



Junos[®] OS

Protocol Family and Interface Address Properties



Published: 2013-08-01

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

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

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

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

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

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

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

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

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

Junos® OS Protocol Family and Interface Address Properties

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	xiii
	Documentation and Release Notes	xiii
	Supported Platforms	xiii
	Using the Examples in This Manual	xiv
	Merging a Full Example	xiv
	Merging a Snippet	xv
	Documentation Conventions	xv
	Documentation Feedback	xvii
	Requesting Technical Support	xvii
	Self-Help Online Tools and Resources	xvii
	Opening a Case with JTAC	xviii
Part 1	Overview	
Chapter 1	Protocol Family and Interface Address Properties	3
	Protocol Family Configuration and Interface Address Statements	3
Part 2	Configuration	
Chapter 2	Protocol Family and Interface Address Properties	9
	Setting the Protocol MTU	9
	Disabling the Removal of Address and Control Bytes	10
	Disabling the Transmission of Redirect Messages on an Interface	10
	Applying Policers	11
	Applying Aggregate Policers	12
	Example: Applying Aggregate Policers	13
	Applying Hierarchical Policers on Enhanced Intelligent Queuing PICs	14
	Hierarchical Policer Overview	14
	Hierarchical Policing Characteristics	15
	Configuring Hierarchical Policers	16
	Configuring a Single-Rate Two-Color Policer	17
	Configuring a Single-Rate Tricolor Policer	17
	Configuring a Two-Rate Tricolor Marker Policer	18
	Applying a Filter to an Interface	19
	Defining Interface Groups in Firewall Filters	21
	Filter-Based Forwarding on the Output Interface	22
	Example: Applying a Filter to an Interface	22
	Enabling Source Class and Destination Class Usage	24
	Examples: Enabling Source Class and Destination Class Usage	28

Configuring the Protocol Family	32
IPv6 Overview	34
IPv4-to-IPv6 Transition	34
VRRP Properties	35
Configuring the Interface Address	35
Configuring Interface IPv4 Addresses	36
Operational Behavior of Interfaces when the Same IPv4 Address is Assigned to Them	37
Configuring Interface IPv6 Addresses	38
Configuring ICCP for MC-LAG	39
Configuring IPCP Options	41
Configuring an IP Address for an Interface	41
Negotiating an IP Address Assignment from the Remote End	42
Configuring an Interface to Be Unnumbered	42
Assigning a Destination Profile to the Remote End	43
Configuring an Unnumbered Interface	43
Configuring an Unnumbered Point-to-Point Interface	43
Example: Configuring an Unnumbered Point-to-Point Interface	44
Configuring an Unnumbered Ethernet or Demux Interface	44
Configuring a Preferred Source Address for Unnumbered Ethernet or Demux Interfaces	46
Configuring Static Routes on Unnumbered Ethernet Interfaces	47
Restrictions for Configuring Unnumbered Ethernet Interfaces	47
Example: Configuring an Unnumbered Ethernet Interface	47
Example: Configuring the Preferred Source Address for an Unnumbered Ethernet Interface	48
Example: Configuring an Unnumbered Ethernet Interface as the Next Hop for a Static Route	48
Configuring Default, Primary, and Preferred Addresses and Interfaces	49
Configuring the Primary Interface for the Router	49
Configuring the Primary Address for an Interface	50
Configuring the Preferred Address for an Interface	50
Configuring Unicast RPF	51
Configuring Unicast RPF Strict Mode	51
Configuring Unicast RPF Loose Mode	52
Configuring Unicast RPF Loose Mode with Ability to Discard Packets	53
Unicast RPF and Default Routes	53
Unicast RPF Behavior with a Default Route	54
Unicast RPF Behavior Without a Default Route	54
Unicast RPF with Routing Asymmetry	54
Configuring Unicast RPF on a VPN	55
Example: Configuring Unicast RPF on a VPN	55

	Example: Configuring Unicast RPF	56
	Configuring Targeted Broadcast	57
	Configuring Targeted Broadcast and Its Options	57
	Display Targeted Broadcast Configuration Options	58
	Forward IP Packets On the Egress Interface and To the Routing Engine	58
	Forward IP Packets On the Egress Interface Only	58
	Understanding Targeted Broadcast	59
	Example: Configuring Unicast Reverse-Path-Forwarding Check	60
	Understanding Unicast Reverse Path Forwarding	60
	Example: Configuring Unicast Reverse-Path-Forwarding Check	60
Chapter 3	Network Interfaces Configuration Statements and Hierarchy	71
	[edit firewall] Hierarchy Level	71
	[edit interfaces] Hierarchy Level	72
	[edit logical-systems] Hierarchy Level	89
	[edit protocols pppoe] Hierarchy Level	93
Chapter 4	Statement Summary	95
	address	96
	aggregate (Hierarchical Policer)	98
	accounting	99
	arp (Interfaces)	100
	bandwidth-limit (Hierarchical Policer)	101
	broadcast	102
	bundle	103
	burst-size-limit (Hierarchical Policer)	104
	cbr	105
	demux0 (Dynamic Interface)	106
	destination (IPCP)	107
	destination (Tunnels)	108
	destination-class-usage	109
	destination-profile	109
	dynamic-profiles	110
	epd-threshold (Logical Interface)	117
	eui-64	118
	family (Dynamic Standard Interface)	119
	family	121
	fast-aps-switch	125
	filter	126
	forward-and-send-to-re	127
	forward-only	127
	hierarchical-policer	128
	if-exceeding	129
	input	129
	input-list	130
	interface-mode	131
	interfaces	132
	interfaces (Static and Dynamic Subscribers)	133
	inverse-arp	137

ipsec-sa	137
keep-address-and-control	138
logical-systems	138
mode (Dynamic Profiles)	139
mode (Interfaces)	139
mtu	140
multicast-only	142
multipoint-destination	143
negotiate-address	144
no-redirects	144
oam-liveness	145
oam-period	146
output	147
output-list	147
policer (Interface)	148
post-service-filter	149
preferred	149
preferred-source-address	150
premium (Hierarchical Policer)	151
primary (Address on Interface)	152
protocols	152
proxy	153
queue-length	154
receive-options-packets	154
receive-ttl-exceeded	155
remote	155
rpf-check (Dynamic Profiles)	156
rpf-check (interfaces)	157
rtvbr	158
sampling (Interfaces)	159
service (Logical Interfaces)	160
service-filter (Interfaces)	161
service-set	161
shaping	162
source-class-usage	163
targeted-broadcast	164
then	165
translate-discard-eligible	165
translate-fecn-and-becn	166
unit (Dynamic Profiles Standard Interface)	167
unit	170
unnumbered-address (Demux)	176
unnumbered-address (Dynamic Profiles)	177
unnumbered-address (Ethernet)	178
unnumbered-address (PPP)	179
vbr	180
vci	181
vlan-id (Logical Port in Bridge Domain)	182
vlan-id-list (Interface in Bridge Domain)	182

Part 3	Administration	
Chapter 5	Monitoring Commands	185
	clear firewall	186
	clear vpls mac-address	188
	clear vpls mac-table	189
	request l2circuit-switchover	190
	show as-path	191
	show as-path domain	195
	show as-path summary	197
	show dynamic-tunnels database	199
	show firewall	202
	show firewall log	209
	show firewall prefix-action-stats	212
	show hfrf profiles	214
	show interfaces (10-Gigabit Ethernet)	216
	show interfaces (Fast Ethernet)	241
	show interfaces (Gigabit Ethernet)	257
	show l2circuit connections	279
	show l2vpn connections	286
	show policer	292
	show route export vrf-target	294
	show route forwarding-table interface-name	296
	show route localization	299
	show route match-prefix	301
	show route range	302
	show route snooping	306
	show route snooping	314
Chapter 6	Command Summaries	323
	ANCP Operational Mode Commands	323
	BFD Operational Mode Commands	324
	BGP Operational Mode Commands	325
	ES-IS Operational Mode Commands	326
	IP Multicast Operational Mode Commands	326
	IPv6 Operational Mode Commands	330
	IS-IS Operational Mode Commands	331
	LLDP Operational Mode Commands	332
	MVRP Operational Mode Commands	333
	OSPF Operational Mode Commands	333
	Protocol-Independent Routing Operational Mode Commands	335
	RIP Operational Mode Commands	337
	RIPng Operational Mode Commands	338
	Firewall Filter Operational Mode Commands	339
	Layer 2 Bridging and Switching Operational Mode Commands	339
	VPN Operational Mode Commands	340

Part 4	Troubleshooting	
Chapter 7	Interface Diagnostics	345
	Interface Diagnostics	345
	Configuring Loopback Testing	345
	Interface Diagnostics	347
	Starting and Stopping a BERT Test	351
	Example: Configuring Bit Error Rate Testing	351
Part 5	Index	
	Index	355

List of Figures

Part 2	Configuration	
Chapter 2	Protocol Family and Interface Address Properties	9
	Figure 1: Hierarchical Policer	15
	Figure 2: Prefix Accounting with Source and Destination Classes	25
	Figure 3: Unicast RPF with Routing Asymmetry	55
	Figure 4: Unicast RPF Sample Topoolgy	62

List of Tables

	About the Documentation	xiii
	Table 1: Notice Icons	xv
	Table 2: Text and Syntax Conventions	xvi
Part 3	Administration	
Chapter 5	Monitoring Commands	185
	Table 3: show as-path Output Fields	191
	Table 4: show as-path domain Output Fields	195
	Table 5: show as-path summary Output Fields	197
	Table 6: show dynamic-tunnels database Output Fields	199
	Table 7: show firewall Output Fields	203
	Table 8: show firewall log Output Fields	209
	Table 9: show firewall prefix-action-stats Output Fields	212
	Table 10: show hfrp profiles Output Fields	214
	Table 11: show interfaces Gigabit Ethernet Output Fields	217
	Table 12: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type	231
	Table 13: show interfaces Fast Ethernet Output Fields	241
	Table 14: show interfaces Gigabit Ethernet Output Fields	258
	Table 15: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type	271
	Table 16: show l2circuit connections Output Fields	280
	Table 17: show l2vpn connections Output Fields	287
	Table 18: show policer Output Fields	292
	Table 19: show route export vrf-target Output Fields	294
	Table 20: show route forwarding-table interface-name Output Fields	296
	Table 21: show route localization Output Fields	299
Chapter 6	Command Summaries	323
	Table 22: ANCP Operational Mode Commands	323
	Table 23: BFD Operational Mode Commands	324
	Table 24: BGP Operational Mode Commands	325
	Table 25: ES-IS Operational Mode Commands	326
	Table 26: IP Multicast Operational Mode Commands	326
	Table 27: IPv6 Operational Mode Commands	331
	Table 28: IS-IS Operational Mode Commands	331
	Table 29: LLDP Operational Mode Commands	332
	Table 30: MVRP Operational Mode Commands	333
	Table 31: OSPF Operational Mode Commands	334
	Table 32: Protocol-Independent Routing Operational Mode Commands	335

