



Junos[®] OS

10-Gigabit Ethernet LAN/WAN PICs Feature Guide for Routing Devices



Published: 2013-08-01

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

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

Juniper Networks, 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.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Junos® OS 10-Gigabit Ethernet LAN/WAN PICs Feature Guide for Routing Devices

Copyright © 2013, 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
Part 1	Overview	
Chapter 1	10-Gigabit Ethernet LAN/WAN PICs	3
	10-port 10-Gigabit Ethernet LAN/WAN PIC Overview	3
	12-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview	7
	24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview	9
Part 2	Configuration	
Chapter 2	10-Gigabit Ethernet LAN/WAN PICs	13
	Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription	13
	Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC	14
	Example: Handling Oversubscription on a 10-Gigabit Ethernet LAN/WAN PIC . . .	17
Chapter 3	Network Interfaces Configuration Statements and Hierarchy	19
	[edit interfaces] Hierarchy Level	19
Chapter 4	Statement Summary	37
	fast-aps-switch	37
	no-pre-classifier	38
	pop	39
	pop-pop	39
	pop-swap	40
	push	41
	push-push	42
	swap	43
	swap-swap	44
	inner-tag-protocol-id	45

	tag-protocol-id (TPID to Rewrite)	46
	vlan-id (VLAN ID to Rewrite)	47
Part 3	Administration	
Chapter 5	Monitoring Commands	51
	show interfaces (10-Gigabit Ethernet)	52
Chapter 6	Command Summary	77
	Ethernet Interface Operational Mode Commands	77
Part 4	Troubleshooting	
Chapter 7	Ethernet	85
	tracertoe ethernet	86
Chapter 8	Interface Diagnostics	89
	Interface Diagnostics	89
	Configuring Loopback Testing	89
	Interface Diagnostics	91
	Starting and Stopping a BERT Test	95
	Example: Configuring Bit Error Rate Testing	95
Part 5	Index	
	Index	99

List of Figures

Part 2	Configuration	
Chapter 2	10-Gigabit Ethernet LAN/WAN PICs	13
	Figure 1: Control Queue Rate Limiter Scenario	14

List of Tables

	About the Documentation	ix
	Table 1: Notice Icons	xi
	Table 2: Text and Syntax Conventions	xi
Part 1	Overview	
Chapter 1	10-Gigabit Ethernet LAN/WAN PICs	3
	Table 3: Capabilities of 10-Gigabit Ethernet LAN/WAN PICs	6
Part 2	Configuration	
Chapter 2	10-Gigabit Ethernet LAN/WAN PICs	13
	Table 4: Handling Oversubscription on 10-Gigabit Ethernet LAN/WAN PICs	17
Part 3	Administration	
Chapter 5	Monitoring Commands	51
	Table 5: show interfaces Gigabit Ethernet Output Fields	53
	Table 6: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type	67
Chapter 6	Command Summary	77
	Table 7: Ethernet Interface Operational Mode Commands	77
Part 4	Troubleshooting	
Chapter 7	Ethernet	85
	Table 8: traceroute ethernet Output Fields	86
Chapter 8	Interface Diagnostics	89
	Table 9: Loopback Modes by Interface Type	90
	Table 10: BERT Capabilities by Interface Type	94

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® 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:

- T4000
- T1600
- T640

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.

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active

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

Convention	Description	Examples
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies book names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS System Basics Configuration Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Enclose optional keywords or variables.	stub <default-metric metric>;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast <i>(string1 string2 string3)</i>
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Enclose a variable for which you can substitute one or more values.	community name members [community-ids]
Indentation and braces ({ })	Identify a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to 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: <https://www.juniper.net/alerts/>

- Join and participate in the Juniper Networks Community Forum:
<http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

- [10-Gigabit Ethernet LAN/WAN PICs on page 3](#)

CHAPTER 1

10-Gigabit Ethernet LAN/WAN PICs

- [10-port 10-Gigabit Ethernet LAN/WAN PIC Overview on page 3](#)
- [12-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview on page 7](#)
- [24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview on page 9](#)

10-port 10-Gigabit Ethernet LAN/WAN PIC Overview

This section describes the main features and caveats of the 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PD-5-10XGE-SFPP) and specifies which routers support this PIC.

The 10-port 10-Gigabit Ethernet LAN/WAN PIC (PD-5-10XGE-SFPP) is supported on Juniper Networks T640 Core Routers, T1600 Core Routers, and T4000 Core Routers. It has the following features:

- Access to all 10-Gigabit Ethernet port counters through SNMP
- Intelligent handling of oversubscribed traffic in applications such as data centers and dense-core uplinks
- Line-rate operation for five 10-Gigabit Ethernet ports from each port group, or a total WAN bandwidth of 100 Gbps with Packet Forwarding Engine bandwidth of 50 Gbps
- Flexible encapsulation, source address and destination address media access control (MAC) filtering, source address MAC learning, MAC accounting, and MAC policing
- Interface encapsulations, such as the following:
 - **ethernet-ccc**—Ethernet cross-connect
 - **vlan-ccc**—802.1Q tagging for a cross-connect
 - **ethernet-tcc**—Ethernet translational cross-connect
 - **vlan-tcc**—Virtual LAN (VLAN) translational cross-connect
 - **extended-vlan-ccc**—Standard Tag Protocol Identifier (TPID) tagging for a cross-connect
 - **ethernet-vpls**—Ethernet virtual private LAN service
 - **vlan-vpls**—VLAN virtual private LAN service
 - **flexible-ethernet-services**—Allows per-unit Ethernet encapsulation configuration

- WAN PHY features, such as the following:
 - WAN PHY mode on a per-port basis
 - Insertion and detection of path trace messages
 - Ethernet WAN Interface Sublayer (WIS) object



NOTE: The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP). Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).

- Single, stacked, and flexible VLAN tagging modes
- Native VLAN configuration to allow untagged frames to be received on the tagged interfaces
- Maximum transmission unit (MTU) size of up to 9192 bytes for Ethernet frames
- Link aggregation group (LAG) on single chassis
- Interoperability with other 10-Gigabit Ethernet PICs in M Series and T Series routers in the LAN PHY and WAN PHY modes
- Interrupt-driven link-down detection mechanism
- Two-to-one oversubscription of traffic across a port group

Traffic from 10 ingress ports to the Packet Forwarding Engine traffic is statically mapped to one of the 5 egress ports. 10 Gbps of bandwidth toward the Packet Forwarding Engine is shared by two ingress ports (called a *port group*), thereby achieving two-to-one oversubscription. This scheme provides two-to-one oversubscription across a port group and not across the entire PIC.

