

Stacking and Rewriting Gigabit Ethernet VLAN Tags



Published: 2013-02-14

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.

Stacking and Rewriting Gigabit Ethernet VLAN Tags
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	xi
	Documentation and Release Notes	xi
	Supported Platforms	xi
	Using the Examples in This Manual	xi
	Merging a Full Example	xii
	Merging a Snippet	xii
	Documentation Conventions	xiii
	Documentation Feedback	xv
	Requesting Technical Support	xv
	Self-Help Online Tools and Resources	xv
	Opening a Case with JTAC	xvi
Part 1	Overview	
Chapter 1	Stacking and Rewriting Gigabit Ethernet VLAN Tags	3
	Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview	3
	Understanding Transparent Tag Operations and IEEE 802.1p Inheritance	4
	Understanding swap-by-poppush	6
Part 2	Configuration	
Chapter 2	Stacking and Rewriting Gigabit Ethernet VLAN Tags	9
	Stacking and Rewriting Gigabit Ethernet VLAN Tags	9
	Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames	12
	Configuring Stacked VLAN Tagging	13
	Configuring Dual VLAN Tags	13
	Configuring Inner and Outer TPIDs and VLAN IDs	14
	Stacking a VLAN Tag	17
	Removing a VLAN Tag	18
	Removing the Outer and Inner VLAN Tags	19
	Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag	19
	Stacking Two VLAN Tags	20
	Rewriting the VLAN Tag on Tagged Frames	21
	Rewriting a VLAN Tag on Untagged Frames	22
	Rewriting a VLAN Tag and Adding a New Tag	25
	Rewriting the Inner and Outer VLAN Tags	25
	Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags	26
Chapter 3	Network Interfaces Configuration Statements and Hierarchy	33
	[edit chassis] Hierarchy Level	33
	[edit interfaces] Hierarchy Level	34

	[edit logical-systems] Hierarchy Level	50
	[edit protocols connections] Hierarchy Level	55
	[edit protocols dot1x] Hierarchy Level	56
	[edit protocols iccp] Hierarchy Level	56
	[edit protocols lacp] Hierarchy Level	57
	[edit protocols oam] Hierarchy Level	57
	[edit protocols ppp] Hierarchy Level	59
	[edit protocols pppoe] Hierarchy Level	59
	[edit protocols protection-group] Hierarchy Level	60
	[edit protocols vrrp] Hierarchy Level	61
Chapter 4	Statement Summary	63
	aggregated-ether-options	64
	ethernet-switch-profile	66
	fast-aps-switch	68
	gether-options	69
	inner-tag-protocol-id	70
	inner-vlan-id	71
	input-vlan-map (Gigabit Ethernet IQ and 10-Gigabit Ethernet SFPP)	72
	input-vlan-map (Aggregated Ethernet)	73
	output-vlan-map (Gigabit Ethernet IQ and 10-Gigabit Ethernet with SFPP)	74
	output-vlan-map (Aggregated Ethernet)	75
	pop	75
	pop-pop	76
	pop-swap	76
	push	77
	push-push	77
	stacked-vlan-tagging	78
	swap	78
	swap-by-poppush	79
	swap-push	79
	swap-swap	80
	tag-protocol-id (TPIDs Expected to Be Sent or Received)	81
	tag-protocol-id (TPID to Rewrite)	82
	unit	83
	vlan-id (VLAN ID to Be Bound to a Logical Interface)	89
	vlan-id (VLAN ID to Rewrite)	90
	vlan-rewrite	91
	vlan-tags (Stacked VLAN Tags)	92
	vlan-tags	94
Part 3	Administration	
Chapter 5	Monitoring Commands	97
	show interfaces (Aggregated Ethernet)	98
	show interfaces (Fast Ethernet)	108
	show interfaces (10-Gigabit Ethernet)	125

Chapter 6	Command Summary	153
	Ethernet Interface Operational Mode Commands	153
	VRRP Operational Mode Commands	158
Part 4	Troubleshooting	
Chapter 7	Ethernet	161
	tracert ethernet	162
	Tracing PPPoE Operations	163
	Configuring the PPPoE Trace Log Filename	164
	Configuring the Number and Size of PPPoE Log Files	164
	Configuring Access to the PPPoE Log File	165
	Configuring a Regular Expression for PPPoE Lines to Be Logged	165
	Configuring the PPPoE Tracing Flags	165
	Troubleshooting PPPoE Service Name Tables	165
	Verifying a PPPoE Configuration	167
Chapter 8	Interface Diagnostics	169
	Interface Diagnostics	169
	Configuring Loopback Testing	169
	Interface Diagnostics	171
	Starting and Stopping a BERT Test	175
	Example: Configuring Bit Error Rate Testing	175
Part 5	Index	
	Index	179

List of Figures

Part 1	Overview
Chapter 1	Stacking and Rewriting Gigabit Ethernet VLAN Tags 3
	Figure 1: swap-push (transparent tag) 5
	Figure 2: swap-push (no transparent tag) 5
	Figure 3: push (transparent tag) 5
	Figure 4: push-push (transparent tag) 6

List of Tables

	About the Documentation	xi
	Table 1: Notice Icons	xiii
	Table 2: Text and Syntax Conventions	xiii
Part 1	Overview	
Chapter 1	Stacking and Rewriting Gigabit Ethernet VLAN Tags	3
	Table 3: VLAN Map Operation and IEEE 802.1p Inheritance	4
Part 2	Configuration	
Chapter 2	Stacking and Rewriting Gigabit Ethernet VLAN Tags	9
	Table 4: Rewrite Operations on Untagged, Single-Tagged, and Dual-Tagged Frames	10
	Table 5: Applying Rewrite Operations to VLAN Maps	11
	Table 6: Rewrite Operations and Statement Usage for Input VLAN Maps	15
	Table 7: Rewrite Operations and Statement Usage for Output VLAN Maps	16
	Table 8: Input VLAN Map Statements Allowed for ethernet-ccc and ethernet-vpls Encapsulations	23
	Table 9: Output VLAN Map Statements Allowed for ethernet-ccc and ethernet-vpls Encapsulations	23
	Table 10: Rules for Applying Rewrite Operations to VLAN Maps	23
Part 3	Administration	
Chapter 5	Monitoring Commands	97
	Table 11: show interfaces (Aggregated Ethernet) Output Fields	98
	Table 12: show interfaces Fast Ethernet Output Fields	108
	Table 13: show interfaces Gigabit Ethernet Output Fields	126
	Table 14: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type	140
Chapter 6	Command Summary	153
	Table 15: Ethernet Interface Operational Mode Commands	153
	Table 16: VRRP Operational Mode Commands	158
Part 4	Troubleshooting	
Chapter 7	Ethernet	161
	Table 17: traceroute ethernet Output Fields	162
Chapter 8	Interface Diagnostics	169

Table 18: Loopback Modes by Interface Type	170
Table 19: BERT Capabilities by Interface Type	174

About the Documentation

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

Documentation and Release Notes

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

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

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

Supported Platforms

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

- M Series
- MX Series
- T Series
- J Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

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

Table 1: Notice Icons

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

Table 2 on page xiii 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.	
J-Web GUI Conventions		
Bold text like this	Represents J-Web graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of J-Web selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. 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

- [Stacking and Rewriting Gigabit Ethernet VLAN Tags on page 3](#)

CHAPTER 1

Stacking and Rewriting Gigabit Ethernet VLAN Tags

- [Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 3](#)
- [Understanding Transparent Tag Operations and IEEE 802.1p Inheritance on page 4](#)
- [Understanding swap-by-poppush on page 6](#)

Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview

Stacking and rewriting VLAN tags allows you to use an additional (outer) VLAN tag to differentiate between customer edge (CE) routers that share one VLAN ID. A frame can be received on an interface, or it can be internal to the system (as a result of the **input-vlan-map** statement).

On IQ2 interfaces, 10-Gigabit Ethernet LAN/WAN PIC, 40-Gigabit Ethernet MIC, 100-Gigabit Ethernet MIC, IQ2-E interfaces, and MX Series interfaces, when a VLAN tag is pushed, the inner VLAN IEEE 802.1p bits are copied to the IEEE bits of the VLAN or VLANs being pushed. If the original packet is untagged, the IEEE bits of the VLAN or VLANs being pushed are set to 0.



NOTE: When swap-by-poppush is configured on the interface, when a VLAN tag is swapped, the inner VLAN IEEE 802.1p bits are copied to the IEEE bits of the VLAN being swapped. If swap-by-poppush is not configured on the interface, the VLAN IEEE 802.1p bits of the of the VLAN being swapped remains same.

You can stack and rewrite VLAN tags on the following interfaces:

- Gigabit Ethernet
- Gigabit Ethernet IQ
- 10-Gigabit Ethernet LAN/WAN PIC
- 40-Gigabit Ethernet MIC
- 100-Gigabit Ethernet MIC
- Gigabit Ethernet IQ2 and IQ2-E

- 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, and MX Series router Gigabit Ethernet Interfaces
- Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces with the VLAN encapsulation type configured to support Layer 2 tunneling protocols such as circuit cross-connect (CCC) or virtual private LAN service (VPLS) (as described in 802.1Q VLANs Overview)

**Related
Documentation**

- 802.1Q VLANs Overview
- [Stacking and Rewriting Gigabit Ethernet VLAN Tags on page 9](#)
- Junos® OS Ethernet Interfaces

Understanding Transparent Tag Operations and IEEE 802.1p Inheritance

When **swap-by-poppush** is configured on IQ2 interfaces, 10-Gigabit Ethernet LAN/WAN PIC, IQ2-E interfaces, and MX Series interfaces, during a swap operation, the inner VLAN IEEE 802.1p bits are copied to the IEEE bits of the tag being swapped. If swap-by-poppush is not configured on the interface, the VLAN IEEE 802.1p bits of the tag being swapped remains same.

When **swap-by-poppush** is configured but the incoming packet has no inner VLAN tag (transparent tag), the IEEE 802.1p bits are set to zero .

[Table 3 on page 4](#) describes the relationship between the VLAN map operation and the inheritance of IEEE 802.1p from the transparent tag. It assumes the presence of the transparent tag in the incoming packet. If the transparent tag is not present, the IEEE 802.1p value is set to 0.

Table 3: VLAN Map Operation and IEEE 802.1p Inheritance

Rewrite Operation	Untagged Logical Interface	Transparent tag IEEE 802.1p Inheritance	Single-tagged Logical Interface	Transparent tag IEEE 802.1p Inheritance	Change in number of tags
push-push	yes	OUTER, INNER	NA	no operation	+2
swap-push	NA	no operation	yes	OUTER, INNER	+1
push	yes	OUTER	yes	*none	+1
swap	NA	NA	yes	OUTER	0

NOTE: *In a **push** operation on a single-tagged logical interface, none of the tags (inner, or outer) inherit the IEEE 802.1p bits from the transparent tag.

The following section shows four different examples of the inheritance of the transparent IEEE 802.1p values into the outer and inner VLAN tags.

Figure 1 on page 5 shows an incoming packet with a transparent tag. A swap-push operation swaps the outer VLAN tag and pushes another VLAN tag. The IEEE 802.1p values are inherited from the transparent tag.

Figure 1: swap-push (transparent tag)

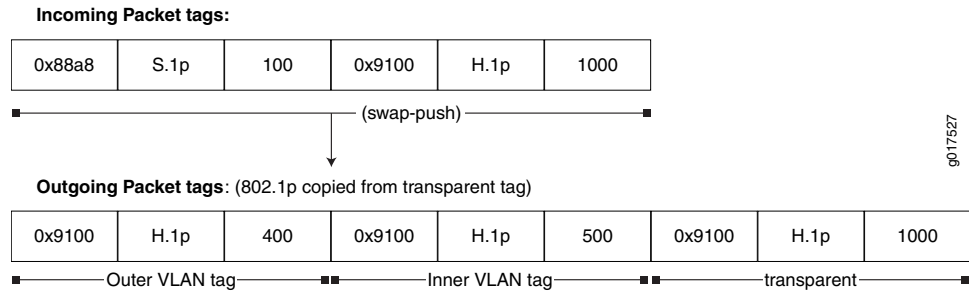


Figure 2 on page 5 shows an incoming packet with no transparent tag. A swap-push operation swaps the outer VLAN tag and pushes another VLAN tag. The IEEE 802.1p value is set to zero, as there is no transparent tag.

Figure 2: swap-push (no transparent tag)




Figure 3 on page 5 shows an incoming packet with a transparent tag. A push operation pushes another VLAN tag. The IEEE 802.1p value is inherited from the transparent tag.

Figure 3: push (transparent tag)

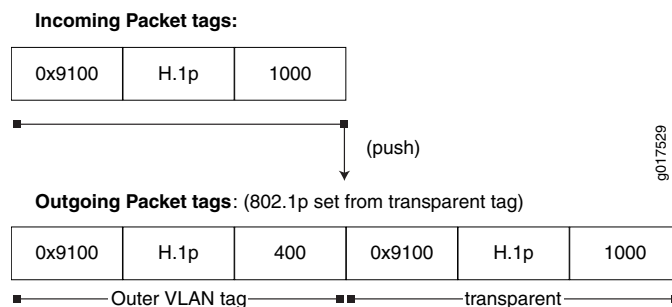
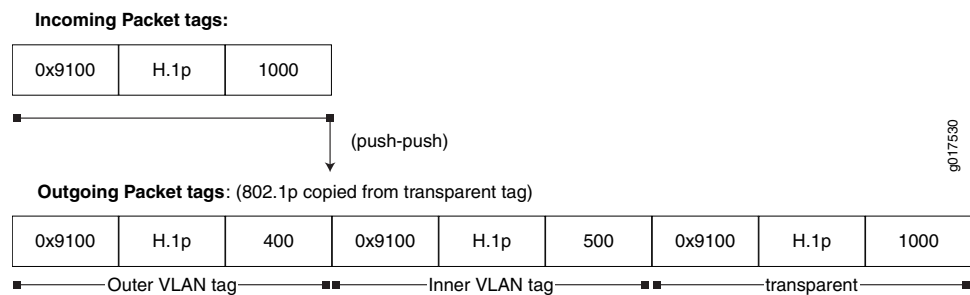


Figure 4 on page 6 shows an incoming packet with a transparent tag. A push-push operation pushes the outer and inner VLAN tags, respectively. The IEEE 802.1p values are inherited from the transparent tag.

Figure 4: push-push (transparent tag)



Related Documentation

- Understanding IEEE 802.1p Inheritance push and swap from a Transparent Tag
- Configuring IEEE 802.1p Inheritance push and swap from the Transparent Tag
- [Understanding swap-by-poppush on page 6](#)
- [swap-by-poppush on page 79](#)
- transparent

Understanding swap-by-poppush

By default, during a swap operation, the IEEE 802.1p bits of the VLAN tag remain unchanged. When the **swap-by-poppush** operation is enabled on a logical interface, the swap operation is treated as a **pop** operation followed by **push** operation. The **pop** operation removes the existing tag and the associated IEEE 802.1p bits and the push operation copies the inner VLAN IEEE 802.1p bits to the IEEE bits of the VLAN or VLANs being pushed. As a result, the IEEE 802.1p bits are inherited from the incoming transparent tag.

In effect, **swap-by-poppush** serves as a VLAN operation property and is used along with a **swap** or **swap-push** VLAN rewrite operation, indicating the nature of the swap operation being performed.

