . . .</b>	<b>59</b>
	address . . . . .	60
	encapsulation (Logical Interface) . . . . .	62
	encapsulation (Physical Interface) . . . . .	66
	family . . . . .	71
	fast-aps-switch . . . . .	75
	flexible-vlan-tagging . . . . .	76
	interface-mode . . . . .	77
	native-vlan-id . . . . .	79
	stacked-vlan-tagging . . . . .	80
	unit . . . . .	81
	vlan-id (VLAN ID to Be Bound to a Logical Interface) . . . . .	87
	vlan-id (VLAN ID to Rewrite) . . . . .	88
	vlan-id-list (Interface in Bridge Domain) . . . . .	89
	vlan-id-range . . . . .	90
	vlan-tagging . . . . .	91
	vlan-tags (Dual-Tagged Logical Interface) . . . . .	92
	vlan-tags (Stacked VLAN Tags) . . . . .	94
	vlan-id-list (Ethernet VLAN Circuit) . . . . .	96
<b>Part 3</b>	<b>Administration</b>	
<b>Chapter 5</b>	<b>Monitoring Commands . . . . .</b>	<b>101</b>
	show interfaces (Aggregated Ethernet) . . . . .	103
	show interfaces demux0 (Demux Interfaces) . . . . .	114
	show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet) . . . . .	123
	show interfaces (Fast Ethernet) . . . . .	144
	show interfaces (10-Gigabit Ethernet) . . . . .	160
	show interfaces interface-set (Ethernet Interface Set) . . . . .	185
	show interfaces interface-set queue . . . . .	188
	show interfaces irb . . . . .	196
	show lacp interfaces . . . . .	202
	show interfaces mac-database (Gigabit Ethernet) . . . . .	206
	show interfaces mc-ae . . . . .	210
	show oam ethernet connectivity-fault-management delay-statistics . . . . .	213
	show oam ethernet connectivity-fault-management forwarding-state . . . . .	217
	show oam ethernet connectivity-fault-management interfaces . . . . .	221
	show oam ethernet connectivity-fault-management linktrace path-database . . . . .	231
	show oam ethernet connectivity-fault-management mep-database . . . . .	234
	show oam ethernet connectivity-fault-management mep-statistics . . . . .	245
	show oam ethernet connectivity-fault-management path-database . . . . .	256

	show oam ethernet evc . . . . .	258
	show oam ethernet link-fault-management . . . . .	259
	show oam ethernet lmi . . . . .	265
	show oam ethernet lmi statistics . . . . .	267
	show protection-group ethernet-ring aps . . . . .	269
	show protection-group ethernet-ring data-channel . . . . .	271
	show protection-group ethernet-ring interface . . . . .	273
	show protection-group ethernet-ring node-state . . . . .	276
	show protection-group ethernet-ring statistics . . . . .	279
	show protection-group ethernet-ring vlan . . . . .	282
<b>Part 4</b>	<b>Troubleshooting</b>	
<b>Chapter 6</b>	<b>Ethernet . . . . .</b>	<b>287</b>
	tracertoe ethernet . . . . .	288
	Tracing PPPoE Operations . . . . .	290
	Configuring the PPPoE Trace Log Filename . . . . .	290
	Configuring the Number and Size of PPPoE Log Files . . . . .	291
	Configuring Access to the PPPoE Log File . . . . .	291
	Configuring a Regular Expression for PPPoE Lines to Be Logged . . . . .	291
	Configuring the PPPoE Tracing Flags . . . . .	291
	Troubleshooting PPPoE Service Name Tables . . . . .	291
	Verifying a PPPoE Configuration . . . . .	293
<b>Chapter 7</b>	<b>Interface Diagnostics . . . . .</b>	<b>295</b>
	Interface Diagnostics . . . . .	295
	Configuring Loopback Testing . . . . .	295
	Interface Diagnostics . . . . .	297
	Starting and Stopping a BERT Test . . . . .	301
	Example: Configuring Bit Error Rate Testing . . . . .	301
<b>Part 5</b>	<b>Index</b>	
	Index . . . . .	305

# List of Tables

	<b>About the Documentation</b> . . . . .	<b>ix</b>
	Table 1: Notice Icons . . . . .	xi
	Table 2: Text and Syntax Conventions . . . . .	xii
<b>Part 2</b>	<b>Configuration</b>	
<b>Chapter 2</b>	<b>Ethernet 802.1Q VLANs</b> . . . . .	<b>7</b>
	Table 3: VLAN ID Range by Interface Type . . . . .	9
	Table 4: Configuration Statements Used to Bind VLAN IDs to Logical Interfaces . . . . .	13
	Table 5: Configuration Statements Used to Associate VLAN IDs to VLAN Demux Interfaces . . . . .	18
	Table 6: Encapsulation Inside Circuits CCC-Connected by VLAN-Bundled Logical Interfaces . . . . .	22
<b>Part 3</b>	<b>Administration</b>	
<b>Chapter 5</b>	<b>Monitoring Commands</b> . . . . .	<b>101</b>
	Table 7: show interfaces (Aggregated Ethernet) Output Fields . . . . .	103
	Table 8: Demux show interfaces Output Fields . . . . .	114
	Table 9: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet DWDM and DWDM OTN PICs . . . . .	124
	Table 10: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics . . . . .	125
	Table 11: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers . . . . .	128
	Table 12: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers . . . . .	130
	Table 13: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers . . . . .	133
	Table 14: show interfaces Fast Ethernet Output Fields . . . . .	144
	Table 15: show interfaces Gigabit Ethernet Output Fields . . . . .	161
	Table 16: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type . . . . .	175
	Table 17: Ethernet show interfaces interface-set Output Fields . . . . .	185
	Table 18: Ethernet show interfaces interface-set queue Output Fields . . . . .	188
	Table 19: show interfaces irb Output Fields . . . . .	196
	Table 20: show lacp interfaces Output Fields . . . . .	202
	Table 21: show interfaces mac-database Output Fields . . . . .	206
	Table 22: show interfaces mc-ae Output Fields . . . . .	210

Table 23: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields . . . . .	214
Table 24: show oam ethernet connectivity-fault-management forwarding-state Output Fields . . . . .	217
Table 25: show oam ethernet connectivity-fault-management interfaces Output Fields . . . . .	222
Table 26: show oam ethernet connectivity-fault-management linktrace path-database Output Fields . . . . .	231
Table 27: show oam ethernet connectivity-fault-management mep-database Output Fields . . . . .	235
Table 28: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields . . . . .	246
Table 29: show oam ethernet connectivity-fault-management path-database Output Fields . . . . .	256
Table 30: show oam ethernet evc Output Fields . . . . .	258
Table 31: show oam ethernet link-fault-management Output Fields . . . . .	259
Table 32: show oam ethernet lmi Output Fields . . . . .	265
Table 33: show oam ethernet lmi statistics Output Fields . . . . .	267
Table 34: show protection-group ethernet-ring aps Output Fields . . . . .	269
Table 35: show protection-group ethernet-ring data-channel Output Fields . . . . .	271
Table 36: MX Series Routers show protection-group ethernet-ring interface Output Fields . . . . .	273
Table 37: show protection-group ethernet-ring node-state Output Fields . . . . .	276
Table 38: show protection-group ethernet-ring statistics Output Fields . . . . .	279
Table 39: show protection-group ethernet-ring vlan Output Fields . . . . .	282

## Part 4

### Chapter 6

### Chapter 7

## Troubleshooting

<b>Ethernet . . . . .</b>	<b>287</b>
Table 40: traceroute ethernet Output Fields . . . . .	288
<b>Interface Diagnostics . . . . .</b>	<b>295</b>
Table 41: Loopback Modes by Interface Type . . . . .	296
Table 42: BERT Capabilities by Interface Type . . . . .	300

# About the Documentation

- [Documentation and Release Notes on page ix](#)
- [Supported Platforms on page ix](#)
- [Using the Examples in This Manual on page ix](#)
- [Documentation Conventions on page xi](#)
- [Documentation Feedback on page xiii](#)
- [Requesting Technical Support on page xiii](#)

## Documentation and Release Notes

---

To obtain the most current version of all Juniper Networks<sup>®</sup> 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:

- [M Series](#)
- [MX Series](#)
- [T Series](#)
- [J 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 xi 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 xii 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>
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</b> <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	<b>broadcast</b>   <b>multicast</b>  ( <i>string1</i>   <i>string2</i>   <i>string3</i> )
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<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> ]
Indentation and braces ( { } )	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

---

## GUI Conventions

---

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

Convention	Description	Examples
<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>
> (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 send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net), or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version (if applicable)

## Requesting Technical Support

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

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

## Self-Help Online Tools and Resources

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

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>

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

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

## Opening a Case with JTAC

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

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

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

## PART 1

# Overview

- [Ethernet 802.1Q VLANs on page 3](#)



## CHAPTER 1

# Ethernet 802.1Q VLANs

- [802.1Q VLANs Overview on page 3](#)

## 802.1Q VLANs Overview

---

For Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet interfaces supporting VPLS, the Junos OS supports a subset of the IEEE 802.1Q standard for channelizing an Ethernet interface into multiple logical interfaces, allowing many hosts to be connected to the same Gigabit Ethernet switch, but preventing them from being in the same routing or bridging domain.

### Related Documentation

- [Configuring Dynamic 802.1Q VLANs on page 8](#)
- [802.1Q VLAN IDs and Ethernet Interface Types on page 8](#)
- [Enabling VLAN Tagging on page 9](#)
- [Binding VLAN IDs to Logical Interfaces on page 12](#)
- [Configuring VLAN Encapsulation on page 18](#)
- [Configuring Extended VLAN Encapsulation on page 20](#)
- [Guidelines for Configuring VLAN ID List-Bundled Logical Interfaces That Connect CCCs on page 21](#)
- [Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface on page 23](#)
- [Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance on page 23](#)
- [Specifying the Interface Over Which VPN Traffic Travels to the CE Router on page 24](#)
- [Specifying the Interface to Handle Traffic for a CCC on page 24](#)
- [Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface on page 27](#)
- [Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance on page 27](#)
- [Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit on page 27](#)

- [Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface on page 29](#)
- [Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface on page 31](#)
- [Configuring a Logical Interface for Access Mode on page 32](#)
- [Configuring a Logical Interface for Trunk Mode on page 33](#)
- [Configuring the VLAN ID List for a Trunk Interface on page 33](#)
- [Configuring a Trunk Interface on a Bridge Network on page 34](#)
- *Ethernet Interfaces*

## PART 2

# Configuration

- [Ethernet 802.1Q VLANs on page 7](#)
- [Network Interfaces Configuration Statements and Hierarchy on page 37](#)
- [Statement Summary on page 59](#)



## CHAPTER 2

# Ethernet 802.1Q VLANs

- [Configuring Dynamic 802.1Q VLANs on page 8](#)
- [802.1Q VLAN IDs and Ethernet Interface Types on page 8](#)
- [Enabling VLAN Tagging on page 9](#)
- [Binding VLAN IDs to Logical Interfaces on page 12](#)
- [Associating VLAN IDs to VLAN Demux Interfaces on page 17](#)
- [Configuring VLAN Encapsulation on page 18](#)
- [Configuring Extended VLAN Encapsulation on page 20](#)
- [Guidelines for Configuring VLAN ID List-Bundled Logical Interfaces That Connect CCCs on page 21](#)
- [Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface on page 23](#)
- [Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance on page 25](#)
- [Specifying the Interface Over Which VPN Traffic Travels to the CE Router on page 25](#)
- [Specifying the Interface to Handle Traffic for a CCC on page 26](#)
- [Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface on page 27](#)
- [Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance on page 28](#)
- [Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit on page 29](#)
- [Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface on page 29](#)
- [Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface on page 31](#)
- [Configuring a Logical Interface for Access Mode on page 32](#)
- [Configuring a Logical Interface for Trunk Mode on page 33](#)
- [Configuring the VLAN ID List for a Trunk Interface on page 33](#)
- [Configuring a Trunk Interface on a Bridge Network on page 34](#)
- [Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers on page 36](#)

## Configuring Dynamic 802.1Q VLANs

---

You can configure the router to dynamically create VLANs when a client accesses an interface and requests a VLAN ID that does not yet exist. When a client accesses a VLAN interface, the router instantiates a VLAN dynamic profile that you have associated with the interface. Using the settings in the dynamic profile, the router extracts information about the client from the incoming packet (for example, the interface and unit values), saves this information in the routing table, and creates a VLAN or stacked VLAN ID for the client from a range of VLAN IDs that you configure for the interface.

Dynamically configuring VLANs or stacked VLANs requires the following general steps:

1. Configure a dynamic profile for dynamic VLAN or dynamic stacked VLAN creation.
2. Associate the VLAN or stacked VLAN dynamic profile with the interface.
3. Specify the Ethernet packet type that the VLAN dynamic profile accepts.
4. Define VLAN ranges for use by the dynamic profile when creating VLAN IDs.

For procedures on how to configure dynamic VLANs and dynamic stacked VLANs for client access, see the *Junos OS Subscriber Management and Services Library*.

- Related Documentation**
- [802.1Q VLANs Overview on page 3](#)
  - [Ethernet Interfaces](#)

## 802.1Q VLAN IDs and Ethernet Interface Types

---

You can partition the router into up to 4095 different VLANs—depending on the router model and the physical interface types—by associating logical interfaces with specific VLAN IDs.

VLAN ID 0 is reserved for tagging the priority of frames. VLAN IDs 1 through 511 are reserved for normal VLANs. VLAN IDs 512 and above are reserved for VLAN circuit cross-connect (CCCs).

For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can configure flexible Ethernet services encapsulation on the physical interface. With flexible Ethernet services encapsulation, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

The maximum number of user-configurable VLANs is 15 on each port of the Dense-FE PIC (8-port/12-port/48-port).

[Table 3 on page 9](#) lists VLAN ID range by interface type.

Table 3: VLAN ID Range by Interface Type

Interface Type	VLAN ID Range
Aggregated Ethernet for Fast Ethernet	1 through 1023
Aggregate Ethernet for Gigabit Ethernet	1 through 4094
4-port, 8-port, and 12-port Fast Ethernet	1 through 1023
48-port Fast Ethernet	1 through 4094
Tri-Rate Ethernet copper	1 through 4094
Gigabit Ethernet	1 through 4094
Gigabit Ethernet IQ	1 through 4094
10-Gigabit Ethernet	1 through 4094
Management and internal Ethernet interfaces	1 through 1023



**NOTE:** For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the built-in Gigabit Ethernet port on the M7i router), VLAN IDs on a single interface can differ from each other.

Because IS-IS has an 8-bit limit for broadcast multiaccess media, you cannot set up more than 255 adjacencies over Gigabit Ethernet using VLAN tagging. For more information, see the *Junos OS Routing Protocols Library for Routing Devices*.

- Related Documentation**
- [802.1Q VLANs Overview on page 3](#)
  - [Ethernet Interfaces](#)

## Enabling VLAN Tagging

You can configure the router to receive and forward single-tag frames, dual-tag frames, or a mixture of single-tag and dual-tag frames. For more information, see the following sections:

- [Configuring Single-Tag Framing on page 10](#)
- [Configuring Dual Tagging on page 10](#)
- [Configuring Mixed Tagging on page 10](#)
- [Configuring Mixed Tagging Support for Untagged Packets on page 11](#)

- [Example: Configuring Mixed Tagging on page 11](#)
- [Example: Configuring Mixed Tagging to Support Untagged Packets on page 12](#)



**NOTE:** If you configure VLAN tagging on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces on M320, M120, and T Series routers, the Junos OS creates an internal logical interface that reserves 50 Kbps of bandwidth from Gigabit Ethernet IQ interfaces and 2 Mbps of bandwidth from Gigabit Ethernet IQ2 and IQ2-E interfaces. As a result, the effective available bandwidth for these interface types is now 999.5 Mbps and 998 Mbps, respectively.

---

## Configuring Single-Tag Framing

To configure the router to receive and forward single-tag frames with 802.1Q VLAN tags, include the **vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]  
vlan-tagging;
```

## Configuring Dual Tagging

To configure the routing platform to receive and forward dual-tag frames with 802.1Q VLAN tags, include the **stacked-vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]  
stacked-vlan-tagging;
```

## Configuring Mixed Tagging

Mixed tagging is supported for Gigabit Ethernet interfaces on Gigabit Ethernet IQ2 and IQ2-E, and IQ or IQE PICs on M Series and T Series routers, for all MX Series router Gigabit and 10-Gigabit Ethernet interfaces, and for aggregated Ethernet interfaces with member links in IQ2 and IQ2-E PICs or in MX Series DPCs. Mixed tagging lets you configure two logical interfaces on the same Ethernet port, one with single-tag framing and one with dual-tag framing.



**NOTE:** Mixed tagging is not supported on Fast Ethernet interfaces or on J Series Services Routers.

To configure mixed tagging, include the **flexible-vlan-tagging** statement at the **[edit interfaces *ge-fpc/pic/port* ]** hierarchy level. You must also include the **vlan-tags** statement with **inner** and **outer** options or the **vlan-id** statement at the **[edit interfaces *ge-fpc/pic/port* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces ge-fpc/pic/port]  
flexible-vlan-tagging;  
unit logical-unit-number {  
    vlan-id number;  
    family family {
```

```

        address address;
    }
}
unit logical-unit-number {
    vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
    family family {
        address address;
    }
}

```



**NOTE:** When you configure the physical interface MTU for mixed tagging, you must increase the MTU to 4 bytes more than the MTU value you would configure for a standard VLAN-tagged interface.

For example, if the MTU value is configured to be 1018 on a VLAN-tagged interface, then the MTU value on a flexible VLAN tagged interface must be 1022—4 bytes more. The additional 4 bytes accommodates the future addition of a stacked VLAN tag configuration on the same physical interface.

If the same physical interface MTU value is configured on both the VLAN and flexible VLAN-tag routers, the L2 circuit configuration does not come up and a MTU mismatch is logged. However, normal traffic flow is unaffected.

For encapsulation type **flexible-ethernet-services**, all VLAN IDs are valid. See “[Configuring VLAN Encapsulation](#)” on page 18.

## Configuring Mixed Tagging Support for Untagged Packets

For 1-, 4-, and 8-port Gigabit Ethernet IQ2 and IQ2-E PICs, for 1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs, for all MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces configured for 802.1Q flexible VLAN tagging, and for aggregated Ethernet interfaces on IQ2 and IQ2-E PICs or MX Series DPCs, you can configure mixed tagging support for untagged packets on a port. Untagged packets are accepted on the same mixed VLAN-tagged port. To accept untagged packets, include the **native-vlan-id** statement and the **flexible-vlan-tagging** statement at the **[edit interfaces interface-name]** hierarchy level:

```

[edit interfaces ge-fpc/pic/port]
flexible-vlan-tagging;
native-vlan-id number;

```

The logical interface on which untagged packets are to be received must be configured with the same native VLAN ID as that configured on the physical interface. To configure the logical interface, include the **vlan-id** statement (matching the **native-vlan-id** statement on the physical interface) at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level.

## Example: Configuring Mixed Tagging

The following example configures mixed tagging. Dual-tag and single-tag logical interfaces are under the same physical interface:

```
[edit interfaces ge-3/0/1]
flexible-vlan-tagging;
unit 0 {
  vlan-id 232;
  family inet {
    address 10.66.1.2/30;
  }
}
unit 1 {
  vlan-tags outer 0x8100.222 inner 0x8100.221;
  family inet {
    address 10.66.1.2/30;
  }
}
```

For information about binding VLAN IDs to logical interfaces, see [“Binding VLAN IDs to Logical Interfaces” on page 12](#). For information about configuring dual VLAN tags using the **vlan-tag** statement, see *Stacking a VLAN Tag*.

### Example: Configuring Mixed Tagging to Support Untagged Packets

The following example configures untagged packets to be mapped to logical unit number 0:

```
[edit interfaces ge-0/2/0]
flexible-vlan-tagging;
native-vlan-id 232;
unit 0 {
  vlan-id 232;
  family inet {
    address 10.66.1.2/30;
  }
}
unit 1 {
  vlan-tags outer 0x8100.222 inner 0x8100.221;
  family inet {
    address 10.66.1.2/30;
  }
}
```

- Related Documentation**
- [802.1Q VLANs Overview on page 3](#)
  - [Ethernet Interfaces](#)

---

## Binding VLAN IDs to Logical Interfaces

The following sections describe how to configure logical interfaces to receive and forward VLAN-tagged frames:

- [Binding VLAN IDs to Logical Interfaces Overview on page 13](#)
- [Binding a VLAN ID to a Logical Interface on page 13](#)
- [Binding a Range of VLAN IDs to a Logical Interface on page 14](#)
- [Binding a List of VLAN IDs to a Logical Interface on page 15](#)

## Binding VLAN IDs to Logical Interfaces Overview

To configure a logical interface to receive and forward VLAN-tagged frames, you must bind a VLAN ID, a range of VLAN IDs, or a list of VLAN IDs to the logical interface.

Table 4 on page 13 lists the configuration statements you use to bind VLAN IDs to logical interfaces, organized by scope of the VLAN IDs used to match incoming packets:

**Table 4: Configuration Statements Used to Bind VLAN IDs to Logical Interfaces**

Scope of VLAN ID Matching	Type of VLAN Framing Supported on the Logical Interface	
	Single-Tag Framing	Dual-Tag Framing
VLAN ID	<code>vlan-id <i>vlan-id</i>;</code>	<code>vlan-tags outer <i>tpid.&lt;vlan-id&gt;</i> inner <i>tpidvlan-id</i>;</code>
VLAN ID Range	<code>vlan-id-range <i>vlan-id-vlan-id</i>;</code>	<code>vlan-tags outer <i>tpid.vlan-id</i> inner-range <i>tpid.vlan-id-vlan-id</i>;</code>
VLAN ID List	<code>vlan-id-list [<i>vlan-id</i> <i>vlan-id-vlan-id</i>];</code>	<code>vlan-tags outer <i>&lt;tpid.&gt;vlan-id</i> inner-list [<i>vlan-id</i> <i>vlan-id-vlan-id</i>];</code>

You can include all of the statements at the following hierarchy levels:

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



**NOTE:** The inner-list option of the `vlan-tags` statement does not support Tag Protocol ID (TPID) values.

## Binding a VLAN ID to a Logical Interface

A logical interface that you have associated (bound) to a particular VLAN ID will receive and forward incoming frames that contain a matching VLAN ID.

### Binding a VLAN ID to a Single-Tag Logical Interface

To bind a VLAN ID to a single-tag logical interface, include the `vlan-id` statement:

```
vlan-id vlan-id;
```

You can include the statement at the following hierarchy levels:

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

To configure an Ethernet interface to support single-tag logical interfaces, include the `vlan-tagging` statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the `flexible-vlan-tagging` statement instead.

### Binding a VLAN ID to a Dual-Tag Logical Interface

---

To bind a VLAN ID to a dual-tag logical interface, include the **vlan-tags** statement:

```
vlan-tags inner <tpid.>vlan-id outer <tpid.>vlan-id;
```

You can include the statement at the following hierarchy levels:

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

To configure an Ethernet interface to support dual-tag logical interfaces, include the **stacked-vlan-tagging** statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

### Binding a Range of VLAN IDs to a Logical Interface

A VLAN range can be used by service providers to interconnect multiple VLANs belonging to a particular customer over multiple sites. Using a VLAN ID range conserves switch resources and simplifies configuration.

#### Binding a Range of VLAN IDs to a Single-Tag Logical Interface

---

To bind a range of VLAN IDs to a single-tag logical interface, include the **vlan-id-range** statement:

```
vlan-id-range vlan-id-vlan-id;
```

You can include the statement at the following hierarchy levels:

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

To configure an Ethernet interface to support single-tag logical interfaces, include the **vlan-tagging** statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

#### Binding a Range of VLAN IDs to a Dual-Tag Logical Interface

---

To bind a range of VLAN IDs to a dual-tag logical interface, include the **vlan-tags** statement. Use the **inner-list** option to specify the VLAN IDs as an inclusive range by separating the starting VLAN ID and ending VLAN ID with a hyphen.

```
vlan-tags inner-list [ vlan-id vlan-id-vlan-id ] outer <tpid.>vlan-id;
```

You can include the statement at the following hierarchy levels:

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

To configure an Ethernet interface to support dual-tag logical interfaces, include the **stacked-vlan-tagging** statement at the `[edit interfaces ethernet-interface-name]` hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

### Example: Binding Ranges VLAN IDs to Logical Interfaces

The following example configures two different ranges of VLAN IDs on two different logical ports:

```
[edit interfaces]
ge-3/0/0 {
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 500-600;
  }
}
ge-3/0/1 {
  flexible-vlan-tagging;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 200-300;
  }
  unit 1 {
    encapsulation vlan-bridge;
    vlan-tags outer 1000 inner-range 100-200;
  }
}
```

### Binding a List of VLAN IDs to a Logical Interface

In Junos OS Release 9.5 and later, on MX Series routers and in Junos OS Release 12.2R2 and later on EX Series switches, you can bind a list of VLAN IDs to a single logical interface, eliminating the need to configure a separate logical interface for every VLAN or VLAN range. A logical interface that accepts packets tagged with any VLAN ID specified in a VLAN ID list is called a *VLAN-bundled* logical interface.

You can use VLAN-bundled logical interfaces to configure circuit cross-connects between Layer 2 VPN routing instances or Layer 2 circuits. Using VLAN-bundled logical interfaces simplifies configuration and reduces use of system resources such as logical interfaces, next hops, and circuits.

As an alternative to configuring multiple logical interfaces (one for each VLAN ID and one for each range of VLAN IDs), you can configure a single VLAN-bundled logical interface based on a list of VLAN IDs.



**NOTE:** The `vlan-id` option is not supported to achieve VLAN normalization on VPLS instances that are configured with `vlan-id-list`. However, you can use the `vlan-maps` option to achieve VLAN normalization.

### Binding a List of VLAN IDs to a Single-Tag Logical Interface

---

To bind a list of VLAN IDs to a single-tag logical interface, include the **vlan-id-list** statement. Specify the VLAN IDs in the list individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both.

```
vlan-id-list [ vlan-id vlan-id-vlan-id ];
```

You can include the statement at the following hierarchy levels:

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

To configure an Ethernet interface to support single-tag logical interfaces, include the **vlan-tagging** statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

### Binding a List of VLAN IDs to a Dual-Tag Logical Interface

---

To bind a list of VLAN IDs to a dual-tag logical interface, include the **vlan-tags** statement. Use the **inner-list** option to specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both:

```
vlan-tags inner-list [ vlan-id vlan-id-vlan-id ] outer <tpid> vlan-id;
```



**NOTE:** The inner-list option of the **vlan-tags** statement does not support Tag Protocol ID (TPID) values.

---

You can include the statement at the following hierarchy levels:

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

To configure an Ethernet interface to support dual-tag logical interfaces, include the **stacked-vlan-tagging** statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

### Example: Binding Lists of VLAN IDs to Logical Interfaces

---

The following example configures two different lists of VLAN IDs on two different logical ports:

```
[edit interfaces]
ge-1/1/0 {
  vlan-tagging; # Only for single-tagging
  encapsulation flexible-ethernet-services;
  unit 10 {
    encapsulation vlan-ccc;
```

```

        vlan-id-list [20 30–40 45];
    }
}
ge-1/1/1 {
    flexible-vlan-tagging; # Only for mixed tagging
    encapsulation flexible-ethernet-services;
    unit 10 {
        encapsulation vlan-ccc;
        vlan-id-list [1 10 20 30–40];
    }
    unit 20 {
        encapsulation vlan-ccc;
        vlan-tags outer 200 inner-list [50–60 80 90–100];
    }
}

```

In the example configuration above, **ge-1/1/0** supports single-tag logical interfaces, and **ge-1/1/1** supports mixed tagging. The single-tag logical interfaces **ge-1/1/0.10** and **ge-1/1/1.20** each bundle lists of VLAN IDs. The dual-tag logical interface **ge-1/1/1.20** bundles lists of inner VLAN IDs.



**TIP:** You can group a range of identical interfaces into an interface range and then apply a common configuration to that interface range. For example, in the above example configuration, both interfaces **ge-1/1/0** and **ge-1/1/1** have the same physical encapsulation type of **flexible-ethernet-services**. Thus you can define an interface range with the interfaces **ge-1/1/0** and **ge-1/1/1** as its members and apply the encapsulation type **flexible-ethernet-services** to that defined interface range. For more information about interface ranges, see *Configuring Interface Ranges*.

- Related Documentation**
- [802.1Q VLANs Overview on page 3](#)
  - *Configuring Interface Ranges*
  - *Ethernet Interfaces*

## Associating VLAN IDs to VLAN Demux Interfaces

The following sections describe how to configure VLAN demux interfaces to receive and forward VLAN-tagged frames:

- [Associating VLAN IDs to VLAN Demux Interfaces Overview on page 17](#)
- [Associating a VLAN ID to a VLAN Demux Interface on page 18](#)

### Associating VLAN IDs to VLAN Demux Interfaces Overview

To configure a VLAN demux interface to receive and forward VLAN-tagged frames, you must associate a VLAN ID or dual tagged (stacked) VLAN ID to the interface.

[Table 5 on page 18](#) shows the configuration statements you use to associate VLAN IDs to VLAN demux interfaces, depending on the VLAN tag framing you use:

Table 5: Configuration Statements Used to Associate VLAN IDs to VLAN Demux Interfaces

	Single-Tag Framing	Dual-Tag Framing
Statement Format	<code>vlan-id <i>vlan-id</i>;</code>	<code>vlan-tags outer <i>tpid</i>.&lt;<i>vlan-id</i>&gt; inner <i>tpid</i><i>vlan-id</i>;</code>

You can include all of the 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]`
- `[edit interfaces demux0 unit logical-unit-number]`

## Associating a VLAN ID to a VLAN Demux Interface

A VLAN demux interface that you have associated to a particular VLAN ID receives and forwards incoming frames that contain a matching VLAN ID. You can associate a VLAN ID to a single-tag logical interface or to a dual-tagged (stacked) logical interface.

1. [Associating a VLAN ID to a Single-Tag VLAN Demux Interface on page 18](#)
2. [Associating a VLAN ID to a Dual-Tag VLAN Demux Interface on page 18](#)

### Associating a VLAN ID to a Single-Tag VLAN Demux Interface

To associate a VLAN ID to a single-tag VLAN demux interface, include the **vlan-id** statement at the `[edit interfaces demux0 unit logical-unit-number]` hierarchy level:

```
vlan-id vlan-id;
```

To configure an interface to support single-tag logical interfaces, you must also include the **vlan-tagging** statement at the `[edit interfaces interface-name]` hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

### Associating a VLAN ID to a Dual-Tag VLAN Demux Interface

To associate a VLAN ID to a dual-tag VLAN demux interface, include the **vlan-tags** statement at the `[edit interfaces demux0 unit logical-unit-number]` hierarchy level:

```
vlan-tags inner <tpid.>vlan-id outer <tpid.>vlan-id;
```

To configure an interface to support dual-tag logical interfaces, include the **stacked-vlan-tagging** statement at the `[edit interfaces interface-name]` hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

## Configuring VLAN Encapsulation

Gigabit Ethernet IQ, Gigabit Ethernet PICs with small form-factor pluggable optics (SFPs), and MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces with VLAN tagging enabled can use flexible Ethernet services, VLAN CCC, or VLAN virtual private LAN service (VPLS) encapsulation.

Aggregated Ethernet interfaces configured for VPLS can use Ethernet VPLS or VLAN VPLS.

To configure the encapsulation on a Gigabit Ethernet IQ or Gigabit Ethernet physical interface, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level, specifying **flexible-ethernet-services**, **vlan-ccc**, or **vlan-vpls**:

```
[edit interfaces interface-name]
encapsulation (flexible-ethernet-services | vlan-ccc | vlan-vpls);
```

To configure the encapsulation on an aggregated Ethernet interface, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level, specifying **flexible-ethernet-services**, **ethernet-vpls**, or **vlan-vpls**:

```
[edit interfaces interface-name]
encapsulation (flexible-ethernet-services | ethernet-vpls | vlan-vpls);
```

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.

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

## 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-ccc;
    unit 0 {
        vlan-id 100;
    }
}
```

- Related Documentation**
- [802.1Q VLANs Overview on page 3](#)
  - *Configuring VPLS Interface Encapsulation*
  - *Ethernet Interfaces*

---

## Configuring Extended VLAN Encapsulation

Gigabit Ethernet, 4-port Fast Ethernet, MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, 10-Gigabit Ethernet, and aggregated Ethernet interfaces with VLAN tagging enabled can use extended VLAN CCC or VLAN VPLS, which allow 802.1Q tagging. To configure the encapsulation on a physical interface, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level, specifying **extended-vlan-ccc** or **extended-vlan-vpls**:

```
[edit interfaces interface-name]
encapsulation (extended-vlan-ccc | extended-vlan-vpls);
```

For extended VLAN CCC and extended VLAN VPLS encapsulation, all VLAN IDs 1 and higher are valid. VLAN ID 0 is reserved for tagging the priority of frames.



**NOTE:** For extended VLAN CCC, the VLAN IDs on ingress and egress interfaces must be the same. For back-to-back connections, all VLAN IDs must be the same.

---

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

Configure extended VLAN CCC encapsulation on Gigabit Ethernet ingress and egress interfaces:

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

```

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

```

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

Configure extended VLAN VPLS encapsulation on an aggregated Ethernet interface:

```

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

```

- Related Documentation**
- [802.1Q VLANs Overview on page 3](#)
  - [Ethernet Interfaces](#)

### Guidelines for Configuring VLAN ID List-Bundled Logical Interfaces That Connect CCCs

For MX Series routers, you can bind a list of VLAN IDs to a logical interface, configure a Layer 2 VPN routing instance or Layer 2 circuit on the logical interface, and then use the logical interface to configure a circuit cross-connect (CCC) to another Layer 2 VPN routing instance or Layer 2 circuit.

A CCC allows you to configure transparent connections between two circuits so that packets from the source circuit are delivered to the destination circuit with, at most, the Layer 2 address being changed. You configure a CCC by connecting circuit interfaces of the same type. For more information, see *Circuit and Translational Cross-Connects Overview*.



**NOTE:** The Junos OS supports binding of Ethernet logical interfaces to lists of VLAN IDs on MX Series routers only. For all other routers, you can bind an Ethernet logical interface to only a single VLAN ID or to a single range of VLAN IDs.

The following configuration guidelines apply to bundling lists of VLAN IDs to Ethernet logical interfaces used to configure CCCs:

- [Guidelines for Configuring Physical Link-Layer Encapsulation to Support CCCs on page 22](#)
- [Guidelines for Configuring Logical Link-Layer Encapsulation to Support CCCs on page 22](#)

## Guidelines for Configuring Physical Link-Layer Encapsulation to Support CCCs

To enable a physical interface to support VLAN-bundled logical interfaces that you will use to configure a CCC, you must specify one of the following physical link-layer encapsulation types as the value of the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level:

**[edit interfaces *interface-name*]**

**encapsulation** (extended-vlan-ccc | flexible-ethernet-services);

- **extended-vlan-ccc**—For Ethernet interfaces with standard TPID tagging.
- **flexible-ethernet-services**—For supported Gigabit Ethernet interfaces for which you want to configure multiple per-unit Ethernet encapsulations.

For more information about configuring the encapsulation on a physical interface, see *Configuring Interface Encapsulation on Physical Interfaces*.

## Guidelines for Configuring Logical Link-Layer Encapsulation to Support CCCs

For VLAN-bundled logical interfaces that you use to configure a CCC, specific logical link-layer encapsulation types are used inside the circuits themselves.

[Table 6 on page 22](#) describes the logical link-layer encapsulation types used within circuits connected using VLAN-bundled logical interfaces of the same type.

**Table 6: Encapsulation Inside Circuits CCC-Connected by VLAN-Bundled Logical Interfaces**

Encapsulation Inside the Circuit	Layer 2 Circuit Joined by Configuring an Interface-to-Interface CCC Connection	
	Layer 2 VPN Routing Instance	Layer 2 Circuit
Syntax	encapsulation-type (ethernet   ethernet-vlan);	encapsulation vlan-ccc;
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> protocols l2vpn], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn]	[edit interfaces <i>ethernet-interface-name</i> unit <i>logical-unit-number</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>ethernet-interface-name</i> unit <i>logical-unit-number</i> ]
Usage Guidelines	See the <i>Junos OS VPNs Library for Routing Devices</i> .	See <i>Configuring Interface Encapsulation on Logical Interfaces, Circuit and Translational Cross-Connects Overview</i> , and <i>Defining the Encapsulation for Switching Cross-Connects</i> .
For a Single-Tag Logical Interface	The MX Series router automatically uses <b>ethernet</b> as the Layer 2 protocol used to encapsulate incoming traffic. Although the connection spans multiple VLANs, the VLANs are bundled and therefore can be encapsulated as a single VLAN.  <b>NOTE:</b> With <b>ethernet</b> encapsulation, the circuit signal processing does not check that the VLAN ID list is the same at both ends of the CCC connection.	Configure the MX Series router to use <b>vlan-ccc</b> as the logical link-layer encapsulation type.

**Table 6: Encapsulation Inside Circuits CCC-Connected by VLAN-Bundled Logical Interfaces (*continued*)**

Encapsulation Inside the Circuit	Layer 2 Circuit Joined by Configuring an Interface-to-Interface CCC Connection	
	Layer 2 VPN Routing Instance	Layer 2 Circuit
For a Dual-Tag Logical Interface	<p>Configure the MX Series router to use <b>ethernet-vlan</b> as the Layer 2 protocol to encapsulate incoming traffic.</p> <p>With <b>ethernet-vlan</b> encapsulation, circuit signal processing checks that the VLAN ID list is the same at both ends of the CCC connection. If a VLAN ID list mismatch is detected, you can view the error condition in the <b>show interfaces</b> command output.</p>	The MX Series router automatically uses <b>vlan-ccc</b> as the logical link-layer encapsulation type, regardless of the value configured.

- Related Documentation**
- [802.1Q VLANs Overview on page 3](#)
  - [Binding VLAN IDs to Logical Interfaces on page 12](#)
  - [Defining the Encapsulation for Switching Cross-Connects](#)

## Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface

This topic describes how to configure a Layer 2 VPN routing instance on a logical interface bound to a list of VLAN IDs.

- [Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance on page 23](#)
- [Specifying the Interface Over Which VPN Traffic Travels to the CE Router on page 24](#)
- [Specifying the Interface to Handle Traffic for a CCC on page 24](#)

## Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance

To configure a VLAN-bundled logical interface, specify the list of VLAN IDs by including the **vlan-id-list** statement or the **vlan-tags** statement on a provider edge (PE) router:

```

interfaces {
  ethernet-interface-name {
    vlan-tagging; # Support single- or dual-tag logical interfaces
    flexible-vlan-tagging; # Support mixed tagging
    encapsulation (extended-vlan-ccc | flexible-ethernet-services);
    unit logical-unit-number {
      vlan-id-list [vlan-id vlan-id-vlan-id]; # For single-tag
      vlan-tags outer <tpid.>vlan-id inner-list [vlan-id vlan-id-vlan-id]; # For dual-tag
    }
    ...
  }
}

```

You can include the statements at the following hierarchy levels:

- **[edit]**

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

## Specifying the Interface Over Which VPN Traffic Travels to the CE Router

To configure a Layer 2 VPN routing instance on a PE router, include the **instance-type** statement and specify the value **l2vpn**. To specify an interface connected to the router, include the **interface** statement and specify the VLAN-bundled logical interface:

```
instance-type l2vpn;  
interface logical-interface-name;
```

You can include the statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

## Specifying the Interface to Handle Traffic for a CCC

To configure the VLAN-bundled logical interface as the interface to handle traffic for a circuit connected to the Layer 2 VPN routing instance, include the following statements:

```
protocols {  
  l2vpn {  
    (control-word | no-control-word);  
    encapsulation-type (ethernet | ethernet-vlan);  
    site site-name {  
      site-identifier identifier;  
      interface logical-interface-name { # VLAN-bundled logical interface  
        ... interface-options ...  
      }  
    }  
  }  
}
```

You can include the statements at the same hierarchy level at which you include the **instance-type l2vpn** and **interface *logical-interface-name*** statements:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

To enable a Layer 2 VPN routing instance on a PE router, include the **l2vpn** statement. For more information, see the *Junos OS VPNs Library for Routing Devices*.

The **encapsulation-type** statement specifies the Layer 2 protocol used for traffic from the customer edge (CE) router. If the Layer 2 VPN routing instance is being connected to a single-tag Layer 2 circuit, specify **ethernet** as the encapsulation type. If the Layer 2 VPN routing instance is being connected to a dual-tag Layer 2 circuit, specify **ethernet-vlan** as the encapsulation type.

To specify the interface to handle traffic for a circuit connected to the Layer 2 VPN routing instance, include the **interface** statement and specify the VLAN-bundled logical interface.

## Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance

To configure a VLAN-bundled logical interface, specify the list of VLAN IDs by including the **vlan-id-list** statement or the **vlan-tags** statement on a provider edge (PE) router:

```
interfaces {
  ethernet-interface-name {
    vlan-tagging; # Support single- or dual-tag logical interfaces
    flexible-vlan-tagging; # Support mixed tagging
    encapsulation (extended-vlan-ccc | flexible-ethernet-services);
    unit logical-unit-number {
      vlan-id-list [vlan-id vlan-id-vlan-id]; # For single-tag
      vlan-tags outer <tpid.>vlan-id inner-list [vlan-id vlan-id-vlan-id]; # For dual-tag
    }
    ...
  }
}
```

You can include the statements at the following hierarchy levels:

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

**Related Documentation**

- [802.1Q VLANs Overview on page 3](#)
- *Ethernet Interfaces*

## Specifying the Interface Over Which VPN Traffic Travels to the CE Router

To configure a Layer 2 VPN routing instance on a PE router, include the **instance-type** statement and specify the value **l2vpn**. To specify an interface connected to the router, include the **interface** statement and specify the VLAN-bundled logical interface:

```
instance-type l2vpn;
interface logical-interface-name;
```

You can include the statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

**Related Documentation**

- [802.1Q VLANs Overview on page 3](#)
- *Ethernet Interfaces*

## Specifying the Interface to Handle Traffic for a CCC

---

To configure the VLAN-bundled logical interface as the interface to handle traffic for a circuit connected to the Layer 2 VPN routing instance, include the following statements:

```
protocols {
  l2vpn {
    (control-word | no-control-word);
    encapsulation-type (ethernet | ethernet-vlan);
    site site-name {
      site-identifier identifier;
      interface logical-interface-name { # VLAN-bundled logical interface
        ... interface-options ...
      }
    }
  }
}
```

You can include the statements at the same hierarchy level at which you include the **instance-type l2vpn** and **interface *logical-interface-name*** statements:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

To enable a Layer 2 VPN routing instance on a PE router, include the **l2vpn** statement. For more information, see the *Junos OS VPNs Library for Routing Devices*.

The **encapsulation-type** statement specifies the Layer 2 protocol used for traffic from the customer edge (CE) router. If the Layer 2 VPN routing instance is being connected to a single-tag Layer 2 circuit, specify **ethernet** as the encapsulation type. If the Layer 2 VPN routing instance is being connected to a dual-tag Layer 2 circuit, specify **ethernet-vlan** as the encapsulation type.

To specify the interface to handle traffic for a circuit connected to the Layer 2 VPN routing instance, include the **interface** statement and specify the VLAN-bundled logical interface.

### Related Documentation

- [802.1Q VLANs Overview on page 3](#)
- [Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface on page 29](#)
- [Ethernet Interfaces](#)

## Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface

This topic describes how to configure a Layer 2 circuit on a logical interface bound to a list of VLAN IDs.

- [Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance on page 27](#)
- [Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit on page 27](#)

## Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance

To configure a VLAN-bundled logical interface, specify the list of VLAN IDs by including the **vlan-id-list** statement or the **vlan-tags** statement:

```
interfaces {
  ethernet-interface-name {
    vlan-tagging; # Support single- or dual-tag logical interfaces
    flexible-vlan-tagging; # Support mixed tagging
    encapsulation (extended-vlan-ccc | flexible-ethernet-services);
    unit logical-unit-number {
      encapsulation vlan-ccc; # Required for single-tag
      vlan-id-list [vlan-id vlan-id-vlan-id]; # For single-tag
      vlan-tags outer tpid.vlan-id inner-list [vlan-id vlan-id-vlan-id]; # For dual-tag
    }
    ...
  }
}
```

You can include the statements at the following hierarchy levels:

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

For a single-tag logical interface, include the **encapsulation** statement and specify **vlan-ccc** so that CCC circuit encapsulation is used inside the Layer 2 circuit.