- Four queues per physical interface on ingress and eight queues per physical interface on egress
- A separate control queue per physical interface to ensure that the control packets are not dropped during oversubscribed traffic. The control queue can be disabled in the CLI.
- Optical diagnostics
- Behavior aggregate (BA) classification (IPv4 DSCP, IPv6 DSCP, Inet precedence, IEEE 802.1P, IEEE 802.1AD, MPLS EXP) and fixed classification
- Weighted round-robin scheduling with two queue priorities (low and strict-high)
- Committed information rate and peak information rate shaping on a per-queue basis
- Excess information rate configuration for allocation of excess bandwidth
- IEEE 802.3ah Operation, Administration, and Maintenance (OAM)-related operations, such as the following:

- Link fault management
- Link discovery
- Graceful Routing Engine Switchover
- IEEE 802.3ag Operation, Administration, and Maintenance (OAM)-related operations, such as the following:
 - Connectivity fault management (CFM)
 - Linktrace
 - Loopback
 - Graceful Routing Engine switchover (GRES)

The 10-port 10-Gigabit Ethernet LAN/WAN PIC has the following caveats:

- Source address and destination address MAC filtering takes place after oversubscription is handled.
- Oversubscription on the PIC operates across a port group of two ports and not at the PIC level.
- Queuing is not supported at the logical interface level.
- Committed information rate and peak information rate configurations are not supported at the physical interface level.
- There is limited packet buffering of 2 MB.
- Delay-bandwidth buffering configuration is not supported.
- Multifield classifiers are not supported at the PIC level.

The multifield classification can be done at the Packet Forwarding Engine using the firewall filters, which overrides the classification done at the PIC level. The multifield classification at the Packet Forwarding Engine occurs after the PIC handles the oversubscribed traffic.

- Egress MAC policer statistics not supported.
- Byte counters are not supported at the queue level.
- Only TPID (0x8100) is supported.
- Line-timing mode is not supported.
- MAC-level Rx VLAN tagged frames counter is not supported.
- OAM unified in-service software upgrade (unified ISSU) is not supported.
- OAM remote loopback is not supported.

The 10-port 10-Gigabit Ethernet LAN/WAN PIC (PD-5-10XGE-SFPP) supports link aggregation. For bandwidth aggregation, load sharing, and link protection, LAG can be enabled. Once aggregated Ethernet is enabled, Link Aggregation Control Protocol (LACP) forms an aggregated bundle of member links.

Only features that are supported across all of the linked devices will be supported in the resulting LAG bundle. The following caveats apply to LAG bundles that involve 10-port 10-Gigabit Ethernet LAN/WAN PIC (PD-5-10XGE-SFPP) ports:

- Non-standard TPID for VLAN tagging is not supported, except for 0x8100.
- The number of user created IFLs is limited to 4065/PIC and 1022/port.
- Classifier tables are limited to 8 for each BA classifier type.
- Forwarding classes are limited to 8.
- The **guaranteed-rate** and **shaping-rate** statements are not supported at the IFD level.
- The **per-unit-scheduler** and **hierarchical-scheduler** statements are not supported.
- Only the **strict-high** and **low** levels of scheduling priorities are supported.
- The **excess-priority** configuration is not supported.
- The **buffer-size** configuration under **schedulers** is not supported.
- WRED is not supported.
- srTCM and trTCM are not supported.
- Shared scheduler mode is not supported.

[Table 3 on page 6](#) 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PD-5-10XGE-SFPP).

Table 3: Capabilities of 10-Gigabit Ethernet LAN/WAN PICs

Capability	Support
Maximum VLANs per PIC	4065
Maximum VLANs per port	1022
MAC learning per port	960
MAC accounting per port	960
MAC filtering per port	960 (64 filters per physical or logical interface) 960 filters across multiple logical interfaces
MAC policers	128 ingress Mac policers 128 egress Mac policers
Classifiers	Eight classifiers per PIC for each BA classifier type

**Related
Documentation**

- [Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription on page 13](#)

- [Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC on page 14](#)
- *Handling Oversubscription on a 10-Port 10-Gigabit Ethernet LAN/WAN PIC*
- *IEEE 802.1ag OAM Connectivity Fault Management Overview*
- *IEEE 802.3ah OAM Link-Fault Management Overview*
- *Ethernet Interfaces*

12-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview

The 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC is a 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number, PF-12XGE-SFPP) on T4000 Core Routers.

The following features are supported on the 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC:

- Access to all 10-Gigabit Ethernet port counters through SNMP.
- Logical interface—level MAC filtering, accounting, policing, and learning for source media access control (MAC).
- Flexible encapsulation.
- Single, stacked, and flexible VLAN tagging modes.
- Native VLAN configuration to allow untagged frames to be received on the tagged interfaces.
- Maximum transmission unit (MTU) size of up to 9192 bytes for Ethernet frames.
- Link aggregation group (LAG) on single chassis.
- Interoperability with other 10-Gigabit Ethernet PICs on M Series and T Series routers in LAN PHY mode.
- Eight queues per physical interface on egress.
- Behavior aggregate (BA) classification (IPv4 DSCP, IPv6 DSCP, Inet precedence, IEEE 802.1P, IEEE 802.1AD, MPLS EXP) and fixed classification.
- Defining the VLAN rewrite operation to be applied to the incoming and outgoing frames on logical interfaces on this PIC.



NOTE: Only the Tag Protocol Identifier (TPID) 0x8100 is supported.

- Interface encapsulations, such as the following:
 - **untagged**—Default encapsulation, when other encapsulation is not configured.
 - You can configure only one logical interface (unit 0) on the port.

- You cannot include the **vlan-id** statement in the configuration of the logical interface.
- **vlan-tagging**—Enable VLAN tagging for all logical interfaces on the physical interface.
- **stacked-vlan-tagging**—Enable stacked VLAN tagging for all logical interfaces on the physical interface.
- **ethernet-ccc**—Ethernet cross-connect.
- **ethernet-tcc**—Ethernet translational cross-connect.
- **vlan-ccc**—802.1Q tagging for a cross-connect.
- **vlan-tcc**—Virtual LAN (VLAN) translational cross-connect.
- **extended-vlan-ccc**—Standard Tag Protocol Identifier (TPID) tagging for a cross-connect.
- **extended-vlan-tcc**—Standard Tag Protocol Identifier (TPID) tagging for an Ethernet translational cross-connect.
- **ethernet-vpls**—Ethernet virtual private LAN service.
- **vlan-vpls**—VLAN virtual private LAN service.
- **flexible-ethernet-services**—Allows per-unit Ethernet encapsulation configuration.
- The following Layer 3 protocols are also supported:
 - IPv4
 - IPv6
 - MPLS
- WAN PHY features, such as the following:
 - WAN PHY mode on a per-port basis.
 - Insertion and detection of path trace messages.
 - Ethernet WAN Interface Sublayer (WIS) object.



.....

NOTE: The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on T4000 routers with 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+.

.....