Related Documentation

- [swap-by-poppush on page 79](#)
- transparent
- Understanding IEEE 802.1p Inheritance push and swap from a Transparent Tag
- Configuring IEEE 802.1p Inheritance push and swap from the Transparent Tag
- [Understanding Transparent Tag Operations and IEEE 802.1p Inheritance on page 4](#)

PART 2

Configuration

- [Stacking and Rewriting Gigabit Ethernet VLAN Tags on page 9](#)
- [Network Interfaces Configuration Statements and Hierarchy on page 33](#)
- [Statement Summary on page 63](#)

CHAPTER 2

Stacking and Rewriting Gigabit Ethernet VLAN Tags

- [Stacking and Rewriting Gigabit Ethernet VLAN Tags on page 9](#)
- [Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames on page 12](#)
- [Configuring Stacked VLAN Tagging on page 13](#)
- [Configuring Dual VLAN Tags on page 13](#)
- [Configuring Inner and Outer TPIDs and VLAN IDs on page 14](#)
- [Stacking a VLAN Tag on page 17](#)
- [Removing a VLAN Tag on page 18](#)
- [Removing the Outer and Inner VLAN Tags on page 19](#)
- [Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag on page 19](#)
- [Stacking Two VLAN Tags on page 20](#)
- [Rewriting the VLAN Tag on Tagged Frames on page 21](#)
- [Rewriting a VLAN Tag on Untagged Frames on page 22](#)
- [Rewriting a VLAN Tag and Adding a New Tag on page 25](#)
- [Rewriting the Inner and Outer VLAN Tags on page 25](#)
- [Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags on page 26](#)

Stacking and Rewriting Gigabit Ethernet VLAN Tags

You can configure rewrite operations to stack (**push**), remove (**pop**), or rewrite (**swap**) tags on single-tagged frames and dual-tagged frames. If a port is not tagged, rewrite operations are not supported on any logical interface on that port.

You can configure the following VLAN rewrite operations:

- **pop**—Remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.
- **pop-pop**—For Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces, remove both the outer and inner VLAN tags of the frame.
- **pop-swap**—For Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces, 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.

- **push**—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.
- **push-push**—For Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces, push two VLAN tags in front of the frame.
- **swap-push**—For Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces, replace the outer VLAN tag of the frame with a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.
- **swap-swap**—For Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces, replace both the inner and the outer VLAN tags of the incoming frame with a user-specified VLAN tag value.

You configure VLAN rewrite operations for logical interfaces in the input VLAN map for incoming frames and in the output VLAN map for outgoing frames. To configure the input VLAN map, include the **input-vlan-map** statement:

```
input-vlan-map {
  ...interface-specific configuration...
}
```

To configure the output VLAN map, include the **output-vlan-map** statement:

```
output-vlan-map {
  ...interface-specific configuration...
}
```

You can include both statements at the following hierarchy levels:

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

The type of VLAN rewrite operation permitted depends upon whether the frame is single-tagged or dual-tagged. [Table 4 on page 10](#) shows supported rewrite operations and whether they can be applied to single-tagged frames or dual-tagged frames. The table also indicates the number of tags being added or removed during the operation.

Table 4: Rewrite Operations on Untagged, Single-Tagged, and Dual-Tagged Frames

Rewrite Operation	Untagged	Single-Tagged	Dual-Tagged	Number of Tags
pop	No	Yes	Yes	– 1
push	Sometimes	Yes	Yes	+1
swap	No	Yes	Yes	0

Table 4: Rewrite Operations on Untagged, Single-Tagged, and Dual-Tagged Frames (*continued*)

Rewrite Operation	Untagged	Single-Tagged	Dual-Tagged	Number of Tags
push-push	Sometimes	Yes	Yes	+2
swap-push	No	Yes	Yes	+1
swap-swap	No	No	Yes	0
pop-pop	No	No	Yes	-2
pop-swap	No	No	Yes	-1

The rewrite operations **push** and **push-push** can be valid in certain circumstances on frames that are not tagged. For example, a single-tagged logical interface (interface 1) and a dual-tagged logical interface (interface 2) have the following configurations:

```

Interface 1 [edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
  pop;
}
output-vlan-map {
  push;
}

Interface 2 [edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
  pop-pop;
}
output-vlan-map {
  push-push;
}

```

When a frame is received on the interface as a result of the **input-vlan-map** operation, the frame is not tagged. As it goes out of the second interface, the **output-vlan-map** operation **push-push** is applied to it. The resulting frame will be dual-tagged at the logical interface output.

Depending on the VLAN rewrite operation, you configure the rewrite operation for the interface in the input VLAN map, the output VLAN map, or in both the input VLAN map and the output VLAN map. [Table 5 on page 11](#) shows what rewrite operation combinations you can configure. “None” means that no rewrite operation is specified for the VLAN map.

Table 5: Applying Rewrite Operations to VLAN Maps

Input VLAN Map	Output VLAN Map								
	none	push	pop	swap	push-push	swap-push	swap-swap	pop-pop	swap-pop
none	Yes	No	No	Yes	No	No	Yes	No	No

Table 5: Applying Rewrite Operations to VLAN Maps (*continued*)

Input VLAN Map	Output VLAN Map								
	none	push	pop	swap	push-push	swap-push	swap-swap	pop-pop	swap-pop
push	No	No	Yes	No	No	No	No	No	No
pop	No	Yes	No	No	No	No	No	No	No
swap	Yes	No	No	Yes	No	No	No	No	No
push-push	No	No	No	No	No	No	No	Yes	No
swap-push	No	No	No	No	No	No	No	No	Yes
swap-swap	Yes	No	No	No	No	No	Yes	No	No
pop-pop	No	No	No	No	Yes	No	No	No	No
pop-swap	No	No	No	No	No	Yes	No	No	No

You must know whether the VLAN rewrite operation is valid and is applied to the input VLAN map or the output VLAN map. You must also know whether the rewrite operation requires you to include statements to configure the inner and outer TPIDs and inner and outer VLAN IDs in the input VLAN map or output VLAN map. For information about configuring inner and outer TPIDs and inner and outer VLAN IDs, see [“Configuring Inner and Outer TPIDs and VLAN IDs” on page 14](#).

Related Documentation

- [Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 3](#)
- [Understanding swap-by-poppush on page 6](#)
- [swap-by-poppush on page 79](#)
- Junos® OS Ethernet Interfaces

Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames

For Gigabit Ethernet IQ interfaces, aggregated Ethernet with Gigabit Ethernet IQ interfaces, Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, you can configure frames with particular TPIDs to be processed as tagged frames. To do this, you specify up to eight IEEE 802.1Q TPID values per port; a frame with any of the specified TPIDs is processed as a tagged frame; however, with IQ2 and IQ2-E interfaces, only the first four IEEE 802.1Q TPID values per port are supported. To configure the TPID values, include the **tag-protocol-id** statement:

```
tag-protocol-id [ tpids ];
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces interface-name gigether-options ethernet-switch-profile]`
- `[edit interfaces interface-name aggregated-ether-options ethernet-switch-profile]`

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

**Related
Documentation**

- [Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 3](#)
- [aggregated-ether-options on page 64](#)
- [ethernet-switch-profile on page 66](#)
- [gigether-options on page 69](#)
- [tag-protocol-id on page 81](#)
- Junos® OS Ethernet Interfaces

Configuring Stacked VLAN Tagging

To configure stacked VLAN tagging for all logical interfaces on a physical interface, include the **stacked-vlan-tagging** statement at the `[edit interfaces interface-name]` hierarchy level:

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

If you include the **stacked-vlan-tagging** statement in the configuration, you must configure dual VLAN tags for all logical interfaces on the physical interface. For more information, see “Stacking a VLAN Tag” on page 17.

**Related
Documentation**

- [stacked-vlan-tagging on page 78](#)
- [Stacking a VLAN Tag on page 17](#)
- Junos® OS Ethernet Interfaces

Configuring Dual VLAN Tags

To configure dual VLAN tags on a logical interface, include the **vlan-tags** statement:

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

You can include this statement at the following hierarchy levels:

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

The outer tag VLAN ID range is from 1 through 511 for normal interfaces, and from 512 through 4094 for VLAN CCC or VLAN VPLS interfaces. The inner tag is not restricted.

You must also include the **stacked-vlan-tagging** statement in the configuration. See [“Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags” on page 26](#).

**Related
Documentation**

- [unit on page 83](#)
- [Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags on page 26](#)
- [Junos® OS Ethernet Interfaces](#)

Configuring Inner and Outer TPIDs and VLAN IDs

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

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

inner-tag-protocol-id *tpid*;

You can include this statement at the following hierarchy levels:

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

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

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

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

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

You can include these statements at the following hierarchy levels:

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

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

All TPIDs you include in input and output VLAN maps must be among those you specify at the [edit interfaces *interface-name* **igether-options ethernet-switch-profile tag-protocol-id [*tpids*]**] hierarchy level or [edit interfaces *interface-name* **aggregated-ether-options ethernet-switch-profile tag-protocol-id [*tpids*]**] hierarchy level. For more information, see [“Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames” on page 12](#).

[Table 6 on page 15](#) and [Table 7 on page 16](#) specify when these statements are required. [Table 6 on page 15](#) indicates valid statement combinations for rewrite operations for the input VLAN map. “No” means the statement must not be included in the input VLAN map for the rewrite operation. “Optional” means the statement may be optionally specified for the rewrite operation in the input VLAN map. “Any” means that you must include the **vlan-id** statement, **tag-protocol-id** statement, **inner-vlan-id** statement, or **inner-tag-protocol-id** statement.

Table 6: Rewrite Operations and Statement Usage for Input VLAN Maps

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

Table 6: Rewrite Operations and Statement Usage for Input VLAN Maps (*continued*)

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

Table 7 on page 16 indicates valid statement combinations for rewrite operations for the output VLAN map. “No” means the statement must not be included in the output VLAN map for the rewrite operation. “Optional” means the statement may be optionally specified for the rewrite operation in the output VLAN map.

Table 7: Rewrite Operations and Statement Usage for Output VLAN Maps

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

The following examples use Table 6 on page 15 and Table 7 on page 16 and show how the **pop-swap** operation can be configured in an input VLAN map and an output VLAN map:

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

	}
Input VLAN Map with inner-tag-protocol-id Statement, Output VLAN Map with Optional inner-tag-protocol-id Statement	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] input-vlan-map { pop-swap; inner-tag-protocol-id <i>tpid</i>; } output-vlan-map { pop-swap; inner-tag-protocol-id <i>tpid</i>; }</pre>
Input VLAN Map with inner-tag-protocol-id and inner-vlan-id Statements	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] input-vlan-map { pop-swap; inner-vlan-id <i>number</i>; inner-tag-protocol-id <i>tpid</i>; }</pre>
Related Documentation	<ul style="list-style-type: none"> • inner-tag-protocol-id on page 70 • input-vlan-map on page 72 • output-vlan-map on page 74 • pop-swap on page 76 • unit on page 83 • Junos® OS Ethernet Interfaces

Stacking a VLAN Tag

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

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

You can include these statements at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number]`
- `[edit interfaces interface-name unit logical-unit-number]`

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

If you include the **push** statement in an interface's input VLAN map, see [Table 5 on page 11](#) for information about permissible rewrite operations,

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

All TPIDs you include in input and output VLAN maps must be among those you specify at the [edit interfaces *interface-name* **giether-options ethernet-switch-profile tag-protocol-id [*tpids*]**] hierarchy level. For more information, see [“Configuring Inner and Outer TPIDs and VLAN IDs” on page 14](#).

**Related
Documentation**

- [tag-protocol-id on page 82](#)
- [unit on page 83](#)
- [Table 5 on page 11](#)
- [802.1Q VLANs Overview](#)
- [Configuring Inner and Outer TPIDs and VLAN IDs on page 14](#)
- [Junos® OS Ethernet Interfaces](#)

Removing a VLAN Tag

To remove a VLAN tag from all tagged frames entering or exiting the interface, include the **pop** statement in the input VLAN map or output VLAN map:

pop;

You can include this statement at the following hierarchy levels:

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

**Related
Documentation**

- [input-vlan-map on page 72](#)
- [output-vlan-map on page 74](#)
- [pop on page 75](#)
- [unit on page 83](#)

- Junos® OS Ethernet Interfaces

Removing the Outer and Inner VLAN Tags

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to remove both the outer and inner VLAN tags of the frame, include the **pop-pop** statement in the input VLAN map or output VLAN map:

pop-pop;

You can include this statement at the following hierarchy levels:

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

See [Table 6 on page 15](#) and [Table 7 on page 16](#) for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Related Documentation

- [input-vlan-map on page 72](#)
- [output-vlan-map on page 74](#)
- [pop-pop on page 76](#)
- [unit on page 83](#)
- See [Table 6 on page 15](#) and [Table 7 on page 16](#) for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.
- Junos® OS Ethernet Interfaces

Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag

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

pop-swap;

The inner tag becomes the outer tag in the final frame.

You can include this statement at the following hierarchy levels:

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

See [Table 6 on page 15](#) and [Table 7 on page 16](#) for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

**Related
Documentation**

- [input-vlan-map on page 72](#)
- [output-vlan-map on page 74](#)
- [pop-swap on page 76](#)
- [unit on page 83](#)
- See [Table 6 on page 15](#) and [Table 7 on page 16](#) for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.
- Junos® OS Ethernet Interfaces

Stacking Two VLAN Tags

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to push two VLAN tags in front of tagged frames entering or exiting the interface, include the **push-push** statement in the input VLAN map or the output VLAN map:

push-push;

You can include this statement at the following hierarchy levels:

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

See [Table 6 on page 15](#) and [Table 7 on page 16](#) for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

**Related
Documentation**

- [input-vlan-map on page 72](#)
- [output-vlan-map on page 74](#)
- [pop on page 75](#)

- [unit on page 83](#)
- See [Table 6 on page 15](#) and [Table 7 on page 16](#) for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.
- Junos® OS Ethernet Interfaces

Rewriting the VLAN Tag on Tagged Frames

To rewrite the VLAN tag on all tagged frames entering the interface to a specified VLAN ID and TPID, include the **swap**, **tag-protocol-id**, and **vlan-id** statements in the input VLAN map:

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

To rewrite the VLAN tag on all tagged frames exiting the interface to a specified VLAN ID and TPID, include the **swap** and **tag-protocol-id** statements in the output VLAN map:

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

You can include these statements at the following hierarchy levels:

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

You cannot include both the **swap** statement and the **vlan-id** statement in the output VLAN map configuration. If you include the **swap** statement in the configuration, the VLAN ID in outgoing frames is rewritten to the VLAN ID bound to the logical interface. For more information about binding a VLAN ID to the logical interface, see [802.1Q VLANs Overview](#).

The swap operation works on the outer tag only, whether or not you include the **stacked-vlan-tagging** statement in the configuration. For more information, see [“Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags” on page 26](#).

Related Documentation

- [input-vlan-map on page 72](#)
- [output-vlan-map on page 74](#)
- [swap on page 78](#)
- [vlan-id on page 90](#)
- [tag-protocol-id on page 82](#)
- [unit on page 83](#)

- For more information about binding a VLAN ID to the logical interface, see [802.1Q VLANs Overview](#).
- For more information about the swap operation, see [Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags on page 26](#).
- Junos® OS Ethernet Interfaces

Rewriting a VLAN Tag on Untagged Frames

On M320, M120, and MX Series routers with Gigabit Ethernet IQ, IQ2, and IQ2E PICs, 10-Gigabit Ethernet IQ, IQ2, and IQ2E PICs, and on MX Series 40-port Gigabit Ethernet R, 40-port Gigabit Ethernet R EQ, 4-port 10-Gigabit Ethernet R, and 4-port 10-Gigabit Ethernet R EQ DPCs, you can rewrite VLAN tags on untagged incoming and outgoing frames under **ethernet-ccc** and **ethernet-vpls** encapsulations. On MX Series routers with IQ2 and IQ2-E PICs, you can perform all rewrite VLAN tag operations. These features provide added flexibility.

Consider a network where two provider edges (PE) are connected by a Layer 2 circuit. PE1 is receiving traffic on an untagged port while the corresponding port on PE2 is tagged. In the normal case, packets coming from PE1 will be dropped at PE2 because it is expecting tagged packets. However, if PE1 can push a VLAN tag on the incoming packet before sending it across to PE2, you can ensure that packets are not dropped. To make it work in both directions, PE1 must strip the VLAN tag from outgoing packets. Therefore, a push on the ingress side is always paired with a pop on the egress side.

The rewrite operations represented by the following statement options are supported under **ethernet-ccc** and **ethernet-vpls** encapsulations:

- **push**—A VLAN tag is added to the incoming untagged frame.
- **pop**—VLAN tag is removed from the outgoing frame.
- **push-push**—An outer and inner VLAN tag are added to the incoming untagged frame.
- **pop-pop**—Both the outer and inner VLAN tags of the outgoing frame are removed.

IQ2 and 10-Gigabit Ethernet PICs support all rewrite operations described above. Details on the possible combinations of usage are explained later in this section.



NOTE: The **push-push** and **pop-pop** operations are not supported on the Gigabit Ethernet IQ PIC.

For the **input-vlan-map** statement, only the **push** and **push-push** options are supported because it does not make sense to remove a VLAN tag from an incoming untagged frame. Similarly, only the **pop** and **pop-pop** options are supported for the **output-vlan-map** statement. Also, with the **push** and **push-push** options, the tag parameters have to be explicitly specified. Apart from this, the other rules for configuring the **input-vlan-map** and **output-vlan-map** statements are the same as for tagged frames. [Table 8 on page 23](#) through [Table 10 on page 23](#) explain the rules in more detail.

For the **input-vlan-map** statement, only the **push** and **push-push** options are supported because it does not make sense to remove a VLAN tag from an incoming untagged frame. Similarly, only the **pop** and **pop-pop** options are supported for the **output-vlan-map** statement. Also, with the **push** and **push-push** options, the **vlan-id** parameters (**vlan-id** for **push** and **vlan-id** or **inner-vlan-id** for **push-push**) have to be explicitly specified. TPID however, is optional and the default value of **0x8100** is set if not configured. Apart from this, the other rules for configuring the **input-vlan-map** and **output-vlan-map** statements are the same as for tagged frames.

Table 8: Input VLAN Map Statements Allowed for ethernet-ccc and ethernet-vpls Encapsulations

Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
push	Yes	Optional	No	Optional
push-push	Yes	Optional	Yes	Optional

Table 9: Output VLAN Map Statements Allowed for ethernet-ccc and ethernet-vpls Encapsulations

Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
pop	No	No	No	No
pop-pop	No	No	No	No

Table 10: Rules for Applying Rewrite Operations to VLAN Maps

Output VLAN Map			
Input VLAN Map	None	pop	pop-pop
None	Yes	No	No
push	No	Yes	No
push-push	No	No	Yes

**Example: push and pop
with Ethernet CCC
Encapsulation**