**NOTE:** In the case of a dual-tag logical interface, the Junos OS automatically uses the **vlan-ccc** encapsulation type.

## Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit

To configure the VLAN-bundled logical interface as the interface to handle traffic for a circuit connected to the Layer 2 circuit, include the following statements:

```
l2circuit {
  neighbor address {
    interface logical-interface-name {
      virtual-circuit-id number;
      no-control-word;
    }
  }
}
```

```
}
}
```

You can include the statements at the following hierarchy levels:

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

To enable a Layer 2 circuit, include the **l2circuit** statement.

To configure the router as a neighbor for a Layer 2 circuit, specify the neighbor address using the **neighbor** statement.

To specify the interface to handle traffic for a circuit connected to the Layer 2 circuit, include the **interface** statement and specify the VLAN-bundled logical interface.

## Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance

To configure a VLAN-bundled logical interface, specify the list of VLAN IDs by including the **vlan-id-list** statement or the **vlan-tags** statement:

```
interfaces {
  ethernet-interface-name {
    vlan-tagging; # Support single- or dual-tag logical interfaces
    flexible-vlan-tagging; # Support mixed tagging
    encapsulation (extended-vlan-ccc | flexible-ethernet-services);
    unit logical-unit-number {
      encapsulation vlan-ccc; # Required for single-tag
      vlan-id-list [vlan-id vlan-id-vlan-id]; # For single-tag
      vlan-tags outer tpid.vlan-id inner-list [vlan-id vlan-id-vlan-id]; # For dual-tag
    }
    ...
  }
}
```

You can include the statements at the following hierarchy levels:

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

For a single-tag logical interface, include the **encapsulation** statement and specify **vlan-ccc** so that CCC circuit encapsulation is used inside the Layer 2 circuit.



**NOTE:** In the case of a dual-tag logical interface, the Junos OS automatically uses the **vlan-ccc** encapsulation type.

### Related Documentation

- [802.1Q VLANs Overview on page 3](#)
- [Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit on page 27](#)

- [Ethernet Interfaces](#)

## Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit

To configure the VLAN-bundled logical interface as the interface to handle traffic for a circuit connected to the Layer 2 circuit, include the following statements:

```
l2circuit {
  neighbor address {
    interface logical-interface-name {
      virtual-circuit-id number;
      no-control-word;
    }
  }
}
```

You can include the statements at the following hierarchy levels:

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

To enable a Layer 2 circuit, include the **l2circuit** statement.

To configure the router as a neighbor for a Layer 2 circuit, specify the neighbor address using the **neighbor** statement.

To specify the interface to handle traffic for a circuit connected to the Layer 2 circuit, include the **interface** statement and specify the VLAN-bundled logical interface.

### Related Documentation

- [802.1Q VLANs Overview on page 3](#)
- [Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface on page 31](#)
- [Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance on page 27](#)
- [Ethernet Interfaces](#)

## Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface

The following configuration shows that the single-tag logical interface **ge-1/0/5.0** bundles a list of VLAN IDs, and the logical interface **ge-1/1/1.0** supports IPv4 traffic using IP address 10.30.1.130 and can participate in an MPLS path.

```
[edit interfaces]
ge-1/0/5 {
  vlan-tagging;
  encapsulation extended-vlan-ccc;
  unit 0 { # VLAN-bundled logical interface
    vlan-id-list [513 516 520-525];
  }
}
```

```
ge-1/1/1 {  
  unit 0 {  
    family inet {  
      address 10.30.1.1/30;  
    }  
    family mpls;  
  }  
}
```

The following configuration shows the type of traffic supported on the Layer 2 VPN routing instance:

```
[edit protocols]  
rsvp {  
  interface all;  
  interface lo0.0;  
}  
mpls {  
  label-switched-path lsp {  
    to 10.255.69.128;  
  }  
  interface all;  
}  
bgp {  
  group g1 {  
    type internal;  
    local-address 10.255.69.96;  
    family l2vpn {  
      signaling;  
    }  
    neighbor 10.255.69.128;  
  }  
}  
ospf {  
  traffic-engineering;  
  area 0.0.0.0 {  
    interface lo0.0;  
    interface ge-1/1/1.0;  
  }  
}
```

The following configuration shows that the VLAN-bundled logical interface is the interface over which VPN traffic travels to the CE router and handles traffic for a CCC to which the VPN connects.

```
[edit routing-instances]  
red {  
  instance-type l2vpn;  
  interface ge-1/0/5.0; # VLAN-bundled logical interface  
  route-distinguisher 10.255.69.96:100;  
  vrf-target target:1:1;  
  protocols {  
    l2vpn {  
      encapsulation-type ethernet; # For single-tag VLAN logical interface  
      site CE_ultima {  
        site-identifier 1;  
        interface ge-1/0/5.0;  
      }  
    }  
  }  
}
```

```

    }
  }
}

```



**NOTE:** Because the VLAN-bundled logical interface supports single-tag frames, Ethernet is the Layer 2 protocol used to encapsulate incoming traffic. Although the connection spans multiple VLANs, the VLANs are bundled and therefore can be encapsulated as a single VLAN.

However, with Ethernet encapsulation, the circuit signal processing does not check that the VLAN ID list is the same at both ends of the CCC connection.

- Related Documentation**
- [802.1Q VLANs Overview on page 3](#)
  - *Ethernet Interfaces*

## Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface

The following configuration shows that the single-tag logical interface **ge-1/0/5.0** bundles a list of VLAN IDs, and the logical interface **ge-1/1/1.0** supports IPv4 traffic using IP address 10.30.1.1/30 and can participate in an MPLS path.

```

[edit interfaces]
ge-1/0/5 {
  vlan-tagging;
  encapsulation extended-vlan-ccc;
  unit 0 { # VLAN-bundled logical interface
    vlan-id-list [513 516 520-525];
  }
}
ge-1/1/1 {
  unit 0 {
    family inet {
      address 10.30.1.1/30;
    }
    family mpls;
  }
}

```

The following configuration shows the type of traffic supported on the Layer 2 VPN routing instance, and shows that the VLAN-bundled logical interface handles traffic for a CCC to which the Layer 2 circuit connects:

```

[edit protocols]
rsvp {
  interface all;
  interface lo0.0;
}
mpls {
  label-switched-path lsp {
    to 10.255.69.128;
  }
}

```

```
    interface all;
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface lo0.0;
      interface ge-1/1/1.0;
    }
  }
  ldp {
    interface ge-1/1/1.0;
    interface ge-1/0/5.0; # VLAN-bundled logical interface
    interface lo0.0;
  }
  l2circuit {
    neighbor 10.255.69.128 {
      interface ge-1/0/5.0 { # VLAN-bundled logical interface
        virtual-circuit-id 3;
        no-control-word;
      }
    }
  }
}
```

- Related Documentation**
- [802.1Q VLANs Overview on page 3](#)
  - [Ethernet Interfaces](#)

---

## 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 bridge domain. A logical interface configured to accept untagged packets is called an *access interface* or *access port*. Access interface configuration is supported on MX Series routers only.

**interface-mode** access;

You can include this statement at the following hierarchy levels:

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

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 bridge domain that is configured with the matching VLAN ID.

### Example: Configuring a Logical Interface for Access Mode

The following example configures a logical interface as an access port with a VLAN ID of 20:

```
[edit interfaces ge-1/2/0]
unit 0 {
  family bridge {
```

```

        interface-mode access;
        vlan-id 20;
    }
}

```

- Related Documentation**
- [802.1Q VLANs Overview on page 3](#)
  - [Ethernet Interfaces](#)

## Configuring a Logical Interface for Trunk Mode

As an alternative to configuring a logical interface for each VLAN, enterprise network administrators can configure a single logical interface to accept untagged packets or packets tagged with any VLAN ID specified in a list of VLAN IDs. Using a VLAN ID list conserves switch resources and simplifies configuration. A logical interface configured to accept packets tagged with any VLAN ID specified in a list is called a *trunk interface* or *trunk port*. Trunk interface configuration is supported on MX Series routers only. Trunk interfaces support integrated routing and bridging (IRB).

To configure a logical interface to accept any packet tagged with a VLAN ID that matches the list of VLAN IDs, include the **interface-mode** statement and specify the **trunk** option:

```
interface-mode trunk;
```

You can include this statement at the following hierarchy levels:

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

- Related Documentation**
- [802.1Q VLANs Overview on page 3](#)
  - [Ethernet Interfaces](#)

## Configuring the VLAN ID List for a Trunk Interface

To configure the list of VLAN IDs to be accepted by the trunk port, include the **vlan-id-list** statement and specify the list of VLAN IDs. You can specify individual VLAN IDs with a space separating the ID numbers, specify a range of VLAN IDs with a dash separating the ID numbers, or specify a combination of individual VLAN IDs and a range of VLAN IDs.

```
vlan-id-list [number number-number];
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family bridge interface-mode trunk]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family bridge interface-mode trunk]

When a packet is received that is tagged with a VLAN ID specified in the trunk interface list of VLAN IDs, the packet is accepted and forwarded within the bridge domain that is configured with the matching VLAN ID.

When a packet is received that is tagged with a VLAN ID not specified in the trunk interface list of VLAN IDs, the native VLAN ID is pushed in front of the existing VLAN tag or tags and the packet is forwarded within the bridge domain that is configured with the matching VLAN ID.

When an untagged packet is received on a trunk interface, the native VLAN ID is added to the packet and the packet is forwarded within the bridge domain that is configured with the matching VLAN ID.

A bridge domain configured with a matching VLAN ID must be configured before the trunk interface is configured. To learn more about configuring bridge domains, see the *Junos Routing Protocols Configuration Guide*.

- Related Documentation**
- [802.1Q VLANs Overview on page 3](#)
  - *Ethernet Interfaces*

---

## Configuring a Trunk Interface on a Bridge Network

---

On MX Series routers, you can configure a trunk interface on a bridge network.

The following output sample shows trunk port configuration on a bridge network:

```
user@host# run show interfaces
ge-0/0/0 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 1;
    }
}
ge-2/0/0 {
    unit 0 {
        family bridge {
            interface-mode trunk;
            vlan-id-list 1-200;
        }
    }
}
ge-2/0/1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 1;
    }
}
```

If you want **igmp-snooping** to be functional for a bridge domain, then you should not configure **interface-mode** and **irb** for that bridge domain. Such a configuration commit

succeeds, but IGMP snooping is not functional, and a message informing the same is displayed as shown after the sample configuration below:

```
user@host# run show configuration
interfaces {
  ge-5/1/1 {
    flexible-vlan-tagging;
    native-vlan-id 1;
    unit 0 {
      family bridge {
        interface-mode trunk;
        vlan-id-list 401;
      }
    }
  }
  irb {
    unit 401 {
      family inet {
        address 192.168.2.2/27;
      }
    }
  }
}
protocols {
  igmp {
    interface all;
  }
}
bridge-domains {
  VLAN-401 {
    vlan-id 401;
    routing-interface irb.401;
    protocols {
      igmp-snooping;
    }
  }
}

user@host# commit
[edit bridge-domains]
'VLAN-401'
IGMP Snooping not supported with IRB and trunk mode interface ge-5/1/1.0
commit complete
```

To achieve IGMP snooping for a bridge domain, you should use such a configuration as shown in the following example:

```
user@host# run show configuration
interfaces {
  ge-0/0/1 {
    flexible-vlan-tagging;
    native-vlan-id 1;
    encapsulation flexible-ethernet-services;
    unit 0 {
      encapsulation vlan-bridge;
      vlan-id 401;
    }
  }
  irb {
    unit 401 {
      family inet {
        address 192.168.2.2/27;
      }
    }
  }
}
```

```
    }
  }
}
protocols {
  igmp {
    interface all;
  }
}
bridge-domains {
  VLAN-401 {
    vlan-id 401;
    interface ge-0/0/1.0;
    routing-interface irb.401;
    protocols {
      igmp-snooping;
    }
  }
}

user@host# commit
commit complete
```

- Related Documentation**
- [802.1Q VLANs Overview on page 3](#)
  - [interface-mode on page 77](#)
  - [Ethernet Interfaces](#)

---

## Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers

---

This topic describes how to configure flexible VLAN tagging on PTX Series Packet Transport Routers. In addition to VLAN tagging and stacked VLAN tagging, you can configure a port for flexible tagging. With flexible VLAN tagging, you can configure two logical interfaces on the same Ethernet port, one with single-tag framing and one with dual-tag framing.

To configure mixed tagging, include the **flexible-vlan-tagging** statement at the **[edit interfaces et-fpc/pic/port]** hierarchy level. You must also include the **vlan-tags** statement with **inner** and **outer** options or the **vlan-id** statement at the **[edit interfaces et-fpc/pic/port unit logical-unit-number]** hierarchy level:

```
[edit interfaces et-fpc/pic/port]
flexible-vlan-tagging;
unit logical-unit-number {
  vlan-id number;
}
unit logical-unit-number {
  vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
}
```

- Related Documentation**
- [Enabling VLAN Tagging on page 9](#)

## CHAPTER 3

# Network Interfaces Configuration Statements and Hierarchy

- [edit interfaces] Hierarchy Level on page 37
- [edit logical-systems] Hierarchy Level on page 53

### [edit interfaces] Hierarchy Level

---

The statements at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level can also be configured at the [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*] hierarchy level.



**NOTE:** The *accounting-profile* statement is an exception to this rule. The *accounting-profile* statement can be configured at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level, but it cannot be configured at the [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*] hierarchy level.

```
interfaces {
  traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable |
      no-world-readable> ;
    flag flag <disable>;
  }
  interface-name {
    accounting-profile name;
    aggregated-ether-options {
      (flow-control | no-flow-control);
      lacp {
        (active | passive);
        link-protection {
          disable;
          (revertive | non-revertive);
          periodic interval;
          system-priority priority;
        }
      }
      link-protection;
      link-speed speed;
    }
  }
}
```

```

(loopback | no-loopback);
mc-ae{
  chassis-id chassis-id;
  mc-ae-id mc-ae-id;
  mode (active-active | active-standby);
  redundancy-group group-id;
  status-control (active | standby);
}
minimum-links number;
source-address-filter {
  mac-address;
}
(source-filtering | no-source-filtering);
}
shared-scheduler;
aggregated-sonet-options {
  link-speed speed | mixed;
  minimum-links number;
}
atm-options {
  cell-bundle-size cells;
  ilmi;
  linear-red-profiles profile-name {
    high-plp-max-threshold percent;
    low-plp-max-threshold percent;
    queue-depth cells high-plp-threshold percent low-plp-threshold percent;
  }
}
mpls {
  pop-all-labels {
    required-depth number;
  }
}
pic-type (atm1 | atm2);
plp-to-clp;
promiscuous-mode {
  vpi vpi-identifier;
}
scheduler-maps map-name {
  forwarding-class class-name {
    epd-threshold cells plp1 cells;
    linear-red-profile profile-name;
    priority (high | low);
    transmit-weight (cells number | percent number);
  }
  vc-cos-mode (alternate | strict);
}
use-null-cw;
vpi vpi-identifier {
  maximum-vcs maximum-vcs;
  oam-liveness {
    down-count cells;
    up-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;
    }
}
clocking clock-source;
damping {
    enable;
    half-life seconds;
    max-suppress seconds;
    reuse number;
    suppress number;
}
data-input (system | interface interface-name);
dce;
serial-options {
    clock-rate rate;
    clocking-mode (dce | internal | loop);
    control-polarity (negative | positive);
    cts-polarity (negative | positive);
    dcd-polarity (negative | positive);
    dce-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    dsr-polarity (negative | positive);
    dte-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    dtr-circuit (balanced | unbalanced);
    dtr-polarity (negative | positive);
    encoding (nrz | nrzi);
    indication-polarity (negative | positive);
    line-protocol protocol;
    loopback mode;
    rts-polarity (negative | positive);
    tm-polarity (negative | positive);
    transmit-clock invert;
}
description text;

```

```
alias alias-name;  
dialer-options {  
    pool pool-name <priority priority>;  
}  
disable;  
ds0-options {  
    bert-algorithm algorithm;  
    bert-error-rate rate;  
    bert-period seconds;  
    byte-encoding (nx56 | nx64);  
    fcs (16 | 32);  
    idle-cycle-flag (flags | ones);  
    invert-data;  
    loopback payload;  
    start-end-flag (filler | shared);  
}  
e1-options {  
    bert-error-rate rate;  
    bert-period seconds;  
    fcs (16 | 32);  
    framing (g704 | g704-no-crc4 | unframed);  
    idle-cycle-flag (flags | ones);  
    invert-data;  
    loopback (local | remote);  
    start-end-flag (filler | shared);  
    timeslots time-slot-range;  
}  
e3-options {  
    atm-encapsulation (direct | plcp);  
    bert-algorithm algorithm;  
    bert-error-rate rate;  
    bert-period seconds;  
    framing feet;  
    compatibility-mode (digital-link | kentrox | larscom) <subrate value>;  
    fcs (16 | 32);  
    framing (g.751 | g.832);  
    idle-cycle-flag (filler | shared);  
    invert-data;  
    loopback (local | remote);  
    (payload-scrambler | no-payload-scrambler);  
    start-end-flag (filler | shared);  
    (unframed | no-unframed);  
}  
encapsulation type;  
es-options {  
    backup-interface es-fpc/pic/port;  
}  
fastether-options {  
    802.3ad aex;  
    (flow-control | no-flow-control);  
    ignore-l3-incompletes;  
    ingress-rate-limit rate;  
    (loopback | no-loopback);  
    mpls {  
        pop-all-labels {  
            required-depth number;
```

```

    }
  }
  source-address-filter {
    mac-address;
  }
  (source-filtering | no-source-filtering);
}
flexible-vlan-tagging;
gigether-options {
  802.3ad aex;
  (asynchronous-notification | no-asynchronous-notification);
  (auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
    local-interface-offline>;
  auto-reconnect seconds;
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  (loopback | no-loopback);
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  no-auto-mdix;
  source-address-filter {
    mac-address;
  }
  (source-filtering | no-source-filtering);
  ethernet-switch-profile {
    (mac-learn-enable | no-mac-learn-enable);
    tag-protocol-id [ tpids ];
    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;
    }
  }
}
}
}
}

```

```
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time up milliseconds down milliseconds;
ima-group-options {
    differential-delay number;
    frame-length (32 | 64 | 128 | 256);
    frame-synchronization {
        alpha number;
        beta number;
        gamma number;
    }
    minimum-links number;
    symmetry (symmetrical-config-and-operation |
        symmetrical-config-asymmetrical-operation);
    test-procedure {
        ima-test-start;
        ima-test-stop;
        interface name;
        pattern number;
        period number;
    }
    transmit-clock (common | independent);
    version (1.0 | 1.1);
}
ima-link-options group-id group-id;
interface-set interface-set-name {
    interface ethernet-interface-name {
        (unit unit-number | vlan-tags-outer vlan-tag);
    }
    interface interface-name {
        (unit unit-number);
    }
}
isdn-options {
    bchannel-allocation (ascending | descending);
    calling-number number;
    pool pool-name <priority priority>;
    spid1 spid-string;
    spid2 spid-string;
    static-tei-val value;
    switch-type (att5e | etsi | nil | ntdms100 | ntt);
    t310 seconds;
    tei-option (first-call | power-up);
}
keepalives <down-count number> <interval seconds> <up-count number>;
link-mode mode;
lmi {
    lmi-type (ansi | itu | c-lmi);
    n391dte number;
    n392dce number;
    n392dte number;
    n393dce number;
    n393dte number;
    t391dte seconds;
    t392dce seconds;
}
lsq-failure-options {
```

```

no-termination-request;
[ trigger-link-failure interface-name ];
}
mac mac-address;
mlfr-uni-nni-bundle-options {
    acknowledge-retries number;
    acknowledge-timer milliseconds;
    action-red-differential-delay (disable-tx | remove-link);
    drop-timeout milliseconds;
    fragment-threshold bytes;
    cisco-interoperability send-lip-remove-link-for-link-reject;
    hello-timer milliseconds;
    link-layer-overhead percent;
    lmi-type (ansi | itu | c-lmi);
    minimum-links number;
    mrru bytes;
    n391 number;
    n392 number;
    n393 number;
    red-differential-delay milliseconds;
    t391 seconds;
    t392 seconds;
    yellow-differential-delay milliseconds;
}
modem-options {
    dialin (console | routable);
    init-command-string initialization-command-string;
}
mtu bytes;
multi-chassis-protection {
    peer a.b.c.d {
        interface interface-name;
    }
}
multiservice-options {
    (core-dump | no-core-dump);
    (syslog | no-syslog);
}
native-vlan-id number;
no-gratuitous-arp-request;
no-keepalives;
no-partition {
    interface-type type;
}
no-vpivci-swapping;
otn-options {
    fec (efec | gfec | none);
    (laser-enable | no-laser-enable);
    (line-loopback | no-line-loopback);
    pass-thru;
    rate (fixed-stuff-bytes | no-fixed-stuff-bytes | pass-thru);
    transmit-payload-type number;
    trigger (oc-lof | oc-lom | oc-los | oc-wavelength-lock | odu-ais | odu-bbe-th | odu-bdi
        | odu-es-th | odu-lck | odu-oci | odu-sd | odu-ses-th | odu-ttim | odu-uas-th |
        opu-ptm | otu-ais | otu-bbe-th | otu-bdi | otu-es-th | otu-fec-deg | otu-fec-exe |
        otu-iae | otu-sd | otu-ses-th | otu-ttim | otu-uas-th);
}

```

```
tti;
}
optics-options {
    wavelength nm;
    alarm alarm-name {
        (syslog | link-down);
    }
    warning warning-name {
        (syslog | link-down);
    }
}
partition partition-number oc-slice oc-slice-range interface-type type;
timeslots time-slot-range;
passive-monitor-mode;
per-unit-scheduler;
ppp-options {
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
    dynamic-profile profile-name;
    no-termination-request;
    pap {
        access-profile name;
        local-name name;
        local-password password;
        compression;
    }
}
psn-vcip psn-vci-identifier;
psn-vpip psn-vpi-identifier;
receive-bucket {
    overflow (discard | tag);
    rate percentage;
    threshold bytes;
}
redundancy-options {
    priority sp-fpc/pic/port;
    secondary sp-fpc/pic/port;
    hot-standby;
}
satop-options {
    payload-size n;
}
schedulers number;
serial-options {
    clock-rate rate;
    clocking-mode (dce | internal | loop);
    control-polarity (negative | positive);
    cts-polarity (negative | positive);
```

```

dcd-polarity (negative | positive);
dce-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
}
dsr-polarity (negative | positive);
dte-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
}
dtr-circuit (balanced | unbalanced);
dtr-polarity (negative | positive);
encoding (nrz | nrzi);
indication-polarity (negative | positive);
line-protocol protocol;
loopback mode;
rts-polarity (negative | positive);
tm-polarity (negative | positive);
transmit-clock invert;
}
services-options {
    inactivity-timeout seconds;
    open-timeout seconds;
    session-limit {
        maximum number;
        rate new-sessions-per-second;
    }
    syslog {
        host hostname {
            facility-override facility-name;
            log-prefix prefix-number;
            services priority-level;
        }
    }
}
shdsl-options {
    annex (annex-a | annex-b);
    line-rate line-rate;
    loopback (local | remote);
    snr-margin {
        current margin;
        snext margin;
    }
}

```

```

    }
  }
  sonet-options {
    aggregate asx;
    aps {
      advertise-interval milliseconds;
      annex-b;
      authentication-key key;
      fast-aps-switch;
      force;
      hold-time milliseconds;
      lockout;
      neighbor address;
      paired-group group-name;
      preserve-interface;
      protect-circuit group-name;
      request;
      revert-time seconds;
      switching-mode (bidirectional | unidirectional);
      working-circuit group-name;
    }
    bytes {
      c2 value;
      e1-quiet value;
      f1 value;
      f2 value;
      s1 value;
      z3 value;
      z4 value;
    }
    fcs (16 | 32);
    loopback (local | remote);
    mpls {
      pop-all-labels {
        required-depth number;
      }
    }
    path-trace trace-string;
    (payload-scrambler | no-payload-scrambler);
    rfc-2615;
    trigger {
      defect ignore;
      hold-time up milliseconds down milliseconds;
    }
    vtmapping (itu-t | klm);
    (z0-increment | no-z0-increment);
  }
  speed (10m | 100m | 1g | oc3 | oc12 | oc48);
  stacked-vlan-tagging;
  switch-options {
    switch-port port-number {
      (auto-negotiation | no-auto-negotiation);
      speed (10m | 100m | 1g);
      link-mode (full-duplex | half-duplex);
    }
  }
}

```

```

t1-options {
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout value;
    byte-encoding (nx56 | nx64);
    crc-major-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5);
    crc-minor-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5 | 5e-6 | 1e-6);
    fcs (16 | 32);
    framing (esf | sf);
    idle-cycle-flag (flags | ones);
    invert-data;
    line-encoding (ami | b8zs);
    loopback (local | payload | remote);
    remote-loopback-respond;
    start-end-flag (filler | shared);
    timeslots time-slot-range;
}
t3-options {
    atm-encapsulation (direct | plcp);
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout feet;
    (cbit-parity | no-cbit-parity);
    compatibility-mode (adtran | digital-link | kentrox | larscom | verilink) <subrate
        value>;
    fcs (16 | 32);
    (feac-loop-respond | no-feac-loop-respond);
    idle-cycle-flag value;
    (long-buildout | no-long-buildout);
    (loop-timing | no-loop-timing);
    loopback (local | payload | remote);
    (mac | no-mac);
    (payload-scrambler | no-payload-scrambler);
    start-end-flag (filler | shared);
}
traceoptions {
    flag flag <flag-modifier> <disable>;
}
transmit-bucket {
    overflow discard;
    rate percentage;
    threshold bytes;
}
(traps | no-traps);
unidirectional;
vlan-tagging;
vlan-vci-tagging;
unit logical-unit-number {
    accept-source-mac {
        mac-address mac-address {
            policer {
                input cos-policer-name;
                output cos-policer-name;
            }
        }
    }
}

```

```
    }
  }
  accounting-profile name;
  advisory-options {
    downstream-rate rate;
    upstream-rate rate;
  }
  allow-any-vci;
  atm-scheduler-map (map-name | default);
  backup-options {
    interface interface-name;
  }
  bandwidth rate;
  cell-bundle-size cells;
  clear-dont-fragment-bit;
  compression {
    rtp {
      f-max-period number;
      maximum-contexts number <force>;
      queues [ queue-numbers ];
      port {
        minimum port-number;
        maximum port-number;
      }
    }
  }
}
compression-device interface-name;
copy-tos-to-outer-ip-header;
demux-destination family;
demux-source family;
demux-options {
  underlying-interface interface-name;
}
description text;
alias alias-name;
interface {
  l2tp-interface-id name;
  (dedicated | shared);
}
dialer-options {
  activation-delay seconds;
  callback;
  callback-wait-period time;
  deactivation-delay seconds;
  dial-string [ dial-string-numbers ];
  idle-timeout seconds;
  incoming-map {
    caller (caller-id | accept-all);
    initial-route-check seconds;
    load-interval seconds;
    load-threshold percent;
    pool pool-name;
    redial-delay time;
    watch-list {
      [ routes ];
    }
  }
}
```

```

    }
}
disable;
disable-mlppp-inner-ppp-pfc;
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;
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;
}
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 {
    down-count cells;
    up-count cells;
}
oam-period (seconds | disable);
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;
}
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;
    pap;
    default-pap-password password;
    local-name name;
    local-password password;
    passive;
}
dynamic-profile profile-name;
lcp-max-conf-req number;
lcp-restart-timer milliseconds;
loopback-clear-timer seconds;
ncp-max-conf-req number;
ncp-restart-timer milliseconds;
}
pppoe-options {
    access-concentrator name;
    auto-reconnect seconds;
    (client | server);
    service-name name;
    underlying-interface interface-name;
}
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;
vci-range start start-vci end end-vci;
vpi vpi-identifier;
vlan-id number;
vlan-id-list [vlan-id vlan-id-vlan-id];
vlan-id-range number-number;
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
vlan-tags-outer tpid.vlan-id inner-list [vlan-id vlan-id-vlan-id];

```

```

family family {
  accounting {
    destination-class-usage;
    source-class-usage {
      direction;
    }
  }
  access-concentrator name;
  address address {
    destination address;
  }
  bundle ml-fpc/pic/port | ls-fpc/pic/port);
  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 ];
    }
  }
  ipsec-sa sa-name;
  keep-address-and-control;
  max-sessions number;
  max-sessions-vs-a-ignore;
  mtu bytes;
  multicast-only;
  negotiate-address;
  no-redirects;
  policer {
    arp policer-template-name;
    disable-arp-policer
    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-names <service-filter filter-name>;
    }
}
service-name-table table-name;
short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
    maximum-seconds>;
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
(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 | preferred-source-address address>;
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 {
        bandwidth-threshold bits-per-second {
            priority;

```

```

    }
    interface 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*
- *Junos OS Network Interfaces Library for Routing Devices*

## [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 {

```





## CHAPTER 4

# Statement Summary

- [address](#) on page 60
- [encapsulation \(Logical Interface\)](#) on page 62
- [encapsulation \(Physical Interface\)](#) on page 66
- [family](#) on page 71
- [fast-aps-switch](#) on page 75
- [flexible-vlan-tagging](#) on page 76
- [interface-mode](#) on page 77
- [native-vlan-id](#) on page 79
- [stacked-vlan-tagging](#) on page 80
- [unit](#) on page 81
- [vlan-id \(VLAN ID to Be Bound to a Logical Interface\)](#) on page 87
- [vlan-id \(VLAN ID to Rewrite\)](#) on page 88
- [vlan-id-list \(Interface in Bridge Domain\)](#) on page 89
- [vlan-id-range](#) on page 90
- [vlan-tagging](#) on page 91
- [vlan-tags \(Dual-Tagged Logical Interface\)](#) on page 92
- [vlan-tags \(Stacked VLAN Tags\)](#) on page 94
- [vlan-id-list \(Ethernet VLAN Circuit\)](#) on page 96

## address

```

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

```

**Hierarchy Level** [edit interfaces *interface-name* unit *logical-unit-number* family *family*],  
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*  
 family *family*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure the interface address.

**Options** *address*—Address of the interface.

- In Junos OS Release 13.3 and later, when you configure an IPv6 host address and an IPv6 subnet address on an interface, the commit operation fails.
- In releases earlier than Junos OS Release 13.3, when you use the same configuration on an interface, the commit operation succeeds, but only one of the IPv6 addresses that was entered is assigned to the interface. The other address is not applied.

The remaining statements are explained separately.



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

**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*
- *Junos OS Administration Library for Routing Devices*
- *family*
- *negotiate-address*
- *unnumbered-address (Ethernet)*

## encapsulation (Logical Interface)

<b>Syntax</b>	encapsulation (atm-ccc-cell-relay   atm-ccc-vc-mux   atm-cisco-nlpid   atm-mlppp-llc   atm-nlpid   atm-ppp-llc   atm-ppp-vc-mux   atm-snap   atm-tcc-snap   atm-tcc-vc-mux   atm-vc-mux   ether-over-atm-llc   ether-vpls-over-atm-llc   ether-vpls-over-fr   ether-vpls-over-ppp   ethernet   ethernet-ccc   ethernet-vpls   ethernet-vpls-fr   frame-relay-ccc   frame-relay-ether-type   frame-relay-ether-type-tcc   frame-relay-ppp   frame-relay-tcc   gre-fragmentation   multilink-frame-relay-end-to-end   multilink-ppp   ppp-over-ether   ppp-over-ether-over-atm-llc   vlan-bridge   vlan-ccc   vlan-vci-ccc   vlan-tcc   vlan-vpls);
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ], [edit interfaces <i>rlsq number</i> <b>unit</b> <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers ( <b>ethernet</b> , <b>vlan-ccc</b> , and <b>vlan-tcc</b> options only). Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Access Routers. Only the <b>atm-ccc-cell-relay</b> and <b>atm-ccc-vc-mux</b> options are supported on ACX Series routers.
<b>Description</b>	Configure a logical link-layer encapsulation type.
<b>Options</b>	<p><b>atm-ccc-cell-relay</b>—Use ATM cell-relay encapsulation.</p> <p><b>atm-ccc-vc-mux</b>—Use ATM virtual circuit (VC) multiplex encapsulation on CCC circuits. When you use this encapsulation type, you can configure the <b>ccc</b> family only.</p> <p><b>atm-cisco-nlpid</b>—Use Cisco ATM network layer protocol identifier (NLPID) encapsulation. When you use this encapsulation type, you can configure the <b>inet</b> family only.</p> <p><b>atm-mlppp-llc</b>—For ATM2 IQ interfaces only, use Multilink Point-to-Point (MLPPP) over AAL5 LLC. For this encapsulation type, your router must be equipped with a Link Services or Voice Services PIC. MLPPP over ATM encapsulation is not supported on ATM2 IQ OC48 interfaces.</p> <p><b>atm-nlpid</b>—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the <b>inet</b> family only.</p> <p><b>atm-ppp-llc</b>—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over AAL5 LLC encapsulation.</p> <p><b>atm-ppp-vc-mux</b>—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over ATM AAL5 multiplex encapsulation.</p> <p><b>atm-snap</b>—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM subnetwork attachment point (SNAP) encapsulation.</p> <p><b>atm-tcc-snap</b>—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.</p>

**atm-tcc-vc-mux**—Use ATM VC multiplex encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

**atm-vc-mux**—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM VC multiplex encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

**ether-over-atm-llc**—(All IP interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) For interfaces that carry IP traffic, use Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces.

**ether-vpls-over-atm-llc**—For ATM2 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.

**ether-vpls-over-fr**—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Frame Relay encapsulation to support Bridged Ethernet over Frame Relay encapsulated TDM interfaces for VPLS applications, per RFC 2427, *Multiprotocol Interconnect over Frame Relay*.



**NOTE:** The SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, and the DS3/E3 MIC do not support Ethernet over Frame Relay encapsulation.

**ether-vpls-over-ppp**—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Point-to-Point Protocol (PPP) encapsulation to support Bridged Ethernet over PPP-encapsulated TDM interfaces for VPLS applications.

**ethernet**—Use Ethernet II encapsulation (as described in RFC 894, *A Standard for the Transmission of IP Datagrams over Ethernet Networks*).

**ethernet-ccc**—Use Ethernet CCC encapsulation on Ethernet interfaces.

**ethernet-vpls**—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values.



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

**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.

**frame-relay-ccc**—Use Frame Relay encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

**frame-relay-ether-type**—Use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay. The physical interface must be configured with flexible-frame-relay encapsulation.

**frame-relay-ether-type-tcc**—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media. The physical interface must be configured with flexible-frame-relay encapsulation.

**frame-relay-ppp**—Use PPP over Frame Relay circuits. When you use this encapsulation type, you can configure the **ppp** family only. J Series routers do not support frame-relay-ppp encapsulation.

**frame-relay-tcc**—Use Frame Relay encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

**gre-fragmentation**—For adaptive services interfaces only, use GRE fragmentation encapsulation to enable fragmentation of IPv4 packets in GRE tunnels. This encapsulation clears the do not fragment (DF) bit in the packet header. If the packet's size exceeds the tunnel's maximum transmission unit (MTU) value, the packet is fragmented before encapsulation.

**multilink-frame-relay-end-to-end**—Use MLFR FRF.15 encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

**multilink-ppp**—Use MLPPP encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces.

**ppp-over-ether**—For underlying Ethernet interfaces on J Series routers, use PPP over Ethernet encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, configure the interface address on the PPP interface. You also use PPP over Ethernet encapsulation to configure an underlying Ethernet interface for a dynamic PPPoE logical interface on M120 and M320 routers with Intelligent Queuing 2 (IQ2) PICs, and on MX Series routers with MPCs.

**ppp-over-ether-over-atm-llc**—(J Series routers and MX Series routers with MPCs using the ATM MIC with SFP only) For underlying ATM interfaces, use PPP over Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, configure the interface address on the PPP interface.

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

**vlan-ccc**—Use Ethernet virtual LAN (VLAN) encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

**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.

**vlan-tcc**—Use Ethernet VLAN encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

**vlan-vpls**—Use Ethernet VLAN encapsulation on VPLS circuits.

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

**Related Documentation**

- *Configuring Layer 2 Switching Cross-Connects Using CCC*
- *Configuring the Encapsulation for Layer 2 Switching TCCs*
- *Configuring Interface Encapsulation on Logical Interfaces*
- *Configuring MPLS LSP Tunnel Cross-Connects Using CCC*
- *Circuit and Translational Cross-Connects Overview*
- *Identifying the Access Concentrator*
- *Configuring ATM Interface Encapsulation*
- [Configuring VLAN Encapsulation on page 18](#)
- [Configuring Extended VLAN Encapsulation on page 20](#)
- *Configuring ISDN Logical Interface Properties*
- *Configuring ATM-to-Ethernet Interworking*
- *Configuring Interface Encapsulation on PTX Series Packet Transport Routers*
- *Configuring CCC Encapsulation for Layer 2 VPNs*
- *Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits*
- *Configuring ATM for Subscriber Access*
- *CoS on ATM IMA Pseudowire Interfaces Overview*
- *Configuring Policing on an ATM IMA Pseudowire*

## encapsulation (Physical Interface)

---

<b>Syntax</b>	<code>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);</code>
<b>Hierarchy Level</b>	<code>[edit interfaces <i>interface-name</i>],</code> <code>[edit interfaces rlsq <i>number:number</i>]</code>
<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>	<b>ppp</b> —Use serial PPP encapsulation.
<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, <i>Multiprotocol Encapsulation over ATM Adaptation Layer 5</i> , 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>	<b>interface</b> —To view this statement in the configuration. <b>interface-control</b> —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 Encapsulation on page 18](#)
- [Configuring Extended VLAN Encapsulation on page 20](#)
- *Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces*
- *Configuring Interfaces for Layer 2 Circuits*
- *Configuring Interface Encapsulation on PTX Series Packet Transport Routers*
- *Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)*
- *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*
- *Configuring Q-in-Q Tunneling (CLI Procedure)*

## 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> <b>unit</b> <i>logical-unit-number</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <b>unit</b> <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*
- *Junos OS Services Interfaces Library for Routing Devices*

## 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>

## flexible-vlan-tagging

---

<b>Syntax</b>	flexible-vlan-tagging;
<b>Hierarchy Level</b>	[edit interfaces <i>aex</i> ], [edit interfaces <i>ge-fpc/pic/port</i> ], [edit interfaces <i>et-fpc/pic/port</i> ], [edit interfaces <i>ps0</i> ], [edit interfaces <i>xe-fpc/pic/port</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1. Support for aggregated Ethernet added in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers. Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
<b>Description</b>	<p>Support simultaneous transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port, and on pseudowire logical interfaces.</p> <p>This statement is supported on M Series and T Series routers, for Fast Ethernet and Gigabit Ethernet interfaces only on Gigabit Ethernet IQ2 and IQ2-E, IQ, and IQE PICs, and for aggregated Ethernet interfaces with member links in IQ2, IQ2-E, and IQ PICs or in MX Series DPCs, or on Ethernet interfaces for PTX Series Packet Transport Routers or 100-Gigabit Ethernet Type 5 PIC with CFP. This statement is supported on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series switches.</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="#">Configuring Mixed Tagging on page 10</a></li><li>• <a href="#">Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers on page 36</a></li><li>• <a href="#">Configuring Q-in-Q Tunneling (CLI Procedure)</a></li></ul>

## interface-mode

<b>Syntax</b>	interface-mode (access   trunk);
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <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> <b>unit</b> <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.

**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](#)” on [page 34](#).

<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>
<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>• <a href="#">Configuring a Logical Interface for Access Mode on page 32</a></li> <li>• <a href="#">Configuring a Logical Interface for Trunk Mode on page 33</a></li> </ul>

- *Example: Connecting Access Switches to a Distribution Switch*

## 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.
<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 <b>flexible-vlan-tagging</b> statement, untagged packets are accepted on the same mixed VLAN-tagged port.</p> <p>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. 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>Options</b>	<p><b><i>number</i></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>	<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>• <a href="#">Configuring Mixed Tagging Support for Untagged Packets on page 11</a></li> <li>• <a href="#">Configuring a Logical Interface for Access Mode on page 32</a></li> <li>• <a href="#">Configuring the Native VLAN Identifier (CLI Procedure)</a></li> <li>• <a href="#">Understanding Bridging and VLANs on EX Series Switches</a></li> <li>• <a href="#">flexible-vlan-tagging on page 76</a></li> </ul>

- *Understanding Q-in-Q Tunneling on EX Series Switches*

## stacked-vlan-tagging

---

<b>Syntax</b>	stacked-vlan-tagging;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
<b>Description</b>	<p>For Gigabit Ethernet IQ interfaces, Gigabit Ethernet, 10-Gigabit Ethernet LAN/WAN PIC, and 100-Gigabit Ethernet Type 5 PIC with CFP, enable stacked VLAN tagging for all logical interfaces on the physical interface.</p> <p>For pseudowire subscriber interfaces, enable stacked VLAN tagging for logical interfaces on the pseudowire service.</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 and Rewriting Gigabit Ethernet VLAN Tags Overview</i></li><li>• <a href="#">vlan-tags (Stacked VLAN Tags) on page 94</a></li></ul>

## unit

```

Syntax  unit logical-unit-number {
    accept-source-mac {
        mac-address mac-address {
            policer {
                input cos-policer-name;
                output cos-policer-name;
            }
        }
    }
    accounting-profile name;
    advisory-options {
        downstream-rate rate;
        upstream-rate rate;
    }
    allow-any-vci;
    atm-scheduler-map (map-name | default);
    backup-options {
        interface interface-name;
    }
    bandwidth rate;
    cell-bundle-size cells;
    clear-dont-fragment-bit;
    compression {
        rtp {
            maximum-contexts number <force>;
            f-max-period number;
            queues [queue-numbers];
            port {
                minimum port-number;
                maximum port-number;
            }
        }
    }
    compression-device interface-name;
    copy-tos-to-outer-ip-header;
    demux-destination family;
    demux-source family;
    demux-options {
        underlying-interface interface-name;
    }
    description text;
    interface {
        l2tp-interface-id name;
        (dedicated | shared);
    }
    dialer-options {
        activation-delay seconds;
        callback;
        callback-wait-period time;
        deactivation-delay seconds;
        dial-string [dial-string-numbers];
        idle-timeout seconds;
    }
  }

```

```

incoming-map {
    caller caller-id | accept-all;
    initial-route-check seconds;
    load-interval seconds;
    load-threshold percent;
    pool pool-name;
    redial-delay time;
    watch-list {
        [routes];
    }
}
}
disable;
disable-mlppp-inner-ppp-pfc;
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;
family family-name {
    ... the family subhierarchy appears after the main [edit interfaces interface-name unit
        logical-unit-number] hierarchy ...
}
fragment-threshold bytes;
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;
}
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 (disable | seconds);

```

```

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;
}
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;
    lcp-restart-timer milliseconds;
    loopback-clear-timer seconds;
    ncp-restart-timer milliseconds;
    pap {
        access-profile name;
        default-pap-password password;
        local-name name;
        local-password password;
        passive;
    }
}
pppoe-options {
    access-concentrator name;
    auto-reconnect seconds;
    (client | server);
    service-name name;
    underlying-interface interface-name;
}
pppoe-underlying-options {
    access-concentrator name;
    direct-connect;
    dynamic-profile profile-name;
    max-sessions number;
}
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;

```

```

targeted-distribution;
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;
vci-range start start-vci end end-vci;
vpi vpi-identifier;
vlan-id number;
vlan-id-range number-number;
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
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;

```