The 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC does not support:

- MAC filtering, accounting, and policing for destination MAC at the logical interface level.



NOTE: Because destination MAC filtering is not supported, the hardware is configured to accept all the multicast packets. This enables the OSPF protocol to work.

- Premium MAC policers at the logical interface level.
- MAC filtering, accounting, and policing at the physical interface level.
- Multiple TPIDs

Capability	Support
Maximum logical interfaces per PIC	32,000
Maximum logical interfaces per port	For IPv4 the limit is 4093. For IPv6 the limit is 1022.
Classifiers	Eight classifiers per PIC for each BA classifier type

- Related Documentation
- *Ethernet Interfaces*
 - [10-port 10-Gigabit Ethernet LAN/WAN PIC Overview on page 3](#)

24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview

This section describes the main features and caveats of the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PF-24XGE-SFPP).

The following major software features are supported on the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PF-24XGE-SFPP):

- Twenty-four 10-Gigabit Ethernet interfaces in two-to-one oversubscription of traffic in oversubscribed mode or 12 ports in line-rate mode. For more information about oversubscribed mode and line-rate mode, see the “[Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription](#)” on page 13.
- Traffic is classified as control traffic or best-effort traffic with non-class-of-service-aware tail drops of best-effort traffic in oversubscribed mode.

The aggregate bandwidth of all the ports together is 120 Gbps. No hard partitioning of bandwidth is done—that is, if one port group is active, it can support 120 Gbps traffic. The bandwidth for best-effort traffic is shared among all the 24 ports.

Note that the preclassification is restricted to two traffic classes, and is not user-configurable.

- All Junos OS configuration commands supported on the existing 10-Gigabit Ethernet LAN/WAN PIC with SFP+.
- The output of the **show interfaces extensive** operational mode command now displays preclassification queue counters.
- Line-rate mode operation of the first 12 ports can be achieved by using the **[set chassis fpc fpc-number pic pic-number linerate-mode]** command. By default, the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ works in oversubscribed mode.
- LAN PHY mode and WAN PHY mode on a per-port basis. WAN PHY mode can be achieved by using the **[set interfaces interface-name framing wan-phy]** command.
- WAN PHY features, such as the following:
 - Insertion and detection of path trace messages.
 - Ethernet WAN Interface Sublayer (WIS) object.
- Aggregated Ethernet is supported only in line-rate mode.
- Link aggregation group (LAG) is supported only in line-rate mode.
- 4000 logical interfaces per physical interface and 32,000 logical interfaces per chassis.
- Access to all 10-Gigabit Ethernet port counters through SNMP.



NOTE: Graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) are now supported on T4000 routers.

**Related
Documentation**

- [Ethernet Interfaces](#)
- [12-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview on page 7](#)
- [Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription on page 13](#)

PART 2

Configuration

- [10-Gigabit Ethernet LAN/WAN PICs on page 13](#)
- [Network Interfaces Configuration Statements and Hierarchy on page 19](#)
- [Statement Summary on page 37](#)

CHAPTER 2

10-Gigabit Ethernet LAN/WAN PICs

- [Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription on page 13](#)
- [Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC on page 14](#)
- [Example: Handling Oversubscription on a 10-Gigabit Ethernet LAN/WAN PIC on page 17](#)

Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription

For 10-Gigabit Ethernet LAN/WAN PICs supporting oversubscription, oversubscribed Ethernet mode is set by default. To configure these PICs in line-rate mode, include the **linerate-mode** statement at the **[edit chassis set fpc fpc-number pic pic-number]** hierarchy level:

```
[edit chassis]
set fpc fpc-number pic pic-number linerate-mode;
```

To return to the default oversubscribed Ethernet mode, delete the **linerate-mode** statement at the **[edit chassis fpc fpc-number pic pic-number]** hierarchy level.



NOTE: When the mode of operation of a PIC is changed, the PIC is taken offline and then brought back online immediately.

The following 10-Gigabit Ethernet LAN/WAN PICs support line-rate mode:

- 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PD-5-10XGE-SFPP)
- 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PF-24XGE-SFPP)

Related Documentation

- [10-port 10-Gigabit Ethernet LAN/WAN PIC Overview on page 3](#)
- [24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview on page 9](#)
- [Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC on page 14](#)
- [Handling Oversubscription on a 10-Port 10-Gigabit Ethernet LAN/WAN PIC](#)

- *Ethernet Interfaces*


Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC

On a 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PD-5-10XGE-SFPP), a control queue is used to queue all control packets received on an ingress port. This ensures that control protocol packets do not get dropped randomly when there is congestion due to oversubscription. The following control protocols are supported:

- OSPF
- OSPF3
- VRRP
- IGMP
- RSVP
- PIM
- BGP
- BFD
- LDP
- IS-IS
- RIP
- RIPV6
- LACP
- ARP
- IPv6 NDP
- Connectivity fault management (CFM)
- Link fault management (LFM)

These control packets can either terminate locally or transit through the router. The control queue has a rate limiter to limit the control traffic to 2 Mbps (fixed, not user-configurable) per port. Hence, if transit control traffic is taking too much bandwidth, then it can cause drops on locally terminating control traffic, as shown in [Figure 1 on page 14](#).

Figure 1: Control Queue Rate Limiter Scenario



If the end users generate a mass of malicious traffic for which the port number is 179 (BGP), the router dispatches that traffic to the ingress control queue. Further, if congestion occurs in this ingress control queue due to this malicious traffic, the provider's network control packets may be affected.

In some applications, this can be perceived as a new vulnerability. To address this concern, you can disable the control queue feature. With the control queue feature disabled, you must take precautions to protect control traffic through other means, such as mapping control packets (using BA classification) to a queue that is marked strict-high or is configured with a high CIR.

You can disable the control queue for all ports on the PIC. To disable the control queue, use the **set chassis fpc *n* pic *n* no-pre-classifier** command. By default, the **no-pre-classifier** statement is not configured and the control queue is operational.

Deleting the **no-pre-classifier** statement re-enables the control queue feature on all ports of the 10-Gigabit Ethernet LAN/WAN PIC.



NOTE:

- This functionality is applicable both in OSE and line-rate modes.
- The control queue feature is enabled by default in both OSE and line-rate modes, which can be overridden by the user configuration.
- When the control queue is disabled, various **show queue** commands will show *control queue* in the output. However, all control queue counters are reported as zeros.
- Changing this configuration (enabling or disabling the control queue feature) results in the PIC being taken offline and brought back online.

Once the control queue is disabled, the Layer 2/Layer 3 control packets are subject to queue selection based on BA classification. However, some control protocol packets will not be classified using BA classification, because they might not have a VLAN, MPLS, or IP header. These are:

- Untagged ARP packets
- Untagged Layer 2 control packets such as LACP or Ethernet OAM
- Untagged IS-IS packets

When the control queue feature is disabled, untagged ARP, IS-IS, and other untagged Layer 2 control packets will go to the restricted queue corresponding to the forwarding class associated with queue 0, as shown in the following two examples.