```
ge-3/1/0 {
  encapsulation ethernet-ccc;
  unit 0 {
    encapsulation ethernet-ccc;
    input-vlan-map {
      push;
      tag-protocol-id 0x8100;
      vlan-id 600;
    }
    output-vlan-map pop;
    family ccc;
  }
}
```

**Example: push-push
and pop-pop with
Ethernet CCC
Encapsulation**

```
ge-3/1/0 {
  encapsulation ethernet-ccc;
  unit 0 {
    encapsulation ethernet-ccc;
    input-vlan-map {
      push-push;
      tag-protocol-id 0x8100;
      inner-tag-protocol-id 0x8100;
      vlan-id 600;
      inner-vlan-id 575;
    }
    output-vlan-map pop-pop;
    family ccc;
  }
}
```

**Example: push and pop
with Ethernet VPLS
Encapsulation**

```
ge-3/1/0 {
  encapsulation ethernet-vpls;
  unit 0 {
    encapsulation ethernet-vpls;
    input-vlan-map {
      push;
      tag-protocol-id 0x8100;
      vlan-id 700;
    }
    output-vlan-map pop;
    family vpls;
  }
}
```

**Example: push-push
and pop-pop with
Ethernet VPLS
Encapsulation**

```
ge-3/1/0 {
  encapsulation ethernet-vpls;
  unit 0 {
    encapsulation ethernet-vpls;
    input-vlan-map {
      push-push;
      tag-protocol-id 0x8100;
      inner-tag-protocol-id 0x8100;
      vlan-id 600;
      inner-vlan-id 575;
    }
    output-vlan-map pop-pop;
    family vpls;
  }
}
```

You can use the **show interface *interface-name*** command to display the status of a modified VLAN map for the specified interface.

**Related
Documentation**

- [input-vlan-map on page 72](#)
- [output-vlan-map on page 74](#)
- [pop on page 75](#)
- [pop-pop on page 76](#)

- [push on page 77](#)
- [push-push on page 77](#)
- [unit on page 83](#)
- Junos® OS Ethernet Interfaces

Rewriting a VLAN Tag and Adding a New Tag

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

swap-push

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

You can include this statement at the following hierarchy levels:

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

See [Table 6 on page 15](#) and [Table 7 on page 16](#) for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Related Documentation

- [input-vlan-map on page 72](#)
- [output-vlan-map on page 74](#)
- [swap-push on page 79](#)
- [unit on page 83](#)
- Junos® OS Ethernet Interfaces

Rewriting the Inner and Outer VLAN Tags

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

`swap-swap;`

You can include this statement at the following hierarchy levels:

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

See [Table 6 on page 15](#) and [Table 7 on page 16](#) for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

**Related
Documentation**

- [input-vlan-map on page 72](#)
- [output-vlan-map on page 74](#)
- [swap-swap on page 80](#)
- [unit on page 83](#)
- Junos® OS Ethernet Interfaces

Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags

Configure a VLAN CCC tunnel in which Ethernet frames enter the tunnel at interface `ge-4/0/0` and exit the tunnel at interface `ge-4/2/0`.

The following examples show how to perform the following tasks:

- [Push a TPID and VLAN ID Pair on Ingress on page 27](#)
- [Stack Inner and Outer VLAN Tags on page 27](#)
- [Swap a VLAN ID on Ingress on page 28](#)
- [Swap a VLAN ID on Egress on page 28](#)
- [Swap a VLAN ID on Both Ingress and Egress on page 29](#)
- [Swap the Outer VLAN Tag and Push a New VLAN Tag on Ingress; Pop the Outer VLAN Tag and Swap the Inner VLAN Tag on Egress on page 30](#)
- [Swap a TPID and VLAN ID Pair for Both VLAN Tags on Ingress and on Egress on page 30](#)
- [Pop the Outer VLAN Tag and Swap the Inner VLAN Tag on Ingress; Swap the Outer VLAN Tag and Push a New VLAN Tag on Egress on page 31](#)
- [Pop a TPID and VLAN ID Pair on Ingress; Push a VLAN ID and TPID Pair on Egress on page 31](#)
- [Pop an Outer VLAN Tag to Connect an Untagged VPLS Interface to Tagged VPLS Interfaces on page 32](#)

Push a TPID and VLAN ID Pair on Ingress

```

[edit interfaces]
ge-4/0/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  gigether-options {
    ethernet-switch-profile {
      tag-protocol-id 0x9909;
    }
  }
}
unit 0 {
  encapsulation vlan-ccc;
  vlan-id 512;
  input-vlan-map {
    push;
    tag-protocol-id 0x9909;
    vlan-id 520;
  }
  output-vlan-map pop;
}
}
ge-4/2/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 515;
    input-vlan-map {
      swap-push;
      vlan-id 520;
      inner-vlan-id 512;
    }
    output-vlan-map {
      pop-swap;
    }
  }
}
[edit protocols]
mpls {
  interface ge-4/0/0.0;
  interface ge-4/2/0.0;
}
connections {
  interface-switch vlan-tag-push {
    interface ge-4/0/0.0;
    interface ge-4/2/0.0;
  }
}
}

```

Stack Inner and Outer VLAN Tags

```

[edit interfaces]
ge-0/2/0 {
  stacked-vlan-tagging;
  mac 00.01.02.03.04.05;
  gigether-options {
    loopback;
  }
}

```

```
        unit 0 {
            vlan-tags outer 0x8100.200 inner 0x8100.200;
        }
    }
}
```

**Swap a VLAN ID on
Ingress**

```
[edit interfaces]
ge-4/0/0 {
    vlan-tagging;
    encapsulation vlan-ccc;
    gigether-options {
        ethernet-switch-profile {
            tag-protocol-id 0x9100;
        }
    }
}
...
unit 1 {
    encapsulation vlan-ccc;
    vlan-id 1000;
    input-vlan-map {
        swap;
        tag-protocol-id 0x9100;
        vlan-id 2000;
    }
}
}
ge-4/2/0 {
    vlan-tagging;
    encapsulation vlan-ccc;
    ...
    unit 1 {
        encapsulation vlan-ccc;
        vlan-id 2000;
        input-vlan-map {
            swap;
            tag-protocol-id 0x9100;
            vlan-id 1000;
        }
    }
}
[edit protocols]
mpls {
    ...
    interface ge-4/0/0.1;
    interface ge-4/2/0.1;
}
connections {
    ...
    interface-switch vlan-tag-swap {
        interface ge-4/2/0.1;
        interface ge-4/0/0.1;
    }
}
}
```

**Swap a VLAN ID on
Egress**

```
[edit interfaces]
ge-4/0/0 {
```

```

vlan-tagging;
encapsulation vlan-ccc;
...
unit 1 {
    encapsulation vlan-ccc;
    vlan-id 1000;
}
}
ge-4/2/0 {
    vlan-tagging;
    encapsulation vlan-ccc;
    gigether-options {
        ethernet-switch-profile {
            tag-protocol-id 0x8800;
        }
    }
}
...
unit 1 {
    encapsulation vlan-ccc;
    vlan-id 2000;
    output-vlan-map {
        swap;
        tag-protocol-id 0x8800;
    }
}
}
[edit protocols]
mpls {
    ...
    interface ge-4/0/0.1;
    interface ge-4/2/0.1;
}
connections {
    ...
    interface-switch vlan-tag-swap {
        interface ge-4/2/0.1;
        interface ge-4/0/0.1;
    }
}
}

```

**Swap a VLAN ID on
Both Ingress and
Egress**

```

[edit interfaces]
ge-4/0/0 {
    vlan-tagging;
    encapsulation vlan-ccc;
    gigether-options {
        ethernet-switch-profile {
            tag-protocol-id [ 0x8800 0x9100 ];
        }
    }
}
...
unit 1 {
    encapsulation vlan-ccc;
    vlan-id 1000;
    input-vlan-map {
        swap;
        tag-protocol-id 0x9100;
    }
}
}

```

```
        vlan-id 2000;
    }
}
ge-4/2/0 {
    vlan-tagging;
    encapsulation vlan-ccc;
    gigether-options {
        ethernet-switch-profile {
            tag-protocol-id [ 0x8800 0x9100 ];
        }
    }
    unit 1 {
        encapsulation vlan-ccc;
        vlan-id 2000;
        output-vlan-map {
            swap;
            tag-protocol-id 0x8800;
        }
    }
}
[edit protocols]
mpls {
    ...
    interface ge-4/0/0.1;
    interface ge-4/2/0.1;
}
connections {
    ...
    interface-switch vlan-tag-swap {
        interface ge-4/2/0.1;
        interface ge-4/0/0.1;
    }
}
```

**Swap the Outer VLAN
Tag and Push a New
VLAN Tag on Ingress;
Pop the Outer VLAN
Tag and Swap the
Inner VLAN Tag on
Egress**

```
[edit interfaces]
ge-1/1/0 {
    unit 1 {
        vlan-id 200;
        input-vlan-map {
            swap-push;
            tag-protocol-id 0x9100;
            vlan-id 400;
            inner-tag-protocol-id 0x9100;
            inner-vlan-id 500;
        }
        output-vlan-map {
            pop-swap;
            inner-tag-protocol-id 0x9100;
        }
    }
}
```

**Swap a TPID and VLAN
ID Pair for Both VLAN**

```
[edit interfaces]
ge-1/1/0 {
```

Tags on Ingress and on
Egress

```
unit 0 {
  vlan-tags {
    inner 0x9100.425;
    outer 0x9200.525;
  }
  input-vlan-map {
    swap-swap;
    tag-protocol-id 0x9100;
    vlan-id 400;
    inner-tag-protocol-id 0x9100;
    inner-vlan-id 500;
  }
  output-vlan-map {
    swap-swap;
    tag-protocol-id 0x9200;
    inner-tag-protocol-id 0x9100;
  }
}
```

Pop the Outer VLAN
Tag and Swap the
Inner VLAN Tag on
Ingress; Swap the
Outer VLAN Tag and
Push a New VLAN Tag
on Egress

```
[edit interfaces]
ge-1/1/0 {
  unit 0 {
    vlan-tags {
      inner 0x9100.425;
      outer 0x9200.525;
    }
    input-vlan-map {
      pop-swap;
      tag-protocol-id 0x9100;
      vlan-id 400;
    }
    output-vlan-map {
      swap-push;
      tag-protocol-id 0x9200;
      inner-tag-protocol-id 0x9100;
    }
  }
}
```

Pop a TPID and VLAN
ID Pair on Ingress; Push
a VLAN ID and TPID
Pair on Egress

```
[edit interfaces]
ge-1/1/0 {
  unit 0 {
    vlan-tags {
      inner 0x9100.425;
      outer 0x9200.525;
    }
    input-vlan-map {
      pop-pop;
    }
    output-vlan-map {
      push-push;
      tag-protocol-id 0x9200;
      inner-tag-protocol-id 0x9100;
    }
  }
}
```

```
    }  
  }  
  
Pop an Outer VLAN Tag to Connect an  
Untagged VPLS Interface to Tagged  
VPLS Interfaces  
[edit interfaces]  
ge-1/1/0 {  
  vlan-tagging;  
  encapsulation extended-vlan-vpls;  
  unit 0 {  
    vlan-id 0;  
    input-vlan-map {  
      push;  
      vlan-id 0;  
    }  
    output-vlan-map pop;  
    family vpls;  
  }  
}
```

**Related
Documentation**

- [input-vlan-map on page 72](#)
- [output-vlan-map on page 74](#)
- [inner-tag-protocol-id on page 70](#)
- [inner-vlan-id on page 71](#)
- [pop on page 75](#)
- [pop-pop on page 76](#)
- [pop-swap on page 76](#)
- [push on page 77](#)
- [push-push on page 77](#)
- [swap on page 78](#)
- [swap-push on page 79](#)
- [swap-swap on page 80](#)
- [unit on page 83](#)
- Junos® OS Ethernet Interfaces

CHAPTER 3

Network Interfaces Configuration Statements and Hierarchy

- [\[edit chassis\] Hierarchy Level on page 33](#)
- [\[edit interfaces\] Hierarchy Level on page 34](#)
- [\[edit logical-systems\] Hierarchy Level on page 50](#)
- [\[edit protocols connections\] Hierarchy Level on page 55](#)
- [\[edit protocols dot1x\] Hierarchy Level on page 56](#)
- [\[edit protocols iccp\] Hierarchy Level on page 56](#)
- [\[edit protocols lacp\] Hierarchy Level on page 57](#)
- [\[edit protocols oam\] Hierarchy Level on page 57](#)
- [\[edit protocols ppp\] Hierarchy Level on page 59](#)
- [\[edit protocols pppoe\] Hierarchy Level on page 59](#)
- [\[edit protocols protection-group\] Hierarchy Level on page 60](#)
- [\[edit protocols vrrp\] Hierarchy Level on page 61](#)

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

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
    }
    sonet {
      device-count number;
    }
  }
  channel-group number {
    ethernet {
      device-count number;
    }
    fpc slot-number {
      pic pic-number {
        adaptive-services {
          service-package (layer-2 | layer-3);
        }
      }
    }
  }
}
```

```

aggregate-ports;
atm-cell-relay-accumulation;
atm-l2circuit-mode (aal5 | cell | trunk trunk);
ce1 {
    e1 link-number {
        channel-group group-number;
        timeslots time-slot-range;
    }
}
channelization;
ct1 {
    t1 link-number {
        channel-group group-number;
        timeslots time-slot-range;
    }
}
ct3 {
    port port-number {
        t1 link-number {
            channel-group group-number;
            timeslots time-slot-range;
        }
    }
    framing sdh;
}
max-queues-per-interface number;
mlfr-uni-nni-bundles num-intf;
no-concatenate;
shdsl {
    pic-mode (1-port-atm | 2-port-atm);
}
vtmapping (klm | itu-t);
}
}
fpc slot-number{
pic pic-number{
    egress-policer-overhead bytes;
    ingress-policer-overhead bytes;
}
}
}