```

mtu bytes;
multicast-only;
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
(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;
        epd-threshold cells <plp1 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 | 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;
    track {
        interface interface-name {
            bandwidth-threshold bits-per-second priority-cost number;
        }
        priority-hold-time seconds;
        route ip-address/prefix-length routing-instance instance-name priority-cost cost;
    }
    virtual-address [addresses];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-interface interface-name;
        active-group group-number;
    }
}
}
}
}

```

**Hierarchy Level** [edit interfaces *interface-name*],  
[edit logical-systems *logical-system-name* interfaces *interface-name*],  
[edit interfaces interface-set *interface-set-name* interface *interface-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.

**Description** Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

**Options** *logical-unit-number*—Number of the logical unit.

**Range:** 0 through 1,073,741,823 for demux and PPPoE static interfaces. 0 through 16,385 for all other static interface types.

The remaining statements are explained separately.

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

- Related Documentation**
- [Configuring Logical Interface Properties](#)
  - [Junos OS Services Interfaces Library for Routing Devices](#)

## vlan-id (VLAN ID to Be Bound to a Logical Interface)

<b>Syntax</b>	<code>vlan-id <i>number</i>;</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.
<b>Description</b>	For Fast Ethernet, Gigabit Ethernet, and Aggregated Ethernet interfaces only, bind a 802.1Q VLAN tag ID to a logical interface.
<b>Options</b>	<p><i>number</i>—A valid VLAN identifier.</p> <p><b>Range:</b> For aggregated Ethernet, 4-port, 8-port, and 12-port Fast Ethernet PICs, and for management and internal Ethernet interfaces, 1 through 1023.</p> <p>For 48-port Fast Ethernet and Gigabit Ethernet PICs, 1 through 4094.</p> <p>VLAN ID 0 is reserved for tagging the priority of frames.</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>• <a href="#">Configuring Mixed Tagging on page 10</a></li> </ul>

## vlan-id (VLAN ID to Rewrite)

---

<b>Syntax</b>	<code>vlan-id <i>number</i>;</code>
<b>Hierarchy Level</b>	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code>input-vlan-map],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code>output-vlan-map]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<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 <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]</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>vlan-id</b> statement that you include at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code> hierarchy level.</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 the VLAN Tag on Tagged Frames</a></li><li>• <a href="#">Binding VLAN IDs to Logical Interfaces on page 12</a></li></ul>

## vlan-id-list (Interface in Bridge Domain)

<b>Syntax</b>	<code>vlan-id-list [ <i>number number-number</i> ];</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge], [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.
<b>Description</b>	Configure a logical interface to forward packets and learn MAC addresses within each bridge domain configured with a VLAN ID that matches a VLAN ID specified in the list. VLAN IDs can be entered individually using a space to separate each ID, entered as an inclusive list separating the starting VLAN ID and ending VLAN ID with a hyphen, or a combination of both.
<b>Options</b>	<p><i>number number</i>—Individual VLAN IDs separated by a space.</p> <p><i>number-number</i>—Starting VLAN ID and ending VLAN ID in an inclusive range.</p> <p><b>Range:</b> 1 through 4095</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>• <a href="#">Configuring a Logical Interface for Trunk Mode on page 33</a></li> <li>• <a href="#">Configuring the VLAN ID List for a Trunk Interface on page 33</a></li> </ul>

## vlan-id-range

---

<b>Syntax</b>	<code>vlan-id-range <i>vlan-id-vlan-id</i></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 in Junos OS Release 8.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Bind a range of VLAN IDs to a logical interface.
<b>Options</b>	<b>number</b> —The first number is the lowest VLAN ID in the range the second number is the highest VLAN ID in the range. <b>Range:</b> 1 through 4094



**NOTE:** Configuring `vlan-id-range` with the entire `vlan-id` range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1-4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.

The following examples illustrate this further:

```
[edit interfaces interface-name]  
vlan-tagging;  
unit number {  
    vlan-id-range 1-4094;  
}  
  
[edit interfaces interface-name]  
unit 0;
```

VLAN ID 0 is reserved for tagging the priority of frames.

<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="#">Binding a Range of VLAN IDs to a Logical Interface on page 14</a></li></ul>

## vlan-tagging

---

<b>Syntax</b>	vlan-tagging;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers. Statement introduced in Junos OS Release 13.2 for PTX Series Routers.
<b>Description</b>	For Fast Ethernet and Gigabit Ethernet interfaces, aggregated Ethernet interfaces configured for VPLS, and pseudowire subscriber interfaces, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.
<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>Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch</i></li> <li>• <i>Example: Configuring BGP Autodiscovery for LDP VPLS</i></li> <li>• <i>Configuring a Layer 3 Subinterface (CLI Procedure)</i></li> <li>• <i>Configuring Tagged Aggregated Ethernet Interfaces</i></li> <li>• <i>Configuring Interfaces for VPLS Routing</i></li> <li>• <i>Enabling VLAN Tagging</i></li> <li>• <a href="#">802.1Q VLANs Overview on page 3</a></li> <li>• <i>vlan-id</i></li> </ul>

## vlan-tags (Dual-Tagged Logical Interface)

<b>Syntax</b>	<code>vlan-tags inner-list [vlan-id vlan-id-vlan-id ] outer &lt;tpid.&gt;vlan-id;</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 in Junos OS Release 9.5.
<b>Description</b>	(MX Series routers only) Binds a dual-tag logical interface to a list of VLAN IDs. Configures the logical interface to receive and forward any dual-tag frame whose inner VLAN ID tag matches the list of VLAN IDs you specify.



### NOTE:

To create a circuit cross-connect (CCC) using VLAN-bundled dual-tag logical interfaces on Layer 2 VPN routing instances, you must include the `encapsulation-type` statement and specify the value `ethernet-vlan` at the one of the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols l2vpn]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols l2vpn]

For more information about the `encapsulation-type` configuration statement and the Layer 2 encapsulation types `ethernet` and `ethernet-vlan`, see the *Junos OS VPNs Library for Routing Devices*.

<b>Options</b>	<p><b>inner-list [vlan-id vlan-id vlan-id-vlan-id]</b>—A list of valid VLAN ID numbers. Specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both.</p> <p><b>Range:</b> 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.</p> <p><b>outer &lt;tpid.&gt;vlan-id</b>—An optional Tag Protocol ID (TPID) and a valid VLAN ID.</p> <p><b>Range:</b> For TPID, specify a hexadecimal value in the format <code>0xnnnn</code>.</p> <p><b>Range:</b> For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.</p>
----------------	--



**NOTE:** Configuring `inner-list` with the entire `vlan-id` range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs of inner tag (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1 through 4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.

The following examples illustrate this further:

```
[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-tags outer vid inner-list 1-4094;
}

[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-id vid;
}
```

---

<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="#">Binding VLAN IDs to Logical Interfaces on page 12</a></li> <li>• <a href="#">encapsulation (Logical Interface) on page 62</a></li> <li>• <a href="#">encapsulation (Physical Interface) on page 66</a></li> <li>• encapsulation-type (Layer 2 VPN routing instance), see the <i>Junos OS VPNs Library for Routing Devices</i>.</li> <li>• <a href="#">flexible-vlan-tagging on page 76</a></li> <li>• <a href="#">vlan-id-list (Ethernet VLAN Circuit) on page 96</a></li> <li>• <a href="#">vlan-tagging on page 91</a></li> </ul>

## vlan-tags (Stacked VLAN Tags)

<b>Syntax</b>	<code>vlan-tags inner <i>tpid.vlan-id</i> inner-range <i>vid1—vid2</i> outer <i>tpid.vlan-id</i>;</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 12.1X48 for PTX Series Packet Transport Routers.
<b>Description</b>	For Gigabit Ethernet IQ and IQE interfaces only, bind TPIDs and 802.1Q VLAN tag IDs to a logical interface.



**NOTE:** The inner-range *vid1—vid2* option is supported on MX Series with IQE PICs only.

<b>Options</b>	<p><b>inner <i>tpid.vlan-id</i></b>—A TPID and a valid VLAN identifier.</p> <p><b>Range:</b> (most routers) For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames. For PTX Series, VLAN ID 0 is not supported.</p> <p><b>inner-range <i>vid1—vid2</i></b>—For MX Series routers with Enhanced IQ (IQE) PICs only; specify a range of VLAN IDs where <i>vid1</i> is the start of the range and <i>vid2</i> is the end of the range.</p> <p><b>Range:</b> For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.</p> <p><b>outer <i>tpid.vlan-id</i></b>—A TPID and a valid VLAN identifier.</p> <p><b>Range:</b> (most routers) For VLAN ID, 1 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces. VLAN ID 0 is reserved for tagging the priority of frames. For PTX Series, VLAN ID 0 is not supported.</p>
----------------	--



**NOTE:** Configuring inner-range with the entire *vlan-id* range consumes system resources and is not a best practice. It should be used only when a subset of VLAN IDs of inner tag (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1–4094), it has the same result as not specifying a range; however, it consumes Packet Forwarding Engine resources such as VLAN lookup table entries, and so on.

The following examples illustrate this further:

```
[edit interfaces interface-name]
  stacked-vlan-tagging;
  unit number {
    vlan-tags outer vid inner-range 1-4094;
  }
```

```
[edit interfaces interface-name]  
vlan-tagging;  
unit number {  
    vlan-id vid;  
}
```

---

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

**Related Documentation**

- [Configuring Dual VLAN Tags](#)
- [Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers on page 36](#)
- [stacked-vlan-tagging on page 80](#)

## vlan-id-list (Ethernet VLAN Circuit)

---

<b>Syntax</b>	<code>vlan-id-list [vlan-id vlan-id–vlan-id];</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 in Junos OS Release 9.5.
<b>Description</b>	Binds a single-tag logical interface to a list of VLAN IDs. Configures a logical interface to receive and forward any tag frame whose VLAN ID tag matches the list of VLAN IDs you specify.



---

**NOTE:**

When you create a circuit cross-connect (CCC) using VLAN-bundled single-tag logical interfaces on Layer 2 VPN routing instances, the circuit automatically uses ethernet encapsulation. For Layer 2 VPN, you need to include the `encapsulation-type` statement and specify the value `ethernet` at either of the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols l2vpn]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols l2vpn]

For more information about the `encapsulation-type` configuration statement and the Layer 2 encapsulation types `ethernet` and `ethernet-vlan`, see the *Junos OS VPNs Library for Routing Devices*.

---

<b>Options</b>	<code>[vlan-id vlan-id–vlan-id]</code> —A list of valid VLAN ID numbers. Specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both. <b>Range:</b> 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.
----------------	--

---



**NOTE:** Configuring `vlan-id-list` with the entire `vlan-id` range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1-4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.

The following examples illustrate this further:

```
[edit interfaces interface-name]  
vlan-tagging;  
unit number {
```

```
        vlan-id-range 1-4094;  
    }  
  
    [edit interfaces interface-name]  
    unit 0;
```

---

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

**Related Documentation**

- [Binding VLAN IDs to Logical Interfaces on page 12](#)
- [encapsulation \(Logical Interface\) on page 62](#)
- [encapsulation \(Physical Interface\) on page 66](#)
- encapsulation-type (Layer 2 VPN routing instance), see the *Junos OS VPNs Library for Routing Devices*
- [flexible-vlan-tagging on page 76](#)
- [vlan-tagging on page 91](#)
- [vlan-tags \(Dual-Tagged Logical Interface\) on page 92](#)



## PART 3

# Administration

- [Monitoring Commands on page 101](#)



## CHAPTER 5

# Monitoring Commands

- `show interfaces` (Aggregated Ethernet)
- `show interfaces demux0` (Demux Interfaces)
- `show interfaces diagnostics optics` (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet)
- `show interfaces` (Fast Ethernet)
- `show interfaces` (10-Gigabit Ethernet)
- `show interfaces interface-set` (Ethernet Interface Set)
- `show interfaces interface-set queue`
- `show interfaces irb`
- `show lacp interfaces`
- `show interfaces mac-database` (Gigabit Ethernet)
- `show interfaces mc-ae`
- `show oam ethernet connectivity-fault-management delay-statistics`
- `show oam ethernet connectivity-fault-management forwarding-state`
- `show oam ethernet connectivity-fault-management interfaces`
- `show oam ethernet connectivity-fault-management linktrace path-database`
- `show oam ethernet connectivity-fault-management mep-database`
- `show oam ethernet connectivity-fault-management mep-statistics`
- `show oam ethernet connectivity-fault-management path-database`
- `show oam ethernet evc`
- `show oam ethernet link-fault-management`
- `show oam ethernet lmi`
- `show oam ethernet lmi statistics`
- `show protection-group ethernet-ring aps`
- `show protection-group ethernet-ring data-channel`
- `show protection-group ethernet-ring interface`
- `show protection-group ethernet-ring node-state`

- `show protection-group ethernet-ring statistics`
- `show protection-group ethernet-ring vlan`

## show interfaces (Aggregated Ethernet)

<b>Syntax</b>	<pre>show interfaces ae <i>number</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>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 14.1 for PTX Series Packet Transport Routers.</p>
<b>Description</b>	(M Series, T Series, MX Series, and PTX Series routers and EX Series switches) Display status information about the specified aggregated Fast Ethernet or Gigabit Ethernet interface.
<b>Options</b>	<p><b>ae <i>number</i></b>—Display standard information about the specified aggregated Fast Ethernet or 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.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information about 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 (Aggregated Ethernet) on page 108</a></p> <p><a href="#">show interfaces brief (Aggregated Ethernet) on page 109</a></p> <p><a href="#">show interfaces detail (Aggregated Ethernet) on page 109</a></p> <p><a href="#">show interfaces extensive (Aggregated Ethernet) on page 110</a></p> <p><a href="#">show interfaces extensive (Aggregated Ethernet with VLAN Stacking) on page 111</a></p>
<b>Output Fields</b>	Table 7 on page 103 lists the output fields for the <b>show interfaces</b> (Aggregated Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 7: show interfaces (Aggregated Ethernet) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface and state of the 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 7: show interfaces (Aggregated Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Interface index</b>	Index number of the physical interface, which reflects its initialization sequence.	All levels
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive</b> none
<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>Flow control</b>	Flow control status: <b>Enabled</b> or <b>Disabled</b> .	All levels
<b>Minimum links needed</b>	Number of child links that must be operational for the aggregate interface to be operational.	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 "Interfaces Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Current address</b>	Configured MAC address.	<b>detail extensive</b>
<b>Hardware address</b>	Hardware MAC address.	<b>detail extensive</b>
<b>Last flapped</b>	Date, time, and how long ago the interface went from down to up or from up to down. 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</b>
<b>Input Rate</b>	Input rate in bits per second (bps) and packets per second (pps).	None specified
<b>Output Rate</b>	Output rate in bps and pps.	None specified
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>

Table 7: show interfaces (Aggregated Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Traffic statistics</b>	<p>Number of and rate at which bytes and packets are received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes and rate, in bps, at which bytes are received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes and rate, in bps, at which bytes are transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets and rate, in pps, at which packets are received on the interface.</li> <li>• <b>Output packets</b>—Number of packets and rate, in pps, at which packets are transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Input errors</b>	<p>Input errors on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of incoming frame aborts and frame check sequence (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 random early detection (RED) mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid 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 were not of interest. Usually, this field reports protocols that 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:</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), then the cable, the far-end system, or the PIC 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>
<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>

Table 7: show interfaces (Aggregated Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Queue counters</b>	CoS queue number and its associated user-configured forwarding class name.  <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>
<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 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 Field" section under <i>Common Output Fields Description</i> .	All levels
<b>VLAN-Tag</b>	Tag Protocol Identifier (TPID) and VLAN identifier.	All levels
<b>Demux</b>	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:  <ul style="list-style-type: none"> <li>• <b>Source Family Inet</b></li> <li>• <b>Destination Family Inet</b></li> </ul>	<b>detail extensive</b> none
<b>Encapsulation</b>	Encapsulation on the logical interface.	All levels

Table 7: show interfaces (Aggregated Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Statistics	<p>Information about the number of packets, packets per second, number of bytes, and bytes per second on this aggregate interface.</p> <ul style="list-style-type: none"> <li>• <b>Bundle</b>—Information about input and output bundle rates.</li> <li>• <b>Link</b>—(<b>detail</b> and <b>extensive</b> only) Information about specific links in the aggregate, including link state and input and output rates.</li> <li>• <b>Adaptive Statistics</b>—(<b>extensive</b> only) Information about adaptive load balancing counter statistics. <ul style="list-style-type: none"> <li>• <b>Adaptive Adjusts</b>—Number of times traffic flow imbalance was corrected by implementation of adaptive load balancing.</li> <li>• <b>Adaptive Scans</b>—Number of times the link utilization on each member link of the AE bundle was scanned by for adaptive load balancing</li> <li>• <b>Adaptive Tolerance</b>—Tolerance level, in percentage, for load imbalance on link utilization on each member link of the AE bundle.</li> <li>• <b>Adaptive Updates</b>—Number of times traffic flow load have been updated on an AE bundle.</li> </ul> </li> <li>• <b>Marker Statistics</b>—(<b>detail</b> and <b>extensive</b> only) Information about 802.3ad marker protocol statistics on the specified links. <ul style="list-style-type: none"> <li>• <b>Marker Rx</b>—Number of valid marker protocol data units (PDUs) received on this aggregation port.</li> <li>• <b>Resp Tx</b>—Number of marker response PDUs transmitted on this aggregation port.</li> <li>• <b>Unknown Rx</b>—Number of frames received that either carry the slow protocols Ethernet type value (43B.4) but contain an unknown PDU, or are addressed to the slow protocols group MAC address (43B.3) but do not carry the slow protocols Ethernet type.</li> <li>• <b>Illegal Rx</b>—Number of frames received that carry the slow protocols Ethernet type value (43B.4) but contain a badly formed PDU or an illegal value of protocol subtype (43B.4).</li> </ul> </li> </ul>	<b>detail extensive</b> none
LACP info	<p>Link Aggregation Control Protocol (LACP) information for each aggregated interface.</p> <ul style="list-style-type: none"> <li>• <b>Role</b> can be one of the following: <ul style="list-style-type: none"> <li>• <b>Actor</b>—Local device participating in LACP negotiation.</li> <li>• <b>Partner</b>—Remote device participating in LACP negotiation.</li> </ul> </li> <li>• <b>System priority</b>—Priority assigned to the system (by management or administrative policy), encoded as an unsigned integer.</li> <li>• <b>System identifier</b>—Actor or partner system ID, encoded as a MAC address.</li> <li>• <b>Port priority</b>—Priority assigned to the port by the actor or partner (by management or administrative policy), encoded as an unsigned integer.</li> <li>• <b>Port number</b>—Port number assigned to the port by the actor or partner, encoded as an unsigned integer.</li> <li>• <b>Port key</b>—Operational key value assigned to the port by the actor or partner, encoded as an unsigned integer.</li> </ul>	

Table 7: show interfaces (Aggregated Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
LACP Statistics	<p>LACP statistics for each aggregated interface.</p> <ul style="list-style-type: none"> <li>• <b>LACP Rx</b>—LACP received counter that increments for each normal hello.</li> <li>• <b>LACP Tx</b>—Number of LACP transmit packet errors logged.</li> <li>• <b>Unknown Rx</b>—Number of unrecognized packet errors logged.</li> <li>• <b>Illegal Rx</b>—Number of invalid packets received.</li> </ul> <p><b>NOTE:</b> For <b>LACP Rx</b> and <b>LACP Tx</b>, Packet count is updated only on snmp timer expiry (30 secs).</p>	
<i>protocol-family</i>	Protocol family configured on the logical interface. Possible values are described in the "Protocol Field" section under <i>Common Output Fields Description</i> .	<b>brief</b>
<b>Protocol</b>	Protocol family configured on the logical interface. 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>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>Flags</b>	Information about protocol family flags. Possible values are described in the "Family Flags Field" section under <i>Common Output Fields Description</i> .	<b>detail extensive</b> none
<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</b> none
<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> none
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive</b> none
<b>Local</b>	IP address of the logical interface.	<b>detail extensive</b> none
<b>Broadcast</b>	Broadcast address 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>

## Sample Output

### show interfaces (Aggregated Ethernet)

```
user@host> show interfaces ae0
```

```
Physical interface: ae0, Enabled, Physical link is Up
Interface index: 153, SNMP ifIndex: 59
Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
Device flags : Present Running
Interface flags: SNMP-Traps 16384
Current address: 00:05:85:8b:bf:f0, Hardware address: 00:05:85:8b:bf:f0
Last flapped : Never
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)
```

```
Logical interface ae0.0 (Index 72) (SNMP ifIndex 60)
Flags: SNMP-Traps 16384 Encapsulation: ENET2
Statistics      Packets      pps      Bytes      bps
Bundle:
  Input :          0          0          0          0
  Output:          0          0          0          0
Protocol inet, MTU: 1500
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.100.1/24, Local: 10.100.1.2, Broadcast: 10.100.1.255
```

#### show interfaces brief (Aggregated Ethernet)

```
user@host> show interfaces ae0 brief
Physical interface: ae0, Enabled, Physical link is Up
Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
Device flags : Present Running
Interface flags: SNMP-Traps 16384

Logical interface ae0.0
Flags: SNMP-Traps 16384 Encapsulation: ENET2
inet 10.100.1.2/24
```

#### show interfaces detail (Aggregated Ethernet)

```
user@host> show interfaces ae0 detail
Physical interface: ae0, Enabled, Physical link is Up
Interface index: 153, SNMP ifIndex: 59, Generation: 36
Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
Device flags : Present Running
Interface flags: SNMP-Traps 16384
Current address: 00:05:85:8b:bf:f0, Hardware address: 00:05:85:8b:bf:f0
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes :          0          0 bps
Output bytes :          0          0 bps
Input packets:          0          0 pps
Output packets:          0          0 pps
Queue counters:      Queued packets  Transmitted packets  Dropped packets

0 best-effort          7375          7375          0
1 expedited-fo          0          0          0
2 assured-forw          0          0          0
3 network-cont        2268          2268          0
```

```

Logical interface ae0.0 (Index 72) (SNMP ifIndex 60) (Generation 18)
Flags: SNMP-Traps 16384 Encapsulation: ENET2
Statistics          Packets          pps          Bytes          bps
Bundle:
  Input :           0           0           0           0
  Output:           0           0           0           0
Link:
  fe-0/1/0.0
    Input :           0           0           0           0
    Output:           0           0           0           0
  fe-0/1/2.0
    Input :           0           0           0           0
    Output:           0           0           0           0
  fe-0/1/3.0
    Input :           0           0           0           0
    Output:           0           0           0           0
Marker Statistics:  Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
fe-0/1/0.0          0          0          0          0
fe-0/1/2.0          0          0          0          0
fe-0/1/3.0          0          0          0          0
Protocol inet, MTU: 1500, Generation: 37, Route table: 0
Flags: Is-Primary, Mac-Validate-Strict
Mac-Validate Failures: Packets: 0, Bytes: 0
  Destination: 10.100.1/24, Local: 10.100.1.2, Broadcast: 10.100.1.255,
  Generation: 49

```

#### show interfaces extensive (Aggregated Ethernet)

```

user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
Interface index: 153, SNMP ifIndex: 59, Generation: 36
Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
Device flags   : Present Running
Interface flags: SNMP-Traps 16384
Current address: 00:05:85:8b:bf:f0, Hardware address: 00:05:85:8b:bf:f0
Last flapped   : Never
Statistics last cleared: Never
Traffic statistics:
  Input bytes :           60           0 bps
  Output bytes :           0           0 bps
  Input packets:           1           0 pps
  Output packets:          0           0 pps
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
Queue counters:      Queued packets  Transmitted packets  Dropped packets

  0 best-effort           7375           7375           0
  1 expedited-fo           0             0             0
  2 assured-forw           0             0             0
  3 network-cont        2268           2268           0

```

```

Logical interface ae0.0 (Index 73) (SNMP ifIndex 563) (Generation 176)
Flags: Up SNMP-Traps 0x4000 Encapsulation: ENET2
Statistics          Packets          pps          Bytes          bps
Bundle:
  Input :           0           0           0           0
  Output:           0           0           0           0
Adaptive Statistics:
  Adaptive Adjusts:           0
  Adaptive Scans  :           0
  Adaptive Updates:           0
Link:
  fe-1/0/3.0
    Input :           0           0           0           0
    Output:           0           0           0           0
LACP info:          Role      System          System      Port      Port      Port
                  priority          identifier  priority  number      key

  fe-1/0/3.0      Actor          127  00:24:dc:85:af:f0          127          2          1
  fe-1/0/3.0      Partner          127  00:23:9c:c3:1f:f0          127          1          1

LACP Statistics:      LACP Rx      LACP Tx      Unknown Rx      Illegal Rx
fe-1/0/3.0           3188          3186           0           0
Marker Statistics:      Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
fe-1/0/3.0           0           0           0           0
Protocol inet, MTU: 1500, Generation: 224, Route table: 0
Flags: Sendbcst-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.40.1.0/30, Local: 10.40.1.1, Broadcast: 10.40.1.3,
Generation: 187
Protocol multiservice, MTU: Unlimited, Generation: 225, Route table: 0
Flags: Is-Primary
Policer: Input: __default_arp_policer__

```

### show interfaces extensive (Aggregated Ethernet with VLAN Stacking)

```

user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
Interface index: 155, SNMP ifIndex: 48, Generation: 186
Link-level type: 52, MTU: 1518, Speed: 2000mbps, Loopback: Disabled, Source
filtering: Disabled,
Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth needed: 0
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Current address: 00:12:1e:19:3f:f0, Hardware address: 00:12:1e:19:3f:f0
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes :           2406875          40152 bps
Output bytes :          1124470          22056 bps
Input packets:           5307           5 pps
Output packets:          13295          21 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, Runt: 0, Giants: 0, Policed discards: 0, Resource errors: 0

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

Ingress queues: 4 supported, 4 in use

Queue counters:	Queued packets	Transmitted packets	Dropped packets
0 best-effort	0	859777	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	1897615	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	0	662505	0

Logical interface ae0.451 (Index 69) (SNMP ifIndex 167) (Generation 601)

Flags: SNMP-Traps VLAN-Tag [ 0x8100.451 ] Encapsulation: VLAN-VPLS

Statistics	Packets	pps	Bytes	bps
------------	---------	-----	-------	-----

Bundle:

Input :	289	0	25685	376
Output:	1698	4	130375	3096

Link:

ge-1/2/0.451				
Input :	289	0	25685	376
Output:	0	0	0	0

ge-1/2/1.451				
Input :	0	0	0	0
Output:	1698	4	130375	3096

Marker Statistics:	Marker	Rx	Resp	Tx	Unknown	Rx	Illegal	Rx
ge-1/2/0.451		0		0		0		0
ge-1/2/1.451		0		0		0		0

Protocol vpls, MTU: 1518, Generation: 849, Route table: 3

Flags: Is-Primary

Logical interface ae0.452 (Index 70) (SNMP ifIndex 170) (Generation 602)

Flags: SNMP-Traps VLAN-Tag [ 0x8100.452 ] Encapsulation: VLAN-VPLS

Statistics	Packets	pps	Bytes	bps
------------	---------	-----	-------	-----

Bundle:

Input :	293	1	26003	1072
Output:	1694	3	130057	2400

Link:

ge-1/2/0.452				
Input :	293	1	26003	1072
Output:	1694	3	130057	2400

ge-1/2/1.452				
Input :	0	0	0	0
Output:	0	0	0	0

```
Marker Statistics:  Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
ge-1/2/0.452       0          0          0          0
ge-1/2/1.452       0          0          0          0
Protocol vpls, MTU: 1518, Generation: 850, Route table: 3
Flags: None
...
```

## show interfaces demux0 (Demux Interfaces)

<b>Syntax</b>	<pre>show interfaces demux0.logical-interface-number &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index snmp-index&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0.
<b>Description</b>	(MX Series and M Series routers only) Display status information about the specified demux interface.
<b>Options</b>	<p><b>none</b>—Display standard information about the specified demux 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 snmp-index</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>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>	<a href="#">show interfaces (Demux) on page 120</a> <a href="#">show interfaces (PPPoE over Aggregated Ethernet) on page 121</a> <a href="#">show interfaces extensive (Targeted Distribution for Aggregated Ethernet Links) on page 121</a> <a href="#">show interfaces demux0 (ACI Interface Set Configured) on page 122</a>
<b>Output Fields</b>	Table 8 on page 114 lists the output fields for the <b>show interfaces</b> (demux interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 8: Demux show interfaces Output Fields

Field Name	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Physical interface</b>	Name of the physical interface.	brief detail extensive none
<b>Interface index</b>	Index number of the physical interface, which reflects its initialization sequence.	brief detail extensive none

Table 8: Demux show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Enabled</b>	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	<b>brief detail extensive</b> none
<b>Physical link</b>	Status of the physical link ( <b>Up</b> or <b>Down</b> ).	<b>detail extensive</b> none
<b>Admin</b>	Administrative state of the interface ( <b>Up</b> or <b>Down</b> ).	<b>terse</b>
<b>Interface index</b>	Index number of the physical interface, which reflects its initialization sequence.	<b>detail extensive</b> none
<b>Link</b>	Status of the physical link ( <b>Up</b> or <b>Down</b> ).	<b>terse</b>
<b>Targeting summary</b>	Status of aggregated Ethernet links that are configured with targeted distribution ( <b>primary</b> or <b>backup</b> )	<b>extensive</b>
<b>Bandwidth</b>	Bandwidth allocated to the aggregated Ethernet links that are configured with targeted distribution.	<b>extensive</b>
<b>Proto</b>	Protocol family configured on the interface.	<b>terse</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive</b> none
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Type</b>	Type of interface. <b>Software-Pseudo</b> indicates a standard software interface with no associated hardware device.	<b>brief detail extensive</b> none
<b>Link-level type</b>	Encapsulation being used on the physical interface.	<b>brief detail extensive</b>
<b>MTU</b>	Maximum transmission unit size on the physical interface.	<b>brief detail extensive</b>
<b>Clocking</b>	Reference clock source: <b>Internal</b> (1) or <b>External</b> (2).	<b>brief detail extensive</b>
<b>Speed</b>	Speed at which the interface is running.	<b>brief detail extensive</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>brief detail extensive</b> none
<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>brief detail extensive</b> none
<b>Link type</b>	Data transmission type.	<b>detail extensive</b> none
<b>Link flags</b>	Information about the link. Possible values are described in the “Link Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive</b> none
<b>Physical info</b>	Information about the physical interface.	<b>detail extensive</b>
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds.	<b>detail extensive</b>

Table 8: Demux show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Current address	Configured MAC address.	detail extensive
Hardware address	Hardware MAC address.	detail extensive
Alternate link address	Backup address of the link.	detail extensive
Last flapped	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> .	detail extensive none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <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> <li>• <b>IPv6 transit statistics</b>—Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled.</li> </ul> <p><b>NOTE:</b> These fields include dropped traffic and exception traffic, as those fields are not separately defined.</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>	detail extensive
Input errors	<p>Input errors on the interface whose definitions are as follows:</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 packet 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>	extensive
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	none

Table 8: Demux show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Output errors</b>	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: <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>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>Output Rate</b>	Output rate in bps and pps.	none
<b>Logical Interface</b>		
<b>Logical interface</b>	Name of the logical interface.	<b>brief detail extensive</b> none
<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</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>brief detail extensive</b> none
<b>Encapsulation</b>	Encapsulation on the logical interface.	<b>brief extensive</b> none
<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 demux interface to create dynamic VLAN subscriber interfaces based on ACI information.	<b>brief detail extensive</b> none
<b>Demux</b>	Specific IP demultiplexing (demux) values: <ul style="list-style-type: none"> <li>• <b>Underlying interface</b>—The underlying interface that the demux interface uses.</li> <li>• <b>Index</b>—Index number of the logical interface.</li> <li>• <b>Family</b>—Protocol family configured on the logical interface.</li> <li>• <b>Source prefixes, total</b>—Total number of source prefixes for the underlying interface.</li> <li>• <b>Destination prefixes, total</b>—Total number of destination prefixes for the underlying interface.</li> <li>• <b>Prefix—in</b>et family prefix.</li> </ul>	<b>detail extensive</b> none

Table 8: Demux show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<i>protocol-family</i>	Protocol family configured on the logical interface.	<b>brief</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.</li> <li>• <b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface set.</li> <li>• <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.</li> </ul> <p><b>NOTE:</b> The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</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>Local statistics</b>	<p>Number of transit bytes and packets received and transmitted on the local 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>Transit statistics</b>	<p>Number and rate of bytes and packets transiting the switch.</p> <p><b>NOTE:</b> The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</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 logical interface if IPv6 statistics tracking is enabled.</p> <p><b>NOTE:</b> The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</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 packets</b>	Number of packets received on the interface.	<b>none</b>

Table 8: Demux show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Output packets</b>	Number of packets transmitted on the interface.	none
<b>Protocol</b>	Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .	detail extensive none
<b>MTU</b>	Maximum transmission unit size on the logical interface.	detail extensive none
<b>Maximum labels</b>	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	detail extensive
<b>Route table</b>	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive
<b>Flags</b>	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
<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.	detail extensive none
<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> .	detail extensive none
<b>Destination</b>	IP address of the remote side of the connection.	detail extensive statistics none
<b>Local</b>	IP address of the logical interface.	detail extensive terse none
<b>Remote</b>	IP address of the remote interface.	terse
<b>Broadcast</b>	Broadcast address of the logical interface.	detail extensive none
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	detail extensive
<b>Link</b>	Name of the physical interfaces for member links in an aggregated Ethernet bundle for a PPPoE over aggregated Ethernet configuration. PPPoE traffic goes out on these interfaces.	detail extensive none
<b>Dynamic-profile</b>	Name of the PPPoE dynamic profile assigned to the underlying interface.	detail extensive none
<b>Service Name Table</b>	Name of the PPPoE service name table assigned to the PPPoE underlying interface.	detail extensive none
<b>Max Sessions</b>	Maximum number of dynamic PPPoE logical interfaces that the router can activate on the underlying interface.	detail extensive none

Table 8: Demux show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Duplicate Protection</b>	State of duplicate protection: <b>On</b> or <b>Off</b> . Duplicate protection prevents the activation of another dynamic PPPoE logical interface on the same underlying interface when a dynamic PPPoE logical interface for a client with the same MAC address is already active on that interface.	<b>detail extensive none</b>
Direct Connect	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>
AC Name	Name of the access concentrator.	<b>detail extensive none</b>

## Sample Output

### show interfaces (Demux)

```

user@host> show interfaces demux0
Physical interface: demux0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 79, Generation: 129
  Type: Software-Pseudo, Link-level type: Unspecified, MTU: 9192, Clocking: 1,
  Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times    : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped  : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :                0                0 bps
    Output bytes  :                0                0 bps
    Input packets :                0                0 pps
    Output packets:                0                0 pps
  IPv6 transit statistics:
    Input bytes   :                0
    Output bytes  :                0
    Input packets :                0
    Output packets:                0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, 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 demux0.0 (Index 87) (SNMP ifIndex 84) (Generation 312)
  Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
  Demux:
    Underlying interface: ge-2/0/1.0 (Index 74)
    Family Inet Source prefixes, total 1
    Prefix: 1.1.1/24
    Traffic statistics:
      Input bytes   :                0

```

```

Output bytes :          1554
Input packets:           0
Output packets:         37
IPv6 transit statistics:
  Input bytes :           0
  Output bytes :           0
  Input packets:          0
  Output packets:         0
Local statistics:
  Input bytes :           0
  Output bytes :         1554
  Input packets:           0
  Output packets:         37
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: 395, Route table: 0
  Flags: Is-Primary, Mac-Validate-Strict
  Mac-Validate Failures: Packets: 0, Bytes: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 11.1.1/24, Local: 11.1.1.1, Broadcast: 11.1.1.255,
    Generation: 434

```

#### show interfaces (PPPoE over Aggregated Ethernet)

```

user@host> show interfaces demux0.100
Logical interface demux0.100 (Index 76) (SNMP ifIndex 61160)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]
  Encapsulation: ENET2
  Demux:
    Underlying interface: ae0 (Index 199)
  Link:
    ge-1/0/0
    ge-1/1/0
  Input packets : 0
  Output packets: 0
  Protocol pppoe
    Dynamic Profile: pppoe-profile,
    Service Name Table: service-table1,
    Max Sessions: 100, Duplicate Protection: On,
    Direct Connect: Off,
    AC Name: pppoe-server-1

```

#### show interfaces extensive (Targeted Distribution for Aggregated Ethernet Links)

```

user@host> show interfaces demux0.1073741824 extensive

Logical interface demux0.1073741824 (Index 75) (SNMP ifIndex 558) (Generation
346)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2
  Demux:
    Underlying interface: ae0 (Index 201)
  Link:
    ge-1/0/0


```

```
ge-1/1/0
ge-2/0/7
ge-2/0/8
Targeting summary:
ge-1/1/0, primary, Physical link is Up
ge-2/0/8, backup, Physical link is Up
Bandwidth: 1000mbps
```

#### show interfaces demux0 (ACI Interface Set Configured)

```
user@host> show interfaces demux0.1073741827
Logical interface demux0.1073741827 (Index 346) (SNMP ifIndex 527)
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1802 0x8100.302 ] Encapsulation:
ENET2
Demux: Source Family Inet
ACI VLAN:
  Dynamic Profile: aci-vlan-set-profile
Demux:
  Underlying interface: ge-1/0/0 (Index 138)
Input packets : 18
Output packets: 16
Protocol inet, MTU: 1500
  Flags: Sendbcst-pkt-to-re, Unnumbered
  Donor interface: lo0.0 (Index 322)
  Preferred source address: 100.20.200.202
  Addresses, Flags: Primary Is-Default Is-Primary
    Local: 10.4.12.119
Protocol 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
```

## show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet)

<b>Syntax</b>	show interfaces diagnostics optics <i>interface-name</i>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for PTX Series routers.
<b>Description</b>	Display diagnostics data, warnings, and alarms for Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, or 100-Gigabit Ethernet interfaces.
<b>Options</b>	<i>interface-name</i> —Interface name. For example: <i>ge-fpc/pic/port</i> <i>et-fpc/pic/port</i> <i>xe-fpc/pic/port</i>
<b>Additional Information</b>	<p>The transceivers are polled in 1-second intervals for diagnostics data, warnings, and alarms. The alarms do not cause the links to go down or the LEDs to change color, nor generate SNMP traps. Changes in alarm and warning status will generate system log messages.</p> <p>Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transceiver vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a device is not working.</p>
	<div>  <p><b>NOTE:</b> Some transceivers do not support all optical diagnostics features described in the output fields.</p> <p>The <code>show interfaces diagnostics optics</code> command for optical interfaces does not report the decibel (dBm) value of the received signal if the received power is zero milliwatts (0.0000 mW).</p> </div>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Supported Network Interface Standards by Transceiver</li> <li>Supported Network Interface Standards by Transceiver for PTX Series Packet Transport Routers</li> </ul>
<b>List of Sample Output</b>	<a href="#">show interfaces diagnostics optics (DWDM and DWDM OTN) on page 136</a> <a href="#">show interfaces diagnostics optics (Bidirectional SFP) on page 136</a> <a href="#">show interfaces diagnostics optics (SFP) on page 137</a> <a href="#">show interfaces diagnostics optics (SFP) on page 137</a>

[show interfaces diagnostics optics \(XFP and CFP Optics\) on page 138](#)

[show interfaces diagnostics optics for 10-Gigabit Ethernet \(PTX 24-10GE-SFPP\) on page 139](#)

[show interfaces diagnostics optics for 40-Gigabit Ethernet on page 140](#)

**Output Fields** Table 9 on page 124 lists the output fields for the **show interfaces diagnostics optics** command for DWDM and DWDM OTN PICs. Output fields are listed in the approximate order in which they appear.

**Table 9: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet DWDM and DWDM OTN PICs**

Field Name	Field Description
Physical interface	Name of the physical interface.
Laser bias current	Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm). This is a software equivalent to the <b>LsPOWMON</b> pin in hardware.
Receiver signal average optical power	Average received optical power, in mW and dBm. This indicator is a software equivalent to the <b>RxPOWMON</b> pin in hardware. Average optical power is vendor-specific.
Laser end-of-life alarm	Laser end-of-life alarm: <b>On</b> or <b>Off</b> .
Laser wavelength alarm	Laser wavelength alarm: <b>On</b> or <b>Off</b> .
Laser bias current alarm	Laser bias current alarm: <b>On</b> or <b>Off</b> .
Laser temperature alarm	Laser temperature alarm: <b>On</b> or <b>Off</b> .
Laser power alarm	Laser power alarm: <b>On</b> or <b>Off</b> .
Modulator temperature alarm	Modulator temperature alarm: <b>On</b> or <b>Off</b> . Transceivers from some vendors do not support this field.
Modulator bias alarm	Modulator bias alarm: <b>On</b> or <b>Off</b> .
Tx multiplexer FIFO error alarm	Transmit multiplexer first in, first out (FIFO) error alarm: <b>On</b> or <b>Off</b> .
Tx loss of PLL lock alarm	Transmit loss of phase-locked loop (PLL) lock alarm: <b>On</b> or <b>Off</b> .
Rx loss of average optical power alarm	Receive loss of average optical power alarm: <b>On</b> or <b>Off</b> .

**Table 9: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet DWDM and DWDM OTN PICs (*continued*)**

Field Name	Field Description
Rx loss of AC power alarm	Receive loss of AC power alarm: <b>On</b> or <b>Off</b> .  Transceivers from some vendors do not support this field.
Rx loss of PLL lock alarm	Receive loss of phase-locked loop (PLL) lock alarm: <b>On</b> or <b>Off</b> .

Table 10 on page 125 lists the output fields for the **show interfaces diagnostics optics** command when the router is operating with bidirectional SFP optics. Output fields are listed in the approximate order in which they appear.

**Table 10: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics**

Field Name	Field Description
Physical interface	Name of the physical interface.
Laser bias current	Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm).
Module temperature	Temperature of the optics module, in Celsius and Fahrenheit.
Module voltage	Internally measured module voltage.
Receiver signal average optical power	Average received optical power, in mW and dBm.
Laser bias current high alarm	Laser bias power setting high alarm. Displays <b>on</b> or <b>off</b> .
Laser bias current low alarm	Laser bias power setting low alarm. Displays <b>on</b> or <b>off</b> .
Laser bias current high warning	Laser bias power setting high warning. Displays <b>on</b> or <b>off</b> .
Laser bias current low warning	Laser bias power setting low warning. Displays <b>on</b> or <b>off</b> .
Laser output power high alarm	Laser output power high alarm. Displays <b>on</b> or <b>off</b> .
Laser output power low alarm	Laser output power low alarm. Displays <b>on</b> or <b>off</b> .

Table 10: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics *(continued)*

Field Name	Field Description
Laser output power high warning	Laser output power high warning. Displays <b>on</b> or <b>off</b> .
Laser output power low warning	Laser output power low warning. Displays <b>on</b> or <b>off</b> .
Module temperature high alarm	Module temperature high alarm. Displays <b>on</b> or <b>off</b> .
Module temperature low alarm	Module temperature low alarm. Displays <b>on</b> or <b>off</b> .
Module temperature high warning	Module temperature high warning. Displays <b>on</b> or <b>off</b> .
Module temperature low warning	Module temperature low warning. Displays <b>on</b> or <b>off</b> .
Module voltage high alarm	Module voltage high alarm. Displays <b>on</b> or <b>off</b> .
Module voltage low alarm	Module voltage low alarm. Displays <b>on</b> or <b>off</b> .
Module voltage high warning	Module voltage high warning. Displays <b>on</b> or <b>off</b> .
Module voltage low warning	Module voltage high warning. Displays <b>on</b> or <b>off</b> .
Laser rx power high alarm	Receive laser power high alarm. Displays <b>on</b> or <b>off</b> .
Laser rx power low alarm	Receive laser power low alarm. Displays <b>on</b> or <b>off</b> .
Laser rx power high warning	Receive laser power high warning. Displays <b>on</b> or <b>off</b> .
Laser rx power low warning	Receive laser power low warning. Displays <b>on</b> or <b>off</b> .
Laser bias current high alarm threshold	Vendor-specified threshold for the laser bias current high alarm: <b>70.000 mA</b> .
Laser bias current low alarm threshold	Vendor-specified threshold for the laser bias current low alarm: <b>0.0002 mA</b> .