	Table 33: RIP Operational Mode Commands	338
	Table 34: RIPng Operational Mode Commands	338
	Table 35: Firewall Filter Operational Mode Commands	339
	Table 36: Layer 2 Bridging and Switching Operational Mode Commands	340
	Table 37: Layer 2 Circuit, Layer 2 VPN, and VPLS Operational Mode Commands	340
Part 4	Troubleshooting	
Chapter 7	Interface Diagnostics	345
	Table 38: Loopback Modes by Interface Type	346
	Table 39: BERT Capabilities by Interface Type	350

About the Documentation

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

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:

- ACX Series
- M Series
- MX Series
- T Series
- J Series
- PTX 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 xv 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 xvi 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
<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 <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Enclose a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identify a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

GUI Conventions

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

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

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

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

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

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

- [Protocol Family and Interface Address Properties on page 3](#)

CHAPTER 1

Protocol Family and Interface Address Properties

- [Protocol Family Configuration and Interface Address Statements on page 3](#)

Protocol Family Configuration and Interface Address Statements

For each logical interface, you must configure one or more protocol families. You can also configure interface address properties. To do this, include the following statements:

```
family family {  
    accounting {  
        destination-class-usage;  
        source-class-usage {  
            direction;  
        }  
    }  
    address address {  
        destination address;  
    }  
    bundle interface-name;  
    filter {  
        dialer filter-name;  
        input filter-name;  
        output filter-name;  
        group filter-group-number;  
    }  
    interface-mode (access | trunk);  
    ipsec-sa sa-name;  
    keep-address-and-control;  
    mtu bytes;  
    multicast-only;  
    negotiate-address;  
    no-redirects;  
    policer {  
        arp policer-template-name;  
        input policer-template-name;  
        output policer-template-name;  
    }  
    primary;  
    protocols [inet iso mpls];  
    proxy inet-address address;
```

```
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check <fail-filter filter-name>;
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>;
    }
}
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;
vlan-id number;
vlan-id-list [number number-number];
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 dlcid dlcid-identifier;
    multipoint-destination address {
        epd-threshold cells;
        inverse-arp;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (disable | seconds);
        shaping {
            (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
                burst length);
            queue-length number;
        }
        vci vpi-identifier.vci-identifier;
    }
    primary;
    preferred;
    (vrrp-group | vrrp-inet6-group) group-number {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        authentication-type authentication;
        authentication-key key;
```

```
fast-interval milliseconds;  
(preempt | no-preempt) {  
    hold-time seconds;  
}  
priority-number number;  
track {  
    priority-cost seconds;  
    priority-hold-time interface-name {  
        interface priority;  
        bandwidth-threshold bits-per-second {  
            priority;  
        }  
    }  
    route ip-address/mask routing-instance instance-name priority-cost cost;  
}  
virtual-address [ addresses ];  
}  
}
```

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*]

For information about interface-specific protocol and address properties, see *Configuring T1 and NxDSO Interfaces*.

PART 2

Configuration

- [Protocol Family and Interface Address Properties on page 9](#)
- [Network Interfaces Configuration Statements and Hierarchy on page 71](#)
- [Statement Summary on page 95](#)

CHAPTER 2

Protocol Family and Interface Address Properties

- [Setting the Protocol MTU on page 9](#)
- [Disabling the Removal of Address and Control Bytes on page 10](#)
- [Disabling the Transmission of Redirect Messages on an Interface on page 10](#)
- [Applying Policers on page 11](#)
- [Applying a Filter to an Interface on page 19](#)
- [Enabling Source Class and Destination Class Usage on page 24](#)
- [Configuring the Protocol Family on page 32](#)
- [Configuring the Interface Address on page 35](#)
- [Configuring ICCP for MC-LAG on page 39](#)
- [Configuring IPCP Options on page 41](#)
- [Configuring an Unnumbered Interface on page 43](#)
- [Configuring Default, Primary, and Preferred Addresses and Interfaces on page 49](#)
- [Configuring Unicast RPF on page 51](#)
- [Configuring Targeted Broadcast on page 57](#)
- [Understanding Targeted Broadcast on page 59](#)
- [Example: Configuring Unicast Reverse-Path-Forwarding Check on page 60](#)

Setting the Protocol MTU

When you initially configure an interface, the protocol maximum transmission unit (MTU) is calculated automatically. If you subsequently change the media MTU, the protocol MTU on existing address families automatically changes.

For a list of default protocol MTU values, see *Configuring the Media MTU*.

To modify the MTU for a particular protocol family, include the **mtu** statement:

mtu *bytes*;

You can include this statement at the following hierarchy levels:

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

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

If you increase the size of the protocol MTU, you must ensure that the size of the media MTU is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. For a list of encapsulation overhead values, see *Configuring the Media MTU*. If you reduce the media MTU size, but there are already one or more address families configured and active on the interface, you must also reduce the protocol MTU size. (You configure the media MTU by including the **mtu** statement at the [edit interfaces *interface-name*] hierarchy level, as discussed in *Configuring the Media MTU*.)



NOTE: Changing the media MTU or protocol MTU causes an interface to be deleted and added again.

The maximum number of data-link connection identifiers (DLCIs) is determined by the MTU on the interface. If you have keepalives enabled, the maximum number of DLCIs is 1000, with the MTU set to 5012.

The actual frames transmitted also contain cyclic redundancy check (CRC) bits, which are not part of the MTU. For example, the default protocol MTU for a Gigabit Ethernet interface is 1500 bytes, but the largest possible frame size is actually 1504 bytes; you need to consider the extra bits in calculations of MTUs for interoperability.

Disabling the Removal of Address and Control Bytes

For Point-to-Point Protocol (PPP) CCC-encapsulated interfaces, the address and control bytes are removed by default before the packet is encapsulated into a tunnel.

You can disable the removal of address and control bytes. To do this, include the **keep-address-and-control** statement:

```
keep-address-and-control;
```

You can include this statement at the following hierarchy levels:

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

Disabling the Transmission of Redirect Messages on an Interface

By default, the interface sends protocol redirect messages. To disable the sending of these messages on an interface, include the **no-redirects** statement:

```
no-redirects;
```

You can include this statement at the following hierarchy levels:

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

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

To disable the sending of protocol redirect messages for the entire router or switch, include the **no-redirects** statement at the [edit system] hierarchy level.

Applying Policers

Policers allow you to perform simple traffic policing on specific interfaces or Layer 2 virtual private networks (VPNs) without configuring a firewall filter. To apply policers, include the **policer** statement:

```
policer {
  arp policer-template-name;
  input policer-template-name;
  output policer-template-name;
}
```

You can include these statements at the following hierarchy levels:

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

In the **family** statement, the protocol family can be **ccc**, **inet**, **inet6**, **mpls**, **tcc**, or **vpls**.

In the **arp** statement, list the name of one policer template to be evaluated when Address Resolution Protocol (ARP) packets are received on the interface. By default, an ARP policer is installed that is shared among all the Ethernet interfaces on which you have configured the **family inet** statement. If you want more stringent or lenient policing of ARP packets, you can configure an interface-specific policer and apply it to the interface. You configure an ARP policer just as you would configure any other policer, at the [edit firewall policer] hierarchy level. If you apply this policer to an interface, the default ARP packet policer is overridden. If you delete this policer, the default policer takes effect again.

In the **input** statement, list the name of one policer template to be evaluated when packets are received on the interface.

In the **output** statement, list the name of one policer template to be evaluated when packets are transmitted on the interface.



NOTE: To use policing on a CCC or TCC interface, you must configure the CCC or TCC protocol family.

You can configure a different policer on each protocol family on an interface, with one input policer and one output policer for each family. When you apply policers, you can configure the family **ccc**, **inet**, **inet6**, **mpls**, **tcc**, or **vpls** only, and one ARP policer for the family **inet** protocol only. Each time a policer is referenced, a separate copy of the policer is installed on the packet forwarding components for that interface.

If you apply both policers and firewall filters to an interface, input policers are evaluated before input firewall filters, and output policers are evaluated after output firewall filters.

If you apply the policer to the interface **lo0**, it is applied to packets received or transmitted by the Routing Engine.

On T Series, M120, and M320 platforms, if the interfaces are on the same FPC, the filters or policers do not act on the sum of traffic entering and exiting the interfaces.

For more information about policers, see the *Routing Policy Feature Guide for Routing Devices*.

This section includes the following topics:

- [Applying Aggregate Policers on page 12](#)
- [Applying Hierarchical Policers on Enhanced Intelligent Queuing PICs on page 14](#)

Applying Aggregate Policers

By default, if you apply a policer to multiple protocol families on the same logical interface, the policer restricts traffic for each protocol family individually. For example, a policer with a 50 Mbps bandwidth limit applied to both IPv4 and IPv6 traffic would allow the interface to accept 50 Mbps of IPv4 traffic and 50 Mbps of IPv6 traffic. If you apply an aggregate policer, the policer would allow the interface to receive only 50 Mbps of IPv4 and IPv6 traffic combined.

To configure an aggregate policer, include the **logical-interface-policer** statement at the **[edit firewall policer *policer-template-name*]** hierarchy level:

```
[edit firewall policer policer-template-name]  
logical-interface-policer;
```

For the policer to be treated as an aggregate, you must apply it to multiple protocol families on a single logical interface by including the **policer** statement:

```
policer {  
  arp policer-template-name;  
  input policer-template-name;  
  output policer-template-name;  
}
```

You can include these statements at the following hierarchy levels:

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

In the **family** statement, the protocol family can be **ccc**, **inet**, **inet6**, **mpls**, **tcc**, or **vpls**.

The protocol families on which you do not apply the policer are not affected by the policer. For example, if you configure a single logical interface to accept MPLS, IPv4, and IPv6 traffic and you apply the logical interface policer **policer1** to only the IPv4 and IPv6 protocol families, MPLS traffic is not subject to the constraints of **policer1**.

If you apply **policer1** to a different logical interface, there are two instances of the policer. This means the Junos OS polices traffic on separate logical interfaces separately, not as an aggregate, even if the same logical-interface policer is applied to multiple logical interfaces on the same physical interface port.

Example: Applying Aggregate Policers

Configure two logical interface policers: **aggregate_police1** and **aggregate_police2**. Apply **aggregate_police1** to IPv4 and IPv6 traffic received on logical interface **fe-0/0/0.0**. Apply **aggregate_police2** to CCC and MPLS traffic received on logical interface **fe-0/0/0.0**. This configuration causes the software to create only one instance of **aggregate_police1** and one instance of **aggregate_police2**.