```

[edit interfaces] Hierarchy Level

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



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

```

interfaces {
  traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable |
      no-world-readable> ;
    flag flag <disable>;
  }
  interface-name {
    accounting-profile name;
    aggregated-ether-options {
      (flow-control | no-flow-control);
      lacp {
        (active | passive);
        link-protection {
          disable;
        }
        (revertive | non-revertive);
        periodic interval;
        system-priority priority;
      }
      link-protection;
      link-speed speed;
      (loopback | no-loopback);
      mc-ae {
        chassis-id chassis-id;
        mc-ae-id mc-ae-id;
        mode (active-active | active-standby);
        redundancy-group group-id;
        status-control (active | standby);
      }
      minimum-links number;
      source-address-filter {
        mac-address;
      }
      (source-filtering | no-source-filtering);
    }
    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);
}
flexible-vlan-tagging;
gigether-options {
  802.3ad aex;
  (asynchronous-notification | no-asynchronous-notification);
  (auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
  local-interface-offline>;
  auto-reconnect seconds;
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  (loopback | no-loopback);
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  no-auto-mdix;
  source-address-filter {
    mac-address;
  }
  (source-filtering | no-source-filtering);
ethernet-switch-profile {
  (mac-learn-enable | no-mac-learn-enable);
  tag-protocol-id [ tpids ];
  ethernet-policer-profile {
    input-priority-map {
      ieee802.1p premium [ values ];
    }
    output-priority-map {
      classifier {
        premium {
          forwarding-class class-name {
            loss-priority (high | low);
          }
        }
      }
    }
  }
  policer cos-policer-name {
    aggregate {
      bandwidth-limit bps;
    }
  }
}
```



```
lmi-type (ansi | itu | c-lmi);
n391dte number;
n392dce number;
n392dte number;
n393dce number;
n393dte number;
t391dte seconds;
t392dce seconds;
}
lsq-failure-options {
  no-termination-request;
  [ trigger-link-failure interface-name ];
}
mac mac-address;
mlfr-uni-nni-bundle-options {
  acknowledge-retries number;
  acknowledge-timer milliseconds;
  action-red-differential-delay (disable-tx | remove-link);
  drop-timeout milliseconds;
  fragment-threshold bytes;
  cisco-interoperability send-lip-remove-link-for-link-reject;
  hello-timer milliseconds;
  link-layer-overhead percent;
  lmi-type (ansi | itu | c-lmi);
  minimum-links number;
  mrru bytes;
  n391 number;
  n392 number;
  n393 number;
  red-differential-delay milliseconds;
  t391 seconds;
  t392 seconds;
  yellow-differential-delay milliseconds;
}
modem-options {
  dialin (console | routable);
  init-command-string initialization-command-string;
}
mtu bytes;
multi-chassis-protection {
  peer a.b.c.d {
    interface interface-name;
  }
}
multiservice-options {
  (core-dump | no-core-dump);
  (syslog | no-syslog);
}
native-vlan-id number;
no-gratuitous-arp-request;
no-keepalives;
no-partition {
  interface-type type;
}
no-vpivci-swapping;
otn-options {
```



```

fec (efec | gfec | none);
(laser-enable | no-laser-enable);
(line-loopback | no-line-loopback);
pass-thru;
rate (fixed-stuff-bytes | no-fixed-stuff-bytes | pass-thru);
transmit-payload-type number;
trigger (oc-lof | oc-lom | oc-los | oc-wavelength-lock | odu-ais | odu-bbe-th | odu-bdi
        | odu-es-th | odu-lck | odu-oci | odu-sd | odu-ses-th | odu-ttim | odu-uas-th |
        opu-ptm | otu-ais | otu-bbe-th | otu-bdi | otu-es-th | otu-fec-deg | otu-fec-exe |
        otu-iae | otu-sd | otu-ses-th | otu-ttim | otu-uas-th);
tti;
}
optics-options {
    wavelength nm;
    alarm alarm-name {
        (syslog | link-down);
    }
    warning warning-name {
        (syslog | link-down);
    }
}
partition partition-number oc-slice oc-slice-range interface-type type;
timeslots time-slot-range;
passive-monitor-mode;
per-unit-scheduler;
ppp-options {
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
    dynamic-profile profile-name;
    no-termination-request;
    pap {
        access-profile name;
        local-name name;
        local-password password;
        compression;
    }
}
psn-vcipsn-vci-identifier;
psn-vpipsn-vpi-identifier;
receive-bucket {
    overflow (discard | tag);
    rate percentage;
    threshold bytes;
}
redundancy-options {
    priority sp-fpc/pic/port;
    secondary sp-fpc/pic/port;
    hot-standby;
}

```

```
}
satop-options {
    payload-size n;
}
schedulers number;
serial-options {
    clock-rate rate;
    clocking-mode (dce | internal | loop);
    control-polarity (negative | positive);
    cts-polarity (negative | positive);
    dcd-polarity (negative | positive);
    dce-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    dsr-polarity (negative | positive);
    dte-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    dtr-circuit (balanced | unbalanced);
    dtr-polarity (negative | positive);
    encoding (nrz | nrzi);
    indication-polarity (negative | positive);
    line-protocol protocol;
    loopback mode;
    rts-polarity (negative | positive);
    tm-polarity (negative | positive);
    transmit-clock invert;
}
services-options {
    inactivity-timeout seconds;
    open-timeout seconds;
    session-limit {
        maximum number;
        rate new-sessions-per-second;
    }
    syslog {
        host hostname {
            facility-override facility-name;
            log-prefix prefix-number;
            services priority-level;
        }
    }
}
```

```

    }
  }
}
shdsl-options {
  annex (annex-a | annex-b);
  line-rate line-rate;
  loopback (local | remote);
  snr-margin {
    current margin;
    snext margin;
  }
}
sonet-options {
  aggregate asx;
  aps {
    advertise-interval milliseconds;
    annex-b;
    authentication-key key;
    fast-aps-switch;
    force;
    hold-time milliseconds;
    lockout;
    neighbor address;
    paired-group group-name;
    preserve-interface;
    protect-circuit group-name;
    request;
    revert-time seconds;
    switching-mode (bidirectional | unidirectional);
    working-circuit group-name;
  }
  bytes {
    c2 value;
    e1-quiet value;
    f1 value;
    f2 value;
    s1 value;
    z3 value;
    z4 value;
  }
  fcs (16 | 32);
  loopback (local | remote);
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  path-trace trace-string;
  (payload-scrambler | no-payload-scrambler);
  rfc-2615;
  trigger {
    defect ignore;
    hold-time up milliseconds down milliseconds;
  }
  vtmapping (itu-t | klm);
  (z0-increment | no-z0-increment);

```

```
}
speed (10m | 100m | 1g | oc3 | oc12 | oc48);
stacked-vlan-tagging;
switch-options {
    switch-port port-number {
        (auto-negotiation | no-auto-negotiation);
        speed (10m | 100m | 1g);
        link-mode (full-duplex | half-duplex);
    }
}
t1-options {
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout value;
    byte-encoding (nx56 | nx64);
    crc-major-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5);
    crc-minor-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5 | 5e-6 | 1e-6);
    fcs (16 | 32);
    framing (esf | sf);
    idle-cycle-flag (flags | ones);
    invert-data;
    line-encoding (ami | b8zs);
    loopback (local | payload | remote);
    remote-loopback-respond;
    start-end-flag (filler | shared);
    timeslots time-slot-range;
}
t3-options {
    atm-encapsulation (direct | plcp);
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout feet;
    (cbit-parity | no-cbit-parity);
    compatibility-mode (adtran | digital-link | kentrox | larscom | verilink) <subrate
        value>;
    fcs (16 | 32);
    (feac-loop-respond | no-feac-loop-respond);
    idle-cycle-flag value;
    (long-buildout | no-long-buildout);
    (loop-timing | no-loop-timing);
    loopback (local | payload | remote);
    (mac | no-mac);
    (payload-scrambler | no-payload-scrambler);
    start-end-flag (filler | shared);
}
traceoptions {
    flag flag <flag-modifier> <disable>;
}
transmit-bucket {
    overflow discard;
    rate percentage;
    threshold bytes;
}
(traps | no-traps);
```

```

unidirectional;
vlan-tagging;
vlan-vci-tagging;
unit logical-unit-number {
  accept-source-mac {
    mac-address mac-address {
      policer {
        input cos-policer-name;
        output cos-policer-name;
      }
    }
  }
  accounting-profile name;
  advisory-options {
    downstream-rate rate;
    upstream-rate rate;
  }
  allow-any-vci;
  atm-scheduler-map (map-name | default);
  backup-options {
    interface interface-name;
  }
  bandwidth rate;
  cell-bundle-size cells;
  clear-dont-fragment-bit;
  compression {
    rtp {
      f-max-period number;
      maximum-contexts number <force>;
      queues [ queue-numbers ];
      port {
        minimum port-number;
        maximum port-number;
      }
    }
  }
  compression-device interface-name;
  copy-tos-to-outer-ip-header;
  demux-destination family;
  demux-source family;
  demux-options {
    underlying-interface interface-name;
  }
  description text;
  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 dlcid-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 dlcid-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-vsa-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](#)
 - [Junos® OS Ethernet Interfaces](#)
 - [Junos® OS Network Interfaces](#)

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

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

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

```

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

```

```

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

```

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

```

```

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

```

```
}
```

**Related
Documentation**

- *Junos OS Hierarchy and RFC Reference*
- Junos® OS Ethernet Interfaces
- Junos® OS Network Interfaces

[edit protocols connections] Hierarchy Level

The following statements can also be configured at the **[edit logical-systems *logical-system-name* protocols connections]** hierarchy level.

```
interface-switch connection-name {  
    interface interface-name.unit-number;  
    interface interface-name.unit-number;  
}
```

- Related Documentation**
- [Junos OS Hierarchy and RFC Reference](#)
 - [Junos® OS Ethernet Interfaces](#)
 - [Junos® OS Network Interfaces](#)

[edit protocols dot1x] Hierarchy Level

```
dot1x {  
  authenticator  
    authentication-profile-name access-profile-name;  
    interface interface-ids {  
      maximum-requests integer;  
      retries integer;  
      quiet-period seconds;  
      transmit-period seconds;  
      reauthentication (disable | interval seconds);  
      server-timeout seconds;  
      supplicant (single);  
      supplicant-timeout seconds;  
    }  
  }  
}
```

- Related Documentation**
- [Junos OS Hierarchy and RFC Reference](#)
 - [Junos® OS Ethernet Interfaces](#)
 - [Junos® OS Network Interfaces](#)

[edit protocols iccp] Hierarchy Level

```
iccp {  
  traceoptions;  
  local-ip-address ip address;  
  session-establishment-hold-time value;  
  authentication-key string;  
  peer ip-address {  
    local-ip-address ip address;  
    session-establishment-hold-time value;  
    authentication-key string;  
    redundancy-group-id-list redundancy-group-id-list;  
    liveness-detection;  
  }  
}
```

- Related Documentation**
- [Junos OS Hierarchy and RFC Reference](#)
 - [Junos® OS Ethernet Interfaces](#)
 - [Junos® OS Network Interfaces](#)

[edit protocols lacp] Hierarchy Level

```

traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <disable>;
}

```

- Related Documentation**
- *Junos OS Hierarchy and RFC Reference*
 - Junos® OS Ethernet Interfaces
 - Junos® OS Network Interfaces

[edit protocols oam] Hierarchy Level

```

ethernet {
  connectivity-fault-management {
    action-profile profile-name {
      default-actions {
        interface-down;
      }
      event {
        adjacency-loss;
        interface-status-tlv (down | lower-layer-down);
        port-status-tlv blocked;
        rdi;
      }
    }
  }
  linktrace {
    age (30m | 10m | 1m | 30s | 10s);
    path-database-size path-database-size;
  }
  maintenance-domain domain-name {
    bridge-domain name;
    routing-instance rl {
      bridge-domain name;
      instance vpls-instance;
      interface (ge | xe) fpc/pic/port.domain;
      level number;
      maintenance-association name{
        mep identifier {
          direction (up | down)
          interface (ge | xe) fpc/pic/port.domain (working | protect );
          auto-discovery;
          lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect |
            rem-err-xcon | xcon);
          priority number;
        }
      }
    }
    mip-half-function (none | default | explicit);
    name-format (character-string | none | dns | mac+2oct);
    short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
    protect-maintenance-association protect-ma-name;
    remote-maintenance-association remote-ma-name;
  }
}

```

```
continuity-check {
  hold-interval minutes;
  interval (10m | 10s | 1m | 1s | 100ms);
  loss-threshold number;
}
maintenance-association ma-name {
  mip-half-function (none | default | explicit);
  mep mep-id {
    auto-discovery;
    direction (up | down);
    interface interface-name (working | protect);
    priority number;
    remote-mep mep-id {
      action-profile profile-name;
      sla-iterator-profile profile-name {
        data-tlv-size bytes;
        iteration-count frames;
        priority priority-value;
      }
    }
  }
}
}
performance-monitoring {
  hardware-assisted-timestamping;
  sla-iterator-profiles {
    profile-name {
      disable;
      calculation-weight {
        delay delay-weight;
        delay-variation delay-variation-weight;
      }
      cycle-time milliseconds;
      iteration-period connections;
      measurement-type (loss | statistical-frame-loss | two-way-delay);
    }
  }
}
link-fault-management {
  action-profile profile-name {
    action {
      syslog;
      link-down;
      send-critical-event;
    }
    event {
      link-adjacency-loss;
      link-event-rate {
        frame-error count;
        frame-period count;
        frame-period-summary count;
        symbol-period count;
      }
      protocol-down;
    }
  }
}
```

```

}
interface interface-name {
  apply-action-profile profile-name;
  event-thresholds {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
  link-discovery (active | passive);
  negotiation-options {
    allow-remote-loopback;
    no-allow-link-events;
  }
  pdu-interval interval;
  pdu-threshold threshold-value;
  remote-loopback;
}
}
fnp {
  interval <100ms | 1s | 10s | 1m | 10m>;
  loss-threshold number
  interface interface name {
    domain-id domain-id
  }
}
}

```

- Related Documentation**
- *Junos OS Hierarchy and RFC Reference*
 - Junos® OS Ethernet Interfaces
 - Junos® OS Network Interfaces

[\[edit protocols ppp\] Hierarchy Level](#)

```

monitor-session (interface-name | all);
traceoptions {
  file filename <files number> <match regular-expression> <size size> <world-readable |
  no-world-readable> ;
  flag flag <disable>;
}

```

- Related Documentation**
- *Junos OS Hierarchy and RFC Reference*
 - Junos® OS Ethernet Interfaces
 - Junos® OS Network Interfaces

[\[edit protocols pppoe\] Hierarchy Level](#)

```

pppoe {
  no-send-pads-error;
  no-send-pads-ac-info
}

```

```
pado-advertise;
service-name-tables table-name {
  service service-name {
    drop;
    delay seconds;
    terminate;
    dynamic-profile profile-name;
    routing-instance routing-instance-name;
    max-sessions number;
    agent-specifier {
      aci circuit-id-string ari remote-id-string {
        drop;
        delay seconds;
        terminate;
        dynamic-profile profile-name;
        routing-instance routing-instance-name;
        static-interface interface-name;
      }
    }
  }
}
}
traceoptions {
  file <filename> <files number> <match regular-expression> <size maximum-file-size>
  <world-readable | no-world-readable>;
  filter {
    aci regular-expression;
    ari regular-expression;
    service-name regular-expression;
    underlying-interface interface-name;
  }
  flag flag;
  level (all | error | info | notice | verbose | warning);
  no-remote-trace;
}
}
```

[edit protocols protection-group] Hierarchy Level

```
ethernet-ring ring-name {
  east-interface {
    control-channel channel-name {
      vlan number;
    }
  }
  guard-interval number;
  node-id mac-address;
  restore-interval number;
  ring-protection-link-owner;
  west-interface {
    control-channel channel-name {
      vlan number;
    }
  }
}
```

- Related Documentation**
- *Junos OS Hierarchy and RFC Reference*
 - Junos® OS Ethernet Interfaces
 - Junos® OS Network Interfaces

[edit protocols vrrp] Hierarchy Level

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

```
protocols {
  vrrp {
    failover-delay milliseconds;
    global-advertisements-threshold advertisement-value;
    skew-timer-disable;
    startup-silent-period seconds;
    traceoptions {
      file <filename> <files number> <match regular-expression> <microsecond-stamp>
        <size maximum-file-size> <world-readable | no-world-readable>;
      flag flag;
      no-remote-trace;
    }
    version-3;
  }
}
```