Table 10: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics *(continued)*

Field Name	Field Description
Laser bias current high warning threshold	Vendor-specified threshold for the laser bias current high warning: <b>65.000 mA</b> .
Laser bias current low warning threshold	Vendor-specified threshold for the laser bias current low warning: <b>0.0002 mA</b> .
Laser output power high alarm threshold	Vendor-specified threshold for the laser output power high alarm: <b>1.0000 mW</b> or <b>0.00 dBm</b> .
Laser output power low alarm threshold	Vendor-specified threshold for the laser output power low alarm: <b>0.0560 mW</b> or <b>-12.52 dBm</b> .
Laser output power high warning threshold	Vendor-specified threshold for the laser output power high warning: <b>0.6300 mW</b> or <b>-2.01 dBm</b> .
Laser output power low warning threshold	Vendor-specified threshold for the laser output power low warning: <b>0.0890 mW</b> or <b>-10.51 dBm</b> .
Module temperature high alarm threshold	Vendor-specified threshold for the module temperature high alarm: <b>100° C</b> or <b>212° F</b> .
Module temperature low alarm threshold	Vendor-specified threshold for the module temperature low alarm: <b>-50° C</b> or <b>-58° F</b> .
Module temperature high warning threshold	Vendor-specified threshold for the module temperature high warning: <b>95 ° C</b> or <b>203 ° F</b> .
Module temperature low warning threshold	Vendor-specified threshold for the module temperature low warning: <b>-48° C</b> or <b>-54° F</b> .
Module voltage high alarm threshold	Module voltage high alarm threshold: <b>3.700 v</b> .
Module voltage low alarm threshold	Module voltage low alarm threshold: <b>2.900 v</b> .
Module voltage high warning threshold	Module voltage high warning threshold: <b>3.7600 v</b> .
Module voltage low warning threshold	Module voltage low warning threshold: <b>3.000 v</b> .
Laser rx power high alarm threshold	Vendor-specified threshold for the laser Rx power high alarm: <b>1.9953 mW</b> or <b>3.00 dBm</b> .
Laser rx power low alarm threshold	Vendor-specified threshold for the laser Rx power low alarm: <b>0.0001 mW</b> or <b>-40.00 dBm</b> .

Table 10: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics (*continued*)

Field Name	Field Description
Laser rx power high warning threshold	Vendor-specified threshold for the laser Rx power high warning: 1.0000 mW or 0.00 dBm.
Laser rx power low warning threshold	Vendor-specified threshold for the laser Rx power low warning: 0.0010 mW or -30.00 dBm.

Table 11 on page 128 lists the output fields for the **show interfaces diagnostics optics** command for SFP transceivers. Output fields are listed in the approximate order in which they appear.

Table 11: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers

Field Name	Field Description
Physical interface	Name of the physical interface.
Laser bias current	Measured laser bias current in uA.
Laser output power	Measured laser output power in mW.
Module temperature	Internally measured module temperature.
Module voltage	Internally measured module voltage.
Laser rx power	Measured receive optical power in mW.
Laser bias current high alarm	Laser bias current high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.
Laser bias current low alarm	Laser bias current low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.
Laser output power high alarm	Laser output power high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.
Laser output power low alarm	Laser output power low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.
Module temp high alarm	Module temperature high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.
Module temp low alarm	Module temperature low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.
Laser rx power high alarm	Laser receive power high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.
Laser rx power low alarm	Laser receive power low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.

Table 11: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers (*continued*)

Field Name	Field Description
Laser bias current high warning	Laser bias current high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.
Laser bias current low warning	Laser bias current low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.
Laser output power high warning	Laser output power high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.
Laser output power low warning	Laser output power low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.
Module temperature high warning	Module temperature high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.
Module temperature low warning	Module temperature low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.
Laser rx power high warning	Laser receive power high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.
Laser rx power low warning	Laser receive power low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.
Laser bias current high alarm threshold	Laser bias current high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser bias current low alarm threshold	Laser bias current low alarm threshold. Alarm threshold ranges are vendor-specific.
Laser bias current high warning threshold	Laser bias current high warning threshold. Warning ranges are vendor-specific.
Laser bias current low warning threshold	Laser bias current low warning threshold. Warning ranges are vendor-specific.
Laser output power high alarm threshold	Laser output power high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser output power low alarm threshold	Laser output power low alarm threshold. Alarm threshold ranges are vendor-specific.
Laser output power high warning threshold	Laser output power high warning threshold. Warning ranges are vendor-specific.
Laser output power low warning threshold	Laser output power low warning threshold. Warning ranges are vendor-specific.

**Table 11: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers (*continued*)**

Field Name	Field Description
Module temperature high alarm threshold	Module temperature high alarm threshold. Alarm threshold ranges are vendor-specific.
Module temperature low alarm threshold	Module temperature low alarm threshold. Alarm threshold ranges are vendor-specific.
Module temperature high warning threshold	Module temperature high warning threshold. Warning ranges are vendor-specific.
Module temperature low warning threshold	Module temperature low warning threshold. Warning ranges are vendor-specific.
Module voltage high alarm threshold	Module voltage high alarm threshold. Alarm ranges are vendor-specific.
Module voltage low alarm threshold	Module voltage low alarm threshold. Alarm ranges are vendor-specific.
Module voltage high warning threshold	Module voltage high warning threshold. Warning ranges are vendor-specific.
Module voltage low warning threshold	Module voltage low warning threshold. Warning ranges are vendor-specific.
Laser rx power high alarm threshold	Laser receive power high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser rx power low alarm threshold	Laser receive power low alarm threshold. Alarm threshold ranges are vendor-specific.
Laser rx power high warning threshold	Laser receive power high warning threshold. Warning threshold ranges are vendor-specific.
Laser rx power high low threshold	Laser receive power high warning threshold. Warning threshold ranges are vendor-specific.

Table 12 on page 130 lists the output fields for the **show interfaces diagnostics optics** command for 10-Gigabit Ethernet transceivers. Output fields are listed in the approximate order in which they appear.

**Table 12: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers**

Field Name	Field Description
Physical interface	Name of the physical interface.
Laser bias current	Measured laser bias current in mA.

Table 12: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers (*continued*)

Field Name	Field Description
Laser output power	Measured laser output power in mW.
Module temperature	Internally measured module temperature.
Laser rx power	Measured receive optical power in mW.
Laser bias current high alarm	Laser bias current high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.
Laser bias current low alarm	Laser bias current low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.
Laser output power high alarm	Laser output power high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.
Laser output power low alarm	Laser output power low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.
Module temp high alarm	Module temperature high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.
Module temp low alarm	Module temperature low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.
Laser rx power high alarm	Laser receive power high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.
Laser rx power low alarm	Laser receive power low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.
Laser bias current high warning	Laser bias current high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.
Laser bias current low warning	Laser bias current low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.
Laser output power high warning	Laser output power high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.
Laser output power low warning	Laser output power low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.
Module temperature high warning	Module temperature high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.
Module temperature low warning	Module temperature low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.

Table 12: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers (*continued*)

Field Name	Field Description
Laser rx power high warning	Laser receive power high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.
Laser rx power low warning	Laser receive power low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.
Laser bias current high alarm threshold	Laser bias current high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser bias current low alarm threshold	Laser bias current low alarm threshold. Alarm threshold ranges are vendor-specific.
Laser output power high alarm threshold	Laser output power high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser output power low alarm threshold	Laser output power low alarm threshold. Alarm threshold ranges are vendor-specific.
Module temperature high alarm threshold	Module temperature high alarm threshold. Alarm threshold ranges are vendor-specific.
Module temperature low alarm threshold	Module temperature low alarm threshold. Alarm threshold ranges are vendor-specific.
Laser rx power high alarm threshold	Laser receive power high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser rx power low alarm threshold	Laser receive power low alarm threshold. Alarm threshold ranges are vendor-specific.
Laser bias current high warning threshold	Laser bias current high warning threshold. Warning ranges are vendor-specific.
Laser bias current low warning threshold	Laser bias current low warning threshold. Warning ranges are vendor-specific.
Laser output power high warning threshold	Laser output power high warning threshold. Warning ranges are vendor-specific.
Laser output power low warning threshold	Laser output power low warning threshold. Warning ranges are vendor-specific.
Module temperature high warning threshold	Module temperature high warning threshold. Warning ranges are vendor-specific.
Module temperature low warning threshold	Module temperature low warning threshold. Warning ranges are vendor-specific.

**Table 12: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers (*continued*)**

Field Name	Field Description
Laser rx power high warning threshold	Laser receive power high warning threshold. Warning threshold ranges are vendor-specific.
Laser rx power low warning threshold	Laser receive power low warning threshold. Warning threshold ranges are vendor-specific.

[Table 13 on page 133](#) lists the output fields for the **show interfaces diagnostics optics** command for XFP transceivers. Output fields are listed in the approximate order in which they appear.

**Table 13: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers**

Field Name	Field Description
Physical interface	Name of the physical interface.
Laser bias current	Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm). This is a software equivalent to the <b>LsPOWMON</b> pin in hardware.
Module temperature	Temperature of the XFP optics module, in Celsius and Fahrenheit.
Laser rx power	Laser received optical power, in mW and dBm.
Laser bias current high alarm	Laser bias power setting high alarm. Displays <b>on</b> or <b>off</b> .
Laser bias current low alarm	Laser bias power setting low alarm. Displays <b>on</b> or <b>off</b> .
Laser bias current high warning	Laser bias power setting high warning. Displays <b>on</b> or <b>off</b> .
Laser bias current low warning	Laser bias power setting low warning. Displays <b>on</b> or <b>off</b> .
Laser output power high alarm	Laser output power high alarm. Displays <b>on</b> or <b>off</b> .
Laser output power low alarm	Laser output power low alarm. Displays <b>on</b> or <b>off</b> .
Laser output power high warning	Laser output power high warning. Displays <b>on</b> or <b>off</b> .

Table 13: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers (*continued*)

Field Name	Field Description
Laser output power low warning	Laser output power low warning. Displays <b>on</b> or <b>off</b> .
Module temperature high alarm	Module temperature high alarm. Displays <b>on</b> or <b>off</b> .
Module temperature low alarm	Module temperature low alarm. Displays <b>on</b> or <b>off</b> .
Module temperature high warning	Module temperature high warning. Displays <b>on</b> or <b>off</b> .
Module temperature low warning	Module temperature low warning. Displays <b>on</b> or <b>off</b> .
Laser rx power high alarm	Receive laser power high alarm. Displays <b>on</b> or <b>off</b> .
Laser rx power low alarm	Receive laser power low alarm. Displays <b>on</b> or <b>off</b> .
Laser rx power high warning	Receive laser power high warning. Displays <b>on</b> or <b>off</b> .
Laser rx power low warning	Receive laser power low warning. Displays <b>on</b> or <b>off</b> .
Module not ready alarm	Module not ready alarm. When <b>on</b> , indicates the module has an operational fault. Displays <b>on</b> or <b>off</b> .
Module power down alarm	Module power down alarm. When <b>on</b> , module is in a limited power mode, low for normal operation. Displays <b>on</b> or <b>off</b> .
Tx data not ready alarm	Any condition leading to invalid data on the transmit path. Displays <b>on</b> or <b>off</b> .
Tx not ready alarm	Any condition leading to invalid data on the transmit path. Displays <b>on</b> or <b>off</b> .
Tx laser fault alarm	Laser fault condition. Displays <b>on</b> or <b>off</b> .
Tx CDR loss of lock alarm	Transmit clock and data recovery (CDR) loss of lock. Loss of lock on the transmit side of the CDR. Displays <b>on</b> or <b>off</b> .
Rx not ready alarm	Any condition leading to invalid data on the receive path. Displays <b>on</b> or <b>off</b> .
Rx loss of signal alarm	Receive Loss of Signal alarm. When <b>on</b> , indicates insufficient optical input power to the module. Displays <b>on</b> or <b>off</b> .
Rx CDR loss of lock alarm	Receive CDR loss of lock. Loss of lock on the receive side of the CDR. Displays <b>on</b> or <b>off</b> .

**Table 13: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers (continued)**

Field Name	Field Description
Laser bias current high alarm threshold	Vendor-specified threshold for the laser bias current high alarm: <b>130.000 mA</b> .
Laser bias current low alarm threshold	Vendor-specified threshold for the laser bias current low alarm: <b>10.000 mA</b> .
Laser bias current high warning threshold	Vendor-specified threshold for the laser bias current high warning: <b>120.000 mA</b> .
Laser bias current low warning threshold	Vendor-specified threshold for the laser bias current low warning: <b>12.000 mA</b> .
Laser output power high alarm threshold	Vendor-specified threshold for the laser output power high alarm: <b>0.8910 mW</b> or <b>-0.50 dBm</b> .
Laser output power low alarm threshold	Vendor-specified threshold for the laser output power low alarm: <b>0.2230 mW</b> or <b>-6.52 dBm</b> .
Laser output power high warning threshold	Vendor-specified threshold for the laser output power high warning: <b>0.7940 mW</b> or <b>-100 dBm</b> .
Laser output power low warning threshold	Vendor-specified threshold for the laser output power low warning: <b>0.2510 mW</b> or <b>-600 dBm</b> .
Module temperature high alarm threshold	Vendor-specified threshold for the module temperature high alarm: <b>90° C</b> or <b>194° F</b> .
Module temperature low alarm threshold	Vendor-specified threshold for the module temperature low alarm: <b>-5° C</b> or <b>23° F</b> .
Module temperature high warning threshold	Vendor-specified threshold for the module temperature high warning: <b>85 ° C</b> or <b>185 ° F</b> .
Module temperature low warning threshold	Vendor-specified threshold for the module temperature low warning: <b>0° C</b> or <b>32° F</b> .
Laser rx power high alarm threshold	Vendor-specified threshold for the laser Rx power high alarm: <b>1.2589 mW</b> or <b>1.00 dBm</b> .
Laser rx power low alarm threshold	Vendor-specified threshold for the laser Rx power low alarm: <b>0.0323 mW</b> or <b>-14.91 dBm</b> .
Laser rx power high warning threshold	Vendor-specified threshold for the laser Rx power high warning: <b>1.1220 mW</b> or <b>0.50 dBm</b> .
Laser rx power low warning threshold	Vendor-specified threshold for the laser Rx power low warning: <b>0.0363 mW</b> or <b>-14.40 dBm</b> .

## Sample Output

### show interfaces diagnostics optics (DWDM and DWDM OTN)

```

user@host> show interfaces diagnostics optics ge-5/0/0
Physical interface: ge-5/0/0
Laser bias current           : 79.938 mA
Laser output power          : 1.592 mW / 2.02 dBm
Receiver signal average optical power : 1.3854 mW / 1.42 dBm
Laser end-of-life alarm     : Off
Laser wavelength alarm      : Off
Laser bias current alarm    : Off
Laser temperature alarm     : Off
Laser power alarm           : Off
Modulator temperature alarm : Off
Modulator bias alarm        : Off
Tx multiplexer FIFO error alarm : Off
Tx loss of PLL lock alarm   : Off
Rx loss of average optical power alarm: Off
Rx loss of AC power alarm   : Off
Rx loss of PLL lock alarm   : Off

```

### show interfaces diagnostics optics (Bidirectional SFP)

```

user@host> show interfaces diagnostics optics ge-3/0/6
Physical interface: ge-3/0/6
Laser bias current           : 13.356 mA
Laser output power          : 0.2210 mW / -6.56 dBm
Module temperature           : 36 degrees C / 96 degrees F
Module voltage               : 3.2180 V
Receiver signal average optical power : 0.2429 mW / -6.15 dBm
Laser bias current high alarm : Off
Laser bias current low alarm  : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm  : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm  : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm     : Off
Module voltage low alarm      : Off
Module voltage high warning   : Off
Module voltage low warning    : Off
Laser rx power high alarm     : Off
Laser rx power low alarm      : Off
Laser rx power high warning   : Off
Laser rx power low warning    : Off
Laser bias current high alarm threshold : 70.000 mA
Laser bias current low alarm threshold : 0.002 mA
Laser bias current high warning threshold : 65.000 mA
Laser bias current low warning threshold : 0.002 mA
Laser output power high alarm threshold : 1.0000 mW / 0.00 dBm
Laser output power low alarm threshold : 0.0560 mW / -12.52 dBm
Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
Laser output power low warning threshold : 0.0890 mW / -10.51 dBm
Module temperature high alarm threshold : 100 degrees C / 212 degrees F
Module temperature low alarm threshold : -50 degrees C / -58 degrees F

```

```

Module temperature high warning threshold : 95 degrees C / 203 degrees F
Module temperature low warning threshold : -48 degrees C / -54 degrees F
Module voltage high alarm threshold      : 3.700 V
Module voltage low alarm threshold       : 2.900 V
Module voltage high warning threshold    : 3.600 V
Module voltage low warning threshold     : 3.000 V
Laser rx power high alarm threshold      : 1.9953 mW / 3.00 dBm
Laser rx power low alarm threshold       : 0.0001 mW / -40.00 dBm
Laser rx power high warning threshold    : 1.0000 mW / 0.00 dBm
Laser rx power low warning threshold     : 0.0010 mW / -30.00 dBm

```

### show interfaces diagnostics optics (SFP)

```
user@host> show interfaces diagnostics optics ge-0/3/0
```

```
Physical interface: ge-0/3/0
```

```

Laser bias current          : 23.408 mA
Laser output power          : 1.479 mW / 1.70 dBm
Module temperature          : 37 degrees C / 99 degrees F
Laser rx power              : 0.121 mW / -9.16 dBm
Laser bias current high alarm : Off
Laser bias current low alarm  : Off
Laser output power high alarm : Off
Laser output power low alarm  : Off
Module temperature high alarm : Off
Module temperature low alarm  : Off
Laser rx power high alarm     : Off
Laser rx power low alarm     : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high warning : Off
Module temperature low warning : Off
Laser rx power high warning   : Off
Laser rx power low warning   : Off
Laser bias current high alarm threshold : 31.000 mA
Laser bias current low alarm threshold  : 10.000 mA
Laser output power high alarm threshold : 6.000 mW / 7.78 dBm
Laser output power low alarm threshold  : 0.100 mW / -10.00 dBm
Module temperature high alarm threshold : 85 degrees C / 185 degrees F
Module temperature low alarm threshold  : 0 degrees C / 32 degrees F
Laser rx power high alarm threshold     : 1.000 mW / 0.00 dBm
Laser rx power low alarm threshold      : 0.001 mW / -30.00 dBm
Laser bias current high warning threshold : 28.000 mA
Laser bias current low warning threshold  : 11.000 mA
Laser output power high warning threshold : 5.000 mW / 6.99 dBm
Laser output power low warning threshold  : 0.500 mW / -3.01 dBm
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold  : 10 degrees C / 50 degrees F
Laser rx power high warning threshold    : 0.501 mW / -3.00 dBm
Laser rx power low warning threshold     : 0.001 mW / -28.86 dBm

```

### show interfaces diagnostics optics (SFP)

```
user@host> show interfaces diagnostics optics ge-1/0/0
```

```
Physical interface: ge-1/0/0
```

```

Laser bias current          : 49.010 mA
Laser output power          : 1.263 mW / 1.01 dBm
Module temperature          : 17 degrees C / 62 degrees F

Module voltage              : 4.21 V

```

```

Laser rx power : 0.060 mW / -12.21 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Module temperature high alarm : Off
Module temperature low alarm : Off
Module voltage high alarm : Off
Module voltage low alarm : Off
Laser rx power high alarm : Off
Laser rx power low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high warning : Off
Module voltage low warning : Off
Laser rx power high warning : Off
Laser rx power low warning : Off
Laser bias current high alarm threshold : 70.000 mA
Laser bias current low alarm threshold : 20.000 mA
Laser bias current high warning threshold : 65.000 mA
Laser bias current low warning threshold : 25.000 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.1990 mW / -7.01 dBm
Laser output power high warning threshold : 1.2580 mW / 1.00 dBm
Laser output power low warning threshold : 0.2230 mW / -6.52 dBm
Module temperature high alarm threshold : 78 degrees C / 172 degrees F

Module temperature low alarm threshold : 13 degrees C / 9 degrees F
Module temperature high warning threshold : 75 degrees C / 167 degrees F

Module temperature low warning threshold : 10 degrees C / 14 degrees F

Module voltage high alarm threshold : 5.71 V
Module voltage low alarm threshold : 2.05 V
Module voltage high warning threshold : 5.20 V
Module voltage low warning threshold : 3.11 V
Laser rx power high alarm threshold : 1.7783 mW / 2.50 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 1.5849 mW / 2.00 dBm
Laser rx power low warning threshold : 0.0158 mW / -18.01 dBm

```

### show interfaces diagnostics optics (XFP and CFP Optics)

```

user@host> show interfaces diagnostics optics xe-2/1/0
Physical interface: xe-2/1/0
Laser bias current : 52.060 mA
Laser output power : 0.5640 mW / -2.49 dBm
Module temperature : 31 degrees C / 88 degrees F
Laser rx power : 0.0844 mW / -10.74 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off

```

```

Module temperature high alarm      : Off
Module temperature low alarm       : Off
Module temperature high warning    : Off
Module temperature low warning     : Off
Laser rx power high alarm          : Off
Laser rx power low alarm           : Off
Laser rx power high warning        : Off
Laser rx power low warning         : Off
Module not ready alarm             : Off
Module power down alarm            : Off
Tx data not ready alarm            : Off
Tx not ready alarm                 : Off
Tx laser fault alarm               : Off
Tx CDR loss of lock alarm          : Off
Rx not ready alarm                 : Off
Rx loss of signal alarm            : Off
Rx CDR loss of lock alarm          : Off
Laser bias current high alarm threshold : 130.000 mA
Laser bias current low alarm threshold : 10.000 mA
Laser bias current high warning threshold : 120.000 mA
Laser bias current low warning threshold : 12.000 mA
Laser output power high alarm threshold : 0.8910 mW / -0.50 dBm
Laser output power low alarm threshold : 0.2230 mW / -6.52 dBm
Laser output power high warning threshold : 0.7940 mW / -1.00 dBm
Laser output power low warning threshold : 0.2510 mW / -6.00 dBm
Module temperature high alarm threshold : 90 degrees C / 194 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 85 degrees C / 185 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Laser rx power high alarm threshold : 1.2589 mW / 1.00 dBm
Laser rx power low alarm threshold : 0.0323 mW / -14.91 dBm
Laser rx power high warning threshold : 1.1220 mW / 0.50 dBm
Laser rx power low warning threshold : 0.0363 mW / -14.40 dBm

```

#### show interfaces diagnostics optics for 10-Gigabit Ethernet (PTX 24-10GE-SFPP)

```

user@host> show interfaces diagnostics optics et-2/0/23
Physical interface: et-2/0/23
Laser bias current                : 8.482 mA
Laser output power                 : 0.5890 mW / -2.30 dBm
Module temperature                 : 51 degrees C / 123 degrees F
Module voltage                     : 3.2970 V
Receiver signal average optical power : 0.5574 mW / -2.54 dBm
Laser bias current high alarm      : Off
Laser bias current low alarm       : Off
Laser bias current high warning    : Off
Laser bias current low warning     : Off
Laser output power high alarm      : Off
Laser output power low alarm       : Off
Laser output power high warning    : Off
Laser output power low warning     : Off
Module temperature high alarm      : Off
Module temperature low alarm       : Off
Module temperature high warning    : Off
Module temperature low warning     : Off
Module voltage high alarm          : Off
Module voltage low alarm           : Off
Module voltage high warning        : Off
Module voltage low warning         : Off
Laser rx power high alarm          : Off
Laser rx power low alarm           : Off

```

```

Laser rx power high warning      : Off
Laser rx power low warning       : Off
Laser bias current high alarm threshold : 11.800 mA
Laser bias current low alarm threshold : 4.000 mA
Laser bias current high warning threshold : 10.800 mA
Laser bias current low warning threshold : 5.000 mA
Laser output power high alarm threshold : 0.8310 mW / -0.80 dBm
Laser output power low alarm threshold : 0.2510 mW / -6.00 dBm
Laser output power high warning threshold : 0.6600 mW / -1.80 dBm
Laser output power low warning threshold : 0.3160 mW / -5.00 dBm
Module temperature high alarm threshold : 93 degrees C / 199 degrees F
Module temperature low alarm threshold : -13 degrees C / 9 degrees F
Module temperature high warning threshold : 88 degrees C / 190 degrees F
Module temperature low warning threshold : -8 degrees C / 18 degrees F
Module voltage high alarm threshold : 3.700 V
Module voltage low alarm threshold : 2.900 V
Module voltage high warning threshold : 3.600 V
Module voltage low warning threshold : 3.000 V
Laser rx power high alarm threshold : 1.0000 mW / 0.00 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.0158 mW / -18.01 dBm

```

#### show interfaces diagnostics optics for 40-Gigabit Ethernet

```
user@host> show interfaces diagnostics optics et-7/1/0
```

```
Physical interface: et-7/1/0
```

```

Module temperature      : 34 degrees C / 94 degrees F
Module voltage          : 3.4720 V
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Module voltage low alarm : Off
Module voltage high warning : Off
Module voltage low warning : Off
Module not ready alarm : Off
Module low power alarm : Off
Module initialization incomplete alarm : Off
Module fault alarm : Off
PLD Flash initialization fault alarm : Off
Power supply fault alarm : Off
Checksum fault alarm : Off
Tx laser disabled alarm : Off
Tx loss of signal functionality alarm : Off
Tx CDR loss of lock alarm : Off
Rx loss of signal alarm : Off
Rx CDR loss of lock alarm : Off
Module temperature high alarm threshold : 80 degrees C / 176 degrees F
Module temperature low alarm threshold : -10 degrees C / 14 degrees F
Module temperature high warning threshold : 75 degrees C / 167 degrees F
Module temperature low warning threshold : -5 degrees C / 23 degrees F
Module voltage high alarm threshold : 3.5990 V
Module voltage low alarm threshold : 3.0000 V
Module voltage high warning threshold : 3.5000 V
Module voltage low warning threshold : 3.0990 V
Laser bias current high alarm threshold : 100.000 mA
Laser bias current low alarm threshold : 10.000 mA
Laser bias current high warning threshold : 80.000 mA

```

```

Laser bias current low warning threshold : 15.000 mA
Laser output power high alarm threshold : 2.8180 mW / 4.50 dBm
Laser output power low alarm threshold : 0.2390 mW / -6.22 dBm
Laser output power high warning threshold : 2.2380 mW / 3.50 dBm
Laser output power low warning threshold : 0.3010 mW / -5.21 dBm
Laser rx power high alarm threshold : 2.5119 mW / 4.00 dBm
Laser rx power low alarm threshold : 0.0316 mW / -15.00 dBm
Laser rx power high warning threshold : 1.9953 mW / 3.00 dBm
Laser rx power low warning threshold : 0.0631 mW / -12.00 dBm
Laser temperature high alarm threshold : 80 degrees C / 176 degrees F
Laser temperature low alarm threshold : -10 degrees C / 14 degrees F
Laser temperature high warning threshold : 75 degrees C / 167 degrees F
Laser temperature low warning threshold : -5 degrees C / 23 degrees F
Lane 0
Laser bias current : 27.829 mA
Laser output power : 0.851 mW / -0.70 dBm
Laser temperature : 34 degrees C / 94 degrees F
Laser receiver power : 0.894 mW / -0.49 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Laser temperature high alarm : Off
Laser temperature low alarm : Off
Laser temperature high warning : Off
Laser temperature low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off
Tx CDR loss of lock alarm : Off
Rx loss of signal alarm : Off
Rx CDR loss of lock alarm : Off
APD supply fault alarm : Off
TEC fault alarm : Off
Wavelength unlocked alarm : Off
Lane 1
Laser bias current : 35.374 mA
Laser output power : 0.896 mW / -0.48 dBm
Laser temperature : 34 degrees C / 94 degrees F
Laser receiver power : 0.707 mW / -1.50 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Laser temperature high alarm : Off
Laser temperature low alarm : Off
Laser temperature high warning : Off
Laser temperature low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off

```

```

Laser receiver power low warning           : Off
Tx loss of signal functionality alarm       : Off
Tx CDR loss of lock alarm                  : Off
Rx loss of signal alarm                    : Off
Rx CDR loss of lock alarm                  : Off
APD supply fault alarm                     : Off
TEC fault alarm                           : Off
Wavelength unlocked alarm                  : Off
Lane 2
Laser bias current                         : 29.173 mA
Laser output power                         : 0.890 mW / -0.51 dBm
Laser temperature                         : 34 degrees C / 94 degrees F
Laser receiver power                      : 0.704 mW / -1.52 dBm
Laser bias current high alarm              : Off
Laser bias current low alarm               : Off
Laser bias current high warning            : Off
Laser bias current low warning             : Off
Laser output power high alarm              : Off
Laser output power low alarm               : Off
Laser output power high warning            : Off
Laser output power low warning             : Off
Laser temperature high alarm               : Off
Laser temperature low alarm                : Off
Laser temperature high warning             : Off
Laser temperature low warning              : Off
Laser receiver power high alarm            : Off
Laser receiver power low alarm             : Off
Laser receiver power high warning          : Off
Laser receiver power low warning           : Off
Tx loss of signal functionality alarm       : Off
Tx CDR loss of lock alarm                  : Off
Rx loss of signal alarm                    : Off
Rx CDR loss of lock alarm                  : Off
APD supply fault alarm                     : Off
TEC fault alarm                           : Off
Wavelength unlocked alarm                  : Off
Lane 3
Laser bias current                         : 36.164 mA
Laser output power                         : 0.899 mW / -0.46 dBm
Laser temperature                         : 34 degrees C / 94 degrees F
Laser receiver power                      : 0.892 mW / -0.50 dBm
Laser bias current high alarm              : Off
Laser bias current low alarm               : Off
Laser bias current high warning            : Off
Laser bias current low warning             : Off
Laser output power high alarm              : Off
Laser output power low alarm               : Off
Laser output power high warning            : Off
Laser output power low warning             : Off
Laser temperature high alarm               : Off
Laser temperature low alarm                : Off
Laser temperature high warning             : Off
Laser temperature low warning              : Off
Laser receiver power high alarm            : Off
Laser receiver power low alarm             : Off
Laser receiver power high warning          : Off
Laser receiver power low warning           : Off
Tx loss of signal functionality alarm       : Off
Tx CDR loss of lock alarm                  : Off
Rx loss of signal alarm                    : Off
Rx CDR loss of lock alarm                  : Off

```

APD supply fault alarm	: Off
TEC fault alarm	: Off
Wavelength unlocked alarm	: Off

## show interfaces (Fast Ethernet)

<b>Syntax</b>	<pre>show interfaces <i>interface-type</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 Fast Ethernet interface.
<b>Options</b>	<p><b><i>interface-type</i></b>—On M Series and T Series routers, the interface type is <b><i>fe-fpc/pic/port</i></b>. On the J Series routers, the interface type is <b><i>fe-pim/O/port</i></b>.</p> <p><b><i>brief   detail   extensive   terse</i></b>—(Optional) Display the specified level of output.</p> <p><b><i>descriptions</i></b>—(Optional) Display interface description strings.</p> <p><b><i>media</i></b>—(Optional) Display media-specific information about network interfaces.</p> <p><b><i>snmp-index snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b><i>statistics</i></b>—(Optional) Display static interface statistics.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show interfaces (Fast Ethernet) on page 157</a></p> <p><a href="#">show interfaces brief (Fast Ethernet) on page 158</a></p> <p><a href="#">show interfaces detail (Fast Ethernet) on page 158</a></p> <p><a href="#">show interfaces extensive (Fast Ethernet) on page 158</a></p>
<b>Output Fields</b>	<p><a href="#">Table 14 on page 144</a> lists the output fields for the <b>show interfaces Fast Ethernet</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 14: show interfaces Fast 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>

Table 14: show interfaces Fast 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>
<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>Link-mode</b>	Type of link connection configured for the physical interface: <b>Full-duplex</b> or <b>Half-duplex</b>	<b>extensive</b>
<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
<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

Table 14: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<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>	(GigabitEthernet 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).	None specified
<b>Output Rate</b>	Output rate in bps and pps.	None specified
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>
<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> <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>

Table 14: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<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 <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 Fast 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>	Total number of egress queues supported on the specified interface.	<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>
<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>	<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>	<b>extensive</b>

Table 14: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Active alarms and Active defects</b>	<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 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>.</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>	<b>detail extensive none</b>
<b>OTN FEC statistics</b>	<p>The forward error correction (FEC) counters provide the following statistics:</p> <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>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>—High bit error rate. Indicates the number of bit errors when the PCS receiver is operating in normal mode.</li> <li>• <b>Errored blocks</b>—Loss of block lock. The number of errored blocks when PCS receiver is operating in normal mode.</li> </ul>	<b>detail extensive</b>

Table 14: show interfaces Fast 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 and 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, Broadcast packets, and 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>	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 Fast 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 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>	<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 Fast 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 14: show interfaces Fast 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 Fast 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 routing device 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 14: show interfaces Fast 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 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
<b>VLAN-Tag</b>	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. <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-pop</b>—An outer VLAN tag is pushed in front of the existing VLAN tag, and the outer VLAN tag of the incoming frame is removed.</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>

Table 14: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Demux:</b>	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: <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>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>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>	Number and rate of bytes and packets received and transmitted on the specified interface set. <ul style="list-style-type: none"> <li><b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface set</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>
<b>Transit statistics</b>	Number and rate of bytes and packets transiting the switch. <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, 0 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>

Table 14: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<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 interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>

## Sample Output

### show interfaces (Fast Ethernet)

```

user@host> show interfaces fe-0/0/0
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
  Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:44 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198)
    Flags: SNMP-Traps Encapsulation: ENET2

```

```
Protocol inet, MTU: 1500
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
```

### show interfaces brief (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0 brief
Physical interface: fe-0/0/0, Enabled, Physical link is Up
Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Logical interface fe-0/0/0.0
Flags: SNMP-Traps Encapsulation: ENET2
inet 10.10.10.1/24
```

### show interfaces detail (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0 detail
Physical interface: fe-0/0/0, Enabled, Physical link is Up
Interface index: 128, SNMP ifIndex: 22, Generation: 5391
Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
CoS queues : 4 supported, 4 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
Last flapped : 2006-01-20 14:50:58 PST (2w4d 00:45 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 42 0 bps
Input packets: 0 0 pps
Output packets: 1 0 pps
Active alarms : None
Active defects : None
Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
Flags: SNMP-Traps Encapsulation: ENET2
Protocol inet, MTU: 1500, Generation: 105, Route table: 0
Flags: Is-Primary, Mac-Validate-Strict
Mac-Validate Failures: Packets: 0, Bytes: 0
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255,
Generation: 136
```

### show interfaces extensive (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0 extensive
Physical interface: fe-0/0/0, Enabled, Physical link is Up
Interface index: 128, SNMP ifIndex: 22, Generation: 5391
Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed:
100mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
CoS queues : 4 supported, 4 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
Last flapped : 2006-01-20 14:50:58 PST (2w4d 00:46 ago)
Statistics last cleared: Never
```

```

Traffic statistics:
Input bytes :          0          0 bps
Output bytes :         42          0 bps
Input packets:         0          0 pps
Output packets:        1          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: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Active alarms : None
Active defects : None
MAC statistics:
Total octets          Receive          Transmit
Total packets         0              1
Unicast packets       0              0
Broadcast packets     0              1
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   1
Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link partner:
  Link partner: Full-duplex, Flow control: None, Remote fault: Ok
Local resolution:
Packet Forwarding Engine configuration:
Destination slot: 0
CoS information:
      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 fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
Flags: SNMP-Traps Encapsulation: ENET2
Protocol inet, MTU: 1500, Generation: 105, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255,
Generation: 136

```

## show interfaces (10-Gigabit Ethernet)

---

<b>Syntax</b>	<code>show interfaces <i>xe-fpc/pic/port</i></code> <code>&lt;brief   detail   extensive   terse&gt;</code> <code>&lt;descriptions&gt;</code> <code>&lt;media&gt;</code> <code>&lt;snmp-index <i>snmp-index</i>&gt;</code> <code>&lt;statistics&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 8.0.
<b>Description</b>	(M320, M120, MX Series, and T Series routers and EX Series switches only) Display status information about the specified 10-Gigabit Ethernet interface.
<b>Options</b>	<p><code><i>xe-fpc/pic/port</i></code>—Display standard information about the specified 10-Gigabit Ethernet interface.</p> <p><code>brief   detail   extensive   terse</code>—(Optional) Display the specified level of output.</p> <p><code>descriptions</code>—(Optional) Display interface description strings.</p> <p><code>media</code>—(Optional) Display media-specific information about network interfaces.</p> <p><code>snmp-index <i>snmp-index</i></code>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><code>statistics</code>—(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 175</a></p> <p><a href="#">show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode) on page 178</a></p> <p><a href="#">show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC) on page 180</a></p> <p><a href="#">show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode) on page 182</a></p> <p><a href="#">show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only) on page 182</a></p> <p><a href="#">show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only) on page 183</a></p>
<b>Output Fields</b>	See <a href="#">Table 15 on page 161</a> for the output fields for the <b>show interfaces</b> (10-Gigabit Ethernet) command.

Table 15: 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 15: 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 15: 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 15 on page 161</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 <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 15: 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>	Total number of egress queues supported on the specified interface.	<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>
<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>	<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>	<b>extensive</b>

Table 15: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Active alarms and Active defects</b>	<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 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>.</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>	<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>	<p>The FECmode configured on the interface.</p> <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>	<p>OTN mode.</p> <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>	<p>The forward error correction (FEC) counters for the DWDM OTN PIC.</p> <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>
<b>OTN FEC alarms</b>	<p>OTN FEC excessive or degraded error alarms triggered on the interface.</p> <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>	<p>OTN OC defects triggered on the interface.</p> <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>

Table 15: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<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>—High bit error rate. Indicates the number of bit errors when the PCS receiver is operating in normal mode.</li> <li><b>Errored blocks</b>—Loss of block lock. The number of errored blocks when PCS receiver is operating in normal mode.</li> </ul>	<b>detail extensive</b>

Table 15: 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 16 on page 175</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 15: 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 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>	<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 15: 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 15: 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 15: 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 15: 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 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 15: 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>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>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 routing device.	<b>extensive</b>

Table 15: 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 15: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

For Gigabit Ethernet IQ PICs, traffic and MAC statistics output varies. [Table 16 on page 175](#) 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 16 on page 175](#), 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 16: 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
Bit errors          0
Errored blocks      0
Seconds

```

```

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 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                    %      bps      %      usec      low
  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 interface-set (Ethernet Interface Set)

<b>Syntax</b>	<code>show interfaces interface-set <i>interface-set-name</i></code> <detail   terse>
<b>Release Information</b>	Command introduced in Junos OS Release 8.5.
<b>Description</b>	<p>Display information about the specified gigabit or 10-Gigabit Ethernet interface set. Supported in MX Series routers with enhanced queuing DPCs or MPCs.</p> <p>You can also use the <b>show interfaces interface-set</b> command to display information about agent circuit identifier (ACI) interface sets configured on MX Series routers with MPCs/MICs.</p>
<b>Options</b>	<p><b>interface-set <i>interface-set-name</i></b>—Display information about the specified Gigabit Ethernet, 10-Gigabit Ethernet, or ACI interface set.</p> <p><b>detail   terse</b>—(Optional) Display the specified level of output.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration</i></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show interfaces interface-set terse on page 186</a></p> <p><a href="#">show interfaces interface-set detail on page 186</a></p> <p><a href="#">show interfaces interface-set (ACI Interface Set) on page 187</a></p>
<b>Output Fields</b>	<a href="#">Table 17 on page 185</a> describes the information for the <b>show interfaces interface-set</b> command.

Table 17: Ethernet show interfaces interface-set Output Fields

Field Name	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Interface set</b>	Name of the interface set or sets.	All levels
<b>Interface set index</b>	<p>Index number of the interface set. For ACI interface sets, the following fields are displayed:</p> <ul style="list-style-type: none"> <li><b>ACI VLAN</b>—ACI interface set that the router uses to create dynamic VLAN subscriber interfaces based on the agent circuit identifier value.</li> <li><b>PPPoE</b>—Dynamic PPPoE subscriber interface that the router creates using the ACI interface set.</li> </ul>	<b>detail none</b>
<b>Agent Circuit ID</b>	For ACI interface sets, string in DHCP or PPPoE control packets that uniquely identifies the subscriber's access node and the DSL line on the access node.	<b>detail none</b>
<b>Max Sessions</b>	For dynamic PPPoE subscriber interfaces, maximum number of PPPoE logical interfaces that that can be activated on the underlying interface.	<b>detail none</b>

Table 17: Ethernet show interfaces interface-set Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Max Sessions VSA Ignore</b>	For dynamic PPPoE subscriber interfaces, whether the router is configured to ignore (clear) the PPPoE maximum session value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks VSA [26-143] and restore the PPPoE maximum session value on the underlying interface to the value configured with the <b>max-sessions</b> statement: <b>Off</b> (default) or <b>On</b> .	<b>detail none</b>
<b>Traffic statistics</b>	Number and rate of bytes and packets received and transmitted on the specified interface set.  <ul style="list-style-type: none"> <li><b>Input bytes, Output bytes</b>—Number of bytes and number of bytes per second received and transmitted on the interface set</li> <li><b>Input packets, Output packets</b>—Number of packets and number of packets per second received and transmitted on the interface set.</li> </ul>	<b>detail</b>
<b>Egress queues supported</b>	Total number of egress queues supported on the specified interface set.	<b>detail</b>
<b>Egress queues in use</b>	Total number of egress queues used on the specified interface set.	<b>detail</b>
<b>Queue counters</b>	<b>Queued packets, Transmitted packets, and Dropped packets</b> statistics for the four forwarding classes.	<b>detail</b>
<b>Members</b>	List of all interface sets or, for ACI interface sets, list of all subscriber interfaces belonging to the specified ACI interface set.	<b>detail none</b>

## Sample Output

### show interfaces interface-set terse

```

user@host> show interfaces interface-set terse
Interface set:
  iflset-xe-11/3/0-0
  ge-1/0/1-0
  ge-1/0/1-2

```

### show interfaces interface-set detail

```

user@host> show interfaces interface-set iflset-xe-11/3/0-0 detail
Interface set: iflset-xe-11/3/0-0
Interface set index: 19
Traffic statistics:
  Output bytes :           751017840           401673504 bps
  Output packets:         11044380           738377 pps
Egress queues: 4 supported, 4 in use
Queue counters:
  Queued packets  Transmitted packets  Dropped packets
0 best-effort    211091327          11044380          199995746
1 expedited-fo           0                0                0
2 assured-forw           0                0                0
3 network-cont          0                0                0
Members:
  xe-11/3/0.0

```

### show interfaces interface-set (ACI Interface Set)