Apply **aggregate_police1** to IPv4 and IPv6 traffic received on another logical interface **fe-0/0/0.1**. This configuration causes the software to create a new instance of **aggregate_police1**, one that applies to unit 0 and another that applies to unit 1.

```
[edit firewall]
policer aggregate_police1 {
  logical-interface-policer;
  if-exceeding {
    bandwidth-limit 100m;
    burst-size-limit 500k;
  }
  then {
    discard;
  }
}
policer aggregate_police2 {
  logical-interface-policer;
  if-exceeding {
    bandwidth-limit 10m;
    burst-size-limit 200k;
  }
  then {
    discard;
  }
}
[edit interfaces fe-0/0/0]
unit 0 {
  family inet {
    policer {
      input aggregate_police1;
    }
  }
  family inet6 {
    policer {
      input aggregate_police1;
    }
  }
  family ccc {
    policer {
      input aggregate_police2;
    }
  }
}
```

```
family mpls {
    policer {
        input aggregate_police2;
    }
}
unit 1 {
    family inet {
        policer {
            input aggregate_police1;
        }
    }
    family inet6 {
        policer {
            input aggregate_police1;
        }
    }
}
```

Applying Hierarchical Policers on Enhanced Intelligent Queuing PICs

M40e, M120, and M320 edge routers and T Series core routers with Enhanced Intelligent Queuing (IQE) PICs support hierarchical policers in the ingress direction and allow you to apply a hierarchical policer for the premium and aggregate (premium plus normal) traffic levels to an interface. Hierarchical policers provide cross-functionality between the configured physical interface and the Packet Forwarding Engine.

Before you begin, there are some general restrictions that apply to hierarchical policers:


- Only one type of policer can be configured for a logical or physical interface. For example, a hierarchical policer and a regular policer in the same direction for the same logical interface is not allowed.
- The chaining of the policers—that is, applying policers to both a port and the logical interfaces of that port—is not allowed.
- There is a limit of 64 policers per interface in case there is no BA classification, providing a single policer per DLCI.
- Only one kind of policer can be applied on a physical or logical interface.
- The policer should be independent of BA classification. Without BA classification, all traffic on an interface will be treated either as EF or non-EF, based on the configuration. With BA classification, an interface can support up to 64 policers. Again, the interface here may be a physical interface or logical interface (for example, DLCI).
- With BA classification, the miscellaneous traffic (the traffic *not* matching with any of the BA classification DSCP/EXP bits) will be policed as non-EF traffic. No separate policers will be installed for this traffic.

Hierarchical Policer Overview

Hierarchical policing uses two token buckets, one for aggregate (non-EF) traffic and one for premium (EF) traffic. Which traffic is EF and which is non-EF is determined by the

class-of-service configuration. Logically, hierarchical policing is achieved by chaining two policers.

Figure 1: Hierarchical Policer



In the example in [Figure 1 on page 15](#), EF traffic is policed by Premium Policer and non EF traffic is policed by Aggregate Policer. What that means is, for EF traffic the out-of-spec action will be the one that is configured for Premium Policer, but the in-spec EF traffic will still consume the tokens from the Aggregate Policer.

But EF traffic will never be submitted to the out-of-spec action of the Aggregate Policer. Also, if the out-of-spec action of the Premium Policer is not set to Discard, those out-of-spec packets will not consume the tokens from the Aggregate Policer. Aggregate Policer only polices the non-EF traffic. As you can see, the Aggregate Policer token bucket can go negative, if all the tokens are consumed by the non-EF traffic and then you get bursts of EF traffic. But that will be for a very short time, and over a period of time it will average out. For example:

- *Premium Policer*: Bandwidth 2 Mbps, OOS Action: Discard
- *Aggregate Policer*: Bandwidth 10 Mbps, OOS Action: Discard

In the above case, EF traffic is guaranteed 2 Mbps and the non-EF traffic will get from 8 Mbps to 10 Mbps, depending on the input rate of the EF traffic.

Hierarchical Policing Characteristics

Hierarchical token bucket features include:

- Ingress traffic is first classified into EF and non-EF traffic prior to applying a policer:
 - Classification is performed by Q-tree lookup
- Channel number selects a shared token bucket policer:
 - Dual token bucket policer is divided into two single bucket policers:
 - Policer1—EF traffic
 - Policer2—non-EF traffic
- Shared token bucket is used to police the traffic as follows:
 - Policer1 is set to EF rate (for example, 2 Mbps)
 - Policer2 is set to aggregate interface policed rate (for example, 10 Mbps).
 - EF traffic gets applied to Policer1.
 - If traffic is in-spec it is allowed to pass and decrement from both Policer1 and Policer2.

- If traffic is out-of-spec it can be discarded or marked with a new FC or loss priority. Policer2 will not do anything with out-of-spec EF traffic.
- Non-EF traffic gets applied only to Policer2.
 - If traffic is in-spec it is allowed to pass through and decremented Policer2.
 - If traffic is out-of-spec it is discarded or marked with a new FC or set with a new drop priority.
- Rate-limit the port speed to a desired rate at Layer 2
- Rate-limit the EF traffic
- Rate-limit the non-EF traffic
- Policing drops counted per color

Configuring Hierarchical Policers

To configure a hierarchical policer, apply the **policing-priority** statement to the proper forwarding class and configure a hierarchical policer for the aggregate and premium level. For more information about class of service, see the *Junos OS Class of Service Library for Routing Devices*.



NOTE: Hierarchical policers can only be configured on SONET physical interfaces hosted on an IQE PIC. Only aggregate and premium levels are supported.

CoS Configuration of Forwarding Classes for Hierarchical Policers

```
[edit class-of-service forwarding-classes]
class fc1 queue-num 0 priority high policing-priority premium;
class fc2 queue-num 1 priority low policing-priority normal;
class fc3 queue-num 2 priority low policing-priority normal;
class fc4 queue-num 3 priority low policing-priority normal;
```

For detailed information on class-of-service configuration and statements, see the *Junos OS Class of Service Library for Routing Devices*.

Firewall Configuration for Hierarchical Policers

```
[edit firewall hierarchical-policer foo]
aggregate {
  if-exceeding {
    bandwidth-limit 70m;
    burst-size-limit 1500;
  }
  then {
    discard;
  }
}
premium {
  if-exceeding {
    bandwidth-limit 50m;
    burst-size-limit 1500;
  }
  then {
    discard;
  }
}
```

```
}  
}
```

You can apply the hierarchical policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]  
input-hierarchical-policer foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]  
input-hierarchical-policer foo;
```

Configuring a Single-Rate Two-Color Policer

You can configure a single-rate two-color policer as follows:

```
[edit firewall policer foo]  
if-exceeding {  
    bandwidth-limit 50m;  
    burst-size-limit 1500;  
}  
then {  
    discard;  
}  
}
```

You can apply the policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]  
input-policer foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]  
input-policer foo;
```

Configuring a Single-Rate Tricolor Policer

This section describes single-rate color blind and color aware policers.

Configuring a Single-Rate Color-Blind Policer

You can configure a single-rate color blind policer as follows:

```
[edit firewall three-color-policer foo]  
single-rate {  
    color-blind;  
    committed-information-rate 50m;  
    committed-burst-size 1500;  
    excess-burst-size 1500;  
}  
}
```

You can apply the single-rate color blind policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]  
input-three-color foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]
```

```
input-three-color foo;
```

Configuring a Single-Rate Color-Aware Policer

You can configure a single-rate color-aware policer as follows:

```
[edit firewall three-color-policer bar]
single-rate {
  color-aware;
  committed-information-rate 50m;
  committed-burst-size 1500;
  excess-burst-size 1500;
}
```

You can apply the single-rate color-aware policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]
input-three-color foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]
input-three-color bar;
```

Configuring a Two-Rate Tricolor Marker Policer

Ingress policing is implemented using a two-rate tricolor marker (trTCM). This is done with a dual token bucket (DTB) that maintains two rates, committed, and a peak. Egress static policing also uses a token bucket.

The token buckets perform the following ingress policing functions:

- (1K) trTCM - Dual token bucket (red, yellow, and green marking)
- Policing is based on Layer 2 packet size:
 - After +/- byte adjust offset
- Marking is color aware and color blind:
 - Color aware needs to have the color set by q-tree lookup based on:
 - ToS
 - EXP
- Programmable marking actions:
 - Color (red, yellow, green)
 - Drop based on color and congestion profile
- Policer is selected based on the arriving channel number:
 - Channel number LUT produces policer index and queue index
 - Multiple channels can share the same policer (LUT produces same policer index)
- Support ingress policing and trTCM at the following levels:

- Queue
- Logical interface (ifl/DLCI)
- Physical interface (ifd)
- Physical port (controller ifd)
- Any combinations of logical interface, physical interface, and port
- Support percentage of interface speed and bits per second

Rate limits may be applied to selected queues on ingress and on predefined queues at egress. The token bucket operates in color aware and color blind modes (specified by RFC 2698).

Configuring a Color-Blind trTCM

```
[edit firewall three-color-policer foo]
two-rate {
  color-blind;
  committed-information-rate 50m;
  committed-burst-size 1500;
  peak-information-rate 100m;
  peak-burst-size 3k;
}
```

You can apply the three-color two-rate color-blind policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]
input-three-color foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]
input-three-color foo;
```

Configuring a Color-Aware trTCM

```
[edit firewall three-color-policer bar]
two-rate {
  color-aware;
  committed-information-rate 50m;
  committed-burst-size 1500;
  peak-information-rate 100m;
  peak-burst-size 3k;
}
```

You can apply the three-color two-rate color-aware policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]
input-three-color bar;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]
input-three-color bar;
```

Applying a Filter to an Interface

To apply firewall filters to an interface, include the **filter** statement:

```
filter {  
    group filter-group-number;  
    input filter-name;  
    input-list [ filter-names ];  
    output filter-name;  
    output-list [ filter-names ];  
}
```

To apply a single filter, include the **input** statement:

```
filter {  
    input filter-name;  
}
```

To apply a list of filters to evaluate packets received on an interface, include the **input-list** statement.

```
filter {  
    input-list [ filter-names ];  
}
```

Up to 16 filter names can be included in an input list.

To apply a list of filters to evaluate packets transmitted on an interface, include the **output-list** statement.

```
filter {  
    output-list [ filter-names ];  
}
```

When you apply filters using the **input-list** statement or the **output-list** statement, a new filter is created with the name *<interface-name>.<unit-direction>*. This filter is exclusively interface-specific.

You can include these statements at the following hierarchy levels:

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

In the **family** statement, the protocol family can be **ccc**, **inet**, **inet6**, **mpls**, or **vpls**.

In the **group** statement, specify the interface group number to associate with the filter.

In the **input** statement, list the name of one firewall filter to be evaluated when packets are received on the interface.

In the **input-list** statement, list the names of filters to evaluate when packets are received on the interface. You can include up to 16 filter names.

In the **output** statement, list the name of one firewall filter to be evaluated when packets are transmitted on the interface.



NOTE: Output filters do not work for broadcast and multicast traffic, including VPLS traffic (except in MX Series routers with MPC/MIC interfaces), as shown in “[Example: Applying a Filter to an Interface](#)” on page 22.



NOTE: On an MX Series router, you cannot apply as an output filter, a firewall filter configured at the `[edit firewall filter family ccc]` hierarchy level. Firewall filters configured for the `family ccc` statement can be applied only as input filters.