- Related Documentation**
- Notational Conventions Used in Junos OS Configuration Hierarchies
 - [edit protocols] Hierarchy Level
 - *Junos OS Hierarchy and RFC Reference*
 - Junos® OS Ethernet Interfaces
 - Junos® OS Network Interfaces

CHAPTER 4

Statement Summary

aggregated-ether-options

```
Syntax aggregated-ether-options {
    ethernet-switch-profile {
        ethernet-policer-profile {
            input-priority-map {
                ieee802.1p premium [ values ];
            }
            output-priority-map {
                classifier {
                    premium {
                        forwarding-class class-name {
                            loss-priority (high | low);
                        }
                    }
                }
            }
            policer cos-policer-name {
                aggregate {
                    bandwidth-limit bps;
                    burst-size-limit bytes;
                }
                premium {
                    bandwidth-limit bps;
                    burst-size-limit bytes;
                }
            }
        }
        (mac-learn-enable | no-mac-learn-enable);
    }
    (flow-control | no-flow-control);
    lacp {
        (active | passive);
        link-protection {
            disable;
            (revertive | non-revertive);
            periodic interval;
            system-priority priority;
            system-id system-id;
        }
        link-protection;
        link-speed speed;
        logical-interface-fpc-redundancy;
        (loopback | no-loopback);
        minimum-links number;
        rebalance-periodic time hour:minute <interval hours>;
        source-address-filter {
            mac-address;
            (source-filtering | no-source-filtering);
        }
    }
}
```

Hierarchy Level [edit interfaces *aex*]

Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure aggregated Ethernet-specific interface properties. The statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Ethernet Interfaces Overview

ethernet-switch-profile

```
Syntax ethernet-switch-profile {
    ethernet-policer-profile {
        input-priority-map {
            ieee802.1p premium [ values ];
        }
        output-priority-map {
            classifier {
                premium {
                    forwarding-class class-name {
                        loss-priority (high | low);
                    }
                }
            }
        }
        policer cos-policer-name {
            aggregate {
                bandwidth-limit bps;
                burst-size-limit bytes;
            }
            premium {
                bandwidth-limit bps;
                burst-size-limit bytes;
            }
        }
        tag-protocol-id tpid;
    }
    (mac-learn-enable | no-mac-learn-enable);
}
```

Hierarchy Level [edit interfaces *interface-name* [gigether-options](#)],
[edit interfaces *interface-name* [aggregated-ether-options](#)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, aggregated Ethernet with Gigabit Ethernet IQ interfaces, and the built-in Gigabit Ethernet port on the M7i router), configure VLAN tag and MAC address accounting and filtering properties.

The statements are explained separately.




NOTE: When you gather interfaces into a bridge domain, the `no-mac-learn-enable` statement at the [edit interfaces *interface-name* [gigether-options](#) ethernet-switch-profile] hierarchy level is not supported. You must use the `no-mac-learning` statement at the [edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name*] hierarchy level to disable MAC learning on an interface in a bridge domain. For information

on disabling MAC learning for a bridge domain, see the *MX Series Layer 2 Configuration Guide*.

.....

Default	If the ethernet-switch-profile statement is not configured, Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router) behave like Gigabit Ethernet interfaces.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Gigabit Ethernet Policers• Configuring MAC Address Filtering• Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 3

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) Reduce the Automatic Protection Switching (APS) switchover time in Layer 2 circuits.
	<div><p>NOTE:</p><ul style="list-style-type: none">• 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.</div>
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">• Reducing APS Switchover Time in Layer 2 Circuits

gigether-options

```
Syntax  gigether-options {
        802.3ad {
            aex (primary | backup);
            lacp {
                port-priority;
            }
        }
        (asynchronous-notification | no-asynchronous-notification);
        (auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
        local-interface-offline>;
        (flow-control | no-flow-control);
        ignore-l3-incompletes;
        (loopback | no-loopback);
        mpls {
            pop-all-labels {
                required-depth number;
            }
        }
        no-auto-mdix
        source-address-filter {
            mac-address;
        }
        (source-filtering | no-source-filtering);
        speed
        ethernet-switch-profile {
            (mac-learn-enable | no-mac-learn-enable);
            tag-protocol-id [ tpids ];
            ethernet-policer-profile {
                input-priority-map {
                    ieee802.1p premium [ values ];
                }
                output-priority-map {
                    classifier {
                        premium {
                            forwarding-class class-name {
                                loss-priority (high | low);
                            }
                        }
                    }
                }
            }
            policer cos-policer-name {
                aggregate {
                    bandwidth-limit bps;
                    burst-size-limit bytes;
                }
                premium {
                    bandwidth-limit bps;
                    burst-size-limit bytes;
                }
            }
        }
    }
```

```
}
```

Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure Gigabit Ethernet specific interface properties. The statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Ethernet Interfaces Overview• gigether-options (ACX Series)

inner-tag-protocol-id

Syntax	inner-tag-protocol-id <i>tpid</i> ;
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.
Description	For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, configure the IEEE 802.1Q TPID value to rewrite for the inner tag. All TPIDs you include in input and output VLAN maps must be among those you specify at the [edit interfaces <i>interface-name</i> gigether-options ethernet-switch-profile tag-protocol-id [<i>tpids</i>]] hierarchy level.
Default	If the inner-tag-protocol-id 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">• Configuring Inner and Outer TPIDs and VLAN IDs on page 14

inner-vlan-id

Syntax	<code>inner-vlan-id <i>number</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map],</p> <p>[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	Statement introduced in Junos OS Release 8.1.
Description	<p>For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, specify the VLAN ID to rewrite for the inner tag of the final packet.</p> <p>You cannot include the inner-vlan-id statement with the swap statement, swap-push statement, push-push statement, or push-swap statement and the inner-vlan-id 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 inner-vlan-id statement you include at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.</p>
Options	<p><i>number</i>—VLAN ID number.</p> <p>Range: 0 through 4094</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"> • Configuring Inner and Outer TPIDs and VLAN IDs on page 14

input-vlan-map (Gigabit Ethernet IQ and 10-Gigabit Ethernet SFPP)

Syntax	<pre>input-vlan-map { (pop pop-pop pop-swap push push-push swap swap-push swap-swap); inner-tag-protocol-id <i>tpid</i>; inner-vlan-id <i>number</i>; tag-protocol-id <i>tpid</i>; vlan-id <i>number</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. pop-pop , pop-swap , push-push , swap-push , and swap-swap statements introduced in Junos OS Release 8.1.
Description	For Gigabit Ethernet IQ and 10-Gigabit Ethernet SFPP interfaces only, define the rewrite profile to be applied to incoming frames on this logical interface. The statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Stacking a VLAN Tag on page 17• output-vlan-map (Gigabit Ethernet IQ and 10-Gigabit Ethernet with SFPP) on page 74

input-vlan-map (Aggregated Ethernet)

Syntax	<pre>input-vlan-map { (pop push swap); tag-protocol-id <i>tpid</i>; vlan-id <i>number</i>; }</pre>
Hierarchy Level	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</pre>
Release Information	Statement introduced in Junos OS Release 8.2.
Description	<p>For aggregated Ethernet interfaces using Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces only, define the rewrite profile to be applied to incoming frames on this logical interface.</p> <p>The statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Stacking a VLAN Tag on page 17 • output-vlan-map (Aggregated Ethernet) on page 75

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

Syntax	<pre>output-vlan-map { (pop pop-pop pop-swap push push-push swap swap-push swap-swap); inner-tag-protocol-id <i>tpid</i>; inner-vlan-id <i>number</i>; tag-protocol-id <i>tpid</i>; vlan-id <i>number</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. pop-pop , pop-swap , push-push , swap-push , and swap-swap statements added in Junos OS Release 8.1.
Description	For Gigabit Ethernet IQ and 10-Port 10-Gigabit Ethernet SFPP interfaces only, define the rewrite operation to be applied to outgoing frames on this logical interface. The statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Stacking and Rewriting Gigabit Ethernet VLAN Tags on page 9• input-vlan-map (Gigabit Ethernet IQ and 10-Gigabit Ethernet SFPP) on page 72

output-vlan-map (Aggregated Ethernet)

Syntax	output-vlan-map { (pop push swap); tag-protocol-id <i>tpid</i> ; vlan-id <i>number</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 8.2.
Description	For aggregated Ethernet interfaces using Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces only, define the rewrite profile to be applied to outgoing frames on this logical interface. The statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Stacking and Rewriting Gigabit Ethernet VLAN Tags on page 9 • input-vlan-map (Aggregated Ethernet) on page 73

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.
Description	For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces, specify the VLAN rewrite operation to remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.
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"> • Removing a VLAN Tag on page 18

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.
Description	For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, 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">• Removing the Outer and Inner VLAN Tags on page 19

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.
Description	For Gigabit Ethernet IQ, IQ2, and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, 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.
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">• Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag on page 19

push

Syntax	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 before Junos OS Release 7.4.
Description	For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces, 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. If you include the push statement in the configuration, you must also include the pop statement at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map] hierarchy level.
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"> • Stacking a VLAN Tag on page 17

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.
Description	For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, specify the VLAN rewrite operation to push two VLAN tags in front 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"> • Stacking Two VLAN Tags on page 20

stacked-vlan-tagging

Syntax	stacked-vlan-tagging;
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Description	For Gigabit Ethernet IQ interfaces, enable stacked VLAN tagging for all logical interfaces on the physical interface. For pseudowire subscriber interfaces, enable stacked VLAN tagging for logical interfaces on the pseudowire service.
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">• Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 3• vlan-tags (Stacked VLAN Tags) on page 92

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.
Description	For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and aggregated Ethernet using Gigabit Ethernet IQ interfaces, 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.
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">• Rewriting the VLAN Tag on Tagged Frames on page 21

swap-by-poppush

Syntax	swap-by-poppush;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 11.2
Description	For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and aggregated Ethernet using Gigabit Ethernet IQ interfaces, specify the VLAN rewrite operation to replace a VLAN tag. Pop original tag, then push an entirely new tag. The swap operation is performed as pop followed by push.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

swap-push

Syntax	swap-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.
Description	For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, specify the VLAN rewrite operation to replace the outer VLAN tag of the frame with a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.
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"> • Rewriting a VLAN Tag and Adding a New Tag on page 25

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.
Description	For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, 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.
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">• Rewriting the Inner and Outer VLAN Tags on page 25

tag-protocol-id (TPIDs Expected to Be Sent or Received)

Syntax	<code>tag-protocol-id [<i>tpids</i>];</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> gigether-options ethernet-switch-profile], [edit interfaces <i>interface-name</i> aggregated-ether-options ethernet-switch-profile]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Description	<p>For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, aggregated Ethernet with Gigabit Ethernet IQ interfaces, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, and the built-in Gigabit Ethernet port on the M7i router), define the TPIDs expected to be sent or received on a particular VLAN. For each Gigabit Ethernet port, you can configure up to eight TPIDs using the tag-protocol-id statement; but only the first four TPIDs are supported on IQ2 and IQ2-E interfaces.</p> <p>For 10-Gigabit Ethernet LAN/WAN PIC interfaces on T Series routers, only the default TPID value (0x8100) is supported.</p>
Options	<i>tpids</i> —TPIDs to be accepted on the VLAN. Specify TPIDs in hexadecimal.
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"> • Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames on page 12

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.
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, value the default TPID value (0x8100) is supported.</p>
Default	If the tag-protocol-id 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">• Configuring Inner and Outer TPIDs and VLAN IDs on page 14

unit

```

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

```

```
incoming-map {
  caller caller-id | accept-all;
  initial-route-check seconds;
  load-interval seconds;
  load-threshold percent;
  pool pool-name;
  redial-delay time;
  watch-list {
    [ routes ];
  }
}
}
disable;
disable-mlppp-inner-ppp-pfc;
dlci dlci-identifier;
drop-timeout milliseconds;
dynamic-call-admission-control {
  activation-priority priority;
  bearer-bandwidth-limit kilobits-per-second;
}
encapsulation type;
epd-threshold cells plp1 cells;
family family-name {
  ... the family subhierarchy appears after the main [edit interfaces interface-name unit
    logical-unit-number] hierarchy ...
}
fragment-threshold bytes;
inner-vlan-id-range start start-id end end-id;
input-vlan-map {
  (pop | pop-pop | pop-swap | push | push-push | swap |
  swap-push | swap-swap);
  inner-tag-protocol-id tpid;
  inner-vlan-id number;
  tag-protocol-id tpid;
  vlan-id number;
}
interleave-fragments;
inverse-arp;
layer2-policer {
  input-policer policer-name;
  input-three-color policer-name;
  output-policer policer-name;
  output-three-color policer-name;
}
link-layer-overhead percent;
minimum-links number;
mrru bytes;
multicast-dlci dlci-identifier;
multicast-vci vpi-identifier.vci-identifier;
multilink-max-classes number;
multipoint;
oam-liveness {
  up-count cells;
  down-count cells;
}
oam-period (disable | seconds);
```

```

output-vlan-map {
  (pop | pop-pop | pop-swap | push | push-push | swap |
  swap-push | swap-swap);
  inner-tag-protocol-id tpid;
  inner-vlan-id number;
  tag-protocol-id tpid;
  vlan-id number;
}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
  chap {
    access-profile name;
    default-chap-secret name;
    local-name name;
    passive;
  }
  compression {
    acfc;
    pfc;
  }
  dynamic-profile profile-name;
  lcp-restart-timer milliseconds;
  loopback-clear-timer seconds;
  ncp-restart-timer milliseconds;
  pap {
    access-profile name;
    default-pap-password password;
    local-name name;
    local-password password;
    passive;
  }
}
pppoe-options {
  access-concentrator name;
  auto-reconnect seconds;
  (client | server);
  service-name name;
  underlying-interface interface-name;
}
pppoe-underlying-options {
  access-concentrator name;
  dynamic-profile profile-name;
  max-sessions number;
}
proxy-arp;
service-domain (inside | outside);
shaping {
  (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst length);
  queue-length number;
}
short-sequence;
targeted-distribution;

```

```
transmit-weight number;  
(traps | no-traps);  
trunk-bandwidth rate;  
trunk-id number;  
tunnel {  
    backup-destination address;  
    destination address;  
    key number;  
    routing-instance {  
        destination routing-instance-name;  
    }  
    source source-address;  
    ttl number;  
}  
vci vpi-identifier.vci-identifier;  
vci-range start start-vci end end-vci;  
vpi vpi-identifier;  
vlan-id number;  
vlan-id-range number-number;  
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;  
family family {  
    accounting {  
        destination-class-usage;  
        source-class-usage {  
            (input | output | input output);  
        }  
    }  
    access-concentrator name;  
    address address {  
        ... the address subhierarchy appears after the main [edit interfaces interface-name unit  
            logical-unit-number family family-name] hierarchy ...  
    }  
    bridge-domain-type (bvlan | svlan);  
    bundle interface-name;  
    core-facing;  
    demux-destination {  
        destination-prefix;  
    }  
    demux-source {  
        source-prefix;  
    }  
    duplicate-protection;  
    dynamic-profile profile-name;  
    filter {  
        group filter-group-number;  
        input filter-name;  
        input-list [ filter-names ];  
        output filter-name;  
        output-list [ filter-names ];  
    }  
    interface-mode (access | trunk);  
    ipsec-sa sa-name;  
    isid-list all-service-groups;  
    keep-address-and-control;  
    mac-validate (loose | strict);  
    max-sessions number;
```

```

mtu bytes;
multicast-only;
no-redirects;
policer {
    arp policer-template-name;
    input policer-template-name;
    output policer-template-name;
}
primary;
protocols [inet iso mpls];
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check {
    fail-filter filter-name
    mode loose;
}
sampling {
    input;
    output;
}
service {
    input {
        post-service-filter filter-name;
        service-set service-set-name <service-filter filter-name>;
    }
    output {
        service-set service-set-name <service-filter filter-name>;
    }
}
service-name-table table-name
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
translate-plp-control-word-de;
unnumbered-address interface-name destination address destination-profile profile-name;
vlan-id number;
vlan-id-list [number number-number];
address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    destination-profile name;
    eui-64;
    master-only;
    multipoint-destination address {
        dlci dlci-identifier;
        epd-threshold cells <plp1 cells>;
        inverse-arp;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (disable | seconds);
        shaping {

```

```

        (cbr rate | rtvbr burst length peak rate sustained rate | vbr burst length peak rate
         sustained rate);
        queue-length number;
    }
    vci vpi-identifier.vci-identifier;
}
preferred;
primary;
(vrrp-group | vrrp-inet6-group) group-number {
    (accept-data | no-accept-data);
    advertise-interval seconds;
    authentication-type authentication;
    authentication-key key;
    fast-interval milliseconds;
    (preempt | no-preempt) {
        hold-time seconds;
    }
    priority number;
    track {
        interface interface-name {
            bandwidth-threshold bits-per-second priority-cost number;
        }
        priority-hold-time seconds;
        route ip-address/prefix-length routing-instance instance-name priority-cost cost;
    }
    virtual-address [ addresses ];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-interface interface-name;
        active-group group-number;
    }
}
}
}
}

```

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

Release Information Statement introduced before Junos OS Release 7.4.