```
user@host> show interfaces interface-set
Interface set: aci-1001-demux0.1073741826
Interface set index: 1
  ACI VLAN:
    Agent Circuit ID: aci-ppp-dhcp-dvlan-60
  PPPoE:
    Max Sessions: 3, Max Sessions VSA Ignore: Off
Members:
  pp0.1073741827
```

## show interfaces interface-set queue

<b>Syntax</b>	show interfaces interface-set queue <i>interface-set-name</i> <aggregate   remaining-traffic> <forwarding-class <i>class-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 8.5.
<b>Description</b>	Display information about the gigabit or 10-Gigabit Ethernet interface set queue. Supported in MX Series routers with enhanced queuing DPCs.
<b>Options</b>	<p><b><i>interface-set-name</i></b>—(Optional) Display information about the specified gigabit or 10-Gigabit Ethernet interface set. Wildcard values can be used in the interface set name.</p> <p><b>aggregate</b>—(Optional) Display the aggregated queuing statistics of all member logical interfaces for interface sets that have traffic-control profiles configured.</p> <p><b>both-ingress-egress</b>—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics.</p> <p><b>egress</b>—(Optional) Display egress queue statistics.</p> <p><b>forwarding-class <i>class-name</i></b>—(Optional) Display queuing statistics for the specified forwarding class.</p> <p><b>ingress</b>—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics.</p> <p><b>remaining-traffic</b>—(Optional) Display the queuing statistics of all member logical interfaces for interface sets that do not have traffic-control profiles configured.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show interfaces interface-set queue (Gigabit Ethernet) on page 189</a></p> <p><a href="#">show interfaces interface-set queue both-ingress-egress (Enhanced DPC) on page 190</a></p> <p><a href="#">show interfaces interface-set queue egress (Enhanced DPC) on page 192</a></p> <p><a href="#">show interfaces interface-set queue forwarding-class (Gigabit Ethernet) on page 193</a></p> <p><a href="#">show interfaces interface-set queue (Enhanced DPC) on page 194</a></p> <p><a href="#">show interfaces interface-set queue remaining-traffic (Gigabit Ethernet) on page 194</a></p>
<b>Output Fields</b>	<a href="#">Table 18 on page 188</a> describes the information for the <b>show interfaces interface-set queue</b> command.

Table 18: Ethernet show interfaces interface-set queue Output Fields

Field Name	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Interface set</b>	Name of the interface set.	All levels
<b>Interface set index</b>	Index number of the interface set.	All levels

Table 18: Ethernet show interfaces interface-set queue Output Fields (*continued*)

Field Name	Field Description	Level of Output
Forwarding classes supported	Total number of forwarding classes supported on the specified interface set.	All levels
Forwarding classes in use	Total number of forwarding classes used on the specified interface set.	All levels
Egress queues supported	Total number of egress queues supported on the specified interface set.	All levels
Egress queues in use	Total number of egress queues used on the specified interface set.	All levels
Ingress queues supported	Total number of ingress queues supported on the specified interface set.	All levels
Ingress queues in use	Total number of ingress queues used on the specified interface set.	All levels
Queue	Egress or ingress queue number for the statistics being displayed.	All levels
Forwarding classes	Forwarding class name for the statistics being displayed.	All levels
Queued	<b>Packet and Byte</b> statistics for the specified queue. <ul style="list-style-type: none"> <li><b>Packets</b>—Number of packets queued and input rate in packets per second.</li> <li><b>Bytes</b>—Number of bytes queued and input rate in bytes per second.</li> </ul>	All levels
Transmitted	<b>Packet and Byte</b> statistics for the specified forwarding class. <ul style="list-style-type: none"> <li><b>Packets</b>—Number of packets transmitted and transmit rate in packets per second.</li> <li><b>Bytes</b>—Number of bytes transmitted and transmit rate in bytes per second.</li> <li><b>Tail-dropped packets</b>—Number of packets tail dropped.</li> <li><b>RED-dropped packets</b>—Number of RED-dropped packets for the <b>low</b>, <b>medium-low</b>, <b>medium-high</b>, and <b>high</b> loss priorities.</li> <li><b>RED-dropped bytes</b>—Number of RED-dropped bytes for the <b>low</b>, <b>medium-low</b>, <b>medium-high</b>, and <b>high</b> loss priorities.</li> </ul>	All levels

## Sample Output

### show interfaces interface-set queue (Gigabit Ethernet)

```

user@host> show interfaces interface-set queue ge-2/2/0-0
Interface set: ge-2/2/0-0
Interface set index: 3
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets           :           3998482           1 pps

```

```

      Bytes          :          271896884          688 bps
Transmitted:
  Packets          :          1077474          1 pps
  Bytes           :          73268340          688 bps
  Tail-dropped packets :          0          0 pps
  RED-dropped packets :          2921008          0 pps
    Low            :          2921008          0 pps
    Medium-low     :          0          0 pps
    Medium-high    :          0          0 pps
    High           :          0          0 pps
  RED-dropped bytes :          198628544          0 bps
    Low            :          198628544          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:
...

```

#### show interfaces interface-set queue both-ingress-egress (Enhanced DPC)

```

user@host> show interfaces interface-set queue ge-2/2/0-0 both-ingress-egress
Interface set: ge-2/2/0-0
Interface set index: 3
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets          :          185968478          473161 pps
  Bytes           :          10042313520          204441336 bps
Transmitted:
  Packets          :          5441673          13780 pps
  Bytes           :          293850342          5952960 bps
  Tail-dropped packets :          0          0 pps
  RED-dropped packets :          180526772          459372 pps
  RED-dropped bytes :          9748446282          198451512 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
  RED-dropped bytes :          0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets          :          522021472          473602 pps
  Bytes           :          28190332480          204599944 bps
Transmitted:
  Packets          :          5791772          4055 pps
  Bytes           :          312755688          1751976 bps
  Tail-dropped packets :          0          0 pps
  RED-dropped packets :          516227139          469546 pps
  RED-dropped bytes :          27876265560          202843872 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
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          :      5417304      13797 pps
Bytes            :     368429508     7506096 bps
Transmitted:
Packets          :      5014996      12769 pps
Bytes            :     341019728     6946560 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :      402189      1028 pps
Low              :      402189      1028 pps
Medium-low      :          0          0 pps
Medium-high     :          0          0 pps
High            :          0          0 pps
RED-dropped bytes :     27348852     559536 bps
Low              :     27348852     559536 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          :      5770534      3963 pps
Bytes            :     396943252     2156144 bps
Transmitted:
Packets          :      3945152      1457 pps
Bytes            :     268270336     792608 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :      1815141      2506 pps
Low              :      1815141      2506 pps
Medium-low      :          0          0 pps
Medium-high     :          0          0 pps
High            :          0          0 pps
RED-dropped bytes :     123429524     1363536 bps
Low              :     123429524     1363536 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 interface-set queue egress (Enhanced DPC)

```

user@host> show interfaces interface-set queue ge-2/2/0-0 egress
Interface set: ge-2/2/0-0
Interface set index: 3
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets              :          3958253          13822 pps
  Bytes               :          269217592        7519712 bps
Transmitted:
  Packets              :          3665035          12729 pps
  Bytes               :          249222380        6924848 bps
  Tail-dropped packets :                0                0 pps
  RED-dropped packets :          293091          1093 pps
    Low               :          293091          1093 pps
    Medium-low        :                0                0 pps
    Medium-high       :                0                0 pps
    High              :                0                0 pps
  RED-dropped bytes   :          19930188        594864 bps
    Low               :          19930188        594864 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             :            5350989            3904 pps
  Bytes               :          368412924          2124048 bps
Transmitted:
  Packets             :            3790469            1465 pps
  Bytes               :          257751892          796960 bps
  Tail-dropped packets :                0                0 pps
  RED-dropped packets :            1550282            2439 pps
    Low               :            1550282            2439 pps
    Medium-low        :                0                0 pps
    Medium-high       :                0                0 pps
    High              :                0                0 pps
  RED-dropped bytes   :          105419176          1327088 bps
    Low               :          105419176          1327088 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 interface-set queue forwarding-class (Gigabit Ethernet)

```

user@host> show interfaces interface-set queue ge-2/2/0-0 forwarding-class best-effort
Interface set: ge-2/2/0-0
Interface set index: 3
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets             :            101857694            1420083 pps
  Bytes               :          6927234456          772532320 bps
Transmitted:
  Packets             :            3984693            55500 pps
  Bytes               :          270959592          30192512 bps
  Tail-dropped packets :                0                0 pps
  RED-dropped packets :            97870952          1364583 pps
    Low               :            97870952          1364583 pps
    Medium-low        :                0                0 pps
    Medium-high       :                0                0 pps
    High              :                0                0 pps

```

RED-dropped bytes	:	6655225776	742339808 bps
Low	:	6655225776	742339808 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

### show interfaces interface-set queue (Enhanced DPC)

```

user@host> show interfaces interface-set queue ge-2/2/0-0 ingress
Interface set: foo
Interface set index: 3
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :      149036817      473711 pps
    Bytes        :      8048003934    204642936 bps
  Transmitted:
    Packets      :      4360749      13891 pps
    Bytes        :      235480446    6000912 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      144676035    459820 pps
    RED-dropped bytes  :      7812506592    198642024 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
    RED-dropped bytes  :      0      0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      :      485089207      473605 pps
    Bytes        :      26195987476    204597576 bps
  Transmitted:
    Packets      :      5480799      3959 pps
    Bytes        :      295963146    1710504 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      479605853    469646 pps
    RED-dropped bytes  :      25898716170    202887072 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
    RED-dropped bytes  :      0      0 bps

```

### show interfaces interface-set queue remaining-traffic (Gigabit Ethernet)

```

user@host> show interfaces interface-set queue ge-2/2/0-0 remaining-traffic
Interface set: ge-2/2/0-0
Interface set index: 12
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use

```

Queue: 0, Forwarding classes: best-effort

Queued:

Packets	:	2201552	0 pps
Bytes	:	149705536	0 bps

Transmitted:

Packets	:	609765	0 pps
Bytes	:	41464020	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	1591787	0 pps
Low	:	1591787	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	108241516	0 bps
Low	:	108241516	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

## show interfaces irb

<b>Syntax</b>	<pre>show interfaces irb &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 8.4.
<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>mac</b>—Display hardware MAC address</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 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 bridging domain 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 200</a></p> <p><a href="#">show interfaces irb snmp-index on page 201</a></p>
<b>Output Fields</b>	<a href="#">Table 19 on page 196</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 19: show interfaces irb 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 physical interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
<b>Proto</b>	Protocol configured on the interface.	<b>terse</b>

Table 19: show interfaces irb Output Fields (*continued*)

Field Name	Field Description	Level of Output
<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 19: 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 19: 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>	<p>Number and rate of bytes and packets received and transmitted on the logical 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 logical 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>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 19: 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, Runt: 0, Giants: 0, Policed discards:
0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0

Logical interface 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 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 lacp interfaces

**Syntax** `show lacp interfaces`  
`<interface-name>`

**Release Information** Command introduced in Junos OS Release 7.6.

**Description** Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet, Fast Ethernet, or Gigabit Ethernet interface.

**Options** **none**—Display LACP information for all interfaces.

**interface-name**—(Optional) Display LACP information for the specified interface:

- Aggregated Ethernet—**aenumber**
- Fast Ethernet—**fe-fpc/pic/port**
- Gigabit Ethernet—**ge-fpc/pic/port**



**NOTE:** The `show lacp interfaces` command returns the following error message if your system is not configured in either active or passive LACP mode:

“Warning: lacp subsystem not running – not needed by configuration”

**Required Privilege Level** view

**List of Sample Output** [show lacp interfaces \(Aggregated Ethernet\) on page 205](#)  
[show lacp interfaces \(Gigabit Ethernet\) on page 205](#)

**Output Fields** [Table 20 on page 202](#) lists the output fields for the `show lacp interfaces` command. Output fields are listed in the approximate order in which they appear.

**Table 20: show lacp interfaces Output Fields**

Field Name	Field Description
Aggregated interface	Aggregated interface value.

Table 20: show lacp interfaces Output Fields (*continued*)

Field Name	Field Description
LACP State	<p>LACP state information for each aggregated interface:</p> <ul style="list-style-type: none"> <li>• <b>Role</b>—Role played by the interface. It can be one of the following: <ul style="list-style-type: none"> <li>• <b>Actor</b>—Local device participating in LACP negotiation.</li> <li>• <b>Partner</b>—Remote device participating in LACP negotiation.</li> </ul> </li> <li>• <b>Exp</b>—Expired state. <b>Yes</b> indicates the actor or partner is in an expired state. <b>No</b> indicates the actor or partner is not in an expired state.</li> <li>• <b>Def</b>—Default. <b>Yes</b> indicates that the actor's receive machine is using the default operational partner information, administratively configured for the partner. <b>No</b> indicates the operational partner information in use has been received in an LACP PDU.</li> <li>• <b>Dist</b>—Distribution of outgoing frames. <b>No</b> indicates distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is <b>Yes</b>.</li> <li>• <b>Col</b>—Collection of incoming frames. <b>Yes</b> indicates collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is <b>No</b>.</li> <li>• <b>Syn</b>—Synchronization. If the value is <b>Yes</b>, the link is considered synchronized. It has been allocated to the correct link aggregation group, the group has been associated with a compatible aggregator, and the identity of the link aggregation group is consistent with the system ID and operational key information transmitted. If the value is <b>No</b>, the link is not synchronized. It is currently not in the right aggregation.</li> <li>• <b>Aggr</b>—Ability of aggregation port to aggregate (<b>Yes</b>) or to operate only as an individual link (<b>No</b>).</li> <li>• <b>Timeout</b>—LACP timeout preference. Periodic transmissions of LACP PDUs occur at either a slow or fast transmission rate, depending upon the expressed LACP timeout preference (<b>Long Timeout</b> or <b>Short Timeout</b>).</li> <li>• <b>Activity</b>—Actor or partner's port activity. <b>Passive</b> indicates the port's preference for not transmitting LAC PDUs unless its partner's control value is <b>Active</b>. <b>Active</b> indicates the port's preference to participate in the protocol regardless of the partner's control value.</li> </ul>

Table 20: show lacp interfaces Output Fields (*continued*)

Field Name	Field Description
LACP Protocol	<p>LACP protocol information for each aggregated interface:</p> <ul style="list-style-type: none"> <li>Link state (active or standby) indicated in parentheses next to the interface when link protection is configured.</li> <li><b>Receive State</b>—One of the following values: <ul style="list-style-type: none"> <li><b>Current</b>—The state machine receives an LACP PDU and enters the <b>Current</b> state.</li> <li><b>Defaulted</b>—If no LACP PDU is received before the timer for the <b>Current</b> state expires a second time, the state machine enters the <b>Defaulted</b> state.</li> <li><b>Expired</b>—If no LACP PDU is received before the timer for the <b>Current</b> state expires once, the state machine enters the <b>Expired</b> state.</li> <li><b>Initialize</b>—When the physical connectivity of a link changes or a Begin event occurs, the state machine enters the <b>Initialize</b> state.</li> <li><b>LACP Disabled</b>—If the port is operating in half duplex, the operation of LACP is disabled on the port, forcing the state to <b>LACP Disabled</b>. This state is similar to the <b>Defaulted</b> state, except that the port is forced to operate as an individual port.</li> <li><b>Port Disabled</b>—If the port becomes inoperable and a Begin event has not occurred, the state machine enters the <b>Port Disabled</b> state.</li> </ul> </li> <li><b>Transmit State</b>—Transmit state of state machine. One of the following values: <ul style="list-style-type: none"> <li><b>Fast Periodic</b>—Periodic transmissions are enabled at a fast transmission rate.</li> <li><b>No Periodic</b>—Periodic transmissions are disabled.</li> <li><b>Periodic Timer</b>—Transitory state entered when the periodic timer expires.</li> <li><b>Slow Periodic</b>—Periodic transmissions are enabled at a slow transmission rate.</li> </ul> </li> <li><b>Mux State</b>—State of the multiplexer state machine for the aggregation port. The state is one of the following values: <ul style="list-style-type: none"> <li><b>Attached</b>—Multiplexer state machine initiates the process of attaching the port to the selected aggregator.</li> <li><b>Collecting—Yes</b> indicates that the receive function of this link is enabled with respect to its participation in an aggregation. Received frames are passed to the aggregator for collection. <b>No</b> indicates the receive function of this link is not enabled.</li> <li><b>Collecting Distributing</b>—Collecting and distributing states are merged together to form a combined state (coupled control). Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution.</li> <li><b>Detached</b>—Process of detaching the port from the aggregator is in progress.</li> <li><b>Distributing—Yes</b> indicates that the transmit function of this link is enabled with respect to its participation in an aggregation. Frames may be passed down from the aggregator's distribution function for transmission. <b>No</b> indicates the transmit function of this link is not enabled.</li> <li><b>Waiting</b>—Multiplexer state machine is in a holding process, awaiting an outcome.</li> </ul> </li> </ul>
LACP Statistics	<p>LACP statistics are returned when the <b>extensive</b> option is used and provides the following information:</p> <ul style="list-style-type: none"> <li><b>LACP Rx</b>—LACP received counter that increments for each normal hello.</li> <li><b>LACP Tx</b>—Number of LACP transmit packet errors logged.</li> <li><b>Unknown Rx</b>—Number of unrecognized packet errors logged.</li> <li><b>Illegal Rx</b>—Number of invalid packets received.</li> </ul>

## Sample Output

### show lacp interfaces (Aggregated Ethernet)

```

user@host> show lacp interfaces ae0 extensive
Aggregated interface: ae0
LACP state:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
ge-1/0/1        Actor  No   Yes  No   No   No   Yes    Fast    Active
ge-1/0/1        Partner No   Yes  No   No   No   Yes    Fast    Passive
ge-1/0/2        Actor  No   Yes  No   No   No   Yes    Fast    Active
ge-1/0/2        Partner No   Yes  No   No   No   Yes    Fast    Passive

LACP protocol:      Receive State      Transmit State      Mux State
ge-1/0/1            CURRENT          Fast periodic      Collecting
distributing
ge-1/0/2            CURRENT          Fast periodic      Collecting
distributing
ge-1/0/1 (active)    CURRENT          Fast periodic      Collecting
distributing
ge-1/0/2 (standby)   CURRENT          Fast periodic      WAITING
LACP Statistics:      LACP Rx      LACP Tx      Unknown Rx      Illegal Rx
ge-1/0/1              0             0             0             0
ge-1/0/2              0             0             0             0

```

### show lacp interfaces (Gigabit Ethernet)

```

user@host> show lacp interfaces ge-0/3/0
Aggregated interface: ae0
LACP State:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
ge-0/3/0        Actor  No   No   Yes  Yes  Yes  Yes    Fast    Active
ge-0/3/0        Partner No   No   Yes  Yes  Yes  Yes    Fast    Active
LACP Protocol:      Receive State      Transmit State      Mux State
ge-0/3/0            Current          Fast periodic      Collecting distributing

```

## show interfaces mac-database (Gigabit Ethernet)

<b>Syntax</b>	<code>show interfaces mac-database (ge-fpc/pic/port   ge-fpc/pic/port.n) &lt;mac-address mac-address&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.
<b>Description</b>	(M Series, T Series, MX Series routers, and PTX Series Packet Transport Routers only) Display media access control (MAC) address information for the specified Gigabit Ethernet interface.
<b>Options</b>	<p><b>ge-fpc/pic/port</b>—Display MAC addresses that have been learned on all logical interfaces on a particular physical interface.</p> <p><b>ge-fpc/pic/port.n</b>—Display MAC addresses that have been learned on a particular logical interface.</p> <p><b>mac-address mac-address</b>—(Optional) Display detailed MAC address statistics, including policer information.</p>
<b>Additional Information</b>	On IQ2 PIC interfaces, the default value for maximum retention of entries in the MAC address table has changed, for cases in which the table is not full. The new holding time is 12 hours. The previous retention time of 3 minutes is still in effect when the table is full.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show interfaces mac-database (All MAC Addresses on a Port) on page 208</a> <a href="#">show interfaces mac-database (All MAC Addresses on a Service) on page 209</a> <a href="#">show interfaces mac-database mac-address on page 209</a>
<b>Output Fields</b>	<a href="#">Table 21 on page 206</a> lists the output fields for the <b>show interfaces mac-database</b> command. Output fields are listed in the approximate order in which they appear.

**Table 21: show interfaces mac-database Output Fields**

Field Name	Field Description
<b>Physical Interface</b>	
<b>Physical interface</b>	Name of the physical interface.
<b>Enabled</b>	State of the physical interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> .
<b>Interface index</b>	Physical interface index number, which reflects its initialization sequence.
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.
<b>Description</b>	Description and name of the interface.

Table 21: show interfaces mac-database Output Fields (*continued*)

Field Name	Field Description
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback: <b>local</b> or <b>remote</b> .
Source filtering	Whether source filtering is configured.
Flow control	Whether flow control is enabled or disabled.
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .
Interface flags	Information about the interface. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .
Link flags	Information about the link. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .
<b>Logical Interface</b>	
Logical interface	Name of the logical interface.
Index	Logical interface index number, which reflects its initialization sequence.
SNMP ifIndex	Logical interface SNMP interface index number.
Flags	Information about the logical interface (possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> ).
Encapsulation	Encapsulation on the logical interface.
MAC address, Input frames, Input bytes, Output frames, Output bytes	MAC address and corresponding number of input frames, input bytes, output frames, and output bytes.
Number of MAC addresses	Number of MAC addresses configured.

Table 21: show interfaces mac-database Output Fields (*continued*)

Field Name	Field Description
<b>Policer Statistics</b>	<p>(Displayed for <b>mac-address</b> option only) Display information about policers applied to a logical interface-MAC pair.</p> <ul style="list-style-type: none"> <li>• <b>Policer type</b>—Type of policer that is out of spec with respect to the configuration. It can be one or more of the following: <ul style="list-style-type: none"> <li>• <b>Input premium</b>—Number of high-priority rating out-of-spec frames or bytes received.</li> <li>• <b>Output premium</b>—Number of high-priority rating out-of-spec frames or bytes sent.</li> <li>• <b>Input aggregate</b>—Total number of out-of-spec frames or bytes received.</li> <li>• <b>Output aggregate</b>—Total number of out-of-spec frames or bytes sent.</li> </ul> </li> <li>• <b>Discarded Frames</b>—Number of discarded frames.</li> <li>• <b>Discarded Bytes</b>—Number of discarded bytes.</li> </ul>

## Sample Output

### show interfaces mac-database (All MAC Addresses on a Port)

```

user@host> show interfaces mac-database xe-0/3/3
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags       : None

Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
MAC address      Input frames  Input bytes  Output frames  Output bytes
00:00:00:00:00:00      1           56           0             0
00:00:c0:01:01:02     7023810     323095260    0             0
00:00:c0:01:01:03     7023810     323095260    0             0
00:00:c0:01:01:04     7023810     323095260    0             0
00:00:c0:01:01:05     7023810     323095260    0             0
00:00:c0:01:01:06     7023810     323095260    0             0
00:00:c0:01:01:07     7023810     323095260    0             0
00:00:c0:01:01:08     7023809     323095214    0             0
00:00:c0:01:01:09     7023809     323095214    0             0
00:00:c0:01:01:0a     7023809     323095214    0             0
00:00:c0:01:01:0b     7023809     323095214    0             0
00:00:c8:01:01:02     30424784    1399540064    37448598      1722635508
00:00:c8:01:01:03     30424784    1399540064    37448598      1722635508
00:00:c8:01:01:04     30424716    1399536936    37448523      1722632058
00:00:c8:01:01:05     30424789    1399540294    37448598      1722635508
00:00:c8:01:01:06     30424788    1399540248    37448597      1722635462
00:00:c8:01:01:07     30424783    1399540018    37448597      1722635462
00:00:c8:01:01:08     30424783    1399540018    37448596      1722635416
00:00:c8:01:01:09      8836796     406492616     8836795       406492570
00:00:c8:01:01:0a     30424712    1399536752    37448521      1722631966
00:00:c8:01:01:0b     30424715    1399536890    37448523      1722632058
Number of MAC addresses : 21

```

**show interfaces mac-database (All MAC Addresses on a Service)**

```

user@host> show interfaces mac-database xe-0/3/3
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address      Input frames    Input bytes    Output frames   Output bytes
00:00:00:00:00:00          1              56              0                0
00:00:c0:01:01:02       7023810       323095260        0                0
00:00:c0:01:01:03       7023810       323095260        0                0
00:00:c0:01:01:04       7023810       323095260        0                0
00:00:c0:01:01:05       7023810       323095260        0                0
00:00:c0:01:01:06       7023810       323095260        0                0
00:00:c0:01:01:07       7023810       323095260        0                0
00:00:c0:01:01:08       7023809       323095214        0                0
00:00:c0:01:01:09       7023809       323095214        0                0
00:00:c0:01:01:0a       7023809       323095214        0                0
00:00:c0:01:01:0b       7023809       323095214        0                0
00:00:c8:01:01:02      31016568      1426762128      38040381      1749857526
00:00:c8:01:01:03      31016568      1426762128      38040382      1749857572
00:00:c8:01:01:04      31016499      1426758954      38040306      1749854076
00:00:c8:01:01:05      31016573      1426762358      38040381      1749857526
00:00:c8:01:01:06      31016573      1426762358      38040381      1749857526
00:00:c8:01:01:07      31016567      1426762082      38040380      1749857480
00:00:c8:01:01:08      31016567      1426762082      38040379      1749857434
00:00:c8:01:01:09       9428580       433714680       9428580       433714680
00:00:c8:01:01:0a      31016496      1426758816      38040304      1749853984
00:00:c8:01:01:0b      31016498      1426758908      38040307      1749854122

```

**show interfaces mac-database mac-address**

```

user@host> show interfaces mac-database xe-0/3/3 mac-address 00:00:c8:01:01:09
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
MAC address: 00:00:c8:01:01:09, Type: Configured,
  Input bytes   : 202324652
  Output bytes  : 202324560
  Input frames  : 4398362
  Output frames : 4398360
Policer statistics:
Policer type    Discarded frames    Discarded bytes
Output aggregate      3992386          183649756

```

## show interfaces mc-ae

<b>Syntax</b>	<b>show interfaces mc-ae</b> <revertive-info> <id <i>identifier</i> unit <i>number</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 9.6. <b>revertive-info</b> statement introduced in Junos OS Release 13.3
<b>Description</b>	On MX Series routers with multichassis aggregated Ethernet ( <b>aeX</b> ) interfaces, displaya information about the <b>aeX</b> interfaces.
<b>Options</b>	<b>revertive-info</b> —(Optional) Display revertive mode information for multichassis aggregated Ethernet interface.  <b>identifier</b> —(Optional) Identifier of the multichassis aggregated Ethernet interface.  <b>number</b> —(Optional) Specify the logical interface by unit number.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Multichassis Link Aggregation</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show interfaces mc-ae on page 211</a> <a href="#">show interfaces mc-ae (Active/Active Bridging and VRRP over IRB on MX Series Routers) on page 211</a> <a href="#">show interfaces mc-ae revertive-info on page 212</a>
<b>Output Fields</b>	Table 22 on page 210 lists the output fields for the <b>show interfaces mc-ae</b> command. Output fields are listed in the approximate order in which they appear.

**Table 22: show interfaces mc-ae Output Fields**

Output Field Name	Field Description
<b>Member Link</b>	Identifiers of the configured multichassis link aggregate interfaces configured interfaces.
<b>Local Status</b>	Status of the local link: <b>active</b> or <b>standby</b> .
<b>Peer Status</b>	Status of the peer link: <b>active</b> or <b>standby</b> .
<b>Local State</b>	Up or down state of the local device.

Table 22: show interfaces mc-ae Output Fields (*continued*)

Output Field Name	Field Description
<b>Peer State</b>	<p>Status of the local and peer links in an <b>active/active</b> bridge or VRRP over integrated routing and bridging (IRB) configuration on MX Series routers, including:</p> <p>Logical Interface—Aggregated Ethernet (AE) aggregate number and unit number.</p> <p>Topology Type—The bridge or VRRP topology type configured on the AE.</p> <p>Local State—Up or down state of the local device.</p> <p>Peer State—Up or down state of the peer device.</p> <p>Peer Ip/ICL-PL/State—Address, interface and state of the peer device.</p>
<b>Logical Interface</b>	Identifier and unit of the multichassis aggregated Ethernet interface.
<b>Core Facing Interface</b>	Label: <b>pseudowire interface</b> or <b>Ethernet interface</b> .
<b>ICL-PL</b>	Label: <b>pseudowire interface</b> or <b>Ethernet interface</b> .
<b>switchover mode</b>	The configured switchover mode for the multichassis aggregated Ethernet interface: <b>revertive</b> or <b>non-revertive</b> .
<b>switchover status</b>	Status of the switchover if the <b>revert-time</b> statement is configured at the <b>[edit interfaces aex mc-ae]</b> hierarchy level.
<b>revert time</b>	Revert time configured for the multichassis aggregated Ethernet interface.
<b>switchover time remaining</b>	Seconds left to trigger the switchover if the switchover is in progress.

## Sample Output

### show interfaces mc-ae

```

user@host> show interfaces mc-ae ae0 unit 512
Member Links      : ae0
Local Status      : active
Peer Status       : active
Logical Interface  : ae0.512
Core Facing Interface : Label Ethernet Interface
ICL-PL            : Label Ethernet Interface

```

### show interfaces mc-ae (Active/Active Bridging and VRRP over IRB on MX Series Routers)

```

user@host# show interfaces mc-ae ge-0/0/0.0
Member Link       : ae0
Current State Machine's State: active

```

```
Local Status      : active
Local State       : up
Peer Status       : active
Peer State        : up
  Logical Interface : ae0.0
  Topology Type     : bridge
  Local State       : up
  Peer State        : up
  Peer Ip/ICL-PL/State : 192.168.100.10 ge-0/0/0.0 up
```

#### show interfaces mc-ae revertive-info

```
user@host> show interfaces mc-ae revertive-info id 2
Member Link      : ae1
Current State Machine's State: mcae active state
Local Status     : active
Local State      : up
Peer Status      : standby
Peer State       : up
Switchover Mode  : Non Revertive
Switchover Status : N/A
Revert Time      : 1 Minutes
Switchover Remaining Time : N/A
  Logical Interface : ae1.1024
  Topology Type     : bridge
  Local State       : up
  Peer State        : up
  Peer Ip/MCP/State : N/A
```

## show oam ethernet connectivity-fault-management delay-statistics

<b>Syntax</b>	<pre>show oam ethernet connectivity-fault-management delay-statistics &lt;count <i>entry-count</i>&gt; &lt;local-mep <i>local-mep-id</i>&gt; maintenance-association <i>ma-name</i> maintenance-domain <i>md-name</i> &lt;remote-mep <i>remote-mep-id</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 11.4 for EX Series switches.</p>
<b>Description</b>	<p>On MX Series routers with Ethernet interfaces on Dense Port Concentrators (DPCs), display ETH-DM delay statistics.</p> <p>On EX Series switches, display delay measurement results.</p>
<b>Options</b>	<p><b>count</b> <i>entry-count</i>—(Optional) Number of entries to display from the statistics table. The range of values is 1 through 100. The default value is 100 entries.</p> <p><b>local-mep</b> <i>local-mep-id</i>—(Optional) Numeric identifier of the local MEP. On MX Series routers, the range of values is 1 through 8192. On EX Series switches, the range of values is 1 through 8191.</p> <p><b>maintenance-association</b> <i>ma-name</i>—Name of an existing CFM maintenance association.</p> <p><b>maintenance-domain</b> <i>md-name</i>—Name of an existing connectivity fault management (CFM) maintenance domain.</p> <p><b>remote-mep</b> <i>remote-mep-id</i>—(Optional) Numeric identifier of the remote MEP. On MX Series routers, the range of values is 1 through 8192. On EX Series switches, the range of values is 1 through 8191.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear oam ethernet connectivity-fault-management statistics</a></li> <li>• <a href="#">clear oam ethernet connectivity-fault-management delay-statistics</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management interfaces on page 221</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management mep-database on page 234</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management mep-statistics on page 245</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show oam ethernet connectivity-fault-management delay-statistics on page 215</a></p> <p><a href="#">show oam ethernet connectivity-fault-management delay-statistics remote-mep on page 215</a></p>
<b>Output Fields</b>	<p>Table 23 on page 214 lists the output fields for the <b>show oam ethernet connectivity-fault-management delay-statistics</b> command and the <b>show oam ethernet</b></p>

**connectivity-fault-management mep-statistics** command. Output fields are listed in the approximate order in which they appear.

**Table 23: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields**

Output Field Name	Field Description
MEP identifier	Maintenance association end point (MEP) numeric identifier.
MAC address	Unicast MAC address configured for the MEP.
Remote MEP count	Number of remote MEPs (unless you specify the <b>remote-mep</b> option).
Remote MEP identifier	Numeric identifier of the remote MEP.
Remote MAC address	Unicast MAC address of the remote MEP.
Index	Index number that corresponds to the ETH-DM entry in the CFM database.
One-way delay (usec)	<p>For a one-way ETH-DM session, the frame delay time, in microseconds, measured at the receiver MEP.</p> <p>For a detailed description of one-way Ethernet frame delay measurement, see the <i>ITU-T Y.1731 Ethernet Service OAM</i> topics in the <i>Junos OS Network Interfaces Library for Routing Devices</i>.</p>
Two-way delay (usec)	<p>For a two-way ETH-DM session, the frame delay time, in microseconds, measured at the initiator MEP.</p> <p>For a detailed description of two-way Ethernet frame delay measurement, see the <i>ITU-T Y.1731 Ethernet Service OAM</i> topics in the <i>Junos OS Network Interfaces Library for Routing Devices</i>.</p>
Average one-way delay	Average one-way frame delay for the statistics displayed.
Average one-way delay variation	Average one-way “frame jitter” for the statistics displayed.
Best-case one-way delay	Lowest one-way frame delay for the statistics displayed.
Worst-case one-way delay	Highest one-way frame delay for the statistics displayed.
Average two-way delay	Average two-way frame delay for the statistics displayed.
Average two-way delay variation	Average two-way “frame jitter” for the statistics displayed.
Best-case two-way delay	Lowest two-way frame delay for the statistics displayed.
Worst-case two-way delay	Highest two-way frame delay calculated in this session.

## Sample Output

show oam ethernet connectivity-fault-  
management  
delay-statistics

```
user@switch> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md6 maintenance-association ma6
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count: 2
Remote MEP identifier: 101
Remote MAC address: 00:05:85:73:39:4a
Delay measurement statistics:
Index  One-way delay  Two-way delay
      (usec)      (usec)
    1      259      519
    2      273      550
    3      287      571
    4      299      610
    5      313      650
Average one-way delay      : 286 usec
Average one-way delay variation: 62 usec
Best case one-way delay    : 259 usec
Worst case one-way delay   : 313 usec
Average two-way delay      : 580 usec
Average two-way delay variation: 26 usec
Best case two-way delay    : 519 usec
Worst case two-way delay   : 650 usec

Remote MEP identifier: 102
Remote MAC address: 00:04:55:63:39:5a
Delay measurement statistics:
Index  One-way delay  Two-way delay
      (usec)      (usec)
    1      29      58
    2      23      59
    3      27      56
    4      29      62
    5      33      68
Average one-way delay      : 28 usec
Average one-way delay variation: 3 usec
Best case one-way delay    : 23 usec
Worst case one-way delay   : 33 usec
Average two-way delay      : 60 usec
Average two-way delay variation: 3 usec
Best case two-way delay    : 56 usec
Worst case two-way delay   : 68 usec
```

show oam ethernet connectivity-fault-  
management delay-statistics remote-mep

```
user@switch> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md6 maintenance-association ma6 remote-mep 101
MEP identifier: 100, MAC address: 00:05:85:73:7b:39

Remote MEP identifier: 101
Remote MAC address: 00:05:85:73:39:4a
Delay measurement statistics:
Index  One-way delay  Two-way delay
      (usec)      (usec)
    1      259      519
```

2	273	550
3	287	571
4	299	610
5	313	650

Average one-way delay : 286 usec  
Average one-way delay variation: 62 usec  
Best case one-way delay : 259 usec  
Worst case one-way delay : 313 usec  
Average two-way delay : 580 usec  
Average two-way delay variation: 26 usec  
Best case two-way delay : 519 usec  
Worst case two-way delay : 650 usec

## show oam ethernet connectivity-fault-management forwarding-state

<b>Syntax</b>	<b>show oam ethernet connectivity-fault-management forwarding-state</b> <b>interface</b> <i>interface-name</i>   <b>instance</b> <i>instance-name</i> <brief   detail   extensive>
<b>Release Information</b>	Command introduced in Junos OS Release 8.4.
<b>Description</b>	On M7i and M10i with the Enhanced CFEB (CFEB-E), M320, MX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management forwarding state information for Ethernet interfaces.
<b>Options</b>	<p><b>interface</b> <i>interface-name</i>—Display forwarding state information for the specified Ethernet interface only.</p> <p><b>instance</b> <i>instance-name</i>—Display forwarding state information for the specified forwarding instance only.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show oam ethernet connectivity-fault-management forwarding-state instance on page 218</a></p> <p><a href="#">show oam ethernet connectivity-fault-management forwarding-state interface on page 218</a></p> <p><a href="#">show oam ethernet connectivity-fault-management forwarding-state interface detail on page 219</a></p> <p><a href="#">show oam ethernet connectivity-fault-management forwarding-state interfaceinterface-name on page 220</a></p>
<b>Output Fields</b>	Table 24 on page 217 lists the output fields for the <b>show oam ethernet connectivity-fault-management forwarding-state</b> command. Output fields are listed in the approximate order in which they appear.

Table 24: show oam ethernet connectivity-fault-management forwarding-state Output Fields

Field Name	Field Description	Level of Output
Interface name	Interface identifier.	All levels
Link (Status)	Local link status.	All levels
Filter action	Filter action for messages at the level.	All levels
Next hop type	Next-hop type.	All levels
Next index	Next-hop index number.	brief
Level	Maintenance domain (MD) level.	detail

Table 24: show oam ethernet connectivity-fault-management forwarding-state Output Fields (*continued*)

Field Name	Field Description	Level of Output
Direction	MEP direction configured.	none
Instance name	Forwarding instance name.	All levels
CEs	Number of customer edge (CE) interfaces.	All levels
VEs	Number of VPN endpoint (VE) interfaces.	All levels

### Sample Output

show oam ethernet  
connectivity-fault-  
management forwarding-  
state instance

```
user@host> show oam ethernet connectivity-fault-management forwarding-state instance
Instance name: __+bd1__
CEs: 3
VEs: 0
Maintenance domain forwarding state:

Level   Direction   Filter action   Nexthop
                     type
0               Drop           none
1               Drop           none
2               Drop           none
3               Drop           none
4               Drop           none
5               Drop           none
6               Drop           none
7               Drop           none
```

show oam ethernet  
connectivity-fault-  
management forwarding-  
state interface

```
user@host> show oam ethernet connectivity-fault-management forwarding-state interface
Interface name: ge-3/0/0.0
Instance name: __+bd1__
Maintenance domain forwarding state:

Level   Direction   Filter action   Nexthop
                     type
0               Drop           none
1               Drop           none
2               Drop           none
3               Drop           none
4               Drop           none
5               Drop           none
6               Drop           none
7       down   Receive        none
```

Interface name: xe-0/0/0.0

Instance name: \_\_+bd1\_\_

Maintenance domain forwarding state:

Level	Direction	Filter action	Nexthop type	Nexthop index
0		Drop	none	
1		Drop	none	
2		Drop	none	
3		Drop	none	
4		Drop	none	
5		Drop	none	
6		Drop	none	
7	down	Receive	none	

**show oam ethernet  
connectivity-fault-  
management forwarding-  
state interface detail**

user@host> **show oam ethernet connectivity-fault-management forwarding-state interface detail**

Interface name: ge-3/0/0.0

Instance name: \_\_+bd1\_\_

Level: 0  
Filter action: Drop  
Nexthop type: none

Level: 1  
Filter action: Drop  
Nexthop type: none

Level: 2  
Filter action: Drop  
Nexthop type: none

Level: 3  
Filter action: Drop  
Nexthop type: none

Level: 4  
Filter action: Drop  
Nexthop type: none

Level: 5  
Filter action: Drop  
Nexthop type: none

Level: 6  
Filter action: Drop  
Nexthop type: none

Level: 7  
Direction: down  
Filter action: Receive  
Nexthop type: none

Interface name: xe-0/0/0.0

Instance name: \_\_+bd1\_\_

```
Level: 0
Filter action: Drop
Nexthop type: none
```

```
Level: 1
Filter action: Drop
Nexthop type: none
```

```
...
```

```
show oam ethernet
connectivity-fault-
management forwarding-
state interface
interface-name
```