Forwarding Untagged Layer2 Control Packets to Queue 3

With this configuration, the forwarding class (FC) associated with queue 0 is "be" (based on the **forwarding-class** statement configuration). "be" maps to restricted-queue number 3 (based on the "restricted-queue" configuration). Hence, with this particular configuration, untagged ARP, IS-IS, and other untagged Layer 2 control packets will go to ingress queue 3 (not to ingress queue 0).

[edit chassis]

```
forwarding-classes {
    queue 0 be;
    queue 1 af-low8;
    queue 2 af-high;
    queue 3 ef;
    queue 4 ops_control;
    queue 5 net_control;
    queue 6 af-low10_12;
}
restricted-queues {
    forwarding-class ef queue-num 0;
    forwarding-class af-low8 queue-num 1;
    forwarding-class af-low10_12 queue-num 1;
    forwarding-class af-high queue-num 2;
    forwarding-class be queue-num 3;
}
```

Forwarding Untagged Layer2 Control Packets to Queue 3

With this configuration, the FC associated with queue 0 is "ef" (based on the **forwarding-class** statement configuration). "ef" maps to restricted-queue number 0 (based on the **restricted-queue** statement configuration). Hence, with this particular configuration, untagged ARP, IS-IS, and other untagged Layer 2 control packets would go to ingress queue 0.

For tagged ARP, IS-IS, or Layer2 control packets, users should configure an explicit dot1p/dot1ad classifier to make sure these packets are directed to the correct queue. Without an explicit dot1p/dot1ad classifier, tagged ARP, IS-IS, or Layer 2 control packets will go to the restricted-queue corresponding to the forwarding class associated with queue 0.

```
[edit chassis]
forwarding-classes {
    queue 0 ef; <<< ef and be are interchanged
    queue 1 af-low8;
    queue 2 af-high;
    queue 3 be; <<< ef and be are interchanged
    queue 4 ops_control;
    queue 5 net_control;
    queue 6 af-low10_12;
}
restricted-queues {
    forwarding-class ef queue-num 0;
    forwarding-class af-low8 queue-num 1;
    forwarding-class af-low10_12 queue-num 1;
    forwarding-class af-high queue-num 2;
    forwarding-class be queue-num 3;
}
```

Related Documentation

- [10-port 10-Gigabit Ethernet LAN/WAN PIC Overview on page 3](#)
- [Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription on page 13](#)
- [no-pre-classifier on page 38](#)
- *Handling Oversubscription on a 10-Port 10-Gigabit Ethernet LAN/WAN PIC*

- *Ethernet Interfaces*

Example: Handling Oversubscription on a 10-Gigabit Ethernet LAN/WAN PIC

Table 4 on page 17 lists the scenarios of handling oversubscription on the 10-port 10-Gigabit Ethernet LAN/WAN PIC for different combinations of port groups and active ports on the PIC.

Table 4: Handling Oversubscription on 10-Gigabit Ethernet LAN/WAN PICs

Number of Port Groups with Two Active Ports (A)	Number of Port Groups with One Active Port (B)	Total Number of Ports Used on PIC ($C = A \times 2 + B$)	Status of Oversubscription and Throughput
0	1	1	Oversubscription is not active. Each port will receive 10 Gbps throughput.
0	2	2	Oversubscription is not active. Each port will receive 10 Gbps throughput.
0	5	5	Oversubscription is not active. Each port will receive 10 Gbps throughput.
1	0	2	Oversubscription is active. Each port will receive 5 Gbps throughput (with default shaper configuration).
1	4	6	<p>Oversubscription is active for the port group that has two active ports. Each port in this port group will receive 5 Gbps throughput (with default shaper configuration).</p> <p>For the remaining four ports, oversubscription is not active. Each port will receive 10 Gbps throughput.</p>
3	0	6	Oversubscription is active. Each port will receive 5 Gbps throughput (with default shaper configuration).
5	0	10	Oversubscription is active on all 10 ports (5 port groups). Each port will receive 5 Gbps throughput (with default shaper configuration).

Related Documentation

- [10-port 10-Gigabit Ethernet LAN/WAN PIC Overview on page 3](#)
- [Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription on page 13](#)
- [Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC on page 14](#)
- *Ethernet Interfaces*

CHAPTER 3

Network Interfaces Configuration Statements and Hierarchy

- [\[edit interfaces\] Hierarchy Level on page 19](#)

[\[edit interfaces\] 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 {
    account-layer2-overhead (Interface Level) {
      value;
      egress bytes;
      ingress bytes;
    }
    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;
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;
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);
```

23

```
    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;
optics-options {
    alarm low-light-alarm {
        (link-down | syslog);
    }
    tx-power dbm;
    warning low-light-warning {
        (link-down | syslog);
    }
    wavelength nm;
}
otn-options {
    bytes transmit-payload-type value;
    fec (efec | gfec | gfec-sdfec | none);
    (is-ma | no-is-ma);
    (laser-enable | no-laser-enable);
    (line-loopback | no-line-loopback);
    (local-loopback | no-local-loopback);
}

```

```
(odu-ttim-action-enable | no-odu-ttim-action-enable);
(otu-ttim-action-enable | no-otu-ttim-action-enable);
odu-delay-management {
    (bypass | no-bypass);
    (monitor-end-point | no-monitor-end-point);
    (number-of-frames | no-number-of-frames);
    (start-measurement | no-start-measurement);
}
(prbs | no-prbs);
preemptive-fast-reroute {
    (backward-frr-enable | no-backward-frr-enable);
    (signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);
}
rate {
    (fixed-stuff-bytes | no-fixed-stuff-bytes);
    otu4;
    (pass-through | no-pass-through);
}
signal-degrade {
    ber-threshold-clear value;
    ber-threshold-signal-degrade value;
    interval value;
}
trigger trigger-identifier;
tti tti-identifier;
}
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;
            }
        }
    }
}
account-layer2-overhead {
    value;
    egress bytes;
    ingress bytes;
}
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;
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;
    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;
        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*

CHAPTER 4

Statement Summary

fast-aps-switch

Syntax	fast-aps-switch;
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options aps]
Release Information	Statement introduced in Junos OS Release 12.1.
Description	(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.
-