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

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

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

The remaining statements are explained separately.

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

- Related Documentation**
- Configuring Logical Interface Properties
 - Example: Configuring E-LINE and E-LAN Services for a PBB Network on MX Series Routers
 - Junos Services Interfaces Configuration Release 12.3

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

Syntax	<code>vlan-id <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For Fast Ethernet, Gigabit Ethernet, and Aggregated Ethernet interfaces only, bind a 802.1Q VLAN tag ID to a logical interface.
Options	<p><i>number</i>—A valid VLAN identifier.</p> <p>Range: For aggregated Ethernet, 4-port, 8-port, and 12-port Fast Ethernet PICs, and for management and internal Ethernet interfaces, 1 through 1023.</p> <p>For 48-port Fast Ethernet and Gigabit Ethernet PICs, 1 through 4094.</p> <p>VLAN ID 0 is reserved for tagging the priority of frames.</p>
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"> • Configuring Mixed Tagging

vlan-id (VLAN ID to Rewrite)

Syntax	<code>vlan-id number;</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.
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 <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]</code> hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the vlan-id statement that you include at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code> hierarchy level.</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">• Rewriting the VLAN Tag on Tagged Frames on page 21• Binding VLAN IDs to Logical Interfaces

vlan-rewrite

Syntax	vlan-rewrite translate (200 500 201 501)
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>number</i> family bridge interface-mode trunk]
Release Information	Statement introduced in Junos OS Release 9.4.
Description	Translates an incoming VLAN to a bridge-domain VLAN, corresponding counter translation at egress. Supports translation of VLAN 200 to VLAN 500 and VLAN 201 to VLAN 501. Other valid VLANs pass through without translation.
Options	<p>translate 200 500—Translates incoming packets with VLAN 200 to 500.</p> <p>translate 201 501—Translates incoming packets with VLAN 201 to 501.</p> <p>translate 202 502—Translates incoming packets with VLAN 202 to 502.</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">• Rewriting a VLAN Tag and Adding a New Tag on page 25

vlan-tags (Stacked VLAN Tags)

Syntax	<code>vlan-tags inner <i>tpid.vlan-id</i> inner-range <i>vid1—vid2</i> outer <i>tpid.vlan-id</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches.
Description	For Gigabit Ethernet IQ and IQE interfaces only, bind TPIDs and 802.1Q VLAN tag IDs to a logical interface.



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

Options	<p>inner <i>tpid.vlan-id</i>—A TPID and a valid VLAN identifier.</p> <p>Range: (most routers) For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.</p> <p>Range: (PTX Series) For VLAN ID, 0 through 4094.</p> <p>inner-range <i>vid1—vid2</i>—For MX Series routers with Enhanced IQ (IQE) PICs only; specify a range of VLAN IDs where <i>vid1</i> is the start of the range and <i>vid2</i> is the end of the range.</p> <p>Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.</p> <p>outer <i>tpid.vlan-id</i>—A TPID and a valid VLAN identifier.</p> <p>Range: (most routers) For VLAN ID, 1 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces. VLAN ID 0 is reserved for tagging the priority of frames.</p> <p>Range: (PTX Series) For VLAN ID, 0 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces.</p>
----------------	---



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

The following examples illustrate this further:

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


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

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

Related Documentation

- [Configuring Dual VLAN Tags on page 13](#)
- [Configuring Flexible VLAN Tagging on PTX Series Packet Transport Switches](#)
- [stacked-vlan-tagging on page 78](#)

vlan-tags

Syntax	<code>vlan-tags outer [<i>tpid</i>].<i>vlan-id</i> [inner [<i>tpid</i>].<i>vlan-id</i>];</code>
Hierarchy Level	<code>[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.5. VLAN demux interface support introduced in Junos OS Release 10.2.
Description	For Gigabit Ethernet IQ and IQE interfaces only, binds TPIDs and 802.1Q VLAN tag IDs to a logical interface. You must include the stacked-vlan-tagging statement at the <code>[edit interfaces <i>interface-name</i>]</code> hierarchy level.
<div> NOTE: The inner-range <i>vid1–vid2</i> option is supported on MX Series routers with IQE PICs only.</div>	
Options	<p>inner [<i>tpid</i>].<i>vlan-id</i>—A TPID (optional) and a valid VLAN identifier in the format <i>tpid.vlan-id</i>. When used in the dynamic-profiles hierarchy, specify the <code>\$junos-vlan-id</code> predefined variable to dynamically obtain the VLAN ID.</p> <p>Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.</p> <p>outer [<i>tpid</i>].<i>vlan-id</i>—A TPID (optional) and a valid VLAN identifier in the format <i>tpid.vlan-id</i>. When used in the dynamic-profiles hierarchy, specify the <code>\$junos-stacked-vlan-id</code> predefined variable.</p> <p>Range: For VLAN ID, 1 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces. VLAN ID 0 is reserved for tagging the priority of frames.</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">• Configuring Dual VLAN Tags on page 13• stacked-vlan-tagging on page 78

PART 3

Administration

- [Monitoring Commands on page 97](#)
- [Command Summary on page 153](#)

CHAPTER 5

Monitoring Commands

show interfaces (Aggregated Ethernet)

Syntax	<pre>show interfaces ae <i>number</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	(M Series, T Series, and MX Series routers and EX Series switches) Display status information about the specified aggregated Fast Ethernet or Gigabit Ethernet interface.
Options	<p>ae <i>number</i>—Display standard information about the specified aggregated Fast Ethernet or Gigabit Ethernet interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information about the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
List of Sample Output	<p>show interfaces (Aggregated Ethernet) on page 103</p> <p>show interfaces brief (Aggregated Ethernet) on page 103</p> <p>show interfaces detail (Aggregated Ethernet) on page 103</p> <p>show interfaces extensive (Aggregated Ethernet) on page 104</p> <p>show interfaces extensive (Aggregated Ethernet with VLAN Stacking) on page 105</p>
Output Fields	Table 11 on page 98 lists the output fields for the show interfaces (Aggregated Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 11: show interfaces (Aggregated Ethernet) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface and state of the interface.	All levels
Enabled	State of the physical interface. Possible values are described in the “Enabled Field” section under Common Output Fields Description.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	All levels

Table 11: show interfaces (Aggregated Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
Flow control	Flow control status: Enabled or Disabled .	All levels
Minimum links needed	Number of child links that must be operational for the aggregate interface to be operational.	All levels
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under Common Output Fields Description.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interfaces Flags" section under Common Output Fields Description.	All levels
Current address	Configured MAC address.	detail extensive
Hardware address	Hardware MAC address.	detail extensive
Last flapped	Date, time, and how long ago the interface went from down to up or from up to down. The format is Last flapped: year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive

Table 11: show interfaces (Aggregated Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number of and rate at which bytes and packets are received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes and rate, in bps, at which bytes are received on the interface. • Output bytes—Number of bytes and rate, in bps, at which bytes are transmitted on the interface. • Input packets—Number of packets and rate, in pps, at which packets are received on the interface. • Output packets—Number of packets and rate, in pps, at which packets are transmitted on the interface. 	detail extensive
Input errors	<p>Input errors on the interface:</p> <ul style="list-style-type: none"> • Errors—Sum of incoming frame aborts and frame check sequence (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 random early detection (RED) mechanism. • Framing errors—Number of packets received with an invalid FCS. • Runts—Number of frames received that are smaller than the runt threshold. • Giants—Number of frames received that are larger than the giant threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or were not of interest. Usually, this field reports protocols that Junos OS does not handle. • Resource errors—Sum of transmit drops. 	detail extensive
Output errors	<p>Output errors on the interface:</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), then the cable, the far-end system, or the PIC 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. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	detail extensive
IPv6 transit statistics	<p>Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive

Table 11: show interfaces (Aggregated Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Queue counters	CoS queue number and its associated user-configured forwarding class name. <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
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 of 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 Field” section under Common Output Fields Description.	All levels
VLAN-Tag	Tag Protocol Identifier (TPID) and VLAN identifier.	All levels
Demux	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: <ul style="list-style-type: none"> Source Family Inet Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
Statistics	Information about the number of packets, packets per second, number of bytes, and bytes per second on this aggregate interface. <ul style="list-style-type: none"> Bundle—Information about input and output bundle rates. Link—(detail and extensive only) Information about specific links in the aggregate, including link state and input and output rates. Marker Statistics—(detail and extensive only) Information about 802.3ad marker protocol statistics on the specified links. <ul style="list-style-type: none"> Marker Rx—Number of valid marker protocol data units (PDUs) received on this aggregation port. Resp Tx—Number of marker response PDUs transmitted on this aggregation port. Unknown Rx—Number of frames received that either carry the slow protocols Ethernet type value (43B.4) but contain an unknown PDU, or are addressed to the slow protocols group MAC address (43B.3) but do not carry the slow protocols Ethernet type. Illegal Rx—Number of frames received that carry the slow protocols Ethernet type value (43B.4) but contain a badly formed PDU or an illegal value of protocol subtype (43B.4). 	detail extensive none
protocol-family	Protocol family configured on the logical interface. Possible values are described in the “Protocol Field” section under Common Output Fields Description.	brief

Table 11: show interfaces (Aggregated Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Protocol	Protocol family configured on the logical interface. Possible values are described in the "Protocol Field" section under Common Output Fields Description.	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
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive
Flags	Information about protocol family flags. Possible values are described in the "Family Flags Field" section under Common Output Fields Description.	detail extensive none
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 address flags. Possible values are described in the "Addresses Flags" section under Common Output Fields Description.	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 interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces (Aggregated Ethernet)

```

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

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

```

show interfaces brief (Aggregated Ethernet)

```

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

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

```

show interfaces detail (Aggregated Ethernet)

```

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

    0 best-effort          7375          7375          0
    1 expedited-fo          0          0          0
    2 assured-forw          0          0          0

```

3 network-cont	2268	2268	0
----------------	------	------	---

Logical interface ae0.0 (Index 72) (SNMP ifIndex 60) (Generation 18)

Flags: SNMP-Traps 16384 Encapsulation: ENET2

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

Bundle:

Input :	0	0	0	0
---------	---	---	---	---

Output:	0	0	0	0
---------	---	---	---	---

Link:

fe-0/1/0.0

Input :	0	0	0	0
---------	---	---	---	---

Output:	0	0	0	0
---------	---	---	---	---

fe-0/1/2.0

Input :	0	0	0	0
---------	---	---	---	---

Output:	0	0	0	0
---------	---	---	---	---

fe-0/1/3.0

Input :	0	0	0	0
---------	---	---	---	---

Output:	0	0	0	0
---------	---	---	---	---

Marker Statistics:	Marker Rx	Resp Tx	Unknown Rx	Illegal Rx
--------------------	-----------	---------	------------	------------

fe-0/1/0.0	0	0	0	0
------------	---	---	---	---

fe-0/1/2.0	0	0	0	0
------------	---	---	---	---

fe-0/1/3.0	0	0	0	0
------------	---	---	---	---

Protocol inet, MTU: 1500, Generation: 37, Route table: 0

Flags: Is-Primary, Mac-Validate-Strict

Mac-Validate Failures: Packets: 0, Bytes: 0

Destination: 10.100.1/24, Local: 10.100.1.2, Broadcast: 10.100.1.255,

Generation: 49

show interfaces extensive (Aggregated Ethernet)

user@host> show interfaces ae0 extensive

Physical interface: ae0, Enabled, Physical link is Up

Interface index: 153, SNMP ifIndex: 59, Generation: 36

Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,

Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1

Device flags : Present Running

Interface flags: SNMP-Traps 16384

Current address: 00:05:85:8b:bf:f0, Hardware address: 00:05:85:8b:bf:f0

Last flapped : Never

Statistics last cleared: Never

Traffic statistics:

Input bytes :	60	0 bps
---------------	----	-------

Output bytes :	0	0 bps
----------------	---	-------

Input packets:	1	0 pps
----------------	---	-------

Output packets:	0	0 pps
-----------------	---	-------

Input errors:

Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,

Policed discards: 0, Resource errors: 0

Output errors:

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

Resource errors: 0

Queue counters:	Queued packets	Transmitted packets	Dropped packets
-----------------	----------------	---------------------	-----------------

0 best-effort	7375	7375	0
---------------	------	------	---

1 expedited-fo	0	0	0
----------------	---	---	---

2 assured-forw	0	0	0
----------------	---	---	---

3 network-cont	2268	2268	0
----------------	------	------	---


```

Logical interface ae0.0 (Index 72) (SNMP ifIndex 60) (Generation 18)
Flags: SNMP-Traps 16384 Encapsulation: ENET2
Statistics          Packets      pps          Bytes          bps
Bundle:
  Input :           1           0           60           0
  Output:           0           0            0           0
Link:
  fe-0/1/0.0
    Input :           0           0            0           0
    Output:           0           0            0           0
  fe-0/1/2.0
    Input :           0           0            0           0
    Output:           0           0            0           0
  fe-0/1/3.0
    Input :           1           0           60           0
    Output:           0           0            0           0
Marker Statistics:  Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
fe-0/1/0.0          0          0          0          0
fe-0/1/2.0          0          0          0          0
fe-0/1/3.0          0          0          0          0
Protocol inet, MTU: 1500, Generation: 37, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.100.1/24, Local: 10.100.1.2, Broadcast: 10.100.1.255,
  Generation: 49

```

show interfaces
extensive (Aggregated

```

user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
Interface index: 155, SNMP ifIndex: 48, Generation: 186

```

Ethernet with VLAN
Stacking)

```

Link-level type: 52, MTU: 1518, Speed: 2000mbps, Loopback: Disabled, Source
filtering: Disabled,
Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth needed: 0
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Current address: 00:12:1e:19:3f:f0, Hardware address: 00:12:1e:19:3f:f0
Last flapped   : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes :          2406875          40152 bps
Output bytes :         1124470          22056 bps
Input packets:           5307           5 pps
Output packets:         13295          21 pps
IPv6 transit statistics:
Input bytes :           0
Output bytes :           0
Input packets:           0
Output packets:          0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
Ingress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort              0              859777              0
  1 expedited-fo             0              0              0
  2 assured-forw             0              0              0
  3 network-cont             0              0              0

Egress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort              0          1897615              0
  1 expedited-fo             0              0              0
  2 assured-forw             0              0              0
  3 network-cont             0          662505              0

Logical interface ae0.451 (Index 69) (SNMP ifIndex 167) (Generation 601)
Flags: SNMP-Traps VLAN-Tag [ 0x8100.451 ] Encapsulation: VLAN-VPLS
Statistics      Packets      pps      Bytes      bps
Bundle:
  Input :          289          0      25685      376
  Output:         1698          4     130375     3096
Link:
  ge-1/2/0.451
    Input :          289          0      25685      376
    Output:           0          0           0           0
  ge-1/2/1.451
    Input :           0          0           0           0
    Output:         1698          4     130375     3096
Marker Statistics:  Marker Rx      Resp Tx      Unknown Rx      Illegal Rx