```
user@host> show oam ethernet connectivity-fault-management forwarding-state interface
interface-name ge-3/0/0/0.0
Interface name: ge-3/0/0.0
Instance name: __+bd1__
Maintenance domain forwarding state:
```

Level	Direction	Filter action	Nexthop type	Nexthop index
0		Drop	none	
1		Drop	none	
2		Drop	none	
3		Drop	none	
4		Drop	none	
5		Drop	none	
6		Drop	none	
7	down	Receive	none	

## show oam ethernet connectivity-fault-management interfaces

<b>Syntax</b>	<pre>show oam ethernet connectivity-fault-management interfaces &lt;ethernet-interface-name&gt; &lt;level md-level&gt; &lt;brief   detail   extensive&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.4.</p> <p>Support for ITU-T Y.1731 frame delay measurement added in Junos OS Release 9.5.</p> <p>Support for ITU-T Y.1731 Ethernet synthetic frame loss measurement (ETH-SLM) added in Junos OS Release 13.2 for ACX Series and MX Series routers.</p>
<b>Description</b>	<p>On M7i and M10i routers with Enhanced CFEB (CFEB-E), and on M320, MX Series, ACX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for Ethernet interfaces.</p> <p>In addition, for Ethernet interfaces on MX Series routers, also display any ITU-T Y.1731 frame delay measurement (ETH-DM) frame counts when <b>detail</b> or <b>extensive</b> mode is specified.</p> <p>For Ethernet interfaces on MX Series routers, display any ITU-T Y.1731 synthetic frame loss measurement (ETH-SLM) statistics and frame counts.</p>
<b>Options</b>	<p><b>brief   detail   extensive</b>—(Optional) Specified level of output.</p> <p><b>ethernet-interface-name</b>—(Optional) CFM information only for CFM entities attached to the specified Ethernet interface.</p> <p><b>level md-level</b>—(Optional) CFM information for CFM identities enclosed within a maintenance domain of the specified level.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear oam ethernet connectivity-fault-management statistics</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management delay-statistics on page 213</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management mep-database on page 234</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management mep-statistics on page 245</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show oam ethernet connectivity-fault-management interfaces on page 226</a></p> <p><a href="#">show oam ethernet connectivity-fault-management interfaces detail on page 226</a></p> <p><a href="#">show oam ethernet connectivity-fault-management interfaces detail (One-Way ETH-DM) on page 227</a></p> <p><a href="#">show oam ethernet connectivity-fault-management interfaces detail (Connection Protection TLV Configured) on page 228</a></p> <p><a href="#">show oam ethernet connectivity-fault-management interfaces extensive on page 229</a></p> <p><a href="#">show oam ethernet connectivity-fault-management interfaces level on page 230</a></p>

[show oam ethernet connectivity-fault-management interfaces \(trunk ports\) on page 230](#)

**Output Fields** Table 25 on page 222 lists the output fields for the **show oam ethernet connectivity-fault-management interfaces** command. Output fields are listed in the approximate order in which they appear.

**Table 25: show oam ethernet connectivity-fault-management interfaces Output Fields**

Field Name	Field Description	Level of Output
<b>Interface</b>	Interface identifier.	All levels
<b>Interface status</b>	Local interface status.	All levels
<b>Link status</b>	Local link status. <b>Up</b> , <b>down</b> , or <b>oam-down</b> .	All levels
<b>Maintenance domain name</b>	Maintenance domain name.	<b>detail extensive</b>
<b>Format (Maintenance domain)</b>	Maintenance domain name format configured.	<b>detail extensive</b>
<b>Level</b>	Maintenance domain level configured.	All levels
<b>Maintenance association name</b>	Maintenance association name.	<b>detail extensive</b>
<b>Format (Maintenance association)</b>	Maintenance association name format configured.	<b>detail extensive</b>
<b>Continuity-check status</b>	Continuity-check status.	<b>detail extensive</b>
<b>Interval</b>	Continuity-check message interval.	<b>detail extensive</b>
<b>Loss-threshold</b>	Lost continuity-check message threshold.	<b>detail extensive</b>
<b>Interface status TLV</b>	Status of the interface status TLV, if configured on the MEP interface: <b>none</b> , <b>up</b> , <b>down</b> , <b>testing</b> , <b>unknown</b> , <b>dormant</b> , <b>notPresent</b> , <b>lowerLayerDown</b>	<b>detail extensive</b>
<b>Port status TLV</b>	Status of the port status TLV, if configured on the MEP interface: <b>none</b> , <b>no</b> , <b>yes</b>	<b>detail extensive</b>
<b>Connection Protection TLV</b>	Status of the connection protection TLV if configured on the MEP interface: <b>no</b> , <b>yes</b>  If <b>yes</b> , then the transmitted connection protection TLV is decoded and the following three fields are displayed: <b>Prefer me</b> , <b>Protection in use</b> , <b>FRR Flag</b>	<b>detail extensive</b>

**Table 25: show oam ethernet connectivity-fault-management interfaces Output Fields (*continued*)**

Field Name	Field Description	Level of Output
<b>Prefer me</b>	If set to <b>yes</b> , the path through which CCM was transmitted is preferred (unless the path fails). It is used for signaling a manual-switch command to the remote side.  Its value can be <b>yes</b> or <b>no</b> .	<b>detail extensive</b>
<b>Protection in use</b>	Used for protection decision coordination. Its value is set to <b>yes</b> if the endpoint transmitting the CCM is currently transmitting the user traffic to protection path.  Its value can be <b>yes</b> or <b>no</b> .	<b>detail extensive</b>
<b>FRR Flag</b>	LSR/LER forwarding the CCM Frame into a bypass tunnel is set.  Its value can be <b>yes</b> or <b>no</b> .	<b>detail extensive</b>
<b>MEP identifier</b>	Maintenance association end point (MEP) identifier.	All levels
<b>Neighbors</b>	Number of MEP neighbors.	All levels
<b>Direction</b>	MEP direction configured.	<b>detail extensive</b>
<b>MAC address</b>	MAC address configured for the MEP.	<b>detail extensive</b>
<b>MEP status</b>	Indicates the status of the connectivity fault management (CFM) protocol running on the MEP: <b>Running</b> , <b>inactive</b> , <b>disabled</b> , or <b>unsupported</b> .	<b>detail extensive</b>
<b>Remote MEP not receiving CCM</b>	Whether the remote MEP is not receiving connectivity check messages (CCMs).	<b>detail extensive</b>
<b>Erroneous CCM received</b>	Whether erroneous CCMs have been received.	<b>detail extensive</b>
<b>Cross-connect CCM received</b>	Whether cross-connect CCMs have been received.	<b>detail extensive</b>
<b>RDI sent by some MEP</b>	Whether the remote defect indication (RDI) bit is set in messages that have been received. The absence of the RDI bit in a CCM indicates that the transmitting MEP is receiving CCMs from all configured MEPs.	<b>detail extensive</b>
<b>CCMs sent</b>	Number of CCMs transmitted.	<b>detail extensive</b>
<b>CCMs received out of sequence</b>	Number of CCMs received out of sequence.	<b>detail extensive</b>
<b>LBMs sent</b>	Number of loopback request messages (LBMs) sent.	<b>detail extensive</b>
<b>Valid in-order LBRs received</b>	Number of loopback response messages (LBRs) received that were valid messages and in sequence.	<b>detail extensive</b>

**Table 25: show oam ethernet connectivity-fault-management interfaces Output Fields (*continued*)**

Field Name	Field Description	Level of Output
<b>Valid out-of-order LBRs received</b>	Number of LBRs received that were valid messages and not in sequence.	<b>detail extensive</b>
<b>LBRs received with corrupted data</b>	Number of LBRs received that were corrupted.	<b>detail extensive</b>
<b>LBRs sent</b>	Number of LBRs transmitted.	<b>detail extensive</b>
<b>LTMs sent</b>	Linktrace messages (LTMs) transmitted.	<b>detail extensive</b>
<b>LTMs received</b>	Linktrace messages received.	<b>detail extensive</b>
<b>LTRs sent</b>	Linktrace responses (LTRs) transmitted.	<b>detail extensive</b>
<b>LTRs received</b>	Linktrace responses received.	<b>detail extensive</b>
<b>Sequence number of next LTM request</b>	Sequence number of next LTM request to be transmitted.	<b>detail extensive</b>
<b>1DMs sent</b>	If the interface is attached to an initiator MEP for a one-way ETH-DM session: Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session.  For all other cases, this field displays 0.	<b>detail extensive</b>
<b>Valid 1DMs received</b>	If the interface is attached to a receiver MEP for a one-way ETH-DM session: Number of valid 1DM frames received.  For all other cases, this field displays 0.	<b>detail extensive</b>
<b>Invalid 1DMs received</b>	If the interface is attached to a receiver MEP for a one-way ETH-DM session: Number of invalid 1DM frames received.  For all other cases, this field displays 0.	<b>detail extensive</b>
<b>Out of sync 1DMs received</b>	If the interface is attached to a receiver MEP for a one-way ETH-DM session: Number of out-of-sync one-way delay measurement request packets received.	<b>detail extensive</b>
<b>DMMs sent</b>	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of Delay Measurement Message (DMM) PDU frames sent to the peer MEP in this session.  For all other cases, this field displays 0.	<b>detail extensive</b>
<b>Valid DMMs received</b>	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of valid two-way delay measurement request packets received.	<b>detail extensive</b>
<b>Invalid DMMs received</b>	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of invalid two-way delay measurement request packets received.	<b>detail extensive</b>

**Table 25: show oam ethernet connectivity-fault-management interfaces Output Fields (*continued*)**

Field Name	Field Description	Level of Output
<b>DMRs sent</b>	If the interface is attached to a responder MEP for a two-way ETH-DM session: Number of delay measurement reply (DMR) frames sent.  For all other cases, this field displays 0.	<b>detail extensive</b>
<b>Valid DMRs received</b>	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of valid DMRs received.  For all other cases, this field displays 0.	<b>detail extensive</b>
<b>Invalid DMRs received</b>	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of invalid DMRs received.  For all other cases, this field displays 0.	<b>detail extensive</b>
<b>LMM sent</b>	If the interface is attached to an initiator MEP for a ETH-LM session: Number of loss measurement message (LMM) PDU frames sent to the peer MEP in this session.	<b>detail extensive</b>
<b>Valid LMM received</b>	If the interface is attached to an initiator MEP for a ETH-LM session: Number of valid loss measurement request packets received.	<b>detail extensive</b>
<b>Invalid LMM received</b>	If the interface is attached to an initiator MEP for a ETH-LM session: Number of invalid loss measurement request packets received.	<b>detail extensive</b>
<b>LMR sent</b>	If the interface is attached to a responder MEP for a ETH-LM session: Number of loss measurement reply (LMR) frames sent.	<b>detail extensive</b>
<b>Valid LMR received</b>	If the interface is attached to an initiator MEP for a ETH-LM session: Number of valid LMR frames received.	<b>detail extensive</b>
<b>Invalid LMR received</b>	If the interface is attached to an initiator MEP for a ETH-LM session: Number of invalid LMR frames received.	<b>detail extensive</b>
<b>SLM sent</b>	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of synthetic loss measurement (SLM) request packets transmitted from the source MEP to the remote or destination MEP in this session.	<b>detail extensive</b>
<b>Valid SLM received</b>	If the interface is attached to a responder MEP for a ETH-SLM session: Number of valid SLM PDUs transmitted from the source MEP to the remote or destination MEP.	<b>detail extensive</b>
<b>Invalid SLM received</b>	If the interface is attached to a responder MEP for a ETH-SLM session: Number of invalid SLM PDUs transmitted from the source MEP to the remote or destination MEP.	<b>detail extensive</b>
<b>SLR sent</b>	If the interface is attached to a responder MEP for a ETH-SLM session: Number detail extensive of synthetic loss reply (SLR) frames sent.	<b>detail extensive</b>

Table 25: show oam ethernet connectivity-fault-management interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Valid SLR received</b>	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of valid SLR PDUs that the source MEP received from the remote or destination MEP.	<b>detail extensive</b>
<b>Invalid SLR received</b>	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of invalid SLR PDUs that the source MEP received from the remote or destination MEP.	<b>detail extensive</b>
<b>Remote MEP count</b>	Number of remote MEPs.	<b>extensive</b>
<b>Identifier</b> (remote MEP)	MEP identifier of the remote MEP.	<b>extensive</b>
<b>MAC address</b> (remote MEP)	MAC address of the remote MEP.	<b>extensive</b>
<b>State</b> (remote MEP)	State of the remote MEP.	<b>extensive</b>
<b>Interface</b> (remote MEP)	Interface of the remote MEP.	<b>extensive</b>

## Sample Output

### show oam ethernet connectivity-fault-management interfaces

```

user@host> show oam ethernet connectivity-fault-management interfaces
Interface      Link      Status      Level      MEP      Neighbors
               Identifier
ge-1/1/0.0     Up        Active      0          2        1
ge-1/1/0.1     Up        Active      0          2        1
ge-1/1/0.10    Up        Active      0          2        1
ge-1/1/0.100   Up        Active      0          2        1
ge-1/1/0.101   Up        Active      0          2        1
ge-1/1/0.102   Up        Active      0          2        1
ge-1/1/0.103   Up        Active      0          2        1
ge-1/1/0.104   Up        Active      0          2        1
ge-1/1/0.105   Up        Active      0          2        1
ge-1/1/0.106   Up        Active      0          2        1
...

```

### show oam ethernet connectivity-fault-management interfaces detail

```

user@host> show oam ethernet connectivity-fault-management interfaces detail
Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
Maintenance domain name: md0, Format: string, Level: 5
Maintenance association name: ma1, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 1, Direction: down, MAC address: 00:90:69:0b:4b:94

```

```

MEP status: running
Defects:
  Remote MEP not receiving CCM           : no
  Erroneous CCM received                 : yes
  Cross-connect CCM received            : no
  RDI sent by some MEP                  : yes
Statistics:
  CCMs sent                             : 76
  CCMs received out of sequence         : 0
  LBMs sent                             : 0
  Valid in-order LBRs received          : 0
  Valid out-of-order LBRs received      : 0
  LBRs received with corrupted data     : 0
  LBRs sent                             : 0
  LTMs sent                             : 0
  LTMs received                         : 0
  LTRs sent                             : 0
  LTRs received                         : 0
  Sequence number of next LTM request   : 0
  1DMs sent                             : 0
  Valid 1DMs received                   : 0
  Invalid 1DMs received                  : 0
  DMMs sent                             : 0
  DMRs sent                             : 0
  Valid DMRs received                   : 0
  Invalid DMRs received                  : 0
  LMM sent                             : 10
  Valid LMM received                    : 20
  Invalid LMM received                  : 0
  LMR sent                             : 20
  Valid LMR received                    : 10
  Invalid LMR received                  : 0
  SLM sent                             : 10
  Valid SLM received                    : 20
  Invalid SLM received                  : 0
  SLR sent                             : 20
  Valid SLR received                    : 10
  Invalid SLR received                  : 0
Remote MEP count: 2
  Identifier  MAC address  State  Interface
  2001       00:90:69:0b:7f:71  ok    ge-5/2/9.0
  4001       00:90:69:0b:09:c5  ok    ge-5/2/9.0

```

### show oam ethernet connectivity-fault-management interfaces detail (One-Way ETH-DM)

```

user@host show oam ethernet connectivity-fault-management interfaces detail
Interface name: ge-0/2/5.0, Interface status: Active, Link status: Up
Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 101, Direction: down, MAC address: 00:90:69:0a:48:57
MEP status: running
Defects:
  Remote MEP not receiving CCM           : no
  Erroneous CCM received                 : no
  Cross-connect CCM received            : no
  RDI sent by some MEP                  : no
Statistics:
  CCMs sent                             : 1590
  CCMs received out of sequence         : 0

```

```

LBMs sent : 0
Valid in-order LBRs received : 0
Valid out-of-order LBRs received : 0
LBRs received with corrupted data : 0
LBRs sent : 0
LTMs sent : 0
LTMs received : 0
LTRs sent : 0
LTRs received : 0
Sequence number of next LTM request : 0
1DMs sent : 10
Valid 1DMs received : 0
Invalid 1DMs received : 0
DMMs sent : 0
DMRs sent : 0
Valid DMRs received : 0
Invalid DMRs received : 0
Remote MEP count: 1
Identifier MAC address State Interface
201 00:90:69:0a:43:94 ok ge-0/2/5.0

```

**show oam ethernet connectivity-fault-  
management interfaces detail  
(Connection Protection TLV Configured)**

user@hostshow oam ethernet connectivity-fault-management interfaces detail

```

Interface name: xe-6/2/0.0 , Interface status: Active, Link status: Up
Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
Interface status TLV: none, Port status TLV: none
Connection Protection TLV: yes
  Prefer me: no, Protection in use: no, FRR Flag: no
MEP identifier: 1, Direction: down, MAC address: 00:19:e2:b1:14:30
MEP status: running
Defects:
  Remote MEP not receiving CCM : no
  Erroneous CCM received : no
  Cross-connect CCM received : no
  RDI sent by some MEP : no
  Some remote MEP's MAC in error state : no
Statistics:
  CCMs sent : 225
  CCMs received out of sequence : 0
  LBMs sent : 0
  Valid in-order LBRs received : 0
  Valid out-of-order LBRs received : 0
  LBRs received with corrupted data : 0
  LBRs sent : 0
  LTMs sent : 0
  LTMs received : 0
  LTRs sent : 0
  LTRs received : 0
  Sequence number of next LTM request : 0
  1DMs sent : 0
  Valid 1DMs received : 0
  Invalid 1DMs received : 0
  Out of sync 1DMs received : 0
  DMMs sent : 0
  Valid DMMs received : 0

```

```

Invalid DMMs received          : 0
DMRs sent                     : 0
Valid DMRs received           : 0
Invalid DMRs received         : 0
LMMs sent                     : 0
Valid LMMs received           : 0
Invalid LMMs received         : 0
LMRs sent                     : 0
Valid LMRs received           : 0
Invalid LMRs received         : 0
Remote MEP count: 1
  Identifier  MAC address      State  Interface
    2        00:90:69:7f:e4:30

```

### show oam ethernet connectivity-fault-management interfaces extensive

```

user@host> show oam ethernet connectivity-fault-management interfaces extensive
Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
Maintenance domain name: md0, Format: string, Level: 5
Maintenance association name: ma1, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
Interface status TLV: none, Port status TLV: none
Connection Protection TLV: no
MEP identifier: 1, Direction: down, MAC address: 00:90:69:0b:4b:94
MEP status: running
Defects:
  Remote MEP not receiving CCM          : no
  Erroneous CCM received                : yes
  Cross-connect CCM received            : no
  RDI sent by some MEP                  : yes
Statistics:
  CCMs sent                            : 76
  CCMs received out of sequence         : 0
  LBMs sent                            : 0
  Valid in-order LBRs received          : 0
  Valid out-of-order LBRs received      : 0
  LBRs received with corrupted data     : 0
  LBRs sent                            : 0
  LTMs sent                            : 0
  LTMs received                        : 0
  LTRs sent                            : 0
  LTRs received                        : 0
  Sequence number of next LTM request   : 0
  1DMs sent                            : 0
  Valid 1DMs received                  : 0
  Invalid 1DMs received                 : 0
  DMMs sent                            : 0
  DMRs sent                            : 0
  Valid DMRs received                  : 0
  Invalid DMRs received                 : 0
  SLM sent                             : 10
  Valid SLM received                   : 20
  Invalid SLM received                 : 0
  SLR sent                             : 20
  Valid SLR received                   : 10
  Invalid SLR received                 : 0
Remote MEP count: 2
  Identifier  MAC address      State  Interface

```

2001	00:90:69:0b:7f:71	ok	ge-5/2/9.0
4001	00:90:69:0b:09:c5	ok	ge-5/2/9.0

### show oam ethernet connectivity-fault-management interfaces level

```
user@host> show oam ethernet connectivity-fault-management interfaces level 7
```

Interface	Link	Status	Level	MEP Identifier	Neighbors
ge-3/0/0.0	Up	Active	7	201	0
xe-0/0/0.0	Up	Active	7	203	1

### show oam ethernet connectivity-fault-management interfaces (trunk ports)

```
user@host> show oam ethernet connectivity-fault-management interfaces
```

Interface	Link	Status	Level	MEP Identifier	Neighbors
ge-4/0/1.0, vlan 100	Up	Active	5	100	0
ge-10/3/10.4091, vlan 4091	Down	Inactive	4	400	0
ge-4/0/0.0	Up	Active	6	200	0

```
user@host> show oam ethernet connectivity-fault-management interfaces ge-4/0/0.0
```

Interface	Link	Status	Level	MEP Identifier	Neighbors
ge-4/0/0.0	Up	Active	6	200	0

```
user@host> show oam ethernet connectivity-fault-management interfaces ge-4/0/1.0 vlan 100
```

Interface	Link	Status	Level	MEP Identifier	Neighbors
ge-4/0/1.0, vlan 100	Up	Active	5	100	0

```
user@host> show oam ethernet connectivity-fault-management interfaces ge-10/3/10.4091
vlan 4091
```

Interface	Link	Status	Level	MEP Identifier	Neighbors
ge-10/3/10.4091, vlan 4091	Down	Inactive	4	400	0

## show oam ethernet connectivity-fault-management linktrace path-database

<b>Syntax</b>	<b>show oam ethernet connectivity-fault-management linktrace path-database mac-address maintenance-association <i>ma-name</i> maintenance-domain <i>md-name</i></b>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0.
<b>Description</b>	On M320, MX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management maintenance linktrace database information.
<b>Options</b>	<p><b>mac-address</b>—Display connectivity fault management path database information for the specified MAC address of the remote host.</p> <p><b>maintenance-association <i>ma-name</i></b>—Display connectivity fault management path database information for the specified maintenance association.</p> <p><b>maintenance-domain <i>md-name</i></b>—Display connectivity fault management path database information for the specified maintenance domain.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show oam ethernet connectivity-fault-management linktrace path-database on page 232</a></p> <p><a href="#">show oam ethernet connectivity-fault-management linktrace path-database (Two traceroute Commands) on page 232</a></p>
<b>Output Fields</b>	Table 26 on page 231 lists the output fields for the <b>show oam ethernet connectivity-fault-management linktrace path-database</b> command. Output fields are listed in the approximate order in which they appear.

**Table 26: show oam ethernet connectivity-fault-management linktrace path-database Output Fields**

Field Name	Field Description
<b>Linktrace to</b>	MAC address of the 802.1ag node to which the linktrace message is targeted.
<b>Interface</b>	Interface used by the local MEP to send the linktrace message (LTM).
<b>Maintenance Domain</b>	Maintenance domain identifier specified in the traceroute command.
<b>Maintenance Association</b>	Maintenance association identifier specified in the traceroute command.
<b>Level</b>	Maintenance domain level configured for the maintenance domain.
<b>Local Mep</b>	MEP identifier of the local MEP originating the linktrace.
<b>Hop</b>	Sequential hop count of the linktrace path.

**Table 26: show oam ethernet connectivity-fault-management linktrace path-database Output Fields (continued)**

Field Name	Field Description
<b>TTL</b>	Number of hops remaining in the linktrace message (LTM). 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>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 responses (LTR), with a previously sent LTM.

## Sample Output

### show oam ethernet connectivity-fault-management linktrace path-database

```

user@host> show oam ethernet connectivity-fault-management linktrace path-database
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

```

### show oam ethernet connectivity-fault-management linktrace path-database (Two traceroute Commands)

```

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:100002
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

Transaction Identifier:100003
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

## [show oam ethernet connectivity-fault-management mep-database](#)

---

<b>Syntax</b>	<b>show oam ethernet connectivity-fault-management mep-database</b> <b>maintenance-domain</b> <i>domain-name</i> <b>maintenance-association</b> <i>ma-name</i> <b>&lt;local-mep</b> <i>local-mep-id</i> <b>&lt;remote-mep</b> <i>remote-mep-id</i>
<b>Release Information</b>	Command introduced in Junos OS Release 8.4. Support for ITU-T Y.1731 frame delay measurement added in Junos OS Release 9.5. Support for ITU-T Y.1731 synthetic frame loss measurement added in Junos OS Release 13.2 for MX Series routers.
<b>Description</b>	<p>On M7i and M10i routers with Enhanced CFEB (CFEB-E), and on M320, M120, MX Series, ACX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.</p> <p>In addition, on M120, M320, and MX series routers, also display port status TLV, interface status TLV, and action profile information.</p> <p>In addition, for Ethernet interfaces on MX Series routers , also display any ITU-T Y.1731 frame delay measurement (ETH-DM) frame counts.</p> <p>For Ethernet interfaces on MX Series routers, display any ITU-T Y.1731 synthetic frame loss measurement (ETH-SLM) statistics and frame counts.</p>
<b>Options</b>	<p><b>maintenance-association</b> <i>ma-name</i>—Name of the maintenance association.</p> <p><b>maintenance-domain</b> <i>domain-name</i>—Name of the maintenance domain.</p> <p><b>local-mep-id</b>—(Optional) Numeric identifier of local MEP.</p> <p><b>remote-mep-id</b>—(Optional) Numeric identifier of the remote MEP.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">clear oam ethernet connectivity-fault-management statistics</a></li><li>• <a href="#">show oam ethernet connectivity-fault-management delay-statistics on page 213</a></li><li>• <a href="#">show oam ethernet connectivity-fault-management interfaces on page 221</a></li><li>• <a href="#">show oam ethernet connectivity-fault-management mep-statistics on page 245</a></li></ul>
<b>List of Sample Output</b>	<p><a href="#">show oam ethernet connectivity-fault-management mep-database on page 239</a></p> <p><a href="#">show oam ethernet connectivity-fault-management mep-database (One-Way ETH-DM) on page 240</a></p> <p><a href="#">show oam ethernet connectivity-fault-management mep-database local-mep remote-mep on page 241</a></p>

[show oam ethernet connectivity-fault-management mep-database remote-mep \(Action Profile Event\) on page 241](#)

[show oam ethernet connectivity-fault-management mep-database \(Connection Protection TLV Configured\) on page 241](#)

[show oam ethernet connectivity-fault-management mep-database on page 242](#)

[show oam ethernet connectivity-fault-management mep-database \(enhanced continuity measurement\) on page 243](#)

**Output Fields** Table 27 on page 235 lists the output fields for the **show oam ethernet connectivity-fault-management mep-database** command. Output fields are listed in the approximate order in which they appear.

**Table 27: show oam ethernet connectivity-fault-management mep-database Output Fields**

Field Name	Field Description
Maintenance domain name	Maintenance domain name.
Format (Maintenance domain)	Maintenance domain name format configured.
Level	Maintenance domain level configured.
Maintenance association name	Maintenance association name.
Format (Maintenance association)	Maintenance association name format configured.
Continuity-check status	Continuity-check status.
Interval	Continuity-check message interval.
Loss-threshold	Lost continuity-check message threshold.
Connection Protection TLV	Status of the connection protection TLV, if configured on the MEP interface: <b>no</b> , <b>yes</b>  If <b>yes</b> , then the transmitted connection protection TLV is decoded and the following three fields are displayed: <b>Prefer me</b> , <b>Protection in use</b> , <b>FRR Flag</b>
Prefer me	If set to <b>yes</b> , the path through which CCM was transmitted is preferred (unless the path fails). It is used for signaling a manual-switch command to remote side.  Its value can be <b>yes</b> or <b>no</b> .
Protection in use	Used for protection decision coordination. Its value is set to <b>yes</b> if the endpoint transmitting the CCM is currently transmitting the user traffic to protection path.  Its value can be <b>yes</b> or <b>no</b> .
FRR Flag	LSR/LER forwarding the CCM Frame into a bypass tunnel is set.  Its value can be <b>yes</b> or <b>no</b> .

Table 27: show oam ethernet connectivity-fault-management mep-database Output Fields (*continued*)

Field Name	Field Description
MEP identifier	Maintenance association end point (MEP) identifier.
Direction	MEP direction configured.
MAC address	MAC address configured for the MEP.
Auto-discovery	Whether automatic discovery is enabled or disabled.
Priority	Priority used for CCMs and linktrace messages transmitted by the MEP.
Interface name	Interface identifier.
Interface status	Local interface status.
Link status	Local link status.
Remote MEP not receiving CCM	Whether the remote MEP is not receiving CCMs.
Erroneous CCM received	Whether erroneous CCMs have been received.
Cross-connect CCM received	Whether cross-connect CCMs have been received.
RDI sent by some MEP	Whether the remote defect indication (RDI) bit is set in messages that have been received. The absence of the RDI bit in a CCM indicates that the transmitting MEP is receiving CCMs from all configured MEPs.
CCMs sent	Number of CCMs transmitted.
CCMs received out of sequence	Number of CCMs received out of sequence.
LBMs sent	Number of loopback messages (LBMs) sent.
Valid in-order LBRs received	Number of loopback response messages (LBRs) received that were valid messages and in sequence.
1DMs sent	If the MEP is an initiator for a one-way ETH-DM session: Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session.  For all other cases, this field displays 0.
Valid 1DMs received	If the MEP is a receiver for a one-way ETH-DM session: Number of valid 1DM frames received.  For all other cases, this field displays 0.

**Table 27: show oam ethernet connectivity-fault-management mep-database Output Fields (*continued*)**

Field Name	Field Description
<b>Invalid 1DMs received</b>	If the MEP is a receiver for a one-way ETH-DM session: Number of invalid 1DM frames received.  For all other cases, this field displays 0.
<b>Out of sync 1DMs received</b>	If the MEP is a receiver for a one-way ETH-DM session: Number of out-of-sync one-way delay measurement request packets received.
<b>DMMs sent</b>	If the MEP is an initiator for a two-way ETH-DM session: Number of Delay Measurement Message (DMM) PDU frames sent to the peer MEP in this session.  For all other cases, this field displays 0.
<b>Valid DMMs received</b>	If the MEP is an initiator for a two-way ETH-DM session: Number of valid two-way delay measurement packets received.
<b>Invalid DMMs received</b>	If the MEP is an initiator for a two-way ETH-DM session: Number of invalid two-way delay measurement packets received.
<b>DMRs sent</b>	If the MEP is a responder for a ETH-DM session: Number of Delay Measurement Reply (DMR) frames sent.  For all other cases, this field displays 0.
<b>Valid DMRs received</b>	If the MEP is an initiator for a two-way ETH-DM session: Number of valid DMRs received.  For all other cases, this field displays 0.
<b>Invalid DMRs received</b>	If the MEP is an initiator for a two-way ETH-DM session: Number of invalid DMRs received.  For all other cases, this field displays 0.
<b>Valid out-of-order LBRs received</b>	Number of LBRs received that were valid messages and not in sequence.
<b>LBRs received with corrupted data</b>	Number of LBRs received that were corrupted.
<b>LBRs sent</b>	Number of LBRs transmitted.
<b>LTMs sent</b>	Linktrace messages (LTMs) transmitted.
<b>LTMs received</b>	Linktrace messages received.
<b>LTRs sent</b>	Linktrace responses (LTRs) transmitted.
<b>LTRs received</b>	Linktrace responses received.
<b>Sequence number of next LTM request</b>	Sequence number of the next linktrace message request to be transmitted.

Table 27: show oam ethernet connectivity-fault-management mep-database Output Fields (*continued*)

Field Name	Field Description
<b>LMM sent</b>	If the interface is attached to an initiator MEP for a ETH-LM session: Number of loss measurement message (LMM) PDU frames sent to the peer MEP in this session.
<b>Valid LMM received</b>	If the interface is attached to an initiator MEP for a ETH-LM session: Number of valid loss measurement request packets received.
<b>Invalid LMM received</b>	If the interface is attached to an initiator MEP for a ETH LM session: Number of invalid loss measurement request packets received.
<b>LMR sent</b>	If the interface is attached to a responder MEP for a ETH-LM session: Number of loss measurement reply (LMR) frames sent.
<b>Valid LMR received</b>	If the interface is attached to an initiator MEP for a ETH LM session: Number of valid LMR frames received.
<b>Invalid LMR received</b>	If the interface is attached to an initiator MEP for a ETH-LM session: Number of invalid LMR frames received.
<b>SLM sent</b>	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of synthetic loss measurement (SLM) request packets transmitted from the source MEP to the remote or destination MEP in this session.
<b>Valid SLM received</b>	If the interface is attached to a responder MEP for a ETH-SLM session: Number of valid SLM PDUs transmitted from the source MEP to the remote or destination MEP.
<b>Invalid SLM received</b>	If the interface is attached to a responder MEP for a ETH-SLM session: Number of invalid SLM PDUs transmitted from the source MEP to the remote or destination MEP.
<b>SLR sent</b>	If the interface is attached to a responder MEP for a ETH-SLM session: Number detail extensive of synthetic loss reply (SLR) frames sent.
<b>Valid SLR received</b>	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of valid SLR PDUs that the source MEP received from the remote or destination MEP.
<b>Invalid SLR received</b>	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of invalid SLR PDUs that the source MEP received from the remote or destination MEP.
<b>Remote MEP identifier</b>	MEP identifier of the remote MEP.
<b>State (remote MEP)</b>	State of the remote MEP: <b>idle</b> , <b>start</b> , <b>ok</b> , or <b>failed</b> .
<b>MAC address</b>	MAC address of the remote MEP.
<b>Type</b>	Whether the remote MEP MAC address was learned using automatic discovery or configured.
<b>Interface</b>	Interface of the remote MEP. A seven-digit number is appended if CFM is configured to run on a routing instance of type VPLS.

Table 27: show oam ethernet connectivity-fault-management mep-database Output Fields (*continued*)

Field Name	Field Description
Last flapped	Date, time, and how long ago the remote MEP 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, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).
Remote defect indication	Whether the remote defect indication (RDI) bit is set in messages that have been received or transmitted.
Port status TLV	<ul style="list-style-type: none"> <li>In the Maintenance domain section, displays the last transmitted port status TLV value.</li> <li>In the Remote MEP section, displays the last value of port status TLV received from the remote MEP.</li> </ul> <p>In the Action profile section, displays, the last occurred event <b>port-status-tlv blocked</b> event. This event occurred due to the reception of <b>blocked</b> value in the port status TLV from remote MEP.</p>
Interface status TLV	<ul style="list-style-type: none"> <li>In the Maintenance domain section, displays the last transmitted interface status TLV value.</li> <li>In the Remote MEP section, displays the last value of interface status TLV received from the remote MEP.</li> </ul> <p>In the Action profile section, if displays, the last occurred event interface-status-tlv event ( either <b>lower-layer-down</b> or <b>down</b>). This event occurred due to the reception of either lower or <b>down</b> value in the interface status TLV from remote MEP.</p>
Action profile	Name of the action profile occurrence associated with a remote MEP.
Last event	When an action profile occurs, displays the last event that triggered it.
Last event cleared	When all the configured and occurred events (under action profile) are cleared, then the action taken gets reverted (such as down interface is made up) and the corresponding time is noted and displayed.
Action	Action taken and the corresponding time of the action occurrence.

## Sample Output

### show oam ethernet connectivity-fault-management mep-database

```

user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain vpls-vlan2000 maintenance-association vpls-vlan200
Maintenance domain name: vpls-vlan2000, Format: string, Level: 5
Maintenance association name: vpls-vlan200, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 200, Direction: up, MAC address: 00:19:e2:b0:74:01
Auto-discovery: enabled, Priority: 0
Interface status TLV: none, Port status TLV: none
Connection Protection TLV: no Interface name: ge-0/0/1.0, Interface status:
Active, Link status: Up
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                      : no
  Cross-connect CCM received                  : no
  RDI sent by some MEP                       : no
Statistics:

```

```

CCMs sent : 1476
CCMs received out of sequence : 0
LBMs sent : 85
Valid in-order LBRs received : 78
Valid out-of-order LBRs received : 0
LBRs received with corrupted data : 0
LBRs sent : 0
LTMs sent : 1
LTMs received : 0
LTRs sent : 0
LTRs received : 1
Sequence number of next LTM request : 1
IDMs sent : 0
Valid IDMs received : 0
Invalid IDMs received : 0
DMMs sent : 0
DMRs sent : 0
Valid DMRs received : 0
Invalid DMRs received : 0
Remote MEP count: 1
  Identifier   MAC address      State   Interface
    100      00:19:e2:b2:81:4b   ok   vt-0/1/10.1049088

```

#### show oam ethernet connectivity-fault- management mep-database (One-Way ETH-DM)

```

user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md6 maintenance-domain ma6
Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 101, Direction: down, MAC address: 00:90:69:0a:48:57
Auto-discovery: enabled, Priority: 0
Interface name: ge-0/2/5.0, Interface status: Active, Link status: Up
Defects:
  Remote MEP not receiving CCM : no
  Erroneous CCM received : no
  Cross-connect CCM received : no
  RDI sent by some MEP : no
Statistics:
  CCMs sent : 1590
  CCMs received out of sequence : 0
  LBMs sent : 0
  Valid in-order LBRs received : 0
  Valid out-of-order LBRs received : 0
  LBRs received with corrupted data : 0
  LBRs sent : 0
  LTMs sent : 0
  LTMs received : 0
  LTRs sent : 0
  LTRs received : 0
  Sequence number of next LTM request : 0
  IDMs sent : 10
  Valid IDMs received : 0
  Invalid IDMs received : 0
  DMMs sent : 0
  DMRs sent : 0
  Valid DMRs received : 0
  Invalid DMRs received : 0
Remote MEP count: 1

```

Identifier	MAC address	State	Interface
201	00:90:69:0a:43:94	ok	ge-0/2/5.0

#### show oam ethernet connectivity-fault-management mep-database local-mep remote-mep

```

user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain vpls-vlan2000 maintenance-association vpls-vlan200 local-mep 200
remote-mep 100
Maintenance domain name: vpls-vlan2000, Format: string, Level: 5
Maintenance association name: vpls-vlan200, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 200, Direction: up, MAC address: 00:19:e2:b0:74:01
Auto-discovery: enabled, Priority: 0
Interface name: ge-0/0/1.0, Interface status: Active, Link status: Up

Remote MEP identifier: 100, State: ok
MAC address: 00:19:e2:b2:81:4b, Type: Learned
Interface: vt-0/1/10.1049088
Last flapped: Never
Remote defect indication: false
Port status TLV: none
Interface status TLV: none

```

#### show oam ethernet connectivity-fault-management mep-database remote-mep (Action Profile Event)

```

user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 remote-mep 200
Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 100, Direction: down, MAC address: 00:05:85:73:e8:ad
Auto-discovery: enabled, Priority: 0
Interface status TLV: none, Port status TLV: none
Interface name: ge-1/0/8.0, Interface status: Active, Link status: Up

Remote MEP identifier: 200, State: ok
MAC address: 00:05:85:73:96:1f, Type: Configured
Interface: ge-1/0/8.0
Last flapped: Never
Remote defect indication: false
Port status TLV: none
Interface status TLV: lower-layer-down
Action profile: juniper
  Last event: Interface-status-tlv lower-layer-down
  Action: Interface-down, Time: 2009-03-27 14:25:10 PDT (00:00:02 ago)

```

#### show oam ethernet connectivity-fault-management mep-database (Connection Protection TLV Configured)

```

user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5

```

If connection-protection is not enabled on down MEPs, but connection-protection TLV is used, MX always sets the protection-in-use flag in connection-protection tlv, while CCMs are sent out. During reversion, this is an indicator to the receiver that protect-path is in use, otherwise the peer (receiver) assumes working is active and reversion does not work as expected. Setting this bit does not affect protection-switching/traffic-loss.

```

Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 1, Direction: down, MAC address: 00:19:e2:b1:14:30
Auto-discovery: enabled, Priority: 0
Interface status TLV: none, Port status TLV: none
Connection Protection TLV: yes
  Prefer me: no, Protection in use: no, FRR Flag: no
Interface name: xe-6/2/0.0, Interface status: Active, Link status: Up
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                      : no
  Cross-connect CCM received                  : no
  RDI sent by some MEP                       : no
  Some remote MEP's MAC in error state        : no
Statistics:
  CCMs sent                                  : 251
  CCMs received out of sequence               : 0
  LBMs sent                                  : 0
  Valid in-order LBRs received                : 0
  Valid out-of-order LBRs received            : 0
  LBRs received with corrupted data           : 0
  LBRs sent                                  : 0
  LTMs sent                                  : 0
  LTMs received                              : 0
  LTRs sent                                  : 0
  LTRs received                              : 0
  Sequence number of next LTM request         : 0
  1DMs sent                                  : 0
  Valid 1DMs received                        : 0
  Invalid 1DMs received                      : 0
  Out of sync 1DMs received                  : 0
  DMMs sent                                  : 0
  Valid DMMs received                       : 0
  Invalid DMMs received                     : 0
  DMRs sent                                  : 0
  Valid DMRs received                       : 0
  Invalid DMRs received                     : 0
  LMMs sent                                  : 0
  Valid LMMs received                       : 0
  Invalid LMMs received                     : 0
  LMRs sent                                  : 0
  Valid LMRs received                       : 0
  Invalid LMRs received                     : 0
Remote MEP count: 1
Identifier    MAC address    State    Interface
  2          00:90:69:7f:e4:30

```

### show oam ethernet connectivity-fault-management mep-database

```

user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5
Maintenance association name: ma1, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 1, Direction: down, MAC address: 00:14:f6:b6:01:fe
Auto-discovery: enabled, Priority: 0
Interface name: ge-1/0/0.0, Interface status: Active, Link status: Up

Defects:
Remote MEP not receiving CCM                : no

```

```

Erroneous CCM received           : no
Cross-connect CCM received       : no
RDI sent by some MEP            : no

Statistics:
CCMs sent                        : 328703
CCMs received out of sequence    : 0
LBMs sent                       : 85
Valid in-order LBRs received     : 78
Valid out-of-order LBRs received : 0
LBRs received with corrupted data : 0
LBRs sent                       : 0
LTMs sent                       : 0
LTMs received                   : 0
LTRs sent                       : 0
LTRs received                   : 0
Sequence number of next LTM request : 0
1DMs sent                       : 10
Valid 1DMs received             : 10
Invalid 1DMs received           : 0
DMMs sent                      : 20
DMRs sent                      : 0
Valid DMRs received            : 10
Invalid DMRs received           : 0
LMM sent                       : 10
Valid LMM received             : 20
Invalid LMM received           : 0
LMR sent                       : 20
Valid LMR received             : 10
Invalid LMR received           : 0
SLM sent                       : 10
Valid SLM received             : 20
Invalid SLM received           : 0
SLR sent                       : 20
Valid SLR received             : 10
Invalid SLR received           : 0

Remote MEP count                 : 1

Identifier    MAC address    State    Interface
  2          00:12:1e:fb:ea:7d    ok      ge-1/0/0.0

```

#### show oam ethernet connectivity-fault- management mep-database (enhanced continuity measurement)

```

user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 local-mep 2001 remote-mep 1001
Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
Auto-discovery: enabled, Priority: 0
Interface status TLV: up, Port status TLV: up
Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

Remote MEP identifier: 1001, State: ok
MAC address   : 00:19:e2:b0:74:00, Type: Learned
Interface     : ge-2/0/0.0
Last flapped  : Never
+ Continuity  : 91%, Admin-enable duration: 2100sec, Oper-down duration: 100sec
Remote defect indication: false

```

Port status TLV: none  
Interface status TLV: none

## show oam ethernet connectivity-fault-management mep-statistics

<b>Syntax</b>	<pre>show oam ethernet connectivity-fault-management mep-statistics maintenance-domain <i>md-name</i> maintenance-association <i>ma-name</i> &lt;mep <i>mep-id</i>&gt; &lt;remote-mep <i>remote-mep-id</i>&gt; &lt;count <i>entry-count</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 11.4 for EX Series switches.</p> <p>Support for ITU-T Y.1731 Ethernet synthetic frame loss measurement (ETH-SLM) added in Junos OS Release 13.2 for MX Series routers.</p>
<b>Description</b>	<p>On MX Series and ACX Series routers and EX Series switches with Ethernet interfaces, display ETH-DM statistics and ETH-DM frame counts.</p> <p>For Ethernet interfaces on MX Series routers, display any ITU-T Y.1731 synthetic frame loss measurement (ETH-SLM) statistics and frame counts.</p>
<b>Options</b>	<p><b>maintenance-domain <i>md-name</i></b>—Name of an existing CFM maintenance domain.</p> <p><b>maintenance-association <i>ma-name</i></b>—Name of an existing CFM maintenance association.</p> <p><b>mep <i>mep-id</i></b>—(Optional) Numeric identifier of the local MEP. The range of values is 1 through 8192. On EX Series switches, the range of values is 1 through 8191.</p> <p><b>remote-mep <i>remote-mep-id</i></b>—(Optional) Numeric identifier of the remote MEP. The range of values is 1 through 8192. On EX Series switches, the range of values is 1 through 8191.</p> <p><b>count <i>entry-count</i></b>—(Optional) Number of entries to display from the statistics table. The range of values is 1 through 100. The default value is 100 entries.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear oam ethernet connectivity-fault-management statistics</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management delay-statistics on page 213</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management interfaces on page 221</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management mep-database on page 234</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show oam ethernet connectivity-fault-management mep-statistics (CIR counters only) on page 247</a></p> <p><a href="#">show oam ethernet connectivity-fault-management mep-statistics (CIR and EIR counters enabled) on page 249</a></p> <p><a href="#">show oam ethernet connectivity-fault-management mep-statistics remote-mep (CIR counters only) on page 250</a></p>

[show oam ethernet connectivity-fault-management mep-statistics remote-mep \(CIR and EIR counters enabled\) on page 251](#)

[show oam ethernet connectivity-fault-management mep-statistics on page 253](#)

[show oam ethernet connectivity-fault-management mep-statistics remote-mep on page 254](#)

**Output Fields** Table 28 on page 246 lists the output fields for the **show oam ethernet connectivity-fault-management mep-statistics** command. Output fields are listed in the approximate order in which they appear.

**Table 28: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields**

Output Field Name	Field Description
MEP identifier	Maintenance association end point (MEP) numeric identifier.
MAC address	Unicast MAC address configured for the MEP.
Remote MEP count	Number of remote MEPs (unless you specify the <b>remote-mep</b> option).
Remote MEP identifier	Numeric identifier of the remote MEP.
Remote MAC address	Unicast MAC address of the remote MEP.
Index	Index number that corresponds to the ETH-DM entry in the CFM database.
One-way delay (usec)	<p>For a one-way ETH-DM session, the frame delay time, in microseconds, measured at the receiver MEP.</p> <p>For a detailed description of one-way Ethernet frame delay measurement, see the <i>ITU-T Y.1731 Ethernet Service OAM</i> topics in the <i>Junos OS Network Interfaces Library for Routing Devices</i>.</p>
Two-way delay (usec)	<p>For a two-way ETH-DM session, the frame delay time, in microseconds, measured at the initiator MEP.</p> <p>For a detailed description of two-way Ethernet frame delay measurement, see the <i>ITU-T Y.1731 Ethernet Service OAM</i> topics in the <i>Junos OS Network Interfaces Library for Routing Devices</i>.</p>
Average one-way delay	Average one-way frame delay for the statistics displayed.
Average one-way delay variation	Average one-way “frame jitter” for the statistics displayed.
Best-case one-way delay	Lowest one-way frame delay for the statistics displayed.
Worst-case one-way delay	Highest one-way frame delay for the statistics displayed.
Average two-way delay	Average two-way frame delay for the statistics displayed.
Average two-way delay variation	Average two-way “frame jitter” for the statistics displayed.

**Table 28: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields (*continued*)**

Output Field Name	Field Description
Best-case two-way delay	Lowest two-way frame delay for the statistics displayed.
Worst-case two-way delay	Highest two-way frame delay calculated in this session.
SLM packets sent	Total number of synthetic loss message (SLM) PDU frames sent from the source MEP to the remote MEP during this ETH-SLM session.
SLM packets received	Total number of synthetic loss message (SLM) PDU frames that the remote MEP received from the source MEP during this ETH-SLM session.
SLR packets sent	Total number of synthetic loss reply (SLR) PDU frames that the remote MEP sent to the source MEP during this measurement session.
SLR packets received	Total number of synthetic loss reply (SLR) PDU frames that the source MEP received from the remote MEP during this measurement session.
Local TXFC1 value	Number of synthetic frames transmitted to the peer MEP for a test ID. A test ID is used to distinguish each synthetic loss measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SLM for the MEG and initiating MEP.
Local RXFC1 value	Number of synthetic frames received from the peer MEP for a test ID. The MEP generates a unique Test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TXFC1 is sent in the packet.
Last Received SLR frame TXFCf(tc)	Value of the local counter TxFC1 at the time of SLM frame transmission.
Last Received SLR frame TXFCb(t)	Value of the local counter RxFC1 at the time of SLR frame transmission.
Frame loss (near-end)	Count of frame loss associated with ingress data frames.
Frame loss (far-end)	Count of frame loss associated with egress data frames.