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

no-pre-classifier

Syntax	no-pre-classifier;
Hierarchy Level	[edit chassis fpc <i>n</i> pic <i>n</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify disabling the control queue for all ports on the 10-Gigabit Ethernet LAN/WAN PIC. Deleting this configuration re-enables the control queue feature on all ports of the 10-Gigabit Ethernet LAN/WAN PIC.
	<div> NOTE: For the 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PD-5-10XGE-SFPP), the control queue has a rate limiter to limit the control traffic to 2 Mbps (fixed, not user-configurable) per port. If the transit control traffic crosses this limit, then it can cause drops on locally terminating control traffic, causing flap of protocols such as BGP and OSPF. To avoid the control traffic being dropped, configure the no-pre-classifier statement to disable the control queue.</div>
Default	The no-pre-classifier statement is not configured and the control queue is operational.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• 10-port 10-Gigabit Ethernet LAN/WAN PIC Overview on page 3• Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC on page 14

pop

Syntax	pop;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces, 10-Gigabit Ethernet SFP interfaces, and 100-Gigabit Ethernet Type 5 PIC with CFP, specify the VLAN rewrite operation to remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Removing a VLAN Tag</i>

pop-pop

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

pop-swap

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

push

Syntax	<code>push;</code>
Hierarchy Level	<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>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Specify the VLAN rewrite operation to add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.</p> <p>You can use this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces, and 100-Gigabit Ethernet Type 5 PIC with CFP.</p> <p>If you include the push statement in the configuration, you must also include the pop statement at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]</code> hierarchy level.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Stacking a VLAN Tag</i>

push-push

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

swap

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

swap-swap

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

inner-tag-protocol-id

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

tag-protocol-id (TPID to Rewrite)

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

vlan-id (VLAN ID to Rewrite)

Syntax	<code>vlan-id <i>number</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<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 vlan-id statement with the swap statement, swap-push statement, push-push statement, or push-swap statement at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map] 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 vlan-id statement that you include at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Rewriting the VLAN Tag on Tagged Frames</i> • <i>Binding VLAN IDs to Logical Interfaces</i>

PART 3

Administration

- [Monitoring Commands on page 51](#)
- [Command Summary on page 77](#)

CHAPTER 5

Monitoring Commands

show interfaces (10-Gigabit Ethernet)

Syntax	<code>show interfaces <i>xe-fpc/pic/port</i></code> <code><brief detail extensive terse></code> <code><descriptions></code> <code><media></code> <code><snmp-index <i>snmp-index</i>></code> <code><statistics></code>
Release Information	Command introduced in Junos OS Release 8.0.
Description	(M320, M120, MX Series, and T Series routers and EX Series switches only) Display status information about the specified 10-Gigabit Ethernet interface.
Options	<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>
Required Privilege Level	view
List of Sample Output	<p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2) on page 67</p> <p>show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode) on page 70</p> <p>show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC) on page 72</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode) on page 74</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only) on page 74</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only) on page 75</p>
Output Fields	See Table 5 on page 53 for the output fields for the show interfaces (10-Gigabit Ethernet) command.

Table 5: show interfaces Gigabit Ethernet Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	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
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces.	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels

Table 5: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Link flags	Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Schedulers	(Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Input Rate	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
Output Rate	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
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Egress accounting overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	detail extensive
Ingress accounting overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	detail extensive

detail
extensive

Table 5: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. 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. • Output bytes—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. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <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 5 on page 53.</p>	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the <code>ignore-l3-incompletes</code> statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 5: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Egress queues	Total number of egress queues supported on the specified interface.	detail extensive
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive
Ingress queues	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	extensive
Queue counters (Ingress)	<p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	extensive

Table 5: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the 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 None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
OTN alarms	Active OTN alarms identified on the interface.	detail extensive
OTN defects	OTN defects received on the interface.	detail extensive
OTN FEC Mode	<p>The FECmode configured on the interface.</p> <ul style="list-style-type: none"> • efec—Enhanced forward error correction (EFEC) is configured to detect and correct bit errors. • gfec—G.709 Forward error correction (GFEC) mode is configured to detect and correct bit errors. • none—FEC mode is not configured. 	detail extensive
OTN Rate	<p>OTN mode.</p> <ul style="list-style-type: none"> • fixed-stuff-bytes—Fixed stuff bytes 11.0957 Gbps. • no-fixed-stuff-bytes—No fixed stuff bytes 11.0491 Gbps. • pass-through—Enable OTN passthrough mode. • no-pass-through—Do not enable OTN passthrough mode. 	detail extensive
OTN Line Loopback	Status of the line loopback, if configured for the DWDM OTN PIC. Its value can be: enabled or disabled .	detail extensive
OTN FEC statistics	<p>The forward error correction (FEC) counters for the DWDM OTN PIC.</p> <ul style="list-style-type: none"> • Corrected Errors—The count of corrected errors in the last second. • Corrected Error Ratio—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits. 	detail extensive
OTN FEC alarms	<p>OTN FEC excessive or degraded error alarms triggered on the interface.</p> <ul style="list-style-type: none"> • FEC Degrade—OTU FEC Degrade defect. • FEC Excessive—OTU FEC Excessive Error defect. 	detail extensive
OTN OC	<p>OTN OC defects triggered on the interface.</p> <ul style="list-style-type: none"> • LOS—OC Loss of Signal defect. • LOF—OC Loss of Frame defect. • LOM—OC Loss of Multiframe defect. • Wavelength Lock—OC Wavelength Lock defect. 	detail extensive

Table 5: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
OTN OTU	OTN OTU defects detected on the interface <ul style="list-style-type: none"> • AIS—OTN AIS alarm. • BDI—OTN OTU BDI alarm. • IAE—OTN OTU IAE alarm. • TTIM—OTN OTU TTIM alarm. • SF—OTN ODU bit error rate fault alarm. • SD—OTN ODU bit error rate defect alarm. • TCA-ES—OTN ODU ES threshold alarm. • TCA-SES—OTN ODU SES threshold alarm. • TCA-UAS—OTN ODU UAS threshold alarm. • TCA-BBE—OTN ODU BBE threshold alarm. • BIP—OTN ODU BIP threshold alarm. • BBE—OTN OTU BBE threshold alarm. • ES—OTN OTU ES threshold alarm. • SES—OTN OTU SES threshold alarm. • UAS—OTN OTU UAS threshold alarm. 	detail extensive
Received DAPI	Destination Access Port Interface (DAPI) from which the packets were received.	detail extensive
Received SAPI	Source Access Port Interface (SAPI) from which the packets were received.	detail extensive
Transmitted DAPI	Destination Access Port Interface (DAPI) to which the packets were transmitted.	detail extensive
Transmitted SAPI	Source Access Port Interface (SAPI) to which the packets were transmitted.	detail extensive
PCS statistics	(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"> • Bit errors—High bit error rate. Indicates the number of bit errors when the PCS receiver is operating in normal mode. • Errored blocks—Loss of block lock. The number of errored blocks when PCS receiver is operating in normal mode. 	detail extensive