```

```

ge-1/2/0.451          0          0          0          0
ge-1/2/1.451          0          0          0          0
Protocol vpls, MTU: 1518, Generation: 849, Route table: 3
Flags: Is-Primary

```

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

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

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

Bundle:

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

Link:

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

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

Marker Statistics:	Marker Rx	Resp Tx	Unknown Rx	Illegal Rx
ge-1/2/0.452	0	0	0	0
ge-1/2/1.452	0	0	0	0

Protocol vpls, MTU: 1518, Generation: 850, Route table: 3

Flags: None

...

show interfaces (Fast Ethernet)

Syntax	<pre>show interfaces <i>interface-type</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display status information about the specified Fast Ethernet interface.
Options	<p><i>interface-type</i>—On M Series and T Series routers, the interface type is <i>fe-fpc/pic/port</i>. On the J Series routers, the interface type is <i>fe-pim/O/port</i>.</p> <p><i>brief detail extensive terse</i>—(Optional) Display the specified level of output.</p> <p><i>descriptions</i>—(Optional) Display interface description strings.</p> <p><i>media</i>—(Optional) Display media-specific information about network interfaces.</p> <p><i>snmp-index snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><i>statistics</i>—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
List of Sample Output	<p>show interfaces (Fast Ethernet) on page 122</p> <p>show interfaces brief (Fast Ethernet) on page 122</p> <p>show interfaces detail (Fast Ethernet) on page 122</p> <p>show interfaces extensive (Fast Ethernet) on page 123</p>
Output Fields	<p>Table 12 on page 108 lists the output fields for the show interfaces Fast Ethernet command. Output fields are listed in the approximate order in which they appear.</p>

Table 12: show interfaces Fast 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 Common Output Fields Description.	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

Table 12: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
Link-mode	Type of link connection configured for the physical interface: Full-duplex or Half-duplex	extensive
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 Common Output Fields Description.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under Common Output Fields Description.	All levels
Link flags	Information about the link. Possible values are described in the "Links Flags" section under Common Output Fields Description.	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels

Table 12: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
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	(GigabitEthernet 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).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface • Output bytes—Number of bytes transmitted on the interface. • 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 31 under the show interfaces (10-Gigabit Ethernet) command.</p>	detail extensive

Table 12: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 ignore-l3-incompletes 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 12: show interfaces Fast 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 12: show interfaces Fast 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 router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value 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 FEC statistics	<p>The forward error correction (FEC) counters provide the following statistics:</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. 	
PCS statistics	<p>(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.</p> <ul style="list-style-type: none"> • 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 12: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • 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 31 under the show interfaces (10-Gigabit Ethernet) command. • 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 12: show interfaces Fast 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 router from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local router (which the router is rejecting). • 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 12: show interfaces Fast 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 12: show interfaces Fast 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 (signal) 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 12: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • 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 router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.</p>	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive

Table 12: show interfaces Fast 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 Common Output Fields Description.	All levels
VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags. <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-pop—An outer VLAN tag is pushed in front of the existing VLAN tag, and then removed. • 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

Table 12: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Demux:	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: <ul style="list-style-type: none"> 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 Common Output Fields Description.	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	Number and rate of bytes and packets received and transmitted on the specified interface set. <ul style="list-style-type: none"> Input bytes, Output bytes—Number of bytes received and transmitted on the interface set 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 router.	extensive
Transit statistics	Number and rate of bytes and packets transiting the switch. <p>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.</p>	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 Common Output Fields Description.	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none

Table 12: show interfaces Fast Ethernet Output Fields (*continued*)

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

Sample Output

show interfaces (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
  Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:44 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198)
  Flags: SNMP-Traps Encapsulation: ENET2
  Protocol inet, MTU: 1500
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
```

show interfaces brief (Fast Ethernet)

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

show interfaces detail (Fast Ethernet)

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

Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255,
Generation: 136

**show interfaces
extensive
(Fast Ethernet)**

```

user@host> show interfaces fe-0/0/0 extensive
Physical interface: fe-0/0/0, Enabled, Physical link is Up
Interface index: 128, SNMP ifIndex: 22, Generation: 5391
Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed:
100mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
CoS queues     : 4 supported, 4 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:46 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   :          0          0 bps
Output bytes  :         42          0 bps
Input packets :          0          0 pps
Output packets:          1          0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Active alarms : None
Active defects : None
MAC statistics:

```

	Receive	Transmit
Total octets	0	64
Total packets	0	1
Unicast packets	0	0
Broadcast packets	0	1
Multicast packets	0	0
CRC/Align errors	0	0
FIFO errors	0	0
MAC control frames	0	0
MAC pause frames	0	0
Oversized frames	0	
Jabber frames	0	
Fragment frames	0	
VLAN tagged frames	0	
Code violations	0	

```

Filter statistics:
Input packet count      0
Input packet rejects    0
Input DA rejects        0
Input SA rejects        0
Output packet count     1
Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link partner:
Link partner: Full-duplex, Flow control: None, Remote fault: Ok
Local resolution:
Packet Forwarding Engine configuration:

```

```
Destination slot: 0
CoS information:
      Bandwidth      Buffer Priority  Limit
              %      bps    %      usec
0 best-effort      95    950000000  95         0    low  none
3 network-control   5     50000000   5         0    low  none
Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
Flags: SNMP-Traps Encapsulation: ENET2
Protocol inet, MTU: 1500, Generation: 105, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255,
  Generation: 136
```

show interfaces (10-Gigabit Ethernet)

Syntax	<pre>show interfaces <i>xe-fpc/pic/port</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre>
Release Information	Command introduced in Junos OS Release 8.0.
Description	(M320, M120, MX Series, and T Series routers only) Display status information about the specified 10-Gigabit Ethernet interface.
Options	<p><i>xe-fpc/pic/port</i>—Display standard information about the specified 10-Gigabit Ethernet interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(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 140</p> <p>show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode) on page 143</p> <p>show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC) on page 145</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode) on page 148</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only) on page 148</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only) on page 149</p>
Output Fields	See Table 13 on page 126 for the output fields for the show interfaces (10-Gigabit Ethernet) command.

Table 13: 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 Common Output Fields Description.	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 Common Output Fields Description.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under Common Output Fields Description.	All levels

Table 13: 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 Common Output Fields Description.	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).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface • Output bytes—Number of bytes transmitted on the interface. • 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 13 on page 126.</p>	detail extensive

Table 13: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 ignore-l3-incompletes 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 13: 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 13: 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 router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value 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 13: 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 13: 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 14 on page 140 • 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 13: 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 router from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local router (which the router is rejecting). • 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 13: 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 13: 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 13: 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 router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.</p>	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive

Table 13: 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 Common Output Fields Description.	All levels

Table 13: 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 Common Output Fields Description.	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 • 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 router.	extensive

Table 13: 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 Common Output Fields Description.	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 Common Output Fields Description.	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 Common Output Fields Description.	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 13: 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 14 on page 140](#) 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 14 on page 140](#), 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 14: 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

```
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
```

**(10-Gigabit Ethernet,
LAN PHY Mode, IQ2)**

```

Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Enabled,
Flow control: Enabled
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags     : None
CoS queues    : 8 supported, 4 maximum usable queues
Schedulers    : 1024
Hold-times    : Up 0 ms, Down 0 ms
Current address: 00:14:f6:b9:f1:f6, Hardware address: 00:14:f6:b9:f1:f6
Last flapped   : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes   :          6970332384          0 bps
Output bytes  :              0          0 bps
Input packets:          81050506          0 pps
Output packets:              0          0 pps
IPv6 transit statistics:
Input bytes   :              0
Output bytes  :              0
Input packets:              0
Output packets:              0
Ingress traffic statistics at Packet Forwarding Engine:
Input bytes   :          6970299398          0 bps
Input packets:          81049992          0 pps
Drop bytes    :              0          0 bps
Drop packets  :              0          0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0,
MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

0 best-effort          81049992          81049992          0

1 expedited-fo              0              0          0

2 assured-forw              0              0          0

3 network-cont              0              0          0

Egress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

0 best-effort          0              0          0

1 expedited-fo          0              0          0

2 assured-forw          0              0          0

3 network-cont          0              0          0

Active alarms : None
Active defects : None
PCS statistics
Bit errors          Seconds
0

```

```

    Errored blocks                                0
MAC statistics:
    Receive
    Transmit
    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
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
    Output packets: 0
    Input bytes : 0 bps
    Output bytes : 0 bps
    Input packets: 0 pps
    Output packets: 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**

```
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
```

(10-Gigabit Ethernet,
WAN PHY Mode)

```

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**

```

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

```

(10-Gigabit Ethernet,
DWDM OTN PIC)

```

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:      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

```

```

Total octets                                0                0
Total packets                              0                0
Unicast packets                            0                0
Broadcast packets                          0                0
Multicast packets                          0                0
CRC/Align errors                           0                0
FIFO errors                                0                0
MAC control frames                         0                0
MAC pause frames                           0                0
Oversized frames                           0
Jabber frames                              0
Fragment frames                            0
VLAN tagged frames                         0
Code violations                             0
OTN alarms : None
OTN defects : None
OTN FEC Mode : GFEC
OTN Rate : Fixed Stuff Bytes 11.0957Gbps
OTN Line Loopback : Enabled
OTN FEC statistics :
    Corrected Errors                                0
    Corrected Error Ratio ( 0 sec average) 0e-0
OTN FEC alarms: Seconds Count State
    FEC Degrade 0 0 OK
    FEC Excessive 0 0 OK
OTN OC: Seconds Count State
    LOS 2 1 OK
    LOF 67164 2 Defect Active
    LOM 67164 71 Defect Active
    Wavelength Lock 0 0 OK
OTN OTU:
    AIS 0 0 OK
    BDI 65919 4814 Defect Active
    IAE 67158 1 Defect Active
    TTIM 7 1 OK
    SF 67164 2 Defect Active
    SD 67164 3 Defect Active
    TCA-ES 0 0 OK
    TCA-SES 0 0 OK
    TCA-UAS 80 40 OK
    TCA-BBE 0 0 OK
    BIP 0 0 OK
    BBE 0 0 OK
    ES 0 0 OK
    SES 0 0 OK
    UAS 587 0 OK
Received DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Received SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
OTN Received Overhead Bytes:
    APS/PCC0: 0x02, APS/PCC1: 0x42, APS/PCC2: 0xa2, APS/PCC3: 0x48
    Payload Type: 0x03
OTN Transmitted Overhead Bytes:
    APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x03
Filter statistics:

```

```

Input packet count          0
Input packet rejects        0
Input DA rejects            0
Input SA rejects            0
Output packet count         0
Output packet pad count     0
Output packet error count   0
CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 7
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit                    %      bps      %      usec      low
0 best-effort            95      9500000000  95      0
none
3 network-control        5      500000000  5      0
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**

```

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,

```

Mode, Unidirectional Mode, Transmit-Only)

```

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
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

```

user@host> show interfaces xe-7/0/0-rx extensive
Physical interface: xe-7/0/0-rx, Enabled, Physical link is Up

```

extensive (10-Gigabit
Ethernet, LAN PHY

Interface index: 174, SNMP ifIndex: 118, Generation: 175
Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Rx-Only

Mode, Unidirectional Mode, Receive-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 153](#)
- [VRRP Operational Mode Commands on page 158](#)

Ethernet Interface Operational Mode Commands

[Table 15 on page 153](#) 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 15: Ethernet Interface Operational Mode Commands

Task	Command
Clear dynamic VLAN interfaces.	clear auto-configuration interfaces
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.	clear auto-configuration interfaces interface-set
Clear Link Aggregation Control Protocol (LACP) statistics.	clear lacp statistics
Clear Link Aggregation Control Protocol (LACP) timeout entries.	clear lacp timeouts
Clear learned MAC addresses from the hardware and MAC database. Static MAC addresses are not cleared.	clear interfaces mac-database
Clear statistics that are collected for every MAC address, including policer statistics, on a given physical or logical interface.	clear interfaces mac-database statistics
Clear statistics that are collected for interface sets.	clear interfaces interface-set statistics
Clear the existing continuity measurement and restart counting the operational uptime.	clear oam ethernet connectivity-fault-management continuity-measurement

Table 15: 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)	clear oam ethernet connectivity-fault-management delay-statistics
Clear Operation, Administration, and Management (OAM) and connectivity fault management (CFM) linktrace database information.	clear oam ethernet connectivity-fault-management linktrace path-database
Clear all loss statistics maintained by CFM for a given maintenance domain and maintenance association.	clear oam ethernet connectivity-fault-management loss-statistics
Clear connectivity-fault-management policer statistics.	clear oam ethernet connectivity-fault-management policer
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).	clear oam ethernet connectivity-fault-management statistics
Clear Operation, Administration, and Management (OAM) link fault management state information and restart the link discovery process on Ethernet interfaces.	clear oam ethernet link-fault-management state
Clear Operation, Administration, and Management (OAM) statistics link fault management statistics for Ethernet interfaces.	clear oam ethernet link-fault-management statistics
Clear the statistics for all Ethernet ring protection groups or a specific Ethernet ring protection group.	clear protection-group ethernet-ring statistics
Check the reachability of a remote IEEE 802.1ag OAM maintenance association end point (MEP) or maintenance association intermediate point (MIP).	ping ethernet
Manually rebalance the subscribers on an aggregated Ethernet bundle with targeted distribution enabled.	request interface rebalance (Aggregated Ethernet for Subscriber Management)
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.	request interface (revert switchover) (Aggregated Ethernet Link Protection)
Force LACP link switchover.	request lacp link-switchover

Table 15: Ethernet Interface Operational Mode Commands (*continued*)

Task	Command
Clear the lockout, force switch, manual switch, exercise, and wait-to-restore states.	request protection-group ethernet-aps clear
Test if APS is operating correctly.	request protection-group ethernet-aps exercise
Force traffic to switch from the active path to the alternate path.	request protection-group ethernet-aps force-switch
Lock the protection path, forcing the use of the working path.	request protection-group ethernet-aps lockout
Force traffic to switch from the active path to the alternate path.	request protection-group ethernet-aps manual-switch
Display status information about aggregated Fast Ethernet or Gigabit Ethernet router interfaces.	show interfaces (Aggregated Ethernet) show interfaces (far-end-interval)
Display status information about Fast Ethernet interfaces.	show interfaces (Fast Ethernet)
Display status information about the specified Gigabit Ethernet interface.	show interfaces (Gigabit Ethernet)
Display status information about 10-Gigabit Ethernet router interfaces.	show interfaces (10-Gigabit Ethernet)
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.	show interfaces extensive
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.	show interfaces interface-set (Ethernet Interface Set)
Display information about Gigabit Ethernet or 10-Gigabit Ethernet router interface set queues.	show interfaces interface-set queue
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.	show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, and 100 Gigabit Ethernet)

Table 15: Ethernet Interface Operational Mode Commands (*continued*)

Task	Command
Display information about integrated routing and bridging interfaces.	show interfaces irb
Display status information about the distribution of subscribers on different links in an aggregated Ethernet bundle.	show interfaces targeting (Aggregated Ethernet for Subscriber Management)
Display Link Aggregation Control Protocol (LACP) information for aggregated, Fast Ethernet, or Gigabit Ethernet router interfaces.	show lacp interfaces
Display Link Aggregation Control Protocol (LACP) statistics.	show lacp statistics
Display Link Aggregation Control Protocol timeout entries.	show lacp timeouts
Display MAC address information for Gigabit Ethernet router interfaces.	show interfaces mac-database (Gigabit Ethernet)
Display information on a specified interface that is part of a multichassis link aggregation configuration.	show interfaces mc-ae
Display ETH-DM statistics for CFM MEPs. (MX Series routers, Ethernet DPCs).	show oam ethernet connectivity-fault-management delay-statistics
Display IEEE 802.1ag OAM connectivity fault management forwarding state information for Ethernet interfaces.	show oam ethernet connectivity-fault-management forwarding-state
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.	show oam ethernet connectivity-fault-management interfaces
Display OAM connectivity fault management linktrace path database information.	show oam ethernet connectivity-fault-management linktrace path-database
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.	show oam ethernet connectivity-fault-management mep-database

Table 15: Ethernet Interface Operational Mode Commands (*continued*)

Task	Command
Display ETH-DM statistics and frame counts for CFM MEPs. (MX Series routers, Ethernet DPCs)	show oam ethernet connectivity-fault-management mep-statistics
Display ETH-LM statistics for on-demand mode only.	show oam ethernet connectivity-fault-management loss-statistics
Display information about maintenance intermediate points (MIPs) for the Ethernet OAM 802.1ag standard for connectivity fault management (CFM).	show oam ethernet connectivity-fault-management mip
Display OAM connectivity fault management path database information for hosts configured with MEP.	show oam ethernet connectivity-fault-management path-database
Displays connectivity-fault-management policer statistics.	show oam ethernet connectivity-fault-management policer
Display OAM Ethernet Virtual Connection (EVC) information for hosts configured with Ethernet Local Management Interface (E-LMI). (MX series only)	show oam ethernet evc
Display OAM fault management statistics for Ethernet interfaces.	show oam ethernet link-fault-management
Display OAM Ethernet Local Management Interface status information for an LMI configured interface. (MX series only)	show oam ethernet lmi
Display OAM Ethernet Local Management Interface statistics for an LMI configured interface. (MX series only)	show oam ethernet lmi statistics
Display protection group Ethernet ring Automatic Protection Switching (APS).	show protection-group ethernet-ring aps
Display data channel information for all Ethernet ring protection groups or for a specific Ethernet ring protection group.	show protection-group ethernet-ring data-channel
Display protection group Ethernet ring interfaces.	show protection-group ethernet-ring interface
Display protection group Ethernet ring nodes.	show protection-group ethernet-ring node-state
Display protection group Ethernet ring statistics.	show protection-group ethernet-ring statistics

Table 15: Ethernet Interface Operational Mode Commands (*continued*)

Task	Command
Display all data channel logical interfaces and the VLAN IDs controlled by a ring instance data channel.	show protection-group ethernet-ring vlan
Trace the path between two Ethernet OAM end points.	traceroute ethernet

VRRP Operational Mode Commands

[Table 16 on page 158](#) summarizes the command-line interface (CLI) commands that you can use to monitor and troubleshoot Virtual Router Redundancy Protocol (VRRP) on Ethernet, Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, and logical tunnel interfaces. Commands are listed in alphabetical order.

Table 16: VRRP Operational Mode Commands

Task	Command
Clear (set to zero) VRRP groups.	clear vrrp
Display VRRP groups.	show vrrp

PART 4

Troubleshooting

- [Ethernet on page 161](#)
- [Interface Diagnostics on page 169](#)

CHAPTER 7

Ethernet

- [Tracing PPPoE Operations on page 163](#)
- [Troubleshooting PPPoE Service Name Tables on page 165](#)
- [Verifying a PPPoE Configuration on page 167](#)

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 seconds>
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 value —Number of hops to use in the linktrace request. The range is 1 to 255 hops. The default is 4. wait seconds —(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 163
Output Fields	Table 17 on page 162 lists the output fields for the traceroute ethernet command. Output fields are listed in the approximate order in which they appear.

Table 17: 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 17: 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

Tracing PPPoE Operations

The Junos OS trace feature tracks PPPoE operations and records events in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

By default, nothing is traced. When you enable the tracing operation, the default tracing behavior is as follows:

1. Important events are logged in a file called **pppoed** located in the **/var/log** directory. You cannot change the directory (**/var/log**) in which trace files are located.
2. When the file **pppoed** reaches 128 kilobytes (KB), it is renamed **pppoed.0**, then **pppoed.1**, and finally **pppoed.2**, until there are three trace files. Then the oldest trace file (**pppoed.2**) is overwritten.

You can optionally specify the number of trace files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB). (For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)

By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure PPPoE tracing operations:

1. Specify that you want to configure tracing options.