## Sample Output

**show oam ethernet connectivity-fault-management mep-statistics (CIR counters only)**

```

user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md1 maintenance-association ma-1 local-mep 3 remote-mep 103 count 3
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count           : 1
CCMs sent                  : 6550
CCMs received out of sequence : 0
LBMs sent                  : 0
Valid in-order LBRs received : 0
Valid out-of-order LBRs received : 0
LBRs received with corrupted data : 0

```

```

LBRs sent : 0
LTMs sent : 0
LTMs received : 0
LTRs sent : 0
LTRs received : 0
Sequence number of next LTM request : 0
1DMs sent : 5
Valid 1DMs received : 0
Invalid 1DMs received : 0
DMMs sent : 5
DMRs sent : 0
Valid DMRs received : 5
Invalid DMRs received : 0
LMM sent : 5
Valid LMM received : 5
Invalid LMM received : 0
LMR sent : 0
Valid LMR received : 5
Invalid LMR received : 0
Remote MEP identifier : 101
Remote MAC address : 00:05:85:73:39:4a

```

## Delay measurement statistics:

Index	One-way delay (usec)	Two-way delay (usec)
1	259	519
2	273	550
3	287	571
4	299	610
5	313	650

```

Average one-way delay : 286 usec
Average one-way delay variation : 62 usec
Best case one-way delay : 259 usec
Average two-way delay : 580 usec
Average two-way delay variation : 26 usec
Best case two-way delay : 519 usec
Worst case two-way delay : 650 usec

```

## Loss measurement statistics:

Index	Near-end Frame loss (CIR)	Far-end Frame loss (CIR)	Near-end Frame loss (EIR)	Far-end Frame loss (EIR)
1	9	9		
2	3	5		
3	7	5		
4	9	6		
5	3	6		

```

Average near-end loss (CIR) : 6.2
Average near-end loss ratio (CIR) : 6.2%
Average far-end loss (CIR) : 6.2
Average far-end loss ratio (CIR) : 6.2%
Near-end best case loss (CIR) : 3
Near-end best case loss ratio (CIR) : 3%
Near-end worst case loss (CIR) : 9
Near-end worst case loss ratio (CIR) : 9%
Far-end best case loss (CIR) : 5
Far-end best case loss ratio (CIR) : 5%
Far-end worst case loss (CIR) : 9
Far-end worst case loss ratio (CIR) : 9%

```

### show oam ethernet connectivity-fault-management mep-statistics (CIR and EIR counters enabled)

```
user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md1 maintenance-association ma-1 local-mep 3 remote-mep 103 count 3
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
```

```
Remote MEP count           : 1
CCMs sent                  : 6550
CCMs received out of sequence : 0
LBMs sent                  : 0
Valid in-order LBRs received : 0
Valid out-of-order LBRs received : 0
LBRs received with corrupted data : 0
LBRs sent                  : 0
LTMs sent                  : 0
LTMs received              : 0
LTRs sent                  : 0
LTRs received              : 0
Sequence number of next LTM request : 0
IDMs sent                  : 5
Valid IDMs received        : 0
Invalid IDMs received      : 0
DMMs sent                  : 5
DMRs sent                  : 0
Valid DMRs received        : 5
Invalid DMRs received      : 0
LMM sent                   : 5
Valid LMM received         : 5
Invalid LMM received       : 0
LMR sent                   : 0
Valid LMR received         : 5
Invalid LMR received       : 0
Remote MEP identifier       : 101
Remote MAC address          : 00:05:85:73:39:4a
```

#### Delay measurement statistics:

Index	One-way delay (usec)	Two-way delay (usec)
1	259	519
2	273	550
3	287	571
4	299	610
5	313	650

```
Average one-way delay           : 286 usec
Average one-way delay variation  : 62 usec
Best case one-way delay          : 259 usec
Average two-way delay            : 580 usec
Average two-way delay variation  : 26 usec
Best case two-way delay          : 519 usec
Worst case two-way delay         : 650 usec
```

#### Loss measurement statistics:

Index	Near-end Frame loss (CIR)	Far-end Frame loss (CIR)	Near-end Frame loss (EIR)	Far-end Frame loss (EIR)
1	9	9	2	4
2	3	5	4	6
3	7	5	0	2
4	9	6	8	2
5	3	6	6	4

Average near-end loss (CIR)	: 6.2
Average near-end loss ratio (CIR)	: 6.2%
Average far-end loss (CIR)	: 6.2
Average far-end loss ratio (CIR)	: 6.2%
Near-end best case loss (CIR)	: 3
Near-end best case loss ratio (CIR)	: 3%
Near-end worst case loss (CIR)	: 9
Near-end worst case loss ratio (CIR)	: 9%
Far-end best case loss (CIR)	: 5
Far-end best case loss ratio (CIR)	: 5%
Far-end worst case loss (CIR)	: 9
Far-end worst case loss ratio (CIR)	: 9%
Average near-end loss (EIR)	: 4
Average near-end loss ratio (EIR)	: 4%
Average far-end loss (EIR)	: 3.4
Average far-end loss ratio (EIR)	: 3.4%
Near-end best case loss (EIR)	: 0
Near-end best case loss ratio (EIR)	: 0%
Near-end worst case loss (EIR)	: 8
Near-end worst case loss ratio (EIR)	: 8%
Far-end best case loss (EIR)	: 2
Far-end best case loss ratio (EIR)	: 2%
Far-end worst case loss (EIR)	: 6
Far-end worst case loss ratio (EIR)	: 6%

#### show oam ethernet connectivity-fault-management mep-statistics remote-mep (CIR counters only)

```

user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md1 maintenance-association ma-1 local-mep 3 remote-mep 103 count 3
remote-mep 101
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
CCMs sent : 7762
CCMs received out of sequence : 0
LBMs sent : 0
Valid in-order LBRs received : 0
Valid out-of-order LBRs received : 0
LBRs received with corrupted data : 0
LBRs sent : 0
LTMs sent : 0
LTMs received : 0
LTRs sent : 0
LTRs received : 0
Sequence number of next LTM request : 0
IDMs sent : 5
Valid IDMs received : 0
Invalid IDMs received : 0
DMMs sent : 5
DMRs sent : 0
Valid DMRs received : 5
Invalid DMRs received : 0
LMM sent : 5
Valid LMM received : 5
Invalid LMM received : 0
LMR sent : 0
Valid LMR received : 5
Invalid LMR received : 0
Remote MEP identifier : 101
Remote MAC address : 00:05:85:73:39:4a

```

## Delay measurement statistics:

Index	One-way delay (usec)	Two-way delay (usec)
1	259	519
2	273	550
3	287	571
4	299	610
5	313	650

Average one-way delay : 286 usec  
 Average one-way delay variation : 62 usec  
 Best case one-way delay : 259 usec  
 Average two-way delay : 580 usec  
 Average two-way delay variation : 26 usec  
 Best case two-way delay : 519 usec  
 Worst case two-way delay : 650 usec

## Loss measurement statistics:

Index	Near-end Frame loss (CIR)	Far-end Frame loss (CIR)	Near-end Frame loss (EIR)	Far-end Frame loss (EIR)
1	9	9		
2	3	5		
3	7	5		
4	9	6		
5	3	6		

Average near-end loss (CIR) : 6.2  
 Average near-end loss ratio (CIR) : 6.2%  
 Average far-end loss (CIR) : 6.2  
 Average far-end loss ratio (CIR) : 6.2%  
 Near-end best case loss (CIR) : 3  
 Near-end best case loss ratio (CIR) : 3%  
 Near-end worst case loss (CIR) : 9  
 Near-end worst case loss ratio (CIR) : 9%  
 Far-end best case loss (CIR) : 5  
 Far-end best case loss ratio (CIR) : 5%  
 Far-end worst case loss (CIR) : 9  
 Far-end worst case loss ratio (CIR) : 9%  
 Average near-end loss (EIR) : 4  
 Average near-end loss ratio (EIR) : 4%  
 Average far-end loss (EIR) : 3.4  
 Average far-end loss ratio (EIR) : 3.4%  
 Near-end best case loss (EIR) : 0  
 Near-end best case loss ratio (EIR) : 0%  
 Near-end worst case loss (EIR) : 8  
 Near-end worst case loss ratio (EIR) : 8%  
 Far-end best case loss (EIR) : 2  
 Far-end best case loss ratio (EIR) : 2%  
 Far-end worst case loss (EIR) : 6  
 Far-end worst case loss ratio (EIR) : 6%

### show oam ethernet connectivity-fault-management mep-statistics remote-mep (CIR and EIR counters enabled)

```

user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md1 maintenance-association ma-1 local-mep 3 remote-mep 103 count 3
remote-mep 101
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
CCMs sent : 7762
CCMs received out of sequence : 0
LBMs sent : 0
  
```

```

Valid in-order LBRs received      : 0
Valid out-of-order LBRs received : 0
LBRs received with corrupted data : 0
LBRs sent                        : 0
LTMs sent                        : 0
LTMs received                    : 0
LTRs sent                        : 0
LTRs received                    : 0
Sequence number of next LTM request : 0
1DMs sent                        : 5
Valid 1DMs received              : 0
Invalid 1DMs received            : 0
DMMs sent                        : 5
DMRs sent                        : 0
Valid DMRs received              : 5
Invalid DMRs received            : 0
LMM sent                         : 5
Valid LMM received               : 5
Invalid LMM received             : 0
LMR sent                         : 0
Valid LMR received               : 5
Invalid LMR received             : 0
Remote MEP identifier            : 101
Remote MAC address               : 00:05:85:73:39:4a

```

## Delay measurement statistics:

Index	One-way delay (usec)	Two-way delay (usec)
1	259	519
2	273	550
3	287	571
4	299	610
5	313	650

```

Average one-way delay      : 286 usec
Average one-way delay variation : 62 usec
Best case one-way delay    : 259 usec
Average two-way delay      : 580 usec
Average two-way delay variation : 26 usec
Best case two-way delay    : 519 usec
Worst case two-way delay   : 650 usec

```

## Loss measurement statistics:

Index	Near-end Frame loss (CIR)	Far-end Frame loss (CIR)	Near-end Frame loss (EIR)	Far-end Frame loss (EIR)
1	10	8	5	12
2	12	7	6	16
3	7	5	0	2
4	9	6	8	2
5	3	6	6	4

```

Average near-end loss (CIR)      : 6.2
Average near-end loss ratio (CIR) : 6.2%
Average far-end loss (CIR)       : 6.2
Average far-end loss ratio (CIR) : 6.2%
Near-end best case loss (CIR)    : 3
Near-end best case loss ratio (CIR) : 3%
Near-end worst case loss (CIR)   : 9
Near-end worst case loss ratio (CIR) : 9%
Far-end best case loss (CIR)     : 5

```

```

Far-end best case loss ratio (CIR)      : 5%
Far-end worst case loss (CIR)           : 9
Far-end worst case loss ratio (CIR)     : 9%
Average near-end loss (EIR)             : 4
Average near-end loss ratio (EIR)       : 4%
Average far-end loss (EIR)              : 3.4
Average far-end loss ratio (EIR)        : 3.4%
Near-end best case loss (EIR)           : 0
Near-end best case loss ratio (EIR)     : 0%
Near-end worst case loss (EIR)          : 8
Near-end worst case loss ratio (EIR)    : 8%
Far-end best case loss (EIR)            : 2
Far-end best case loss ratio (EIR)      : 2%
Far-end worst case loss (EIR)           : 6
Far-end worst case loss ratio (EIR)     : 6%

```

### show oam ethernet connectivity-fault-management mep-statistics

```

user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md1 maintenance-association ma-1

```

```
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
```

```
Remote MEP count: 1
```

```

CCMs sent                               : 6550
CCMs received out of sequence           : 0
LBMs sent                               : 0
Valid in-order LBRs received            : 0
Valid out-of-order LBRs received        : 0
LBRs received with corrupted data       : 0
LBRs sent                               : 0
LTMs sent                               : 0
LTMs received                           : 0
LTRs sent                               : 0
LTRs received                           : 0
Sequence number of next LTM request     : 0
1DMs sent                               : 5
Valid 1DMs received                     : 0
Invalid 1DMs received                   : 0
DMMs sent                               : 5
DMRs sent                               : 0
Valid DMRs received                     : 5
Invalid DMRs received                   : 0
SLM sent                                : 10
Valid SLM received                      : 20
Invalid SLM received                    : 0
SLR sent                                : 20
Valid SLR received                      : 10
Invalid SLR received                    : 0

```

```
Remote MEP identifier: 101
```

```
Remote MAC address: 00:05:85:73:39:4a
```

```
Delay measurement statistics:
```

Index	One-way delay (usec)	Two-way delay (usec)
1	259	519
2	273	550
3	287	571
4	299	610
5	313	650

```
Average one-way delay : 286 usec
```

```
Average one-way delay variation: 62 usec
```

```

Best case one-way delay      : 259 usec
Worst case one-way delay    : 313 usec
Average two-way delay       : 580 usec
Average two-way delay variation: 26 usec
Best case two-way delay     : 519 usec
Worst case two-way delay    : 650 usec
Synthetic Loss measurement
statistics:
  SLM packets sent          : 100
  SLM packets received     : 0
  SLR packets sent         : 100
  SLR packets received     : 0
  Accumulated SLM statistics:
    Local TXFC1 value      : 100
    Local RXFC1 value      : 100
    Last Received SLR frame TXFCftc : 100
    Last Received SLR frame TXFCbtc : 100
  SLM Frame Loss:
    Frame Loss (far-end)   : 0 (0.00 %)
    Frame Loss (near-end)  : 0 (0.00 %)

```

#### show oam ethernet connectivity-fault-management mep-statistics remote-mep

```

user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md1 maintenance-association ma1 remote-mep 101
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
  CCMs sent                  : 7762
  CCMs received out of sequence : 0
  LBMs sent                  : 0
  Valid in-order LBRs received : 0
  Valid out-of-order LBRs received : 0
  LBRs received with corrupted data : 0
  LBRs sent                  : 0
  LTMs sent                  : 0
  LTMs received              : 0
  LTRs sent                  : 0
  LTRs received              : 0
  Sequence number of next LTM request : 0
  1DMs sent                  : 5
  Valid 1DMs received        : 0
  Invalid 1DMs received      : 0
  DMMs sent                  : 5
  DMRs sent                  : 0
  Valid DMRs received        : 5
  Invalid DMRs received      : 0
  SLM sent                   : 10
  Valid SLM received         : 20
  Invalid SLM received       : 0
  SLR sent                   : 20
  Valid SLR received         : 10
  Invalid SLR received       : 0

Remote MEP identifier: 101
Remote MAC address: 00:05:85:73:39:4a
Delay measurement statistics:
  Index  One-way delay  Two-way delay
         (usec)        (usec)
  1      259           519
  2      273           550
  3      287           571
  4      299           610

```

```

      5      313      650
Average one-way delay      : 286 usec
Average one-way delay variation: 62 usec
Best case one-way delay    : 259 usec
Worst case one-way delay   : 313 usec
Average two-way delay      : 580 usec
Average two-way delay variation: 26 usec
Best case two-way delay    : 519 usec
Worst case two-way delay   : 650 usec
Synthetic Loss measurement
statistics:
  SLM packets sent          : 100
  SLM packets received      : 0
  SLR packets sent          : 100
  SLR packets received      : 0
  Accumulated SLM statistics:
  Local TXFC1 value         : 100
  Local RXFC1 value         : 100
  Last Received SLR frame TXFCftc : 100
  Last Received SLR frame TXFCbtc : 100
  SLM Frame Loss:
  Frame Loss (far-end)      : 0 (0.00 %)
  Frame Loss (near-end)     : 0 (0.00 %)
```

## show oam ethernet connectivity-fault-management path-database

<b>Syntax</b>	show oam ethernet connectivity-fault-management path-database <host-mac-address> <maintenance-association <i>ma-name</i> > <maintenance-domain <i>domain-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 8.4.
<b>Description</b>	On M7i and M10i with Enhanced CFEB (CFEB-E), M320, MX Series, ACX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management path database information for a host configured with an MEP.
<b>Options</b>	<p><b>host-mac-address</b>—(Optional) Display connectivity fault management path database information for a specified Ethernet host.</p> <p><b>maintenance-association <i>ma-name</i></b>—(Optional) Display connectivity fault management path database information for the specified maintenance association.</p> <p><b>maintenance-domain <i>domain-name</i></b>—(Optional) Display connectivity fault management path database information for the specified maintenance domain.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show oam ethernet connectivity-fault-management path-database on page 257</a>
<b>Output Fields</b>	<a href="#">Table 29 on page 256</a> lists the output fields for the <b>show oam ethernet connectivity-fault-management path-database</b> command. Output fields are listed in the approximate order in which they appear.

**Table 29: show oam ethernet connectivity-fault-management path-database Output Fields**

Field Name	Field Description
Linktrace to	MAC address of the remote MEPs in the path.
Interface	Interface identifier.
Maintenance domain name	Maintenance domain name.
Format (Maintenance domain)	Maintenance domain name format configured.
Level	Maintenance domain level configured.
Maintenance association name	Maintenance association name.

Table 29: show oam ethernet connectivity-fault-management path-database Output Fields (*continued*)

Field Name	Field Description
Local Mep	Local MEP identifier.

### Sample Output

show oam ethernet  
connectivity-fault-  
management  
path-database

```
user@host> show oam ethernet connectivity-fault-management path-database
maintenance-domain md1 maintenance-association ma1 00:05:85:79:39:ef
Linktrace to 00:05:85:79:39:ef, Interface : ge-3/0/0
  Maintenance Domain: md1, Level: 7
  Maintenance Association: ma1, Local Mep: 201
```

## show oam ethernet evc

<b>Syntax</b>	<b>show oam ethernet evc &lt;evc-id&gt;</b>
<b>Release Information</b>	Command introduced in Junos OS Release 9.5.
<b>Description</b>	On MX Series routers with OAM Ethernet Virtual Connection (EVC) configurations, displays the EVC configuration and status information.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	View
<b>Output Fields</b>	<a href="#">Table 30 on page 258</a> lists the output fields for the <b>show oam ethernet evc</b> command. Output fields are listed in the approximate order in which they appear.

**Table 30: show oam ethernet evc Output Fields**

Field Name	Field Description
<b>EVC identifier</b>	Header for the EVC information showing the EVC name, configuration, and active/inactive status.
<b>UNI count</b>	Number of configured and active UNIs.
<b>Protocol</b>	Protocol configured between the UNIs.
<b>Local UNIs</b>	Heading for the list of local UNIs
<b>UNI Identifier</b>	Name of the UNI.
<b>Interface</b>	Interface type-dpc/pic/port.unit-number.
<b>Status</b>	Status operational or not operational.

## Sample Output

### show oam ethernet evc

```

user@host> show oam ethernet evc
EVC identifier: evc1, Point-to-Point, Active
UNI count: Configured(2), Active(2)
Protocol: cfm, Management domain: md, Management association: ma
Local UNIs:
  UNI Identifier      Interface      Status
  uni1                ge-1/1/1      Operational
  uni2                ge-1/1/1      Not Operational

```

## show oam ethernet link-fault-management

<b>Syntax</b>	show oam ethernet link-fault-management <brief   detail> <interface-name>
<b>Release Information</b>	Command introduced in Junos OS Release 8.2.
<b>Description</b>	On EX Series switches and M320, M120, MX Series, T320, and T640 routers, display Operation, Administration, and Management (OAM) link fault management information for Ethernet interfaces.
<b>Options</b>	<b>brief   detail</b> —(Optional) Display the specified level of output.  <b>interface-name</b> —(Optional) Display link fault management information for the specified Ethernet interface only.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show oam ethernet link-fault-management brief on page 263</a> <a href="#">show oam ethernet link-fault-management detail on page 263</a>
<b>Output Fields</b>	<a href="#">Table 31 on page 259</a> lists the output fields for the <b>show oam ethernet link-fault-management</b> command. Output fields are listed in the approximate order in which they appear.

**Table 31: show oam ethernet link-fault-management Output Fields**

Field Name	Field Description	Level of Output
<b>Status</b>	Indicates the status of the established link.  <ul style="list-style-type: none"> <li>• <b>Fail</b>—A link fault condition exists.</li> <li>• <b>Running</b>—A link fault condition does not exist.</li> <li>• <b>ISSU</b>—The local end is in ISSU.</li> </ul>	All levels
<b>Discovery state</b>	State of the discovery mechanism:  <ul style="list-style-type: none"> <li>• <b>Passive Wait</b></li> <li>• <b>Send Any</b></li> <li>• <b>Send Local Remote</b></li> <li>• <b>Send Local Remote Ok</b></li> <li>• <b>Fault</b></li> </ul>	All levels
<b>Peer address</b>	Address of the OAM peer.	All levels

Table 31: show oam ethernet link-fault-management Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Flags</b>	<p>Information about the interface. Possible values are described in the “Link Flags” section under <i>Common Output Fields Description</i>.</p> <ul style="list-style-type: none"> <li>• <b>Remote-Stable</b>—Indicates remote OAM client acknowledgment of and satisfaction with local OAM state information. <b>False</b> indicates that remote DTE either has not seen or is unsatisfied with local state information. <b>True</b> indicates that remote DTE has seen and is satisfied with local state information.</li> <li>• <b>Local-Stable</b>—Indicates local OAM client acknowledgment of and satisfaction with remote OAM state information. <b>False</b> indicates that local DTE either has not seen or is unsatisfied with remote state information. <b>True</b> indicates that local DTE has seen and is satisfied with remote state information.</li> <li>• <b>Remote-State-Valid</b>—Indicates the OAM client has received remote state information found within Local Information TLVs of received Information OAM PDUs. <b>False</b> indicates that OAM client has not seen remote state information. <b>True</b> indicates that the OAM client has seen remote state information.</li> </ul>	All levels
<b>Remote loopback status</b>	Indicates the remote loopback status. An OAM entity can put its remote peer into loopback mode using the Loopback control OAM PDU. In loopback mode, every frame received is transmitted back on the same port (except for OAM PDUs, which are needed to maintain the OAM session).	All levels
<b>Remote entity information</b>	<p>Remote entity information.</p> <ul style="list-style-type: none"> <li>• <b>Remote MUX action</b>—Indicates the state of the multiplexer functions of the OAM sublayer. Device is forwarding non-OAM PDUs to the lower sublayer or discarding non-OAM PDUs.</li> <li>• <b>Remote parser action</b>—Indicates the state of the parser function of the OAM sublayer. Device is forwarding non-OAM PDUs to higher sublayer, looping back non-OAM PDUs to the lower sublayer, or discarding non-OAM PDUs.</li> <li>• <b>Discovery mode</b>—Indicates whether discovery mode is active or inactive.</li> <li>• <b>Unidirectional mode</b>—Indicates the ability to operate a link in a unidirectional mode for diagnostic purposes.</li> <li>• <b>Remote loopback mode</b>—Indicates whether remote loopback is supported or unsupported.</li> <li>• <b>Link events</b>—Indicates whether interpreting link events is supported or unsupported on the remote peer.</li> <li>• <b>Variable requests</b>—Indicates whether variable requests are supported. The Variable Request OAM PDU, is used to request one or more MIB variables from the remote peer. Also indicates if the remote end is in ISSU.</li> </ul>	All levels
<b>OAM Receive Statistics</b>		
<b>Information</b>	The total number of information PDUs received.	<b>detail</b>
<b>Event</b>	The total number of loopback control PDUs received.	<b>detail</b>
<b>Variable request</b>	The total number of variable request PDUs received.	<b>detail</b>
<b>Variable response</b>	The total number of variable response PDUs received.	<b>detail</b>

Table 31: show oam ethernet link-fault-management Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Loopback control</b>	The total number of loopback control PDUs received.	<b>detail</b>
<b>Organization specific</b>	The total number of vendor organization specific PDUs received.	<b>detail</b>
<b>OAM Transmit Statistics</b>		
<b>Information</b>	The total number of information PDUs transmitted.	<b>detail</b>
<b>Event</b>	The total number of event notification PDUs transmitted.	<b>detail</b>
<b>Variable request</b>	The total number of variable request PDUs transmitted.	<b>detail</b>
<b>Variable response</b>	The total number of variable response PDUs transmitted.	<b>detail</b>
<b>Loopback control</b>	The total number of loopback control PDUs transmitted.	<b>detail</b>
<b>Organization specific</b>	The total number of vendor organization specific PDUs transmitted.	<b>detail</b>
<b>OAM Received Symbol Error Event information</b>		
<b>Events</b>	The number of symbol error event TLVs that have been received since the OAM sublayer was reset.	<b>detail</b>
<b>Window</b>	The symbol error event window in the received PDU.  The protocol default value is the number of symbols that can be received in one second on the underlying physical layer.	<b>detail</b>
<b>Threshold</b>	The number of errored symbols in the period required for the event to be generated.	<b>detail</b>
<b>Errors in period</b>	The number of symbol errors in the period reported in the received event PDU.	<b>detail</b>
<b>Total errors</b>	The number of errored symbols that have been reported in received event TLVs since the OAM sublayer was reset.  Symbol errors are coding symbol errors.	<b>detail</b>
<b>OAM Received Frame Error Event Information</b>		
<b>Events</b>	The number of errored frame event TLVs that have been received since the OAM sublayer was reset.	<b>detail</b>
<b>Window</b>	The duration of the window in terms of the number of 100 ms period intervals.	<b>detail</b>
<b>Threshold</b>	The number of detected errored frames required for the event to be generated.	<b>detail</b>
<b>Errors in period</b>	The number of detected errored frames in the period.	<b>detail</b>

Table 31: show oam ethernet link-fault-management Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Total errors</b>	The number of errored frames that have been reported in received event TLVs since the OAM sublayer was reset.  A frame error is any frame error on the underlying physical layer.	<b>detail</b>
<b>OAM Received Frame Period Error Event Information</b>		
<b>Events</b>	The number of frame seconds errors event TLVs that have been received since the OAM sublayer was reset.	<b>detail</b>
<b>Window</b>	The duration of the frame seconds window.	<b>detail</b>
<b>Threshold</b>	The number of frame seconds errors in the period.	<b>detail</b>
<b>Errors in period</b>	The number of frame seconds errors in the period.	<b>detail</b>
<b>Total errors</b>	The number of frame seconds errors that have been reported in received event TLVs since the OAM sublayer was reset.	<b>detail</b>
<b>OAM Transmitted Symbol Error Event Information</b>		
<b>Events</b>	The number of symbol error event TLVs that have been transmitted since the OAM sublayer was reset.	<b>detail</b>
<b>Window</b>	The symbol error event window in the transmitted PDU.	<b>detail</b>
<b>Threshold</b>	The number of errored symbols in the period required for the event to be generated.	<b>detail</b>
<b>Errors in period</b>	The number of symbol errors in the period reported in the transmitted event PDU.	<b>detail</b>
<b>Total errors</b>	The number of errored symbols reported in event TLVs that have been transmitted since the OAM sublayer was reset.	<b>detail</b>
<b>OAM Current Symbol Error Event Information</b>		
<b>Events</b>	The number of symbol error TLVs that have been generated regardless of whether the threshold for sending event TLVs has been crossed.	<b>detail</b>
<b>Window</b>	The symbol error event window in the transmitted PDU.	<b>detail</b>
<b>Threshold</b>	The number of errored symbols in the period required for the event to be generated.	<b>detail</b>
<b>Errors in period</b>	The total number of symbol errors in the period reported.	<b>detail</b>
<b>Total errors</b>	The number of errored symbols reported in event TLVs that have been generated regardless of whether the threshold for sending event TLVs has been crossed.	<b>detail</b>
<b>OAM Transmitted Frame Error Event Information</b>		

Table 31: show oam ethernet link-fault-management Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Events</b>	The number of errored frame event TLVs that have been transmitted since the OAM sublayer was reset.	<b>detail</b>
<b>Window</b>	The duration of the window in terms of the number of 100 ms period intervals.	<b>detail</b>
<b>Threshold</b>	The number of detected errored frames required for the event to be generated.	<b>detail</b>
<b>Errors in period</b>	The number of detected errored frames in the period.	<b>detail</b>
<b>Total errors</b>	The number of errored frames that have been detected since the OAM sublayer was reset.	<b>detail</b>
<b>OAM Current Frame Error Event Information</b>		
<b>Events</b>	The number of errored frame event TLVs that have been generated regardless of whether the threshold for sending event TLVs has been crossed.	<b>detail</b>
<b>Window</b>	The duration of the window in terms of the number of 100 ms period intervals.	<b>detail</b>
<b>Threshold</b>	The number of detected errored frames required for the event to be generated.	<b>detail</b>
<b>Errors in period</b>	The number of errored frames in the period.	<b>detail</b>
<b>Total errors</b>	The number of errored frames detected regardless of whether the threshold for transmitting event TLVs has been crossed.	<b>detail</b>

## Sample Output

### show oam ethernet link-fault-management brief

```

user@host> show oam ethernet link-fault-management brief
Interface: ge-3/1/3
Status: Running, Discovery state: Send Any, ISSU
Peer address: 00:90:69:72:2c:83
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
Remote loopback status: Disabled on local port, Enabled on peer port
Remote entity information:
  Remote MUX action: discarding, Remote parser action: loopback
  Discovery mode: active, Unidirectional mode: unsupported
  Remote loopback mode: supported, Link events: supported
  Variable requests: unsupported, Remote in ISSU

```

### show oam ethernet link-fault-management detail

```

user@host> show oam ethernet link-fault-management detail
Interface: ge-6/1/0
Status: Running, Discovery state: Send Any, ISSU
Peer address: 00:90:69:0a:07:14
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
OAM receive statistics:
  Information: 186365, Event: 0, Variable request: 0, Variable response: 0
  Loopback control: 0, Organization specific: 0

```

OAM transmit statistics:  
Information: 186347, Event: 0, Variable request: 0, Variable response: 0  
Loopback control: 0, Organization specific: 0  
OAM received symbol error event information:  
Events: 0, Window: 0, Threshold: 0  
Errors in period: 0, Total errors: 0  
OAM received frame error event information:  
Events: 0, Window: 0, Threshold: 0  
Errors in period: 0, Total errors: 0  
OAM received frame period error event information:  
Events: 0, Window: 0, Threshold: 0  
Errors in period: 0, Total errors: 0  
OAM transmitted symbol error event information:  
Events: 0, Window: 0, Threshold: 1  
Errors in period: 0, Total errors: 0  
OAM current symbol error event information:  
Events: 0, Window: 0, Threshold: 1  
Errors in period: 0, Total errors: 0  
OAM transmitted frame error event information:  
Events: 0, Window: 0, Threshold: 1  
Errors in period: 0, Total errors: 0  
OAM current frame error event information:  
Events: 0, Window: 0, Threshold: 1  
Errors in period: 0, Total errors: 0  
Remote entity information:  
Remote MUX action: forwarding, Remote parser action: forwarding  
Discovery mode: active, Unidirectional mode: unsupported  
Remote loopback mode: supported, Link events: supported  
Variable requests: unsupported, Remote in ISSU

## show oam ethernet lmi

<b>Syntax</b>	<b>show oam ethernet lmi</b> ( <i>interface &lt;interface-name&gt;</i> )
<b>Release Information</b>	Command introduced in Junos OS Release 9.5.
<b>Description</b>	On MX Series routers with Gigabit Ethernet, Fast Ethernet, or aggregated Ethernet, and OAM Ethernet Local Management Interface (LMI) configuration, display the LMI information for the configured interfaces or optionally for a specified interface.
<b>Options</b>	<p><b>interface</b>—(Optional) Display LMI information for a specified interface.</p> <p><b>interface-name</b>—(Optional) Display Ethernet LMI information for the specified interface only.</p>
<b>Required Privilege Level</b>	View
<b>Output Fields</b>	Table 32 on page 265 lists the output fields for the <b>show oam ethernet lmi</b> command. Output fields are listed in the approximate order in which they appear.

**Table 32: show oam ethernet lmi Output Fields**

Field Name	Field Description
Physical Interface	Header for the EVC information showing the Ethernet virtual circuit (EVC) name, configuration, and active/inactive status.
UNI Identifier	Name of the UNI.
EVC map type	EVC configuration.
Polling verification timer	Polling verification timer status.
E-LMI state	Operational status of the E-LMI configuration in the interfaces or specified interface.
Priority/Untagged VLAN ID	To be provided.
Default EVC	The EVC set as the default EVC.
Associated EVCs	Heading for the list of configured EVCs.
EVC Identifier	EVC name.
Reference ID	To be provided.
Status	Status active or not active.
CE VLAN IDs	Customer edge VLAN ID numbers.

## Sample Output

### show oam ethernet lmi interface

```
user@host> show oam ethernet lmi interface ge-1/1/1
Physical interface: ge-1/1/1, Physical link is Up
UNI identifier: uni-ce1, EVC map type: Bundling
Polling verification timer: Enabled, E-LMI state: Operational
Priority/Untagged VLAN ID: 20, Default EVC: evc1
Associated EVCs:
  EVC          Reference      Status          CE VLAN IDs
  Identifier   ID
  evc1         1      Active (New)    1-2048
  evc2         2      Not Active      2049-4096
```

## show oam ethernet lmi statistics

<b>Syntax</b>	<code>show oam ethernet lmi statistics &lt;interface <i>interface-name</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.5.
<b>Description</b>	On MX Series routers with Gigabit Ethernet, Fast Ethernet, or aggregated Ethernet PICs, displays OAM Ethernet Local Management Interface (LMI) statistics.
<b>Options</b>	<p><b>interface</b>—(Optional) Display LMI statistics for a specified interface.</p> <p><b>interface-name</b>—(Optional) Display Ethernet LMI information for the specified Ethernet interface only.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show oam ethernet lmi statistics on page 267</a>
<b>Output Fields</b>	Table 33 on page 267 lists the output fields for the <code>show oam ethernet lmi statistics</code> command. Output fields are listed in the approximate order in which they appear.

**Table 33: show oam ethernet lmi statistics Output Fields**

Field Name	Field Description
Physical interface	Name of the interface for the displayed statistics.
Reliability errors	Number of E-LMI reliability errors logged.
Protocol errors	Number of E-LMI protocol errors.
Status check received	Number of E-LMI status check receive errors.
Status check sent	Number of E-LMI status check sent errors.
Full status received	Number of E-LMI full status receive errors.
Full status sent	Number of E-LMI full status sent errors.
Full status continued received	Number of E-LMI status continued received errors.
Full status continued sent	Number of E-LMI full status continued sent errors.
Asynchronous status sent	Number of E-LMI asynchronous status sent errors.

## Sample Output

### show oam ethernet lmi statistics

```
user@host> show oam ethernet lmi statistics interface ge-1/1/1
```

Physical interface: ge-1/1/1	
Reliability errors	4 Protocol errors
0	
Status check received	0 Status check sent
0	
Full status received	694 Full status sent
694	
Full status continued received	0 Full status continued sent
0	
Asynchronous status sent	0

## show protection-group ethernet-ring aps

<b>Syntax</b>	<b>show protection-group ethernet-ring aps</b>
<b>Release Information</b>	Command introduced in Junos OS Release 9.4. Command introduced in Junos OS Release 12.1 for EX Series switches.
<b>Description</b>	Display the status of the Automatic Protection Switching (APS) and Ring APS (RAPS) messages on an Ethernet ring.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show protection-group ethernet-ring data-channel on page 271</a></li> <li>• <a href="#">show protection-group ethernet-ring interface on page 273</a></li> <li>• <a href="#">show protection-group ethernet-ring node-state on page 276</a></li> <li>• <a href="#">show protection-group ethernet-ring statistics on page 279</a></li> <li>• <a href="#">show protection-group ethernet-ring vlan on page 282</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show protection-group ethernet-ring aps (EX Switches) on page 270</a> <a href="#">show protection-group ethernet-ring aps (Owner Node, Normal Operation on MX Routers) on page 270</a> <a href="#">show protection-group ethernet-ring aps (Ring Node, Normal Operation on MX Routers) on page 270</a> <a href="#">show protection-group ethernet-ring aps (Owner Node, Failure Condition on MX Routers) on page 270</a> <a href="#">show protection-group ethernet-ring aps (Ring Node, Failure Condition on MX Routers) on page 270</a>
<b>Output Fields</b>	<a href="#">Table 34 on page 269</a> lists the output fields for the <b>show protection-group ethernet-ring aps</b> command. Output fields are listed in the approximate order in which they appear.

**Table 34: show protection-group ethernet-ring aps Output Fields**

Field Name	Field Description
<b>Ethernet Ring Name</b>	Name configured for the Ethernet ring.
<b>Request/State</b>	Status of the Ethernet ring RAPS messages. <ul style="list-style-type: none"> <li>• <b>NR</b>—Indicates there is no request for APS on the ring.</li> <li>• <b>SF</b>—Indicates there is a signal failure on the ring.</li> </ul>
<b>No Flush</b>	State of the ring flushing: <b>No</b> (normal) or <b>Yes</b> (failure).
<b>Ring Protection Link Blocked</b>	Blocking on the ring protection link: <b>Yes</b> or <b>No</b> .

Table 34: show protection-group ethernet-ring aps Output Fields (*continued*)

Field Name	Field Description
Originator	Whether this node is the ring originator: <b>Yes</b> or <b>No</b> .
Remote Node ID	Identifier (in MAC address format) of the remote node.

## Sample Output

### show protection-group ethernet-ring aps (EX Switches)

```
user@switch>> show protection-group ethernet-ring aps
Ring Name    Request/state  No Flush  RPL Blocked  Originator  Remote Node ID
erp1         NR             no        yes          no          00:1F:12:30:B8:81
```

## Sample Output

### show protection-group ethernet-ring aps (Owner Node, Normal Operation on MX Routers)

```
user@host> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg101              NR             No        Yes

Originator  Remote Node ID
Yes
```

### show protection-group ethernet-ring aps (Ring Node, Normal Operation on MX Routers)

```
user@host> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg102              NR             No        Yes

Originator  Remote Node ID
No          00:01:01:00:00:01
```

### show protection-group ethernet-ring aps (Owner Node, Failure Condition on MX Routers)

```
user@host> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg101              SF             No        No

Originator  Remote Node ID
No          00:01:02:00:00:01
```

### show protection-group ethernet-ring aps (Ring Node, Failure Condition on MX Routers)

```
user@host> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg102              SF             No        Yes

Originator  Remote Node ID
Yes         00:00:00:00:00:00
```

## show protection-group ethernet-ring data-channel

<b>Syntax</b>	show protection-group ethernet-ring data-channel <brief   detail> <group-name <i>group-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 10.2.
<b>Description</b>	On MX Series routers, display data channel information for all Ethernet ring protection groups or for a specific Ethernet ring protection group.
<b>Options</b>	<b>brief   detail</b> —(Optional) Display the specified level of output.  <b>group-name</b> —(Optional) Protection group for which to display statistics. If you omit this optional field, all protection group statistics for configured groups will be displayed.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show protection-group ethernet-ring aps on page 269</a></li> <li>• <a href="#">show protection-group ethernet-ring interface on page 273</a></li> <li>• <a href="#">show protection-group ethernet-ring node-state on page 276</a></li> <li>• <a href="#">show protection-group ethernet-ring statistics on page 279</a></li> <li>• <a href="#">show protection-group ethernet-ring vlan on page 282</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show protection-group ethernet-ring data-channel on page 272</a>
<b>Output Fields</b>	<a href="#">Table 35 on page 271</a> lists the output fields for the <b>show protection-group ethernet-ring data-channel</b> command. Output fields are listed in the approximate order in which they appear.

**Table 35: show protection-group ethernet-ring data-channel Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of the interface configured for the Ethernet ring.
<b>STP index</b>	The Spanning Tree Protocol (STP) index number used by each interface in an Ethernet ring. The STP index controls the forwarding behavior for a set of VLANs on a data channel on an Ethernet ring port. For multiple Ethernet ring instances on a physical ring port, there are multiple STP index numbers. Different ring instances will have different STP index numbers and may have different forwarding behavior.
<b>Forward State</b>	Forwarding state on the Ethernet ring. <ul style="list-style-type: none"> <li>• <b>fowarding</b>—Indicates packets are being forwarded.</li> <li>• <b>discarding</b>—Indicates packets are being discarded.</li> </ul>

## Sample Output

### show protection-group ethernet-ring data-channel

```
user@host> show protection-group ethernet-ring data-channel
Ethernet ring data channel information for protection group pg301
Interface    STP index    Forward State
ge-1/0/3     71           forwarding
ge-1/0/4     82           forwarding

Ethernet ring data channel information for protection group pg302
Interface    STP index    Forward State
ge-1/0/3     52           forwarding
ge-1/0/4     91           forwarding
```

## show protection-group ethernet-ring interface

<b>Syntax</b>	<b>show protection-group ethernet-ring interface</b>
<b>Release Information</b>	Command introduced in Junos OS Release 9.4.
<b>Description</b>	Displays the status of the Automatic Protection Switching (APS) interfaces on an Ethernet ring.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show protection-group ethernet-ring data-channel on page 271</a></li> <li>• <a href="#">show protection-group ethernet-ring aps on page 269</a></li> <li>• <a href="#">show protection-group ethernet-ring node-state on page 276</a></li> <li>• <a href="#">show protection-group ethernet-ring statistics on page 279</a></li> <li>• <a href="#">show protection-group ethernet-ring vlan on page 282</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show protection-group ethernet-ring interface (EX Series Switch Owner Node) on page 274</a> <a href="#">show protection-group ethernet-ring interface (Owner Node MX Series Router ) on page 274</a> <a href="#">show protection-group ethernet-ring interface (EX Series Switch Ring Node) on page 274</a> <a href="#">show protection-group ethernet-ring interface (MX Series Router Ring Node) on page 274</a>
<b>Output Fields</b>	Table 36 on page 273 lists the output fields for both the EX Series switch and the MX Series router <b>show protection-group ethernet-ring interface</b> commands. Output fields are listed in the approximate order in which they appear.

**Table 36: MX Series Routers show protection-group ethernet-ring interface Output Fields**

Field Name	Field Description
Ethernet ring port parameters for protection group <i>group-name</i>	Output is organized by configured protection group.
Interface	Physical interfaces configured for the Ethernet ring.
Control Channel	(MX Series router only) Logical unit configured on the physical interface. <ul style="list-style-type: none"> <li>• <b>NR</b>—Indicates there is no request for APS on the ring.</li> <li>• <b>SF</b>—Indicates there is a signal failure on the ring.</li> </ul>
Forward State	State of the ring forwarding on the interface: <b>discarding</b> or <b>forwarding</b> .

**Table 36: MX Series Routers show protection-group ethernet-ring interface Output Fields (*continued*)**

Field Name	Field Description
Ring Protection Link End	Whether this interface is the end of the ring: <b>Yes</b> or <b>No</b> .
Signal Failure	Whether there a signal failure exists on the link: <b>Clear</b> or <b>Set</b> .
Admin State	State of the interface: For EX switches, <b>ready</b> , <b>ifl ready</b> , or <b>waiting</b> . For MX routers, <b>IFF ready</b> or <b>IFF disabled</b> .

## Sample Output

### show protection-group ethernet-ring interface (EX Series Switch Owner Node)

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101
```

Interface	Forward State	RPL End	Signal Failure	Admin State
ge-0/0/3.0	discarding	Yes	Clear	ready
ge-0/0/9.0	forwarding	No	Clear	ready

### show protection-group ethernet-ring interface (Owner Node MX Series Router )

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101
```

Interface	Control Channel	Forward State	Ring Protection Link End
ge-1/0/1	ge-1/0/1.1	discarding	Yes
ge-1/2/4	ge-1/2/4.1	forwarding	No

Signal Failure	Admin State
Clear	IFF ready
Clear	IFF ready

### show protection-group ethernet-ring interface (EX Series Switch Ring Node)

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102
```

Ethernet ring port parameters for protection group pg101

Interface	Forward State	RPL End	Signal Failure	Admin State
ge-0/0/3.0	discarding	Yes	Clear	ready
ge-0/0/9.0	forwarding	No	Clear	ready

### show protection-group ethernet-ring interface (MX Series Router Ring Node)

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102
```

Interface	Control Channel	Forward State	Ring Protection Link End
ge-1/2/1	ge-1/2/1.1	forwarding	No
ge-1/0/2	ge-1/0/2.1	forwarding	No

Signal	Failure	Admin	State
Clear		IFF	ready
Clear		IFF	ready

## show protection-group ethernet-ring node-state

<b>Syntax</b>	<b>show protection-group ethernet-ring node-state</b>
<b>Release Information</b>	Command introduced in Junos OS Release 9.4. Command introduced in Junos OS Release 12.1 for EX Series switches.
<b>Description</b>	Display the status of the Automatic Protection Switching (APS) nodes on an Ethernet ring.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show protection-group ethernet-ring data-channel on page 271</a></li> <li>• <a href="#">show protection-group ethernet-ring aps on page 269</a></li> <li>• <a href="#">show protection-group ethernet-ring interface on page 273</a></li> <li>• <a href="#">show protection-group ethernet-ring statistics on page 279</a></li> <li>• <a href="#">show protection-group ethernet-ring vlan on page 282</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show protection-group ethernet-ring node-state (EX Series Switch) on page 277</a> <a href="#">show protection-group ethernet-ring node-state (MX Series Router - Owner Node, Normal Operation ) on page 277</a> <a href="#">show protection-group ethernet-ring node-state (MX Series Router - Ring Node, Normal Operation) on page 277</a> <a href="#">show protection-group ethernet-ring node-state (MX Series Router - Owner Node, Remote Signal Failure Condition) on page 278</a> <a href="#">show protection-group ethernet-ring node-state (MX Series Router - Ring Node, Local Signal Failure Condition) on page 278</a> <a href="#">show protection-group ethernet-ring node-state detail (MX Series Router - Node state at RPL-owner after signal failure condition is cleared in the ring and before reversion) on page 278</a>
<b>Output Fields</b>	<a href="#">Table 37 on page 276</a> lists the output fields for the <b>show protection-group ethernet-ring node-state</b> command. Output fields are listed in the approximate order in which they appear.