In the **output-list** statement, list the names of filters to evaluate when packets are transmitted on the interface. You can include up to 16 filter names.

You can use the same filter one or more times. On M Series routers (except the M320 and M120 routers), if you apply a firewall filter or policer to multiple interfaces, the filter or policer acts on the sum of traffic entering or exiting those interfaces.

On T Series, M120, and M320 routers, interfaces are distributed among multiple packet forwarding components. Therefore, on these routers, if you apply a firewall filter or policer to multiple interfaces, the filter or policer acts on the traffic stream entering or exiting each interface, regardless of the sum of traffic on the multiple interfaces.

For more information on Understanding Ethernet Frame Statistics, see the *MX Series Layer 2 Configuration Guide*.

If you apply the filter to the interface **lo0**, it is applied to packets received or transmitted by the Routing Engine. You cannot apply MPLS filters to the management interface (**fxp0** or **em0**) or the loopback interface (**lo0**).

Filters applied at the `[set interfaces lo0 unit 0 family any filter input]` hierarchy level are not installed on T4000 Type 5 FPCs.

For more information about firewall filters, see the *Routing Policy Feature Guide for Routing Devices*. For more information about MPLS filters, see the *Junos OS MPLS Applications Library for Routing Devices*.

See also the following sections:

- [Defining Interface Groups in Firewall Filters on page 21](#)
- [Filter-Based Forwarding on the Output Interface on page 22](#)
- [Example: Applying a Filter to an Interface on page 22](#)

Defining Interface Groups in Firewall Filters

When applying a firewall filter, you can define an interface to be part of an *interface group*. Packets received on that interface are tagged as being part of the group. You can then match these packets using the **interface-group** match statement, as described in the *Routing Policy Feature Guide for Routing Devices*.

To define the interface to be part of an interface group, include the **group** statement:

```
group filter-group-number;
```

You can include this statement at the following hierarchy levels:

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



NOTE: The number 0 is not a valid interface group number.

Filter-Based Forwarding on the Output Interface

If port-mirrored packets are to be distributed to multiple monitoring or collection interfaces, based on patterns in packet headers, it is helpful to configure a filter-based forwarding (FBF) filter on the port-mirroring egress interface.

When an FBF filter is installed as an output filter, a packet that is forwarded to the filter has already undergone at least one route lookup. After the packet is classified at the egress interface by the FBF filter, it is redirected to another routing table for additional route lookup. To avoid packet looping inside the Packet Forwarding Engine, the route lookup in the latter routing table (designated by an FBF routing instance) must result in a different next hop from any next hop specified in a table that has already been applied to the packet.

If an input interface is configured for FBF, the source lookup is disabled for those packets headings to a different routing instance, since the routing table is not set up to handle the source lookup.

For more information about FBF configuration, see the *Junos OS Routing Protocols Library for Routing Devices*. For more information about port mirroring, see the *Junos OS Services Interfaces Library for Routing Devices*.

Example: Applying a Filter to an Interface

Input Filter for VPLS Traffic

For M Series and T Series routers only, apply an input filter to VPLS traffic. Output filters do not work for broadcast and multicast traffic, including VPLS traffic. Note that on MX Series routers with MPC/MIC interfaces, the VPLS filters on the egress is applicable to broadcast, multicast, and unknown unicast traffic.



NOTE:

```
[edit interfaces]
fe-2/2/3 {
  vlan-tagging;
  encapsulation vlan-vpls;
  unit 601 {
    encapsulation vlan-vpls;
    vlan-id 601;
```



```

family vpls {
  filter {
    input filter1; # Works for multicast destination MAC address
    output filter1; # Does not work for multicast destination MAC address
  }
}
}
[edit firewall]
family vpls {
  filter filter1 {
    term 1 {
      from {
        destination-mac-address {
          01:00:0c:cc:cc:cd/48;
        }
      }
      then {
        discard;
      }
    }
    term 2 {
      then {
        accept;
      }
    }
  }
}
}

```

Filter-Based Forwarding at the Output Interface

The following example illustrates the configuration of filter-based forwarding at the output interface. In this example, the packet flow follows this path:

1. A packet arrives at interface **fe-1/2/0.0** with source and destination addresses **10.50.200.1** and **10.50.100.1** respectively.
2. The route lookup in routing table **inet.0** points to the egress interface **so-0/0/3.0**.
3. The output filter installed at **so-0/0/3.0** redirects the packet to routing table **fbf.inet.0**.
4. The packet matches the entry **10.50.100.0/25** in the **fbf.inet.0** table, and finally leaves the router from interface **so-2/0/0.0**.

```

[edit interfaces]
so-0/0/3 {
  unit 0 {
    family inet {
      filter {
        output fbf;
      }
      address 10.50.10.2/25;
    }
  }
}
fe-1/2/0 {
  unit 0 {
    family inet {

```

```
        address 10.50.50.2/25;
    }
}
so-2/0/0 {
    unit 0 {
        family inet {
            address 10.50.20.2/25;
        }
    }
}
[edit firewall]
filter fbf {
    term 0 {
        from {
            source-address {
                10.50.200.0/25;
            }
        }
        then routing-instance fbf;
    }
    term d {
        then count d;
    }
}
[edit routing-instances]
fbf {
    instance-type forwarding;
    routing-options {
        static {
            route 10.50.100.0/25 next-hop so-2/0/0.0;
        }
    }
}
[edit routing-options]
interface-routes {
    rib-group inet fbf-group;
}
static {
    route 10.50.100.0/25 next-hop 10.50.10.1;
}
rib-groups {
    fbf-group {
        import-rib [inet.0 fbf.inet.0];
    }
}
```

Enabling Source Class and Destination Class Usage

For interfaces that carry IPv4, IPv6, MPLS, or peer AS billing traffic, you can maintain packet counts based on the entry and exit points for traffic passing through your network. Entry and exit points are identified by source and destination prefixes grouped into disjoint sets defined as *source classes* and *destination classes*. You can define classes based on a variety of parameters, such as routing neighbors, autonomous systems, and route filters.

Source class usage (SCU) counts packets sent to customers by performing lookup on the IP source address. SCU makes it possible to track traffic originating from specific prefixes on the provider core and destined for specific prefixes on the customer edge. You must enable SCU accounting on both the inbound and outbound physical interfaces, and the route for the source of the packet must be in located in the forwarding table.



NOTE: SCU and DCU accounting do not work with directly connected interface routes. Source class usage does not count packets coming from sources with direct routes in the forwarding table because of software architecture limitations.

Destination class usage (DCU) counts packets from customers by performing lookup of the IP destination address. DCU makes it possible to track traffic originating from the customer edge and destined for specific prefixes on the provider core router.



NOTE: SCU and DCU accounting are supported on the J Series router only for IPv4 and IPv6 traffic.

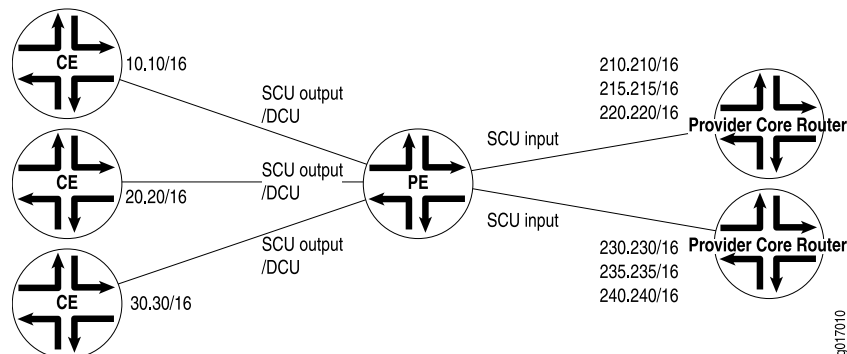


NOTE: We recommend that you stop the network traffic on an interface before you modify the DCU or SCU configuration for that interface. Modifying the DCU or SCU configuration without stopping the traffic might corrupt the DCU or SCU statistics. Before you restart the traffic after modifying the configuration, enter the `clear interfaces statistics` command.

Figure 2 on page 25 illustrates an Internet service provider (ISP) network. In this topology, you can use DCU to count packets customers send to specific prefixes. For example, you can have three counters, one per customer, that count the packets destined for prefix 210.210/16 and 220.220/16.

You can use SCU to count packets the provider sends from specific prefixes. For example, you can count the packets sent from prefix 210.210/16 and 215.215/16 and transmitted on a specific output interface.

Figure 2: Prefix Accounting with Source and Destination Classes



You can configure up to 126 source classes and 126 destination classes. For each interface on which you enable destination class usage and source class usage, the Junos OS maintains an interface-specific counter for each corresponding class up to the 126 class limit.



NOTE: For transit packets exiting the router through the tunnel, forwarding path features, such as RPF, forwarding table filtering, source class usage, and destination class usage are not supported on the interfaces you configure as the output interface for tunnel traffic. For firewall filtering, you must allow the output tunnel packets through the firewall filter applied to input traffic on the interface that is the next-hop interface towards the tunnel destination.



NOTE:

Performing DCU accounting when an output service is enabled produces inconsistent behavior in the following configuration:

- Both SCU input and DCU are configured on the packet input interface.
- SCU output is configured on the packet output interface.
- Interface services is enabled on the output interface.

For an incoming packet with source and destination prefixes matching the SCU and DCU classes respectively configured in the router, both SCU and DCU counters will be incremented. This behavior is not harmful or negative. However, it is inconsistent with non-serviced packets, in that only the SCU count will be incremented (because the SCU class ID will override the DCU class ID in this case).

To enable packet counting on an interface, include the **accounting** statement:

```
accounting {  
  destination-class-usage;  
  source-class-usage {  
    direction;  
  }  
}
```

direction can be one of the following:

- **input**—Configure at least one expected ingress point.
- **output**—Configure at least one expected egress point.
- **input output**—On a single interface, configure at least one expected ingress point and one expected egress point.

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6 | mpls)]

- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family (inet | inet6 | mpls)]`

For SCU to work, you must configure at least one input interface and at least one output interface.

The ability to count a single packet for both SCU and DCU accounting depends on the underlying physical interface.

- For traffic over interfaces on Junos Trio chipset-based Packet Forwarding Engine, a single incoming packet is counted for both SCU and DCU accounting if both SCU and DCU are configured. To ensure the outgoing packet is counted, include the **source-class-usage output** statements in the configuration of the outgoing interface.
- For traffic over other interfaces, an incoming packet is counted only once, and SCU takes priority over DCU. This means that when a packet arrives on an interface on which you include the **source-class-usage input** and **destination-class-usage** statements in the configuration, and when the source and destination both match accounting prefixes, the Junos OS associates the packet with the source class only.

For traffic over interfaces on Junos Trio chipset-based Packet Forwarding Engine, SCU and DCU accounting is performed after output filters are evaluated. If a packet matches a firewall filter match condition, the packet is included in SCU or DCU accounting except in the case where the action of the matched term is **discard**.