Table 5: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 6 on page 67 • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—Number of frames that exceed 1518 octets. • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	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 5: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—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). • Input SA rejects—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. • Output packet count—Number of packets that the filter has given to the MAC hardware. • Output packet pad count—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. • Output packet error count—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. • CAM destination filters, CAM source filters—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. 	extensive
PMA PHY	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • PHY Lock—Phase-locked loop • PHY Light—Loss of optical signal 	extensive

Table 5: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
WIS section	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOL—Loss of light • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section) 	extensive
WIS line	<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"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. State other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line) 	extensive

Table 5: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
WIS path	<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"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path) 	extensive

Table 5: 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"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner status—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the attached Ethernet device, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is None. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Local resolution—Information from the link partner: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive). 	extensive
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"> • Destination slot—FPC slot number. 	extensive

Table 5: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> .	All levels

Table 5: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—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. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. • pop-swap—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. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none
Demux:	<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"> • Source Family Inet • Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—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. • Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the routing device.	extensive

Table 5: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: 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 Output bytes and Output packets 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.	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	detail extensive none
Input Filters	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.	detail extensive
Output Filters	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.	detail extensive
Mac-Validate Failures	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
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about address flag (possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interlace.	detail extensive none

Table 5: 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 6 on page 67](#) 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 6 on page 67](#), 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 6: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound physical interface	show interfaces ge-0/3/0 extensive	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	show interfaces ge-0/3/0.50 extensive	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	
Outbound physical interface	show interfaces ge-0/0/0 extensive	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	show interfaces ge-0/0/0.50 extensive	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, Runt: 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          Seconds
                    0
Errored blocks      0

```

```

MAC statistics:
Total octets          6970332384
Total packets        81050506
Unicast packets      81050000
Broadcast packets    506
Multicast packets    0
CRC/Align errors    0
FIFO errors          0
MAC control frames   0
MAC pause frames     0
Oversized frames     0
Jabber frames        0
Fragment frames      0
VLAN tagged frames   0
Code violations       0

Filter statistics:
Input packet count    81050506
Input packet rejects  506
Input DA rejects      0
Input SA rejects      0
Output packet count   0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0

Packet Forwarding Engine configuration:
Destination slot: 5

CoS information:
Direction : Output
CoS transmit queue   Bandwidth      Buffer Priority Limit
                        %      bps      %      usec
0 best-effort        95    950000000  95      0      low  none
3 network-control    5     50000000   5       0      low  none

Direction : Input
CoS transmit queue   Bandwidth      Buffer Priority Limit
                        %      bps      %      usec
0 best-effort        95    950000000  95      0      low  none
3 network-control    5     50000000   5       0      low  none