```
[edit protocols pppoe]
user@host# edit traceoptions
```
2. (Optional) Configure the name for the file used for the trace output.
3. (Optional) Configure the number and size of the log files.
4. (Optional) Configure access to the log file.
5. (Optional) Configure a regular expression to filter logging events.
6. (Optional) Configure flags to filter the operations to be logged.

Optional PPPoE traceoptions operations are described in the following sections:

- [Configuring the PPPoE Trace Log Filename on page 164](#)
- [Configuring the Number and Size of PPPoE Log Files on page 164](#)
- [Configuring Access to the PPPoE Log File on page 165](#)
- [Configuring a Regular Expression for PPPoE Lines to Be Logged on page 165](#)
- [Configuring the PPPoE Tracing Flags on page 165](#)

Configuring the PPPoE Trace Log Filename

By default, the name of the file that records trace output for PPPoE is **pppoed**. You can specify a different name with the **file** option.

Configuring the Number and Size of PPPoE Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

The archived files are differentiated by a suffix in the format **.number.gz**. The newest archived file is **.0.gz** and the oldest archived file is **.(maximum number)-1.gz**. When the current trace log file reaches the maximum size, it is compressed and renamed, and any

existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation, *filename*, reaches 2 MB, *filename* is compressed and renamed *filename.0.gz*, and a new file called *filename* is created. When the new *filename* reaches 2 MB, *filename.0.gz* is renamed *filename.1.gz* and *filename* is compressed and renamed *filename.0.gz*. This process repeats until there are 20 trace files. Then the oldest file, *filename.19.gz*, is simply overwritten when the next oldest file, *filename.18.gz* is compressed and renamed to *filename.19.gz*.

Configuring Access to the PPPoE Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

Configuring a Regular Expression for PPPoE Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

You can refine the output by including regular expressions to be matched.

Configuring the PPPoE Tracing Flags

By default, no events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged, configure the flags:

- `[edit protocols pppoe traceoptions]`
`user@host# set flag authentication`

Related Documentation

- PPPoE Overview
- Junos® OS Ethernet Interfaces

Troubleshooting PPPoE Service Name Tables

Problem A misconfiguration of a PPPoE service name table can prevent PPPoE services from being properly activated. Configuration options for PPPoE service name tables are simple, which should simplify discovering where a misconfiguration exists. PPPoE clients cannot connect if the service name table contains no match for the service name tag carried in the PADI packet.

The symptom of a service name table misconfiguration is that the client connection process stops at the negotiation stage and the PADI packets are ignored. You can use the **show pppoe statistics** command to examine the PPPoE packet counts for a problem.

When the service name table is properly configured, packets sent and received increment symmetrically. The following sample output shows a PADO sent count equal to the PADI received count, and PADS sent count equal to the PADR received count. This output

indicates that the PPPoE negotiation is proceeding successfully and that the service name table is not misconfigured.

```
user@host> show pppoe statistics ge-2/0/3.1
```

```
Active PPPoE sessions: 2
```

PacketType	Sent	Received
PADI	0	16
PADO	16	0
PADR	0	16
PADS	16	0
PADT	0	0
Service name error	0	0
AC system error	0	0
Generic error	0	0
Malformed packets	0	0
Unknown packets	0	0

When the service name table is misconfigured, the output of the **show pppoe statistics** command indicates that the number of PADI packets received on the underlying interface is increasing, but the number of PADO packets sent remains at zero. The following sample output shows a PADI count of 100 and a PADO count of 0.

```
user@host> show pppoe statistics ge-2/0/3.1
```

```
Active PPPoE sessions: 0
```

PacketType	Sent	Received
PADI	0	100
PADO	0	0
PADR	0	0
PADS	0	0
PADT	0	0
Service name error	0	0
AC system error	0	0
Generic error	0	0
Malformed packets	0	0
Unknown packets	0	0

When you believe a misconfiguration exists, use the **monitor traffic interface** command on the underlying interface to determine which service name is being requested by the PPPoE client. The following sample output shows that the client is requesting Service1 in the service name tag.

```
user@host> monitor traffic interface ge-2/0/3.1 print-hex print-ascii
```

```
Listening on ge-2/0/3.1, capture size 96 bytes
```

```
11:49:41.436682 In PPPoE PADI [Service-Name "Service1"] [Host-Uniq UTF8]
[TAG-0x120 UTF8] [Vendor-Specific UTF8]
0x0000 ffff ffff ffff 0090 1a42 0ac1 8100 029a .....B.....
0x0010 8863 1109 0000 00c9 0101 0008 5365 7276 .c.....Serv
0x0020 6963 6531 0103 0004 1200 9c43 0120 0002 ice1.....C....
0x0030 044a 0105 00ab 0000 0de9 0124 783a 3132 .J.....$x:12
0x0040 3030 3963                                009c
```

You can then use the **show pppoe service-name-tables** command to determine whether you have misspelled the name of the service or perhaps not configured the service at all.

Cause Typical misconfigurations appear in the service name table configurations.

Solution Use the appropriate statements to correct the misconfiguration.

- Related Documentation**
- Configuring PPPoE Service Name Tables
 - show pppoe service-name-tables
 - show pppoe statistics
 - show pppoe underlying-interfaces
 - PPPoE Overview
 - Junos® OS Ethernet Interfaces

Verifying a PPPoE Configuration

Purpose You can use show commands to display and verify the PPPoE configuration.

Action To verify a PPPoE configuration, you can issue the following operational mode commands:

- show interfaces at-*fpc/pic/port* extensive
- show interfaces pp0
- show pppoe interfaces
- show pppoe version
- show pppoe service-name-tables
- show pppoe sessions
- show pppoe statistics
- show pppoe underlying-interfaces

For more information about these operational mode commands, see Junos OS Operational Mode Commands.

- Related Documentation**
- PPPoE Overview
 - Junos® OS Ethernet Interfaces

CHAPTER 8

Interface Diagnostics

- [Interface Diagnostics on page 169](#)

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 169](#)
- [Interface Diagnostics on page 171](#)

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, NxDSO, 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 18 on page 170 shows the loopback modes supported on the various interface types.

Table 18: Loopback Modes by Interface Type

Interface	Loopback Modes	Usage Guidelines
Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet	Local	Configuring Ethernet Loopback Capability
Circuit Emulation E1	Local and remote	Configuring E1 Loopback Capability
Circuit Emulation T1	Local and remote	Configuring T1 Loopback Capability
E1 and E3	Local and remote	Configuring E1 Loopback Capability and Configuring E3 Loopback Capability
NxDSO	Payload	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
Serial (V.35 and X.21)	Local and remote	Configuring Serial Loopback Capability
Serial (EIA-530)	DCE local, DCE remote, local, and remote	Configuring Serial Loopback Capability
SONET/SDH	Local and remote	Configuring SONET/SDH Loopback Capability

Table 18: Loopback Modes by Interface Type (*continued*)

Interface	Loopback Modes	Usage Guidelines
T1 and T3	Local, payload, and remote	Configuring T1 Loopback Capability and Configuring T3 Loopback Capability See also Configuring the T1 Remote Loopback Response

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* gige-ether-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 19 on page 174 shows the BERT capabilities for various interface types.

Table 19: 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 179](#)

Index

Symbols

#, comments in configuration statements.....	xiv
(), in syntax descriptions.....	xiv
802.1Q VLANs	
VLAN IDs.....	89
< >, in syntax descriptions.....	xiv
[], in configuration statements.....	xiv
{ }, in configuration statements.....	xiv
(pipe), in syntax descriptions.....	xiv

A

aggregated Ethernet interfaces	
status information, displaying.....	98
aggregated-ether-options statement.....	64

B

BERT	
configuring interface diagnostics.....	171
bert-algorithm statement	
usage guidelines.....	171
bert-error-rate statement	
usage guidelines.....	171
bert-period statement	
usage guidelines.....	171
bit error rate test See BERT	
braces, in configuration statements.....	xiv
brackets	
angle, in syntax descriptions.....	xiv
square, in configuration statements.....	xiv

C

comments, in configuration statements.....	xiv
connections	
configuration statements.....	55, 59
conventions	
text and syntax.....	xiii
curly braces, in configuration statements.....	xiv
customer support.....	xv
contacting JTAC.....	xv

D

documentation	
comments on.....	xv
dot1x	
configuration statements.....	56
dynamic profiles statements	
vlan-tags.....	94

E

Ethernet interfaces	
status information, displaying	
aggregated.....	98
Fast Ethernet.....	108
Gigabit Ethernet.....	125
VLAN IDs.....	89
Ethernet Ring Protection	
configuration statements.....	60
ethernet-switch-profile statement.....	66
usage guidelines.....	12

F

Fast Ethernet interfaces	
status information, displaying.....	108
VLAN IDs.....	89
fast-aps-switch statement.....	68
font conventions.....	xiii

G

Gigabit Ethernet interfaces	
status information, displaying.....	125
VLAN IDs.....	89
Gigabit Ethernet IQ PIC	
traffic and MAC statistics.....	125
gether-options statement.....	69

I

iccp	
configuration statements.....	56
IEEE 802.1p inheritance	
and transparent tag.....	4
inner-tag-protocol-id statement.....	70
usage guidelines.....	14
inner-vlan-id statement.....	71
usage guidelines.....	14
interfaces	
configuration statements.....	34

L

linktrace	
tracing.....	162
log files	
number of PPPoE.....	164
size of PPPoE.....	164
logical interfaces	
VLAN IDs.....	89
logical systems	
configuration statements.....	50
loopback testing.....	169

M

manuals	
comments on.....	xv

O

OAM	
configuration statements.....	57
output-vlan-map statement	
Aggregated Ethernet.....	75
Gigabit Ethernet IQ.....	74

P

parentheses, in syntax descriptions.....	xiv
pop statement	
Gigabit Ethernet IQ interfaces.....	75
usage guidelines.....	18
pop-pop statement	
Gigabit Ethernet IQ interfaces.....	76
usage guidelines.....	19
pop-swap statement	
Gigabit Ethernet IQ interfaces.....	76
usage guidelines.....	19
PPPoE	
flags for tracing operations.....	165
log file access for tracing operations.....	165
log file size and number.....	164
log filenames for tracing operations.....	164
regular expressions for tracing operations.....	165
service name tables	
configuration troubleshooting.....	165
verifying the configuration.....	167
tracing operations.....	163
protection-group	
configuration statements.....	60
protocols connections	
configuration statements.....	55, 59

protocols dot1x	
configuration statements.....	56
protocols Ethernet Ring Protection	
configuration statements.....	60
protocols iccp	
configuration statements.....	56
protocols OAM	
configuration statements.....	57
protocols VRRP	
configuration statements.....	57
push statement	
Gigabit Ethernet IQ interfaces.....	77
usage guidelines.....	17
push-push statement	
Gigabit Ethernet IQ interfaces.....	77
usage guidelines.....	20

R

rewrite VLAN tag on untagged frame	
usage guidelines.....	22

S

service name tables	
PPPoE	
configuration troubleshooting.....	165
verifying the configuration.....	167
show interfaces (10-Gigabit Ethernet)	
command.....	125
show interfaces (Aggregated Ethernet)	
command.....	98
show interfaces (Fast Ethernet) command.....	108
stacked VLAN-tag framing	
VLAN ID list.....	92
stacked-vlan-tagging statement.....	78
support, technical See technical support	
swap statement	
Gigabit Ethernet IQ interfaces.....	78
swap-by-poppush statement	
Gigabit Ethernet IQ2 and IQ2E interfaces.....	79
swap-push statement	
Gigabit Ethernet IQ interfaces.....	79
usage guidelines.....	25
swap-swap statement	
Gigabit Ethernet IQ interfaces.....	80
usage guidelines.....	25
syntax conventions.....	xiii

T

tag-protocol-id statement	
TPID to rewrite.....	82
TPIDs expected to be sent or received.....	81
usage guidelines.....	14
technical support	
contacting JTAC.....	xv
traceroute ethernet.....	162
traceroute ethernet command.....	162
tracing operations	
PPPoE.....	163
transparent tag operations.....	4

U

unit statement.....	83
---------------------	----

V

Virtual Router Redundancy Protocol	See VRRP
vlan-id statement	
802.1Q VLANs.....	89
Ethernet interfaces.....	89
rewriting at ingress or egress.....	90
vlan-rewrite statement.....	91
vlan-tags statement	
dynamic profiles.....	94
stacked VLAN tags.....	92
usage guidelines.....	13
VRRP	
configuration statements.....	57