On T Series, M120, and M320 routers, the source class and destination classes are not carried across the router fabric. The implications of this are as follows:

- On T Series, M120, and M320 routers, SCU and DCU accounting is performed before the packet enters the fabric.
- On M7i, M10i, M120, and M320 routers, on MX Series routers with non-Trio chipset-based Packet Forwarding Engine, and on T Series routers, SCU and DCU accounting is performed before output filters are evaluated. Consequently, if a packet matches a firewall filter match condition, the packet is included in SCU or DCU accounting; the packet is counted for any term action (including the **discard** action).
- On M120, M320, and T Series routers, the **destination-class** and **source-class** statements are supported at the `[edit firewall family family-name filter filter-name term term-name from]` hierarchy level only for the filter applied to the forwarding table. On M7i, M10i, and MX Series routers, these statements are supported.

Once you enable accounting on an interface, the Junos OS maintains packet counters for that interface, with separate counters for **inet**, **inet6**, and **mpls** protocol families. You must then configure the source class and destination class attributes in policy action statements, which must be included in forwarding-table export policies.



NOTE: When configuring policy action statements, you can configure only one source class for each matching route. In other words, more than one source class cannot be applied to the same route.

In Junos OS Release 9.3 and later, you can configure SCU accounting for Layer 3 VPNs configured with the **vrf-table-label** statement. Include the **source-class-usage** statement at the **[edit routing-instances routing-instance-name vrf-table-label]** hierarchy level. The **source-class-usage** statement at this hierarchy level is supported only for the virtual routing and forwarding (VRF) instance type.



NOTE: DCU counters cannot be enabled on the label-switched interface (LSI) that is created dynamically when the **vrf-table-label** statement is configured within a VRF. For more information, see the *Junos OS VPNs Library for Routing Devices*.

For a complete discussion about source and destination class accounting profiles, see the *Network Management Administration Guide for Routing Devices*. For more information about MPLS, see the *Junos OS MPLS Applications Library for Routing Devices*.

Examples: Enabling Source Class and Destination Class Usage

Configure DCU and SCU output on one interface:

```
[edit]
interfaces {
  so-6/1/0 {
    unit 0 {
      family inet {
        accounting {
          destination-class-usage;
          source-class-usage {
            output;
          }
        }
      }
    }
  }
}
```

Complete SCU Configuration

Source routers A and B use loopback addresses as the prefixes to be monitored. Most of the configuration tasks and actual monitoring occur on transit Router SCU.

The loopback address on Router A contains the origin of the prefix that is to be assigned to source class A on Router SCU. However, no SCU processing happens on this router. Therefore, configure Router A for basic OSPF routing and include your loopback interface and interface **so-0/0/2** in the OSPF process.

Router A

```
[edit]
interfaces {
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.255.50.2/24;
      }
    }
  }
  lo0 {
```

```

    unit 0 {
        family inet {
            address 10.255.192.10/32;
        }
    }
}
protocols {
    ospf {
        area 0.0.0.0 {
            interface so-0/0/2.0;
            interface lo0.0;
        }
    }
}

```

Router SCU Last, apply the policy to the forwarding table.

Router SCU handles the bulk of the activity in this example. On Router SCU, enable source class usage on the inbound and outbound interfaces at the **[edit interfaces *interface-name* unit *unit-number* family inet accounting]** hierarchy level. Make sure you specify the expected traffic: input, output, or, in this case, both.

Next, configure a route filter policy statement that matches the prefixes of the loopback addresses from routers A and B. Include statements in the policy that classify packets from Router A in one group named **scu-class-a** and packets from Router B in a second class named **scu-class-b**. Notice the efficient use of a single policy containing multiple terms.

```

[edit]
interfaces {
    so-0/0/1 {
        unit 0 {
            family inet {
                accounting {
                    source-class-usage {
                        input;
                        output;
                    }
                }
            }
            address 10.255.50.1/24;
        }
    }
}
so-0/0/3 {
    unit 0 {
        family inet {
            accounting {
                source-class-usage {
                    input;
                    output;
                }
            }
            address 10.255.10.3/24;
        }
    }
}

```

```
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.6.111/32;
    }
  }
}
}
protocols {
  ospf {
    area 0.0.0.0 {
      interface so-0/0/1.0;
      interface so-0/0/3.0;
    }
  }
}
routing-options {
  forwarding-table {
    export scu-policy;
  }
}
policy-options {
  policy-statement scu-policy {
    term 0 {
      from {
        route-filter 10.255.192.0/24 orlonger;
      }
      then source-class scu-class-a;
    }
    term 1 {
      from {
        route-filter 10.255.165.0/24 orlonger;
      }
      then source-class scu-class-b;
    }
  }
}
```

Router B Just as Router A provides a source prefix, Router B's loopback address matches the prefix assigned to **scu-class-b** on Router SCU. Again, no SCU processing happens on this router, so configure Router B for basic OSPF routing and include your loopback interface and interface **so-0/0/4** in the OSPF process.

```
interfaces {
  so-0/0/4 {
    unit 0 {
      family inet {
        address 10.255.10.4/24;
      }
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.165.226/32;
    }
  }
}
```



```

    }
  }
}
protocols {
  ospf {
    area 0.0.0.0 {
      interface so-0/0/4.0;
      interface lo0.0;
    }
  }
}

```

Enabling Packet Counting for Layer 3 VPNs

You can use SCU and DCU to count packets on Layer 3 VPNs. To enable packet counting for Layer 3 VPN implementations at the egress point of the MPLS tunnel, you must configure a virtual loopback tunnel interface (**vt**) on the PE router, map the virtual routing and forwarding (VRF) instance type to the virtual loopback tunnel interface, and send the traffic received from the VPN out the source class output interface, as shown in the following example:

Configure a virtual loopback tunnel interface on a provider edge router equipped with a tunnel PIC:

```

[edit interfaces]
vt-0/3/0 {
  unit 0 {
    family inet {
      accounting {
        source-class-usage {
          input;
        }
      }
    }
  }
}

```

Map the VRF instance type to the virtual loopback tunnel interface.

In Junos OS Release 9.3 and later, you can configure SCU accounting for Layer 3 VPNs configured with the **vrf-table-label** statement. Include the **source-class-usage** statement at the **[edit routing-instances routing-instance-name vrf-table-label]** hierarchy level. The **source-class-usage** statement at this hierarchy level is supported only for the virtual routing and forwarding (VRF) instance type. DCU is not supported when the **vrf-table-label** statement is configured. For more information, see the *Junos OS VPNs Library for Routing Devices*.

```

[edit routing-instances]
VPN-A {
  instance-type vrf;
  interface at-2/1/1.0;
  interface vt-0/3/0.0;
  route-distinguisher 10.255.14.225:100;
  vrf-import import-policy-A;
  vrf-export export-policy-A;
  protocols {
    bgp {

```

```
group to-r4 {  
    local-address 10.27.253.1;  
    peer-as 400;  
    neighbor 10.27.253.2;  
}  
}  
}
```

Send traffic received from the VPN out the source class output interface:

```
[edit interfaces]  
at-2/1/0 {  
    unit 0 {  
        family inet {  
            accounting {  
                source-class-usage {  
                    output;  
                }  
            }  
        }  
    }  
}
```

For more information about VPNs, see the *Junos OS VPNs Library for Routing Devices*. For more information about virtual loopback tunnel interfaces, see the *Junos OS Services Interfaces Library for Routing Devices*.

**Related
Documentation**

- [accounting on page 99](#)
- [destination-classes](#)
- [family on page 121](#)
- [forward-and-send-to-re on page 127](#)
- [source-classes](#)
- [targeted-broadcast on page 164](#)
- [unit on page 170](#)

Configuring the Protocol Family

For each logical interface, you can configure one or more of the following protocols that run on the interface:

- **any**—Protocol-independent family used for Layer 2 packet filtering. This option is not supported on J Series routers and on T4000 Type 5 FPCs.
- **bridge**—(M Series and T Series routers only) Configure only when the physical interface is configured with **ethernet-bridge** type encapsulation or when the logical interface is configured with **vlan-bridge** type encapsulation. You can optionally configure this protocol family for the logical interface on which you configure VPLS.

- **ccc**—Circuit cross-connect (CCC). You can configure this protocol family for the logical interface of CCC physical interfaces. When you use this encapsulation type, you can configure the **ccc** family only.
- **inet**—IP. You must configure this protocol family for the logical interface to support IP protocol traffic, including Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), Internet Control Message Protocol (ICMP), and Internet Protocol Control Protocol (IPCP).
- **inet6**—IP version 6 (IPv6). You must configure this protocol family for the logical interface to support IPv6 protocol traffic, including Routing Information Protocol for IPv6 (RIPng), Intermediate System-to-Intermediate System (IS-IS), BGP, and Virtual Router Redundancy Protocol for IPv6 (VRRP). For more information about IPv6, see [“IPv6 Overview” on page 34](#).
- **iso**—International Organization for Standardization (ISO). You must configure this protocol family for the logical interface to support IS-IS traffic.
- **mlfr-uni-nni**—Multilink Frame Relay (MLFR) FRF.16 user-to-network network-to-network (UNI NNI). You must configure this protocol or **mlfr-end-to-end** for the logical interface to support link services and voice services bundling.
- **mlfr-end-to-end**—Multilink Frame Relay end-to-end. You must configure this protocol or multilink Point-to-Point Protocol (MLPPP) for the logical interface to support multilink bundling.
- **mlppp**—MLPPP. You must configure this protocol (or **mlfr-end-to-end**) for the logical interface to support multilink bundling.
- **mpls**—Multiprotocol Label Switching (MPLS). You must configure this protocol family for the logical interface to participate in an MPLS path.
- **tcc**—Translational cross-connect (TCC). You can configure this protocol family for the logical interface of TCC physical interfaces.
- **tnp**—Trivial Network Protocol. This protocol is used to communicate between the Routing Engine and the router's packet forwarding components. The Junos OS automatically configures this protocol family on the router's internal interfaces only, as discussed in *Understanding Internal Ethernet Interfaces*.
- **vpls**—M Series and T Series routers support Virtual Private LAN service (VPLS). You can optionally configure this protocol family for the logical interface on which you configure VPLS. VPLS provides an Ethernet-based point-to-multipoint Layer 2 VPN to connect customer edge (CE) routers across an MPLS backbone. When you configure a VPLS encapsulation type, the **family vpls** statement is assumed by default.

MX Series routers support dynamic profiles for VPLS pseudowires, VLAN identifier translation, and automatic bridge domain configuration.

For more information about VPLS, see the *Junos OS VPNs Library for Routing Devices* and the *Junos OS Feature Guides*.

To configure the logical interface's protocol family, include the **family** statement, specifying the selected family. To configure more than one protocol **family** on a logical interface, include multiple **family** statements. Following is the minimum configuration:

```
family family {  
  mtu size;  
  multicast-only;  
  no-redirects;  
  primary;  
  address address {  
    destination address;  
    broadcast address;  
    preferred;  
    primary;  
  }  
}
```

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*]

IPv6 Overview

IP version 4 (IPv4) has been widely deployed and used to network the Internet today. With the rapid growth of the Internet, enhancements to IPv4 are needed to support the influx of new subscribers, Internet-enabled devices, and applications. IPv6 is designed to enable the global expansion of the Internet.