Logical interface xe-5/0/0.0 (Index 71) (SNMP ifIndex 95) (Generation 195)
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ] Encapsulation: ENET2
Egress accounting overhead: 100
Ingress accounting 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
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 accounting overhead: 100
Ingress accounting 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__

```

CHAPTER 6

Command Summary

- [Ethernet Interface Operational Mode Commands on page 77](#)

Ethernet Interface Operational Mode Commands

Table 7 on page 77 summarizes the command-line interface (CLI) commands that you can use to monitor and troubleshoot aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces. Commands are listed in alphabetical order.

Table 7: Ethernet Interface Operational Mode Commands

Task	Command
Clear dynamic VLAN interfaces.	<i>clear auto-configuration interfaces</i>
Clear a specified dynamic agent circuit identifier (ACI) interface set configured on the router. You can clear only those ACI interface sets that have no subscriber interface members.	<i>clear auto-configuration interfaces interface-set</i>
Clear Link Aggregation Control Protocol (LACP) statistics.	<i>clear lacp statistics</i>
Clear Link Aggregation Control Protocol (LACP) timeout entries.	<i>clear lacp timeouts</i>
Clear learned MAC addresses from the hardware and MAC database. Static MAC addresses are not cleared.	<i>clear interfaces mac-database</i>
Clear statistics that are collected for every MAC address, including policer statistics, on a given physical or logical interface.	<i>clear interfaces mac-database statistics</i>
Clear statistics that are collected for interface sets.	<i>clear interfaces interface-set statistics</i>
Clear the existing continuity measurement and restart counting the operational uptime.	<i>clear oam ethernet connectivity-fault-management continuity-measurement</i>

Table 7: Ethernet Interface Operational Mode Commands (*continued*)

Task	Command
Clear ITU-T Y.1731 Ethernet frame delay measurement (ETH-DM) delay statistics and ETH-DM frame counts. (MX Series routers)	<i>clear oam ethernet connectivity-fault-management delay-statistics</i>
Clear Operation, Administration, and Management (OAM) and connectivity fault management (CFM) linktrace database information.	<i>clear oam ethernet connectivity-fault-management linktrace path-database</i>
Clear all loss statistics maintained by CFM for a given maintenance domain and maintenance association.	<i>clear oam ethernet connectivity-fault-management loss-statistics</i>
Clear connectivity-fault-management policer statistics.	<i>clear oam ethernet connectivity-fault-management policer</i>
Clear the Ethernet OAM service-level agreement (SLA) iterator statistics.	<i>clear oam ethernet connectivity-fault-management sla-iterator-statistics</i>
Clear all statistics maintained by CFM. (Routers that support IEEE 802.1ag OAM CFM) In addition, for interfaces that support ITU-T Y.1731 Ethernet frame delay measurement (ETH-DM), also clear any ETH-DM statistics and frame counts for CFM maintenance association end points (MEPs).	<i>clear oam ethernet connectivity-fault-management statistics</i>
Clear ITU-T Y.1731 Ethernet synthetic loss measurement (ETH-SLM) delay statistics and ETH-SLM frame counts. (MX Series routers, Modular Port Concentrators only)	<i>clear oam ethernet connectivity-fault-management synthetic-loss-measurement</i>
Clear Operation, Administration, and Management (OAM) link fault management state information and restart the link discovery process on Ethernet interfaces.	<i>clear oam ethernet link-fault-management state</i>
Clear Operation, Administration, and Management (OAM) statistics link fault management statistics for Ethernet interfaces.	<i>clear oam ethernet link-fault-management statistics</i>
Clear the statistics for all Ethernet ring protection groups or a specific Ethernet ring protection group.	<i>clear protection-group ethernet-ring statistics</i>
Check the reachability of a remote IEEE 802.1ag OAM maintenance association end point (MEP) or maintenance association intermediate point (MIP).	<i>ping ethernet</i>

Table 7: Ethernet Interface Operational Mode Commands (*continued*)

Task	Command
Manually rebalance the subscribers on an aggregated Ethernet bundle with targeted distribution enabled.	<i>request interface rebalance (Aggregated Ethernet for Subscriber Management)</i>
Manually revert egress traffic from the designated backup link to the designated primary link of an aggregated Ethernet interface for which link protection is enabled, or manually switch egress traffic from the primary link to the backup link.	<i>request interface (revert switchover) (Aggregated Ethernet Link Protection)</i>
Force LACP link switchover.	<i>request lacp link-switchover</i>
Clear the lockout, force switch, manual switch, exercise, and wait-to-restore states.	<i>request protection-group ethernet-aps clear</i>
Test if APS is operating correctly.	<i>request protection-group ethernet-aps exercise</i>
Force traffic to switch from the active path to the alternate path.	<i>request protection-group ethernet-aps force-switch</i>
Lock the protection path, forcing the use of the working path.	<i>request protection-group ethernet-aps lockout</i>
Force traffic to switch from the active path to the alternate path.	<i>request protection-group ethernet-aps manual-switch</i>
Display status information about aggregated Fast Ethernet or Gigabit Ethernet router interfaces.	<i>show interfaces (Aggregated Ethernet)</i> <i>show interfaces (far-end-interval)</i>
Display status information about Fast Ethernet interfaces.	<i>show interfaces (Fast Ethernet)</i>
Display status information about the specified Gigabit Ethernet interface.	<i>show interfaces (Gigabit Ethernet)</i>
Display status information about 10-Gigabit Ethernet router interfaces.	<i>show interfaces (10-Gigabit Ethernet)</i>

Table 7: Ethernet Interface Operational Mode Commands (*continued*)

Task	Command
Display IPv6 interface statistics for IPv6 traffic traversing through the IQ2 and IQ2E PICs on standalone T640 routers and on T640 routers in a TX Matrix or in a TXP Matrix.	<i>show interfaces extensive</i>
Display IPv6 interface statistics for IPv6 traffic traversing through the IQ2 PICs on M10i and M120 routers.	
Display IPv6 interface statistics for IPv6 traffic traversing through the IQ2E PICs on M10i, M120, and M320 routers.	
Display information about Gigabit Ethernet or 10-Gigabit Ethernet router interface sets.	<i>show interfaces interface-set (Ethernet Interface Set)</i>
Display information about Gigabit Ethernet or 10-Gigabit Ethernet router interface set queues.	<i>show interfaces interface-set queue</i>
Display the transceiver temperature, laser bias current, laser output power, receive optical power, and related alarms for 10-Gigabit Ethernet dense wavelength-division multiplexing (DWDM) interfaces.	<i>show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet)</i>
Display information about integrated routing and bridging interfaces.	<i>show interfaces irb</i>
Display status information about the distribution of subscribers on different links in an aggregated Ethernet bundle.	<i>show interfaces targeting (Aggregated Ethernet for Subscriber Management)</i>
Display Link Aggregation Control Protocol (LACP) information for aggregated, Fast Ethernet, or Gigabit Ethernet router interfaces.	<i>show lacp interfaces</i>
Display Link Aggregation Control Protocol (LACP) statistics.	<i>show lacp statistics</i>
Display Link Aggregation Control Protocol timeout entries.	<i>show lacp timeouts</i>
Display MAC address information for Gigabit Ethernet router interfaces.	<i>show interfaces mac-database (Gigabit Ethernet)</i>
Display information on a specified interface that is part of a multichassis link aggregation configuration.	<i>show interfaces mc-ae</i>
Display ETH-DM statistics for CFM MEPs. (MX Series routers, Ethernet DPCs).	<i>show oam ethernet connectivity-fault-management delay-statistics</i>

Table 7: Ethernet Interface Operational Mode Commands (*continued*)

Task	Command
Display IEEE 802.1ag OAM connectivity fault management forwarding state information for Ethernet interfaces.	<i>show oam ethernet connectivity-fault-management forwarding-state</i>
Display OAM connectivity fault management information for Ethernet interfaces. For interfaces that support ETH-DM, also display any ETH-DM frame counts when the detail or extensive option is included. In all other cases, ETH-DM frame counts are zero.	<i>show oam ethernet connectivity-fault-management interfaces</i>
Display OAM connectivity fault management linktrace path database information.	<i>show oam ethernet connectivity-fault-management linktrace path-database</i>
Display OAM connectivity fault management maintenance association end point (MEP) database information. For interfaces that support ETH-DM, also display any ETH-DM frame counts. In all other cases, ETH-DM frame counts are zero.	<i>show oam ethernet connectivity-fault-management mep-database</i>
Display ETH-DM statistics and frame counts for CFM MEPs. (MX Series routers, Ethernet DPCs)	<i>show oam ethernet connectivity-fault-management mep-statistics</i>
Display ETH-LM statistics for on-demand mode only.	<i>show oam ethernet connectivity-fault-management loss-statistics</i>
Display information about maintenance intermediate points (MIPs) for the Ethernet OAM 802.1ag standard for connectivity fault management (CFM).	<i>show oam ethernet connectivity-fault-management mip</i>
Display OAM connectivity fault management path database information for hosts configured with MEP.	<i>show oam ethernet connectivity-fault-management path-database</i>
Displays connectivity-fault-management policer statistics.	<i>show oam ethernet connectivity-fault-management policer</i>
Display the Ethernet OAM service-level agreement (SLA) iterator statistics.	<i>show oam ethernet connectivity-fault-management sla-iterator-statistics</i>
Display ETH-SLM statistics for CFM MEPs (on-demand mode only). (MX Series routers, Ethernet MPCs)."	<i>show oam ethernet connectivity-fault-management synthetic-loss-statistics</i>

Table 7: Ethernet Interface Operational Mode Commands (*continued*)

Task	Command
Display OAM Ethernet Virtual Connection (EVC) information for hosts configured with Ethernet Local Management Interface (E-LMI). (MX series only)	<i>show oam ethernet evc</i>
Display OAM fault management statistics for Ethernet interfaces.	<i>show oam ethernet link-fault-management</i>
Display OAM Ethernet Local Management Interface status information for an LMI configured interface. (MX series only)	<i>show oam ethernet lmi</i>
Display OAM Ethernet Local Management Interface statistics for an LMI configured interface. (MX series only)	<i>show oam ethernet lmi statistics</i>
Display protection group Ethernet ring Automatic Protection Switching (APS).	<i>show protection-group ethernet-ring aps</i>
Display data channel information for all Ethernet ring protection groups or for a specific Ethernet ring protection group.	<i>show protection-group ethernet-ring data-channel</i>
Display protection group Ethernet ring interfaces.	<i>show protection-group ethernet-ring interface</i>
Display protection group Ethernet ring nodes.	<i>show protection-group ethernet-ring node-state</i>
Display protection group Ethernet ring statistics.	<i>show protection-group ethernet-ring statistics</i>
Display all data channel logical interfaces and the VLAN IDs controlled by a ring instance data channel.	<i>show protection-group ethernet-ring vlan</i>
Trace the path between two Ethernet OAM end points.	<i>traceroute ethernet</i>

PART 4

Troubleshooting

- [Ethernet on page 85](#)
- [Interface Diagnostics on page 89](#)

CHAPTER 7

Ethernet

traceroute ethernet

Syntax	traceroute ethernet (<i>mac-address</i> <i>mep-id</i>) maintenance-association <i>ma-name</i> maintenance-domain <i>md-name</i> ttl <i>value</i> < wait <i>seconds</i> >
Release Information	Command introduced in Junos OS Release 9.0. mep-id option introduced in Junos OS Release 9.1.
Description	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 show oam ethernet connectivity-fault-management path-database command. Before using the traceroute command, you can verify the remote MEP's MAC address using the show oam ethernet connectivity-fault-management path-database command.
Options	mac-address —Destination unicast MAC address of the remote maintenance point. mep-id —MEP identifier of the remote maintenance point. The range of values is 1 through 8191. maintenance-association <i>ma-name</i> —Specifies an existing maintenance association from the set of configured maintenance associations. maintenance-domain <i>md-name</i> —Specifies an existing maintenance domain from the set of configured maintenance domains. ttl <i>value</i> —Number of hops to use in the linktrace request. The range is 1 to 255 hops. The default is 4. wait <i>seconds</i> —(Optional) Maximum time to wait for a response to the traceroute request. The range is 1 to 255 seconds. The default is 5.
Required Privilege Level	network
List of Sample Output	traceroute ethernet on page 87
Output Fields	Table 8 on page 86 lists the output fields for the traceroute ethernet command. Output fields are listed in the approximate order in which they appear.

Table 8: 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 8: traceroute ethernet Output Fields (*continued*)

Field Name	Field Description
Maintenance Domain	Maintenance domain specified in the traceroute command.
Level	Maintenance domain level configured.
Maintenance Association	Maintenance association specified in the traceroute command.
Local Mep	The local maintenance end point identifier.
Transaction Identifier	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.
Hop	Sequential hop count of the linktrace path.
TTL	Number of hops remaining in the linktrace message. The time to live (TTL) is decremented at each hop.
Source MAC address	MAC address of the 802.1ag maintenance point that is sending the linktrace message.
Next-hop MAC address	MAC address of the 802.1ag node that is the next hop in the LTM path.

Sample Output

traceroute ethernet