**Table 37: show protection-group ethernet-ring node-state Output Fields**

Field Name	Field Description
Ring Name/Ethernet Ring	Name configured for the Ethernet ring.

Table 37: show protection-group ethernet-ring node-state Output Fields (*continued*)

Field Name	Field Description
APS State	State of the Ethernet ring APS. <ul style="list-style-type: none"> <li><b>idle</b>—Indicates that the ring is working in normal condition and no protection-switching request active or pending in the ring. When the ring is in idle state, it is blocked at RPL link.</li> <li><b>protected</b>—Indicates that there is a protection switch on the ring due to signal failure condition on the ring link.</li> </ul>
Event	Events on the ring. <ul style="list-style-type: none"> <li><b>NR-RB</b>—Indicates there is no APS request and the ring link is blocked on the ring owner node.</li> <li><b>NR</b>—Indicates there is no APS request pending in the ring.</li> <li><b>local SF</b>—Indicates there is signal failure on one or both the ring links of the node.</li> <li><b>remote SF</b>—Indicates there is signal failure on ring links of any other node of the ring.</li> <li><b>WTR running</b>—Indicates wait to restore timer is running on RPL-owner.</li> </ul>
RPL Owner / Ring Protection Link Owner	Whether this node is the ring owner: <b>Yes</b> or <b>No</b> .
WTR Timer / Restore Timer	Restoration timer: <b>running</b> or <b>disabled</b> .
Guard Timer	Guard timer: <b>running</b> or <b>disabled</b> .
Op state / Operational State	State of the node: <b>Operational</b> or <b>any internal wait state</b> .

## Sample Output

### show protection-group ethernet-ring node-state (EX Series Switch)

```

user@switch> show protection-group ethernet-ring node-state
Ring Name APS State Event RPL Owner WTR Timer Guard Timer Op State
erp1      idle      NR-RB  yes    disabled disabled operational

```

### show protection-group ethernet-ring node-state (MX Series Router - Owner Node, Normal Operation )

```

user@host> show protection-group ethernet-ring node-state
Ethernet ring   APS State   Event       Ring Protection Link Owner
pg101          idle        NR-RB       Yes

Restore Timer  Guard Timer  Operation state
disabled       disabled    operational

```

### show protection-group ethernet-ring node-state (MX Series Router - Ring Node, Normal Operation)

```

user@host> show protection-group ethernet-ring node-state

```

Ethernet ring	APS State	Event	Ring Protection Link Owner
pg102	idle	NR-RB	No

Restore Timer	Guard Timer	Operation state
disabled	disabled	operational

**show protection-group ethernet-ring node-state (MX Series Router - Owner Node, Remote Signal Failure Condition)**

```
user@host> show protection-group ethernet-ring node-state
```

Ethernet ring	APS State	Event	Ring Protection Link Owner
pg101	protected	remote SF	Yes

Restore Timer	Guard Timer	Operation state
disabled	disabled	operational

**show protection-group ethernet-ring node-state (MX Series Router - Ring Node, Local Signal Failure Condition)**

```
user@host> show protection-group ethernet-ring node-state
```

Ethernet ring	APS State	Event	Ring Protection Link Owner
pg102	protected	local SF	No

Restore Timer	Guard Timer	Operation state
disabled	disabled	operational

**show protection-group ethernet-ring node-state detail (MX Series Router - Node state at RPL-owner after signal failure condition is cleared in the ring and before reversion)**

```
user@host> show protection-group ethernet-ring node-state detail
```

Ethernet-Ring name	: pg_major
APS State	: protected
Event	: WTR running
Ring Protection Link Owner	: Yes
Restore Timer	: running (time to expire: 269 sec)
Guard Timer	: disabled
Operation state	: operational

## show protection-group ethernet-ring statistics

<b>Syntax</b>	<b>show protection-group ethernet-ring statistics</b> <group-name <i>group-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 9.4. Command introduced in Junos OS Release 12.1 for EX Series switches.
<b>Description</b>	Display statistics regarding Automatic Protection Switching (APS) protection groups on an Ethernet ring.
<b>Options</b>	<b>group-name</b> —Protection group for which to display statistics. In you omit this optional field, all protection group statistics for configured groups will be displayed.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show protection-group ethernet-ring data-channel on page 271</a></li> <li>• <a href="#">show protection-group ethernet-ring aps on page 269</a></li> <li>• <a href="#">show protection-group ethernet-ring node-state on page 276</a></li> <li>• <a href="#">show protection-group ethernet-ring interface on page 273</a></li> <li>• <a href="#">show protection-group ethernet-ring vlan on page 282</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show protection-group ethernet-ring statistics (EX Switch) on page 280</a> <a href="#">show protection-group ethernet-ring statistics (Owner Node, Normal Operation on MX Router) on page 280</a> <a href="#">show protection-group ethernet-ring statistics (Ring Node, Normal Operation on MX Router) on page 280</a> <a href="#">show protection-group ethernet-ring statistics (Owner Node, Failure Condition on MX Router) on page 280</a> <a href="#">show protection-group ethernet-ring statistics (Ring Node, Failure Condition on MX Router) on page 281</a>
<b>Output Fields</b>	<a href="#">Table 38 on page 279</a> lists the output fields for the <b>show protection-group ethernet-ring statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 38: show protection-group ethernet-ring statistics Output Fields**

Field Name	Field Description
<b>Ethernet Ring Statistics for PG</b>	Name of the protection group for which statistics are displayed.
<b>RAPS sent</b>	Number of Ring Automatic Protection Switching (RAPS) messages sent. (On MX Series switches only)
<b>RAPS received</b>	Number of RAPS messages received. (On MX Series switches only)

Table 38: show protection-group ethernet-ring statistics Output Fields (*continued*)

Field Name	Field Description
<b>Local SF</b>	Number of times a signal failure (SF) has occurred locally.
<b>Remote SF</b>	Number of times a signal failure (SF) has occurred anywhere else on the ring.
<b>NR event</b>	Number of times a No Request (NR) event has occurred on the ring.
<b>NR-RB event</b>	Number of times a No Request, Ring Blocked (NR-RB) event has occurred on the ring.

## Sample Output

### show protection-group ethernet-ring statistics (EX Switch)

```

user@switch> show protection-group ethernet-ring statistics
Ring Name Local SF Remote SF NR Event NR-RB Event
erp1      2      1      2      3

```

### show protection-group ethernet-ring statistics (Owner Node, Normal Operation on MX Router)

```

user@host> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent                : 1
RAPS received             : 0
Local SF happened         : 0
Remote SF happened        : 0
NR event happened         : 0
NR-RB event happened      : 1

```

### show protection-group ethernet-ring statistics (Ring Node, Normal Operation on MX Router)

```

user@host> show protection-group ethernet-ring statistics group-name pg102
Ethernet Ring statistics for PG pg102
RAPS sent                : 0
RAPS received             : 1
Local SF happened         : 0
Remote SF happened        : 0
NR event happened         : 0
NR-RB event happened      : 1

```

### show protection-group ethernet-ring statistics (Owner Node, Failure Condition on MX Router)

```

user@host> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent                : 1
RAPS received             : 1
Local SF happened         : 0
Remote SF happened        : 1
NR event happened         : 0
NR-RB event happened      : 1

```

**show protection-group ethernet-ring statistics (Ring Node, Failure Condition on MX Router)**

```
user@host> show protection-group ethernet-ring statistics group-name pg102
Ethernet Ring statistics for PG pg102
RAPS sent                : 1
RAPS received            : 1
Local SF happened        : 1
Remote SF happened       : 0
NR event happened        : 0
NR-RB event happened     : 1
```

## show protection-group ethernet-ring vlan

<b>Syntax</b>	show protection-group ethernet-ring vlan <brief   detail> <group-name <i>group-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 10.2.
<b>Description</b>	On MX Series routers, display all data channel logical interfaces and the VLAN IDs controlled by a ring instance data channel.
<b>Options</b>	<b>brief   detail</b> —(Optional) Display the specified level of output.  <b>group-name</b> —(Optional) Protection group for which to display details such as data channel interfaces, vlan, and bridge-domain. If you omit this optional field, details for all configured protection groups will be displayed.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show protection-group ethernet-ring aps on page 269</a></li> <li>• <a href="#">show protection-group ethernet-ring data-channel on page 271</a></li> <li>• <a href="#">show protection-group ethernet-ring interface on page 273</a></li> <li>• <a href="#">show protection-group ethernet-ring node-state on page 276</a></li> <li>• <a href="#">show protection-group ethernet-ring statistics on page 279</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show protection-group ethernet-ring vlan on page 283</a> <a href="#">show protection-group ethernet-ring vlan brief on page 283</a> <a href="#">show protection-group ethernet-ring vlan detail on page 283</a> <a href="#">show protection-group ethernet-ring vlan group-name vkm01 on page 283</a>
<b>Output Fields</b>	Table 39 on page 282 lists the output fields for the <b>show protection-group ethernet-ring vlan</b> command. Output fields are listed in the approximate order in which they appear.

**Table 39: show protection-group ethernet-ring vlan Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of the interface configured for the Ethernet protection ring.
<b>Vlan</b>	Name of the VLAN associated with the interface configured for the Ethernet protection ring.
<b>STP index</b>	The Spanning Tree Protocol (STP) index number used by each interface in an Ethernet ring. The STP index controls the forwarding behavior for a set of VLANs on a data channel on an Ethernet ring port. For multiple Ethernet ring instances on a physical ring port, there are multiple STP index numbers. Different ring instances will have different STP index numbers and may have different forwarding behavior.

**Table 39: show protection-group ethernet-ring vlan Output Fields (*continued*)**

Field Name	Field Description
Bridge Domain	Name of the bridge domain that is associated with the VLAN configured for the Ethernet protection ring.

## Sample Output

### show protection-group ethernet-ring vlan

```
user@host> show protection-group ethernet-ring vlan
Ethernet ring IFBD parameters for protection group vkm01

Interface  Vlan    STP Index  Bridge Domain
ge-2/0/8   100     130       default-switch/bd100
ge-2/0/4   100     126       default-switch/bd100
```

### show protection-group ethernet-ring vlan brief

```
user@host> show protection-group ethernet-ring vlan brief
Ethernet ring IFBD parameters for protection group vkm01

Interface  Vlan    STP Index  Bridge Domain
ge-2/0/8   100     130       default-switch/bd100
ge-2/0/4   100     126       default-switch/bd100
```

### show protection-group ethernet-ring vlan detail

```
user@host> show protection-group ethernet-ring vlan detail
Ethernet ring IFBD parameters for protection group vkm01

Interface name      : ge-2/0/8
Vlan                 : 100
STP index            : 130
Bridge Domain        : default-switch/bd100

Interface name      : ge-2/0/4
Vlan                 : 100
STP index            : 126
Bridge Domain        : default-switch/bd100
```

### show protection-group ethernet-ring vlan group-name vkm01

```
user@host> show protection-group ethernet-ring vlan vkm01
Ethernet ring IFBD parameters for protection group vkm01

Interface  Vlan    STP Index  Bridge Domain
ge-2/0/8   100     130       default-switch/bd100
ge-2/0/4   100     126       default-switch/bd100
```



## PART 4

# Troubleshooting

- [Ethernet on page 287](#)
- [Interface Diagnostics on page 295](#)



## CHAPTER 6

# Ethernet

- [traceroute ethernet](#)
- [Tracing PPPoE Operations on page 290](#)
- [Troubleshooting PPPoE Service Name Tables on page 291](#)
- [Verifying a PPPoE Configuration on page 293](#)

## traceroute ethernet

<b>Syntax</b>	<b>traceroute ethernet</b> ( <i>mac-address</i>   <i>mep-id</i> ) <b>maintenance-association</b> <i>ma-name</i> <b>maintenance-domain</b> <i>md-name</i> <b>ttl</b> <i>value</i> < <b>wait seconds</b> >
<b>Release Information</b>	Command introduced in Junos OS Release 9.0. <b>mep-id</b> option introduced in Junos OS Release 9.1.
<b>Description</b>	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.  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.
<b>Options</b>	<b>mac-address</b> —Destination unicast MAC address of the remote maintenance point.  <b>mep-id</b> —MEP identifier of the remote maintenance point. The range of values is 1 through 8191.  <b>maintenance-association</b> <i>ma-name</i> —Specifies an existing maintenance association from the set of configured maintenance associations.  <b>maintenance-domain</b> <i>md-name</i> —Specifies an existing maintenance domain from the set of configured maintenance domains.  <b>ttl value</b> —Number of hops to use in the linktrace request. The range is 1 to 255 hops. The default is 4.  <b>wait seconds</b> —(Optional) Maximum time to wait for a response to the traceroute request. The range is 1 to 255 seconds. The default is 5.
<b>Required Privilege Level</b>	network
<b>List of Sample Output</b>	<a href="#">traceroute ethernet on page 289</a>
<b>Output Fields</b>	<a href="#">Table 40 on page 288</a> lists the output fields for the <b>traceroute ethernet</b> command. Output fields are listed in the approximate order in which they appear.

**Table 40: traceroute ethernet Output Fields**

Field Name	Field Description
Linktrace to	MAC address of the destination maintenance point.
Interface	Local interface used to send the linktrace message (LTM).

Table 40: traceroute ethernet Output Fields (*continued*)

Field Name	Field Description
<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.

## 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

## Tracing PPPoE Operations

---

The Junos OS trace feature tracks PPPoE operations and records events in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

By default, nothing is traced. When you enable the tracing operation, the default tracing behavior is as follows:

1. Important events are logged in a file called **pppoed** located in the **/var/log** directory. You cannot change the directory (**/var/log**) in which trace files are located.
2. When the file **pppoed** reaches 128 kilobytes (KB), it is renamed **pppoed.0**, then **pppoed.1**, and finally **pppoed.2**, until there are three trace files. Then the oldest trace file (**pppoed.2**) is overwritten.

You can optionally specify the number of trace files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB). (For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)

By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure PPPoE tracing operations:

1. Specify that you want to configure tracing options.  

```
[edit protocols pppoe]
user@host# edit traceoptions
```
2. (Optional) Configure the name for the file used for the trace output.
3. (Optional) Configure the number and size of the log files.
4. (Optional) Configure access to the log file.
5. (Optional) Configure a regular expression to filter logging events.
6. (Optional) Configure flags to filter the operations to be logged.

Optional PPPoE traceoptions operations are described in the following sections:

- [Configuring the PPPoE Trace Log Filename on page 290](#)
- [Configuring the Number and Size of PPPoE Log Files on page 291](#)
- [Configuring Access to the PPPoE Log File on page 291](#)
- [Configuring a Regular Expression for PPPoE Lines to Be Logged on page 291](#)
- [Configuring the PPPoE Tracing Flags on page 291](#)

### Configuring the PPPoE Trace Log Filename

By default, the name of the file that records trace output for PPPoE is **pppoed**. You can specify a different name with the **file** option.

## Configuring the Number and Size of PPPoE Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

The archived files are differentiated by a suffix in the format *.number.gz*. The newest archived file is *.0.gz* and the oldest archived file is *.(maximum number)-1.gz*. When the current trace log file reaches the maximum size, it is compressed and renamed, and any existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation, *filename*, reaches 2 MB, *filename* is compressed and renamed *filename.0.gz*, and a new file called *filename* is created. When the new *filename* reaches 2 MB, *filename.0.gz* is renamed *filename.1.gz* and *filename* is compressed and renamed *filename.0.gz*. This process repeats until there are 20 trace files. Then the oldest file, *filename.19.gz*, is simply overwritten when the next oldest file, *filename.18.gz* is compressed and renamed to *filename.19.gz*.

## Configuring Access to the PPPoE Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

## Configuring a Regular Expression for PPPoE Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

You can refine the output by including regular expressions to be matched.

## Configuring the PPPoE Tracing Flags

By default, no events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged, configure the flags:

- `[edit protocols pppoe traceoptions]`  
`user@host# set flag authentication`

- Related Documentation**
- [PPPoE Overview](#)
  - [Ethernet Interfaces](#)

## Troubleshooting PPPoE Service Name Tables

- Problem** A misconfiguration of a PPPoE service name table can prevent PPPoE services from being properly activated. Configuration options for PPPoE service name tables are simple, which should simplify discovering where a misconfiguration exists. PPPoE clients cannot

connect if the service name table contains no match for the service name tag carried in the PADI packet.

The symptom of a service name table misconfiguration is that the client connection process stops at the negotiation stage and the PADI packets are ignored. You can use the **show pppoe statistics** command to examine the PPPoE packet counts for a problem.

When the service name table is properly configured, packets sent and received increment symmetrically. The following sample output shows a PADO sent count equal to the PADI received count, and PADS sent count equal to the PADR received count. This output indicates that the PPPoE negotiation is proceeding successfully and that the service name table is not misconfigured.

```
user@host> show pppoe statistics ge-2/0/3.1
```

```
Active PPPoE sessions: 2
```

PacketType	Sent	Received
PADI	0	16
PADO	16	0
PADR	0	16
PADS	16	0
PADT	0	0
Service name error	0	0
AC system error	0	0
Generic error	0	0
Malformed packets	0	0
Unknown packets	0	0

When the service name table is misconfigured, the output of the **show pppoe statistics** command indicates that the number of PADI packets received on the underlying interface is increasing, but the number of PADO packets sent remains at zero. The following sample output shows a PADI count of 100 and a PADO count of 0.

```
user@host> show pppoe statistics ge-2/0/3.1
```

```
Active PPPoE sessions: 0
```

PacketType	Sent	Received
PADI	0	100
PADO	0	0
PADR	0	0
PADS	0	0
PADT	0	0
Service name error	0	0
AC system error	0	0
Generic error	0	0
Malformed packets	0	0
Unknown packets	0	0

When you believe a misconfiguration exists, use the **monitor traffic interface** command on the underlying interface to determine which service name is being requested by the PPPoE client. The following sample output shows that the client is requesting Service1 in the service name tag.

```
user@host> monitor traffic interface ge-2/0/3.1 print-hex print-ascii
Listening on ge-2/0/3.1, capture size 96 bytes
```

```
11:49:41.436682 In PPPoE PADI [Service-Name "Service1"] [Host-Uniq UTF8]
[Tag-0x120 UTF8] [Vendor-Specific UTF8]
```

```

0x0000  ffff ffff ffff 0090 1a42 0ac1 8100 029a  ....B.....
0x0010  8863 1109 0000 00c9 0101 0008 5365 7276  .c.....Serv
0x0020  6963 6531 0103 0004 1200 9c43 0120 0002  ice1.....C....
0x0030  044a 0105 00ab 0000 0de9 0124 783a 3132  .J.....$x:12
0x0040  3030 3963                                     009c

```

You can then use the **show pppoe service-name-tables** command to determine whether you have misspelled the name of the service or perhaps not configured the service at all.

**Cause** Typical misconfigurations appear in the service name table configurations.

**Solution** Use the appropriate statements to correct the misconfiguration.

**Related Documentation**

- *Configuring PPPoE Service Name Tables*
- *show pppoe service-name-tables*
- *show pppoe statistics*
- *show pppoe underlying-interfaces*
- *PPPoE Overview*
- *Ethernet Interfaces*

## Verifying a PPPoE Configuration

**Purpose** You can use show commands to display and verify the PPPoE configuration.

**Action** To verify a PPPoE configuration, you can issue the following operational mode commands:

- **show interfaces at-*fpc/pic/port* extensive**
- **show interfaces pp0**
- **show pppoe interfaces**
- **show pppoe version**
- **show pppoe service-name-tables**
- **show pppoe sessions**
- **show pppoe statistics**
- **show pppoe underlying-interfaces**

For more information about these operational mode commands, see [CLI Explorer](#).

**Related Documentation**

- *PPPoE Overview*
- *Ethernet Interfaces*



## CHAPTER 7

# Interface Diagnostics

- [Interface Diagnostics on page 295](#)

## Interface Diagnostics

---

You can use two diagnostic tools to test the physical layer connections of interfaces: loopback testing and bit error rate test (BERT) testing. Loopback testing enables you to verify the connectivity of a circuit. BERT testing enables you to identify poor signal quality on a circuit. This section contains the following topics:

- [Configuring Loopback Testing on page 295](#)
- [Interface Diagnostics on page 297](#)

## Configuring Loopback Testing

Loopback testing allows you to verify the connectivity of a circuit. You can configure any of the following interfaces to execute a loopback test: Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, E1, E3, NxDS0, serial, SONET/SDH, T1, and T3.

The physical path of a network data circuit usually consists of segments interconnected by devices that repeat and regenerate the transmission signal. The transmit path on one device connects to the receive path on the next device. If a circuit fault occurs in the form of a line break or a signal corruption, you can isolate the problem by using a loopback test. Loopback tests allow you to isolate segments of the circuit and test them separately.

To do this, configure a *line loopback* on one of the routers. Instead of transmitting the signal toward the far-end device, the line loopback sends the signal back to the originating router. If the originating router receives back its own data link layer packets, you have verified that the problem is beyond the originating router. Next, configure a line loopback farther away from the local router. If this originating router does not receive its own data link layer packets, you can assume the problem is on one of the segments between the local router and the remote router's interface card. In this case, the next troubleshooting step is to configure a line loopback closer to the local router to find the source of the problem.

There are several types of loopback testing supported by the Junos OS, as follows:

- DCE local—Loops packets back on the local DCE.
- DCE remote—Loops packets back on the remote DCE.

- **Local**—Useful for troubleshooting physical PIC errors. Configuring local loopback on an interface allows transmission of packets to the channel service unit (CSU) and then to the circuit toward the far-end device. The interface receives its own transmission, which includes data and timing information, on the local router's PIC. The data received from the CSU is ignored. To test a local loopback, issue the **show interfaces *interface-name*** command. If PPP keepalives transmitted on the interface are received by the PIC, the **Device Flags** field contains the output **Loop-Detected**.
- **Payload**—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A payload loopback loops data only (without clocking information) on the remote router's PIC. With payload loopback, overhead is recalculated.
- **Remote**—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A remote loopback loops packets, including both data and timing information, back on the remote router's interface card. A router at one end of the circuit initiates a remote loopback toward its remote partner. When you configure a remote loopback, the packets received from the physical circuit and CSU are received by the interface. Those packets are then retransmitted by the PIC back toward the CSU and the circuit. This loopback tests all the intermediate transmission segments.

Table 41 on page 296 shows the loopback modes supported on the various interface types.

**Table 41: Loopback Modes by Interface Type**

Interface	Loopback Modes	Usage Guidelines
Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet	Local	<i>Configuring Ethernet Loopback Capability</i>
Circuit Emulation E1	Local and remote	<i>Configuring E1 Loopback Capability</i>
Circuit Emulation T1	Local and remote	<i>Configuring T1 Loopback Capability</i>
E1 and E3	Local and remote	<i>Configuring E1 Loopback Capability and Configuring E3 Loopback Capability</i>
NxDSO	Payload	<i>Configuring Channelized E1 IQ and IQE Interfaces, Configuring T1 and NxDSO Interfaces, Configuring Channelized OC12/STM4 IQ and IQE Interfaces (SONET Mode), Configuring Channelized STM1 IQ and IQE Interfaces, and Configuring Channelized T3 IQ Interfaces</i>
Serial (V.35 and X.21)	Local and remote	<i>Configuring Serial Loopback Capability</i>
Serial (EIA-530)	DCE local, DCE remote, local, and remote	<i>Configuring Serial Loopback Capability</i>
SONET/SDH	Local and remote	<i>Configuring SONET/SDH Loopback Capability</i>

Table 41: Loopback Modes by Interface Type (*continued*)

Interface	Loopback Modes	Usage Guidelines
T1 and T3	Local, payload, and remote	<i>Configuring T1 Loopback Capability and Configuring T3 Loopback Capability</i>  <i>See also Configuring the T1 Remote Loopback Response</i>

To configure loopback testing, include the **loopback** statement:

**loopback mode;**

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* aggregated-ether-options]
- [edit interfaces *interface-name* ds0-options]
- [edit interfaces *interface-name* e1-options]
- [edit interfaces *interface-name* e3-options]
- [edit interfaces *interface-name* fastether-options]
- [edit interfaces *interface-name* gigether-options]
- [edit interfaces *interface-name* serial-options]
- [edit interfaces *interface-name* sonet-options]
- [edit interfaces *interface-name* t1-options]
- [edit interfaces *interface-name* t3-options]

## Interface Diagnostics

BERT allows you to troubleshoot problems by checking the quality of links. You can configure any of the following interfaces to execute a BERT when the interface receives a request to run this test: E1, E3, T1, T3; the channelized DS3, OC3, OC12, and STM1 interfaces; and the channelized DS3 IQ, E1 IQ, and OC12 IQ interfaces.

A BERT test requires a line loop to be in place on either the transmission devices or the far-end router. The local router generates a known bit pattern and sends it out the transmit path. The received pattern is then verified against the sent pattern. The higher the bit error rate of the received pattern, the worse the noise is on the physical circuit. As you move the position of the line loop increasingly downstream toward the far-end router, you can isolate the troubled portion of the link.

To configure BERT, you must configure the duration of the test, the bit pattern to send on the transmit path, and the error rate to monitor when the inbound pattern is received.

To configure the duration of the test, the pattern to send in the bit stream, and the error rate to include in the bit stream, include the **bert-period**, **bert-algorithm**, and **bert-error-rate** statements, respectively, at the [edit interfaces *interface-name* *interface-type*-options] hierarchy level:

```
[edit interfaces interface-name interface-type-options]
bert-algorithm algorithm;
bert-error-rate rate;
bert-period seconds;
```

By default, the BERT period is 10 seconds. You can configure the BERT period to last from 1 through 239 seconds on some PICs and from 1 through 240 seconds on other PICs.

**rate** is the bit error rate. This can be an integer from 0 through 7, which corresponds to a bit error rate from  $10^{-0}$  (1 error per bit) to  $10^{-7}$  (1 error per 10 million bits).

**algorithm** is the pattern to send in the bit stream. For a list of supported algorithms, enter a ? after the **bert-algorithm** statement; for example:

```
[edit interfaces t1-0/0/0 t1-options]
user@host# set bert-algorithm ?
Possible completions:
pseudo-2e11-o152      Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151      Pattern is 2^15 - 1 (per 0.152 standard)
pseudo-2e20-o151      Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e20-o153      Pattern is 2^20 - 1 (per 0.153 standard)
...
```

For specific hierarchy information, see the individual interface types.



**NOTE:** The 4-port E1 PIC supports only the following algorithms:

pseudo-2e11-o152	Pattern is $2^{11} - 1$ (per 0.152 standard)
pseudo-2e15-o151	Pattern is $2^{15} - 1$ (per 0.151 standard)
pseudo-2e20-o151	Pattern is $2^{20} - 1$ (per 0.151 standard)
pseudo-2e23-o151	Pattern is $2^{23}$ (per 0.151 standard)

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.



**NOTE:** The 12-port T1/E1 Circuit Emulation (CE) PIC supports only the following algorithms:

```
all-ones-repeating    Repeating one bits
all-zeros-repeating   Repeating zero bits
alternating-double-ones-zeros Alternating pairs of ones and zeros
alternating-ones-zeros Alternating ones and zeros
pseudo-2e11-o152      Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151      Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151      Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e7            Pattern is 2^7 - 1
pseudo-2e9-o153       Pattern is 2^9 - 1 (per 0.153 standard)
repeating-1-in-4       1 bit in 4 is set
repeating-1-in-8       1 bit in 8 is set
repeating-3-in-24      3 bits in 24 are set
```

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.



**NOTE:** The IQE PICs support only the following algorithms:

```
all-ones-repeating    Repeating one bits
all-zeros-repeating   Repeating zero bits
alternating-double-ones-zeros Alternating pairs of ones and zeros
alternating-ones-zeros Alternating ones and zeros
pseudo-2e9-o153       Pattern is 2^9 - 1 (per 0.153 (511 type) standard)
pseudo-2e11-o152      Pattern is 2^11 - 1 (per 0.152 and 0.153 (2047 type)
standards)
pseudo-2e15-o151      Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151      Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e20-o153      Pattern is 2^20 - 1 (per 0.153 standard)
pseudo-2e23-o151      Pattern is 2^23 - 1 (per 0.151 standard)
repeating-1-in-4       1 bit in 4 is set
repeating-1-in-8       1 bit in 8 is set
repeating-3-in-24      3 bits in 24 are set
```

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.



**NOTE:** BERT is supported on the PDH interfaces of the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP and the DS3/E3 MIC. The following BERT algorithms are supported:

all-ones-repeating	Repeating one bits
all-zeros-repeating	Repeating zero bits
alternating-double-ones-zeros	Alternating pairs of ones and zeros
alternating-ones-zeros	Alternating ones and zeros
repeating-1-in-4	1 bit in 4 is set
repeating-1-in-8	1 bit in 8 is set
repeating-3-in-24	3 bits in 24 are set
pseudo-2e9-o153	Pattern is $2^9 - 1$ (per 0.153 standard)
pseudo-2e11-o152	Pattern is $2^{11} - 1$ (per 0.152 standard)
pseudo-2e15-o151	Pattern is $2^{15} - 1$ (per 0.151 standard)
pseudo-2e20-o151	Pattern is $2^{20} - 1$ (per 0.151 standard)
pseudo-2e20-o153	Pattern is $2^{20} - 1$ (per 0.153 standard)
pseudo-2e23-o151	Pattern is $2^{23} - 1$ (per 0.151 standard)

Table 42 on page 300 shows the BERT capabilities for various interface types.

**Table 42: BERT Capabilities by Interface Type**

Interface	T1 BERT	T3 BERT	Comments
12-port T1/E1 Circuit Emulation	Yes (ports 0–11)		<ul style="list-style-type: none"> <li>Limited algorithms</li> </ul>
4-port Channelized OC3/STM1 Circuit Emulation	Yes (port 0–3)		<ul style="list-style-type: none"> <li>Limited algorithms</li> </ul>
E1 or T1	Yes (port 0–3)	Yes (port 0–3)	<ul style="list-style-type: none"> <li>Single port at a time</li> <li>Limited algorithms</li> </ul>
E3 or T3	Yes (port 0–3)	Yes (port 0–3)	<ul style="list-style-type: none"> <li>Single port at a time</li> </ul>
Channelized OC12	N/A	Yes (channel 0–11)	<ul style="list-style-type: none"> <li>Single channel at a time</li> <li>Limited algorithms</li> <li>No bit count</li> </ul>
Channelized STM1	Yes (channel 0–62)	N/A	<ul style="list-style-type: none"> <li>Multiple channels</li> <li>Only one algorithm</li> <li>No error insert</li> <li>No bit count</li> </ul>
Channelized T3 and Multichannel T3	Yes (channel 0–27)	Yes (port 0–3 on channel 0)	<ul style="list-style-type: none"> <li>Multiple ports and channels</li> <li>Limited algorithms for T1</li> <li>No error insert for T1</li> <li>No bit count for T1</li> </ul>

These limitations do not apply to channelized IQ interfaces. For information about BERT capabilities on channelized IQ interfaces, see *Channelized IQ and IQE Interfaces Properties*.

### Starting and Stopping a BERT Test

Before you can start the BERT test, you must disable the interface. To do this, include the **disable** statement at the **[edit interfaces *interface-name*]** hierarchy level:

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

After you configure the BERT properties and commit the configuration, begin the test by issuing the **test interface *interface-name* *interface-type*-bert-start** operational mode command:

```
user@host> test interface interface-name interface-type-bert-start
```

The test runs for the duration you specify with the **bert-period** statement. If you wish to terminate the test sooner, issue the **test interface *interface-name* *interface-type*-bert-stop** command:

```
user@host> test interface interface-name interface-type-bert-stop
```

For example:

```
user@host> test interface t3-1/2/0 t3-bert-start
user@host> test interface t3-1/2/0 t3-bert-stop
```

To view the results of the BERT test, issue the **show interfaces extensive | find BERT** command:

```
user@host> show interfaces interface-name extensive | find BERT
```

For more information about running and evaluating the results of the BERT procedure, see the [CLI Explorer](#).



**NOTE:** To exchange BERT patterns between a local router and a remote router, include the **loopback remote** statement in the interface configuration at the remote end of the link. From the local router, issue the **test interface** command.

### Example: Configuring Bit Error Rate Testing

Configure a BERT test on a T3 interface. In this example, the run duration lasts for 120 seconds. The configured error rate is 0, which corresponds to a bit error rate of  $10^{-0}$  (1 error per bit). The configured bit pattern of **all-ones-repeating** means that every bit the interface sends is a set to a value of 1.

```
[edit interfaces]
t3-1/2/0 {
  t3-options {
    bert algorithm all-ones-repeating;
    bert-error-rate 0;
    bert-period 120;
```

```
}  
}
```

## PART 5

# Index

- [Index on page 305](#)



# Index

## Symbols

#, comments in configuration statements.....	xii
( ), in syntax descriptions.....	xii
10-Gigabit Ethernet interfaces	
DWDM	
data and alarms, displaying.....	123
status information, displaying.....	185, 188
100-Gigabit Ethernet interfaces.....	123
40-Gigabit Ethernet interfaces.....	123
802.1Q VLANs	
mixed VLAN tagging.....	11, 76, 79
VLAN IDs.....	87, 90
values, listed by Ethernet interface	
type.....	8
VLAN tagging.....	91
< >, in syntax descriptions.....	xii
[ ], in configuration statements.....	xii
{ }, in configuration statements.....	xii
(pipe), in syntax descriptions.....	xii

## A

access interface	
interface-mode statement.....	32
ACI (agent circuit identifier) interface sets	
status information, displaying.....	185
address statement.....	60
aggregated Ethernet interfaces	
status information, displaying.....	103
VLAN IDs.....	8

## B

BERT	
configuring interface diagnostics.....	297
bert-algorithm statement	
usage guidelines.....	297
bert-error-rate statement	
usage guidelines.....	297
bert-period statement	
usage guidelines.....	297
bit error rate test See BERT	
braces, in configuration statements.....	xii

## brackets

angle, in syntax descriptions.....	xii
square, in configuration statements.....	xii

## bridge network

trunk interface.....	34
----------------------	----

## C

### CCC

encapsulation	
VLAN-bundled dual-tag logical	
interfaces.....	92
VLAN-bundled single-tag logical	
interfaces.....	96

### CFM statistics

displaying for CFM interfaces.....	234
displaying for interfaces.....	221

### circuit cross-connect (CCC)

encapsulation	
VLAN-bundled dual-tag logical	
interfaces.....	92
VLAN-bundled single-tag logical	
interfaces.....	96

### comments, in configuration statements.....

### Configuring a Layer 2 Circuit on a VLAN-Bundled

Logical Interface.....	27
------------------------	----

### Configuring a VLAN-Bundled Logical

Interface.....	23, 25, 27, 28
----------------	----------------

### Configuring Logical Link-Layer Encapsulation to

Support CCCs.....	22
-------------------	----

### Configuring VLAN ID List-Bundled Logical Interfaces

That Connect CCCs.....	21
------------------------	----

### conventions

text and syntax.....	xi
----------------------	----

### curly braces, in configuration statements.....

### customer support.....

contacting JTAC.....	xiii
----------------------	------

## D

### documentation

comments on.....	xiii
------------------	------

### dual-tag framing

VLAN ID list.....	92
-------------------	----

### DWDM 10-Gigabit Ethernet interfaces

data and alarms, displaying.....	123
----------------------------------	-----

## E

### encapsulation

extended VLAN CCC.....	20
------------------------	----

encapsulation statement	
logical interfaces.....	62
physical interface.....	66
ETH-DM frame counts (with CFM statistics)	
displaying for MEPs by enclosing CFM.....	234
displaying for MEPs by interface or domain level.....	221
ETH-DM statistics (only)	
displaying.....	213
ETH-DM statistics and frame counts	
displaying.....	245
ETH-SLM frame counts (with CFM statistics)	
displaying for MEPs by enclosing CFM.....	234
displaying for MEPs by interface or domain level.....	221
Ethernet interface set	
status information, displaying.....	185
Ethernet interface set queue	
status information, displaying.....	188
Ethernet interfaces	
CFM statistics	
displaying for CFM interfaces.....	234
displaying for interfaces.....	221
ETH-DM frame counts (with CFM statistics)	
displaying for MEPs by enclosing CFM.....	234
displaying for MEPs by interface or domain level.....	221
ETH-DM statistics (only)	
displaying.....	213
ETH-DM statistics and frame counts	
displaying.....	245
ETH-SLM frame counts (with CFM statistics)	
displaying for MEPs by enclosing CFM.....	234
displaying for MEPs by interface or domain level.....	221
mixed VLAN tagging.....	11, 76, 79
status information, displaying	
aggregated.....	103
Fast Ethernet.....	144
Gigabit Ethernet.....	160
VLAN IDs.....	87
VLAN tagging.....	91
ethernet ring	
protection group APS.....	269
protection group data channel.....	271, 282
protection group interface.....	273
protection group node state.....	276
protection group statistics.....	279
Ethernet VLAN circuit	
VLAN ID list.....	96
extended VLAN	
CCC	
applying.....	20
example configuration.....	20, 21
<b>F</b>	
family bridge	
VLAN ID list.....	89
family statement.....	71
Fast Ethernet interfaces	
status information, displaying.....	144
VLAN IDs.....	8, 87
VLAN tagging.....	91
fast-aps-switch statement.....	75
flexible-vlan-tagging statement.....	76
font conventions.....	xi
<b>G</b>	
Gigabit Ethernet interfaces	
demultiplexing interface information,	
displaying.....	114
diagnostics information, displaying.....	123
LACP, displaying.....	202
MAC database, displaying.....	206
status information, displaying.....	160, 185, 188
VLAN IDs.....	8, 87, 90
VLAN tagging.....	91
Gigabit Ethernet IQ PIC	
traffic and MAC statistics.....	160
<b>I</b>	
integrated routing and bridging interfaces	
integrated routing and bridging interfaces,	
displaying.....	196
interface statement	
usage guidelines.....	34
interface-mode statement.....	77
usage guidelines.....	32, 33
interfaces	
configuration statements.....	37
mixed VLAN tagging.....	11, 76, 79
<b>L</b>	
LACP	
displaying.....	202

- Layer 2.5 VPNs
  - Ethernet.....20
- linktrace
  - database, displaying.....231
  - tracing.....288
- log files
  - number of PPPoE.....291
  - size of PPPoE.....291
- logical interfaces
  - VLAN IDs.....87, 90
  - VLAN-bundled
    - dual-tag.....92
    - single-tag.....96
- logical systems
  - configuration statements.....53
- loopback testing.....295
- M**
  - MAC database, displaying.....206
  - manuals
    - comments on.....xiii
  - mixed VLAN tagging.....11, 76, 79
  - multi-chassis link aggregation interface
    - displaying.....210
- N**
  - native-vlan-id statement.....79
- P**
  - parentheses, in syntax descriptions.....xii
  - physical interfaces
    - mixed VLAN tagging.....11, 76, 79
    - VLAN tagging.....91
  - PPPoE
    - flags for tracing operations.....291
    - log file access for tracing operations.....291
    - log file size and number.....291
    - log filenames for tracing operations.....290
    - regular expressions for tracing operations.....291
    - service name tables
      - configuration troubleshooting.....291
      - verifying the configuration.....293
    - tracing operations.....290
  - protection group
    - ethernet ring APS.....269
    - ethernet ring interface.....273
    - ethernet ring node state.....276
    - ethernet ring statistics.....279
    - ethernet ring VLANs.....271, 282
  - pseudowires
    - mixed VLAN tagging.....76
- R**
  - remote statement
    - usage guidelines.....20
- S**
  - service name tables
    - PPPoE
      - configuration troubleshooting.....291
      - verifying the configuration.....293
  - show interfaces (10-Gigabit Ethernet)
    - command.....160
  - show interfaces (Aggregated Ethernet)
    - command.....103
  - show interfaces (Fast Ethernet) command.....144
  - show interfaces demux0 (Demux Interfaces)
    - command.....114
  - show interfaces diagnostics optics (Gigabit Ethernet) command.....123
  - show interfaces interface-set command.....185
  - show interfaces interface-set queue
    - command.....188
  - show interfaces irb command.....196
  - show interfaces mac-database (Gigabit Ethernet)
    - command.....206
  - show interfaces mc-ae.....210
  - show lacp interfaces command.....202
  - show oam ethernet connectivity-fault-management
    - delay-statistics command.....213
    - forwarding-state command.....217
    - interfaces command.....221
    - linktrace path-database command.....231
    - mep-database command.....234
    - mep-statistics command.....245
    - path-database command.....256
  - show oam ethernet link-fault-management
    - command.....259
  - show protection-group ethernet-ring aps
    - command.....269
  - show protection-group ethernet-ring data-channel
    - command.....271

show protection-group ethernet-ring interface command.....	273
show protection-group ethernet-ring node-state command.....	276
show protection-group ethernet-ring statistics command.....	279
show protection-group ethernet-ring vlan command.....	282
Specifying the Interface Over Which VPN Traffic Travels to the CE Router.....	24, 25
Specifying the Interface to Handle Traffic for a CCC.....	24, 26
Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit.....	27, 29
stacked VLAN-tag framing	
VLAN ID list.....	94
stacked-vlan-tagging statement.....	80
support, technical See technical support	
syntax conventions.....	xi

## T

technical support	
contacting JTAC.....	xiii
traceroute ethernet.....	288
traceroute ethernet command.....	288
tracing operations	
PPPoE.....	290
trunk interface.....	77
interface-mode statement.....	33
usage guidelines.....	34
vlan-id-list statement.....	33
trunk port.....	33, 77
VLAN ID list.....	89

## U

unit statement.....	81
---------------------	----

## V

VLAN	
Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface.....	27
Configuring a VLAN-Bundled Logical Interface.....	27, 28
Configuring Logical Link-Layer Encapsulation to Support CCCs.....	22
Configuring VLAN ID List-Bundled Logical Interfaces That Connect CCCs.....	21

Specifying the Interface Over Which VPN Traffic Travels to the CE Router.....	24, 25
Specifying the Interface to Handle Traffic for a CCC.....	24, 26
VLAN CCC	
configuration guidelines.....	18
example configuration.....	19, 20
VLAN IDs.....	90
values, listed by Ethernet interface type.....	8
VLAN tagging.....	91
VLAN tagging on PTX Series Packet Transport Routers	
configuring.....	36
VLAN VPLS	
configuration guidelines.....	18
example configuration.....	19, 20
vlan-id statement	
802.1Q VLANs.....	87
Ethernet interfaces.....	87
Gigabit Ethernet interfaces	
usage guidelines.....	12, 17
rewriting at ingress or egress.....	88
vlan-id-list statement	
bridge domain.....	89
Ethernet VLAN circuit.....	96
Gigabit Ethernet interfaces	
usage guidelines.....	12
usage guidelines.....	33
vlan-id-range statement.....	90
Ethernet interfaces.....	90
Gigabit Ethernet interfaces	
usage guidelines.....	12
vlan-tagging statement.....	91
vlan-tags statement	
dual-tag framing.....	92
Gigabit Ethernet interfaces	
usage guidelines.....	17
stacked VLAN tags.....	94