IPv6 builds upon the functionality of IPv4, providing improvements to addressing, configuration and maintenance, and security.

IPv6 is defined in the following documents:

- RFC 2373, *IP Version 6 Addressing Architecture*
- RFC 2460, *Internet Protocol, Version 6 (IPv6)*

IPv4-to-IPv6 Transition

Implementing IPv6 requires a transition mechanism to allow interoperability between IPv6 nodes (both routers and hosts) and IPv4 nodes. The transition mechanism is the key factor in the successful deployment of IPv6. Because millions of IPv4 nodes already exist, upgrading every node to IPv6 at the same time is not feasible.

As a result, transition from IPv4 to IPv6 happens gradually, allowing nodes to be upgraded independently and without disruption to other nodes. While a gradual upgrade occurs, compatibility between IPv6 and IPv4 nodes becomes a requirement. Otherwise, an IPv6 node would not be able to communicate with an IPv4 node.

Transition mechanisms allow IPv6 and IPv4 nodes to coexist together in the same network, and make gradual upgrading possible. The transition mechanism supported by the Junos OS is tunneling. Tunnels allow IPv6 packets to be encapsulated into IPv4 headers and sent across an IPv4 infrastructure. For more information about configuring tunnels to support IPv4-to-IPv6 transition, see the *Junos OS Services Interfaces Library for Routing Devices*.

VRRP Properties

The Virtual Router Redundancy Protocol (VRRP) provides a much faster switchover to a backup router when the default router fails. Using VRRP, a backup router can take over a failed default router within a few seconds. This is done with minimum amount of VRRP traffic and without any interactions with the hosts.

For more information on VRRP properties, see the *Junos OS High Availability Library for Routing Devices*.

Related Documentation

- [Understanding Internal Ethernet Interfaces](#)

Configuring the Interface Address

You assign an address to an interface by specifying the address when configuring the protocol family. For the **inet** or **inet6** family, configure the interface IP address. For the **iso** family, configure one or more addresses for the loopback interface. For the **ccc**, **ethernet-switching**, **tcc**, **mpls**, **tnp**, and **vpls** families, you never configure an address.



NOTE: The point-to-point (PPP) address is taken from the loopback interface address that has the primary attribute. When the loopback interface is configured as an unnumbered interface, it takes the primary address from the donor interface.

To assign an address to an interface, include the **address** statement:

```
address address {
  broadcast address;
  destination address;
  destination-profile name;
  eui-64;
  preferred;
  primary;
}
```

You can include these statements at the following hierarchy levels:

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

In the **address** statement, specify the network address of the interface.

For each address, you can optionally configure one or more of the following:

- Broadcast address for the interface subnet—Specify this in the **broadcast** statement; this applies only to Ethernet interfaces, such as the management interface **fxp0**, **em0**, or **me0** the Fast Ethernet interface, and the Gigabit Ethernet interface.
- Address of the remote side of the connection (for point-to-point interfaces only)—Specify this in the **destination** statement.
- PPP properties to the remote end—Specify this in the **destination-profile** statement. You define the profile at the **[edit access group-profile name ppp]** hierarchy level (for point-to-point interfaces only).
- Whether the router or switch automatically generates the host number portion of interface addresses—The **eui-64** statement applies only to interfaces that carry IPv6 traffic, in which the prefix length of the address is 64 bits or less, and the low-order 64 bits of the address are zero. This option does not apply to the loopback interface (**lo0**) because IPv6 addresses configured on the loopback interface must have a 128-bit prefix length.
- Whether this address is the preferred address—Each subnet on an interface has a preferred local address. If you configure more than one address on the same subnet, the preferred local address is chosen by default as the source address when you originate packets to destinations on the subnet.

By default, the preferred address is the lowest-numbered address on the subnet. To override the default and explicitly configure the preferred address, include the **preferred** statement when configuring the address.

- Whether this address is the primary address—Each interface has a primary local address. If an interface has more than one address, the primary local address is used by default as the source address when you send packets from an interface where the destination provides no information about the subnet (for example, some **ping** commands).

By default, the primary address on an interface is the lowest-numbered non-127 (in other words, non-loopback) preferred address on the interface. To override the default and explicitly configure the preferred address, include the **primary** statement when configuring the address.

- [Configuring Interface IPv4 Addresses on page 36](#)
- [Configuring Interface IPv6 Addresses on page 38](#)

Configuring Interface IPv4 Addresses

You can configure router or switch interfaces with a 32-bit IP version 4 (IPv4) address and optionally with a destination prefix, sometimes called a *subnet mask*. An IPv4 address utilizes a 4-octet dotted decimal address syntax (for example, **192.16.1.1**). An IPv4 address with destination prefix utilizes a 4-octet dotted decimal address syntax with a destination prefix appended (for example, **192.16.1.1/30**).

To configure an IPv4 address on routers and switches running Junos OS, use the **edit interface *interface-name* unit *number* family inet address *a.b.c.d/nn*** statement at the **[edit interfaces]** hierarchy level.



NOTE: Juniper Networks routers and switches support /31 destination prefixes when used in point-to-point Ethernet configurations; however, they are not supported by many other devices, such as hosts, hubs, routers, or switches. You must determine if the peer system also supports /31 destination prefixes before configuration.

Operational Behavior of Interfaces when the Same IPv4 Address is Assigned to Them

You can configure the same IPv4 address on multiple physical interfaces. When you assign the same IPv4 address to multiple physical interfaces, the operational behavior of those interfaces differs, depending on whether they are implicitly or explicitly point-to-point.



NOTE: By default, all interfaces are assumed to be point-to-point (PPP) interfaces. For all interfaces except aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet, you can explicitly configure an interface to be a point-to-point connection.

The following examples show the sample configuration of assigning the same IPv4 address to implicitly and explicitly point-to-point interfaces, and their corresponding **show interfaces terse** command outputs to see their operational status.

Configuring same IPv4 address on implicitly PPP interfaces:

```
[edit]
user@host# show
ge-0/1/0 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}
ge-3/0/1 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}
```

The sample output shown below for the above configuration reveals that only **ge-0/1/0.0** was assigned the same IPv4 address **200.1.1.1/24** and its **link** state was **up**, while **ge-3/0/1.0** was not assigned the IPv4 address, though its **link** state was up, which means that it will be operational only when it gets a unique IPv4 address other than **200.1.1.1/24**.

```
user@host> show interfaces terse ge*
```

Interface	Admin	Link	Proto	Local	Remote
ge-0/1/0		up	up		
ge-0/1/0.0		up	up	inet 200.1.1.1/24	
			multiservice		
ge-0/1/1		up	down		
ge-3/0/0		up	down		
ge-3/0/1		up	up		
ge-3/0/1.0		up	up	inet	
			multiservice		

Configuring same IPv4 address on explicitly PPP interfaces:

```
[edit]
user@host# show
so-0/0/0 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}
so-0/0/3 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}
```

The sample output shown below for the above configuration reveals that both **so-0/0/0.0** and **so-0/0/3.0** were assigned the same IPv4 address **200.1.1.1/24** and that their link states were down, which means that to make them operational at least one of them will have to be configured with a unique IPv4 address other than **200.1.1.1/24**.

```
user@host> show interfaces terse so*
Interface      Admin Link Proto  Local      Remote
so-0/0/0       up    up
so-0/0/0.0     up    down inet     200.1.1.1/24
so-0/0/1       up    up
so-0/0/2       up    down
so-0/0/3       up    up
so-0/0/3.0     up    down inet     200.1.1.1/24
so-1/1/0       up    down
so-1/1/1       up    down
so-1/1/2       up    up
so-1/1/3       up    up
so-2/0/0       up    up
so-2/0/1       up    up
so-2/0/2       up    up
so-2/0/3       up    down
```

Configuring Interface IPv6 Addresses



NOTE: IPv6 is not currently supported for the QFX Series.

You represent IP version 6 (IPv6) addresses in hexadecimal notation using a colon-separated list of 16-bit values.

You assign a 128-bit IPv6 address to an interface by including the **address** statement:

```
address aaaa:bbbb:...:zzzz/nn;
```



NOTE: You cannot configure a subnet zero IPv6 address because RFC 2461 reserves the subnet-zero address for anycast addresses, and Junos OS complies with the RFC.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet6]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6]

The double colon (::) represents all bits set to 0, as shown in the following example:

```
interfaces fe-0/0/1 {
  unit 0 {
    family inet6 {
      address fec0:1:1::2/64;
    }
  }
}
```



NOTE: You must manually configure the router or switch advertisement and advertise the default prefix for autoconfiguration to work on a specific interface.

Related Documentation

- [Configuring IPCP Options on page 41](#)
- [Configuring Default, Primary, and Preferred Addresses and Interfaces on page 49](#)

Configuring ICCP for MC-LAG

For multichassis link aggregation (MC-LAG), you must configure Inter-Control Center Communications Protocol (ICCP) to exchange information between two MC-LAG peers.

To enable ICCP, include the **iccp** statement at the [edit protocols] hierarchy level:

```
[edit protocols]
iccp {
  authentication-key string;
  local-ip-addr ipv4-address;
  peer ip-address {
    authentication-key string;
    liveness-detection {
```

```
    detection-time {  
        threshold milliseconds;  
    }  
    minimum-interval milliseconds;  
    minimum-receive-interval milliseconds;  
    multiplier number;  
    no-adaptation;  
    transmit-interval {  
        minimum-interval milliseconds;  
        threshold milliseconds;  
    }  
    version (1 | automatic);  
}  
local-ip-addr ipv4-address;  
redundancy-group-id-list [ redundancy-groups ];  
session-establishment-hold-time value;  
}  
session-establishment-hold-time value;  
traceoptions;  
}
```

The **local-ip-address** statement sets the source address. This could be a specified address or interface address. The **session-establishment-hold-time** statement determines whether a chassis takes over as the master at the ICCP session.

The **authentication-key** statement is provided by TCP Message Digest 5 (md5) option for an ICCP TCP session. The **redundancy-group-id-list** statement specifies the redundancy groups between ICCP peers and the **liveness-detection** hierarchy configures Bidirectional Forwarding Detection (BFD) protocol options.



NOTE: ICCP is based on TCP and it uses IP routes to reach the MC-LAG peer. To ensure that the ICCP session is as resilient as possible, we recommend that you configure alternative routes between the ICCP end-point IP addresses. Alternatively, configure a LAG interface that has two or more interfaces between the MC-LAG pairs to prevent session failure when there are no alternative routes.

For Inter-Control Center Communications Protocol (ICCP) in a multichassis link aggregation group (MC-LAG) configured in an active-active bridge domain, you must ensure that you configure the same peer IP address hosting the MC-LAG by including the **peer ip-address** statement at the **[edit protocols iccp]** hierarchy level and the **multi-chassis-protection peer ip-address** statement at the **[edit interfaces interface-name]** hierarchy level. Multichassis protection reduces the configuration at the logical interface level for MX Series routers with multichassis aggregated Ethernet (MC-AE) interfaces. If the ICCP is UP and the interchassis data link (ICL) comes UP, the router configured as standby will bring up the MC-AE interfaces shared with the peer active-active node specified by the **peer** statement.