```
user@host> traceroute ethernet maintenance-domain md1 maintenance-association ma1
00:90:69:7e:01:ff
```

```
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:bb:bb:bb:bb
2	62	00:00:bb:bb:bb:bb	00:00:cc:cc:cc:cc
3	61	00:00:cc:cc:cc:cc	00:01:02:03:04:05
4	60	00:01:02:03:04:05	00:00:00:00:00:00

CHAPTER 8

Interface Diagnostics

- [Interface Diagnostics on page 89](#)

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 89](#)
- [Interface Diagnostics on page 91](#)

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 9 on page 90 shows the loopback modes supported on the various interface types.

Table 9: 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 9: Loopback Modes by Interface Type (*continued*)

Interface	Loopback Modes	Usage Guidelines
T1 and T3	Local, payload, and remote	<i>Configuring T1 Loopback Capability</i> and <i>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 10 on page 94 shows the BERT capabilities for various interface types.

Table 10: 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"> Limited algorithms
4-port Channelized OC3/STM1 Circuit Emulation	Yes (port 0–3)		<ul style="list-style-type: none"> Limited algorithms
E1 or T1	Yes (port 0–3)	Yes (port 0–3)	<ul style="list-style-type: none"> Single port at a time Limited algorithms
E3 or T3	Yes (port 0–3)	Yes (port 0–3)	<ul style="list-style-type: none"> Single port at a time
Channelized OC12	N/A	Yes (channel 0–11)	<ul style="list-style-type: none"> Single channel at a time Limited algorithms No bit count
Channelized STM1	Yes (channel 0–62)	N/A	<ul style="list-style-type: none"> Multiple channels Only one algorithm No error insert No bit count
Channelized T3 and Multichannel T3	Yes (channel 0–27)	Yes (port 0–3 on channel 0)	<ul style="list-style-type: none"> Multiple ports and channels Limited algorithms for T1 No error insert for T1 No bit count for T1

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 *Junos OS Operational Mode Commands*.



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 99](#)

Index

Symbols

#, comments in configuration statements.....	xii
(), in syntax descriptions.....	xii
10-Gigabit Ethernet LAN/WAN PIC	
control queue disable.....	14
handling oversubscription.....	17
line rate mode.....	13
no-pre-classifier statement.....	38
oversubscribed Ethernet mode	
control queue disable.....	14, 38
oversubscribed mode.....	13
< >, in syntax descriptions.....	xii
[], in configuration statements.....	xii
{ }, in configuration statements.....	xii
(pipe), in syntax descriptions.....	xii

B

BERT	
configuring interface diagnostics.....	91
bert-algorithm statement	
usage guidelines.....	91
bert-error-rate statement	
usage guidelines.....	91
bert-period statement	
usage guidelines.....	91
bit error rate test See BERT	
braces, in configuration statements.....	xii
brackets	
angle, in syntax descriptions.....	xii
square, in configuration statements.....	xii

C

comments, in configuration statements.....	xii
conventions	
text and syntax.....	xi
curly braces, in configuration statements.....	xii
customer support.....	xiii
contacting JTAC.....	xiii

D

documentation	
comments on.....	xiii

E

Ethernet interfaces	
status information, displaying	
Gigabit Ethernet.....	52

F

fast-aps-switch statement.....	37
font conventions.....	xi

G

Gigabit Ethernet interfaces	
status information, displaying.....	52
Gigabit Ethernet IQ PIC	
traffic and MAC statistics.....	52

I

inner-tag-protocol-id statement.....	45
interfaces	
configuration statements.....	19

L

linktrace	
tracing.....	86
loopback testing.....	89

M

manuals	
comments on.....	xiii

P

parentheses, in syntax descriptions.....	xii
pop statement	
Gigabit Ethernet IQ interfaces.....	39
pop-pop statement	
Gigabit Ethernet IQ interfaces.....	39
pop-swap statement	
Gigabit Ethernet IQ interfaces.....	40
push statement	
Gigabit Ethernet IQ interfaces.....	41
push-push statement	
Gigabit Ethernet IQ interfaces.....	42

S

show interfaces (10-Gigabit Ethernet)	
command.....	52
support, technical See technical support	

swap statement	
Gigabit Ethernet IQ interfaces.....	43
swap-swap statement	
Gigabit Ethernet IQ interfaces.....	44
syntax conventions.....	xi

T

tag-protocol-id statement	
TPID to rewrite.....	46
technical support	
contacting JTAC.....	xiii
traceroute ethernet.....	86
traceroute ethernet command.....	86

V

vlan-id statement	
rewriting at ingress or egress.....	47