For example, the following statements illustrate how the same peer IP address can be configured for both the ICCP peer and multichassis protection link:

```
set interfaces ae1 unit 0 multi-chassis-protection 10.255.34.112 interface ae0.0
set protocols iccp peer 10.255.34.112 redundancy-group-id-list 1
```

Although you can commit an MC-LAG configuration with various parameters defined for it, you can configure multichassis protection between two peers without configuring the ICCP peer address. You can also configure multiple ICCP peers and commit such a configuration.

Related Documentation

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

Configuring IPCP Options

For interfaces with PPP encapsulation, you can configure IPCP to negotiate IP address assignments and to pass network-related information such as Windows Name Service (WINS) and Domain Name System (DNS) servers, as defined in RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*.



NOTE: The Junos OS does not request name servers from the remote end; the software does, however, send name servers to the remote end if requested.

On the logical interface, the following PPP encapsulation types are supported:

- `atm-mlppp-llc`
- `atm-ppp-llc`
- `atm-ppp-vc-mux`
- `multilink-ppp`

When you enable a PPP interface, you can configure an IP address, enable the interface to negotiate an IP address assignment from the remote end, or allow the interface to be unnumbered. You can also assign a destination profile to the remote end. The destination profile includes PPP properties, such as primary and secondary DNS and NetBIOS Name Servers (NBNSs). These options are described in the following sections:

- [Configuring an IP Address for an Interface on page 41](#)
- [Negotiating an IP Address Assignment from the Remote End on page 42](#)
- [Configuring an Interface to Be Unnumbered on page 42](#)
- [Assigning a Destination Profile to the Remote End on page 43](#)

Configuring an IP Address for an Interface

You can configure an IP address for the interface by including the **address** statement in the configuration. For more information, see [“Configuring the Interface Address” on page 35](#).

If you include the **address** statement in the configuration, you cannot include the **negotiate-address** or **unnumbered-address** statement in the configuration.

When you include the **address** statement in the interface configuration, you can assign PPP properties to the remote end, as shown in [“Assigning a Destination Profile to the Remote End” on page 43](#).



NOTE: The option to negotiate an IP address is not allowed in MLFR and MFR encapsulations.

Negotiating an IP Address Assignment from the Remote End

To enable the interface to obtain an IP address from the remote end, include the **negotiate-address** statement:

```
negotiate-address;
```

You can include this statement at the following hierarchy levels:

- [edit [interfaces](#) *interface-name* [unit](#) *logical-unit-number* [family](#) inet]
- [edit [logical-systems](#) *logical-system-name* [interfaces](#) *interface-name* [unit](#) *logical-unit-number* [family](#) inet]

If you include the **negotiate-address** statement in the configuration, you cannot include the **address** or **unnumbered-address** statement in the configuration.

Configuring an Interface to Be Unnumbered

To configure an interface to be unnumbered, include the **unnumbered-address** and **destination** statements in the configuration:

```
unnumbered-address interface-name destination address;
```

The **unnumbered-address** statement enables the local address to be derived from the specified interface. The interface name must include a logical unit number and must have a configured address (see [“Configuring the Interface Address” on page 35](#)). Specify the IP address of the remote interface with the **destination** statement.

You can include these statements at the following hierarchy levels:

- [edit [interfaces](#) *interface-name* [unit](#) *logical-unit-number* [family](#) inet]
- [edit [logical-systems](#) *logical-system-name* [interfaces](#) *interface-name* [unit](#) *logical-unit-number* [family](#) inet]

If you include the **unnumbered-address** statement in the configuration, you cannot include the **address** or **negotiate-address** statement in the interface configuration.

When you include the **unnumbered-address** statement in the interface configuration, you can assign PPP properties to the remote end, as shown in [“Assigning a Destination Profile to the Remote End” on page 43](#).

Assigning a Destination Profile to the Remote End

When you include the **address** or **unnumbered-address** statement in the interface configuration, you can assign PPP properties to the remote end. To do this, include the **destination-profile** statement:

```
destination-profile name;
```

You can include this statement at the following hierarchy levels:

- [edit **interfaces** *interface-name* **unit** *logical-unit-number* **family** *inet* **address** *address*]
- [edit **interfaces** *interface-name* **unit** *logical-unit-number* **family** *inet* **unnumbered-address** *interface-name*]
- [edit **logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number* **family** *inet* **address** *address*]
- [edit **logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number* **family** *inet* **unnumbered-address** *interface-name*]

The profile name is a PPP group profile. You define the profile by including the following statements at the [edit **access group-profile** *name* **ppp**] hierarchy level:

```
[edit access group-profile name ppp]
framed-pool pool-id;
interface-id interface-id;
primary-dns primary-dns;
primary-wins primary-win-server;
secondary-dns secondary-dns;
secondary-wins secondary-wins;
```

For more information about PPP group profiles, see the *Junos OS Administration Library for Routing Devices*.

Configuring an Unnumbered Interface

When you need to conserve IP addresses, you can configure unnumbered interfaces. Setting up an unnumbered interface enables IP processing on the interface without assigning an explicit IP address to the interface. For IPv6, in which conserving addresses is not a major concern, you can configure unnumbered interfaces to share the same subnet across multiple interfaces. IPv6 unnumbered interfaces are only supported on Ethernet interfaces. The statements you use to configure an unnumbered interface depend on the type of interface you are configuring: a point-to-point interface or an Ethernet interface:

- [Configuring an Unnumbered Point-to-Point Interface on page 43](#)
- [Configuring an Unnumbered Ethernet or Demux Interface on page 44](#)

Configuring an Unnumbered Point-to-Point Interface

To configure an unnumbered point-to-point interface, configure the protocol family, but do not include the **address** statement:

family *family*;

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*]



NOTE: For interfaces with PPP encapsulation, you can configure an unnumbered interface by including the **unnumbered-interface** statement in the configuration. For more information, see [“Configuring IPCP Options” on page 41](#).

When configuring unnumbered interfaces, you must ensure that a source address is configured on some interface in the router. This address is the default address. We recommend that you do this by assigning an address to the loopback interface (**lo0**), as described in *Configuring the Loopback Interface*. If you configure an address (other than a martian) on the **lo0** interface, that address is always the default address, which is preferable because the loopback interface is independent of any physical interfaces and therefore is always accessible.

Example: Configuring an Unnumbered Point-to-Point Interface

Configure an unnumbered point-to-point interface:

```
[edit]
interfaces {
  so-6/1/0 {
    unit 0 {
      family inet;
      family iso;
    }
  }
}
```

Configuring an Unnumbered Ethernet or Demux Interface

To configure an unnumbered Ethernet or demultiplexing interface, include the **unnumbered-address** statement in the configuration:

unnumbered-address *interface-name*;

You can include this statement at the following hierarchy levels:

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

For dynamic profiles, include the **unnumbered-address** statement at the following hierarchy levels:

- [edit **dynamic-profiles** *profile-name* **interfaces** *interface-name* **unit** *logical-unit-number* **family** *family*]
- [edit **dynamic-profiles** *profile-name* **interfaces** **demux0** **unit** *logical-unit-number* **family** *family*]

The **unnumbered-address** statement currently supports configuration of unnumbered demux interfaces only for the IPv4 address family. You can configure unnumbered Ethernet interfaces for both IPv4 and IPv6 address families.

The interface that you configure to be unnumbered *borrow*s an assigned IP address from another interface, and is referred to as the *borrower interface*. The interface from which the IP address is borrowed is referred to as the *donor interface*. In the **unnumbered-address** statement, *interface-name* specifies the donor interface. For an unnumbered Ethernet interface, the donor interface can be an Ethernet, ATM, SONET, or loopback interface that has a logical unit number and configured IP address and is not itself an unnumbered interface. For an unnumbered IP demultiplexing interface, the donor interface can be an Ethernet or loopback interface that has a logical unit number and configured IP address and is not itself an unnumbered interface. In addition, for either Ethernet or demux, the donor interface and the borrower interface must be members of the same routing instance and the same logical system.

When you configure an unnumbered Ethernet or demux interface, the IP address of the donor interface becomes the source address in packets generated by the unnumbered interface.

You can configure a host route that points to an unnumbered Ethernet or demux interface. For information about host routes, see the *Junos OS MPLS Applications Library for Routing Devices*.

For more information, see the following sections:

- [Configuring a Preferred Source Address for Unnumbered Ethernet or Demux Interfaces on page 46](#)
- [Configuring Static Routes on Unnumbered Ethernet Interfaces on page 47](#)
- [Restrictions for Configuring Unnumbered Ethernet Interfaces on page 47](#)
- [Example: Configuring an Unnumbered Ethernet Interface on page 47](#)
- [Example: Configuring the Preferred Source Address for an Unnumbered Ethernet Interface on page 48](#)
- [Example: Configuring an Unnumbered Ethernet Interface as the Next Hop for a Static Route on page 48](#)

For additional information about dynamic-profiles, see *Dynamic Profiles Overview*.

Configuring a Preferred Source Address for Unnumbered Ethernet or Demux Interfaces

When a loopback interface with multiple secondary IP addresses is configured as the donor interface for an unnumbered Ethernet or demux interface, you can optionally specify any one of the loopback interface's secondary addresses as the preferred source address for the unnumbered Ethernet or demux interface. This feature enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet or demux interfaces in your network.

To configure a secondary address on a loopback donor interface as the preferred source address for an unnumbered Ethernet or demux interface, include the **preferred-source-address** option in the **unnumbered-address** statement:

```
unnumbered-address interface-name <preferred-source-address address>;
```

You can include this statement at the following hierarchy levels:

- [edit **interfaces** *interface-name* **unit** *logical-unit-number* **family** *family*]
- [edit **logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number* **family** *family*]
- [edit **dynamic-profiles** *profile-name* **interfaces** *interface-name* **unit** *logical-unit-number* **family** *family*]
- [edit **dynamic-profiles** *profile-name* **interfaces** **demux0** **unit** *logical-unit-number* **family** *family*]

The following considerations apply when you configure a preferred source address on an unnumbered Ethernet or demux interface:

- The **unnumbered-address** statement currently supports the configuration of a preferred source address only for the IPv4 address family for demux interfaces, and for IPv4 and IPv6 address families for Ethernet interfaces.
- If you do not specify the preferred source address, the router uses the default primary IP address of the donor interface.
- You cannot delete an address on a donor loopback interface while it is being used as the preferred source address for an unnumbered Ethernet or demux interface.
- The router uses the preferred source address, if configured for an unnumbered Ethernet or demux interface, in ARP requests and replies. ARP requests must match the preferred source address.

For a configuration example that illustrates this feature, see “[Example: Configuring the Preferred Source Address for an Unnumbered Ethernet Interface](#)” on page 48.

To display the preferred source address for an unnumbered Ethernet or demux interface, use the **show interfaces** operational mode command. For information about using this command, see the *Junos OS Operational Mode Commands*.

Configuring Static Routes on Unnumbered Ethernet Interfaces

You can configure static routes on an unnumbered Ethernet interface. To do so, you use the **qualified-next-hop** statement to specify the unnumbered Ethernet interface as the next-hop interface for a configured static route. This feature enables you to specify independent preferences and metrics for static routes on a next-hop basis.

For a configuration example that illustrates this feature, see [“Example: Configuring an Unnumbered Ethernet Interface as the Next Hop for a Static Route” on page 48](#).

For information about how to specify an independent preference for a static route, see the *Junos OS Routing Protocols Library for Routing Devices*.

Restrictions for Configuring Unnumbered Ethernet Interfaces

The following restrictions apply when you configure unnumbered Ethernet interfaces:

- The **unnumbered-address** statement currently supports the configuration of unnumbered Ethernet interfaces for IPv4 and IPv6 address families.
- You cannot assign an IP address to an Ethernet interface that is already configured as an unnumbered interface.
- The donor interface for an unnumbered Ethernet interface must have one or more configured IP addresses.
- The donor interface for an unnumbered Ethernet interfaced cannot be configured as unnumbered.
- An unnumbered Ethernet interface does not support configuration of the following **address** statement options: **arp**, **broadcast**, **primary**, **preferred**, and **vrrp-group**. For information about these options, see [“Configuring the Interface Address” on page 35](#).
- Running IGMP and PIM are supported only on unnumbered Ethernet interfaces that directly face the host and have no downstream PIM neighbors. IGMP and PIM are not supported on unnumbered Ethernet interfaces that act as upstream interfaces in a PIM topology.
- Running OSPF and IS-IS on unnumbered Ethernet interfaces is not supported. However, you can run OSPF over unnumbered Ethernet interfaces configured as a Point-to-Point connection.

Example: Configuring an Unnumbered Ethernet Interface

In this example, **ge-1/0/0** is the unnumbered interface and **ge-0/0/0** is the donor interface from which **ge-1/0/0** “borrows” an IP address.

```

interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 4.4.4.1/24;
      }
    }
  }
}

```

```
ge-1/0/0 {
  unit 0 {
    family inet {
      unnumbered-address ge-0/0/0.0;
    }
  }
}
```

Example: Configuring the Preferred Source Address for an Unnumbered Ethernet Interface

In this example, loopback interface **lo0** is the donor interface from which unnumbered Ethernet interface **ge-4/0/0** “borrows” an IP address. The example also configures one of the loopback interface’s secondary addresses, 3.3.3.1, as the preferred source address for the unnumbered Ethernet interface.

```
interfaces {
  lo0 {
    unit 0 {
      family inet {
        address 2.2.2.1/32;
        address 3.3.3.1/32;
      }
    }
  }
}
interfaces {
  ge-4/0/0 {
    unit 0 {
      family inet {
        unnumbered-address lo0.0 preferred-source-address 3.3.3.1;
      }
    }
  }
}
```

Example: Configuring an Unnumbered Ethernet Interface as the Next Hop for a Static Route

In this example, **ge-0/0/0** is the unnumbered interface and a loopback interface, **lo0**, is the donor interface from which **ge-0/0/0** “borrows” an IP address. The example also configures a static route to **7.7.7.1/32** with a next hop through unnumbered interface **ge-0/0/0.0**.

```
interfaces {
  lo0 {
    unit 0 {
      family inet {
        address 5.5.5.1/32;
        address 6.6.6.1/32;
      }
    }
  }
}
```

```

interfaces
  ge-0/0/0 {
    unit 0 {
      family inet {
        unnumbered-address lo0.0;
      }
    }
  }
}
routing-options {
  static {
    route 7.7.7.1/32 {
      qualified next-hop ge-0/0/0.0;
    }
  }
}

```

Configuring Default, Primary, and Preferred Addresses and Interfaces

The router has a default address and a primary interface, and interfaces have primary and preferred addresses.

The *default address* of the router is used as the source address on unnumbered interfaces. The routing protocol process tries to pick the default address as the router ID, which is used by protocols, including OSPF and internal BGP (IBGP).

The *primary interface* for the router is the interface that packets go out when no interface name is specified and when the destination address does not imply a particular outgoing interface.

An interface's *primary address* is used by default as the local address for broadcast and multicast packets sourced locally and sent out the interface. An interface's *preferred address* is the default local address used for packets sourced by the local router to destinations on the subnet.

The default address of the router is chosen using the following sequence:

1. The primary address on the loopback interface **lo0** that is not **127.0.0.1** is used.
2. The primary address on the primary interface is used.

To configure these addresses and interfaces, you can do the following:

- [Configuring the Primary Interface for the Router on page 49](#)
- [Configuring the Primary Address for an Interface on page 50](#)
- [Configuring the Preferred Address for an Interface on page 50](#)

Configuring the Primary Interface for the Router

The *primary interface* for the router has the following characteristics:

- It is the interface that packets go out when you type a command such as ping 255.255.255.255—that is, a command that does not include an interface name (there

is no interface **type-0/0/0.0** qualifier) and where the destination address does not imply any particular outgoing interface.

- It is the interface on which multicast applications running locally on the router, such as Session Announcement Protocol (SAP), do group joins by default.
- It is the interface from which the default local address is derived for packets sourced out an unnumbered interface if there are no non-127 addresses configured on the loopback interface, lo0.

By default, the multicast-capable interface with the lowest-index address is chosen as the primary interface. If there is no such interface, the point-to-point interface with the lowest index address is chosen. Otherwise, any interface with an address could be picked. In practice, this means that, on the router, the **fxp0** or **em0** interface is picked by default.

To configure a different interface to be the primary interface, include the **primary** statement:

primary;

You can include this statement at the following hierarchy levels:

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

Configuring the Primary Address for an Interface

The *primary address* on an interface is the address that is used by default as the local address for broadcast and multicast packets sourced locally and sent out the interface. For example, the local address in the packets sent by a **ping interface so-0/0/0.0 255.255.255.255** command is the primary address on interface **so-0/0/0.0**. The primary address flag also can be useful for selecting the local address used for packets sent out unnumbered interfaces when multiple non-127 addresses are configured on the loopback interface, lo0. By default, the primary address on an interface is selected as the numerically lowest local address configured on the interface.

To set a different primary address, include the **primary** statement:

primary;

You can include this statement at the following hierarchy levels:

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

Configuring the Preferred Address for an Interface

The *preferred address* on an interface is the default local address used for packets sourced by the local router to destinations on the subnet. By default, the numerically lowest local address is chosen. For example, if the addresses **172.16.1.1/12**, **172.16.1.2/12**, and **172.16.1.3/12**

are configured on the same interface, the preferred address on the subnet (by default, 172.16.1.1) would be used as a local address when you issue a **ping 172.16.1.5** command.

To set a different preferred address for the subnet, include the **preferred** statement:

```
preferred;
```

You can include this statement at the following hierarchy levels:

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

Configuring Unicast RPF

For interfaces that carry IPv4 or IPv6 traffic, you can reduce the impact of denial of service (DoS) attacks by configuring unicast reverse path forwarding (RPF). Unicast RPF helps determine the source of attacks and rejects packets from unexpected source addresses on interfaces where unicast RPF is enabled.



NOTE: If you want to configure unicast RPF, your router must be equipped with the Internet Processor II application-specific integrated circuit (ASIC).

If you enable unicast RPF on live traffic, some packets are dropped while the packet forwarding components are updating.

For transit packets exiting the router through the tunnel, forwarding path features, such as RPF, forwarding table filtering, source class usage, and destination class usage are not supported on the interfaces you configure as the output interface for tunnel traffic. For firewall filtering, you must allow the output tunnel packets through the firewall filter applied to input traffic on the interface that is the next-hop interface towards the tunnel destination.

The following sections describe unicast RPF in detail:

- [Configuring Unicast RPF Strict Mode on page 51](#)
- [Configuring Unicast RPF Loose Mode on page 52](#)
- [Unicast RPF and Default Routes on page 53](#)
- [Unicast RPF with Routing Asymmetry on page 54](#)
- [Configuring Unicast RPF on a VPN on page 55](#)
- [Example: Configuring Unicast RPF on page 56](#)

Configuring Unicast RPF Strict Mode

In strict mode, unicast RPF checks whether the incoming packet has a source address that matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix.

If the incoming packet fails the unicast RPF check, the packet is not accepted on the interface. When a packet is not accepted on an interface, unicast RPF counts the packet and sends it to an optional fail filter. If the fail filter is not configured, the default action is to silently discard the packet.

The optional fail filter allows you to apply a filter to packets that fail the unicast RPF check. You can define the fail filter to perform any filter operation, including accepting, rejecting, logging, sampling, or policing.

When unicast RPF is enabled on an interface, Bootstrap Protocol (BOOTP) packets and Dynamic Host Configuration Protocol (DHCP) packets are not accepted on the interface. To allow the interface to accept BOOTP packets and DHCP packets, you must apply a fail filter that accepts all packets with a source address of **0.0.0.0** and a destination address of **255.255.255.255**. For a configuration example, see [“Example: Configuring Unicast RPF” on page 56](#).

For more information about unicast RPF, see the *Junos OS Routing Protocols Library for Routing Devices*. For more information about defining fail filters, see the *Routing Policy Feature Guide for Routing Devices*.

To configure unicast RPF, include the **rpf-check** statement:

```
rpf-check <fail-filter filter-name>;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6)]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (inet | inet6)]**

Using unicast RPF can have several consequences when implemented with traffic filters:

- RPF fail filters are evaluated after input filters and before output filters.
- If you configure a filter counter for packets dropped by an input filter, and you want to know the total number of packets dropped, you must also configure a filter counter for packets dropped by the RPF check.
- To count packets that fail the RPF check and are accepted by the RPF fail filter, you must configure a filter counter.
- If an input filter forwards packets anywhere other than the inet.0 or inet6.0 routing tables, the unicast RPF check is not performed.
- If an input filter forwards packets anywhere other than the routing instance the input interface is configured for, the unicast RPF check is not performed.

Configuring Unicast RPF Loose Mode

By default, unicast RPF uses strict mode. Unicast RPF loose mode is similar to unicast RPF strict mode and has the same configuration restrictions. The only check in loose mode is whether the packet has a source address with a corresponding prefix in the routing table; loose mode does not check whether the interface expects to receive a

packet with a specific source address prefix. If a corresponding prefix is not found, unicast RPF loose mode does not accept the packet. As in strict mode, loose mode counts the failed packet and optionally forwards it to a fail filter, which either accepts, rejects, logs, samples, or polices the packet.

To configure unicast RPF loose mode, include the **mode**:

```
mode loose;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) rpf-check <fail-filter *filter-name*>]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) rpf-check <fail-filter *filter-name*>]

Configuring Unicast RPF Loose Mode with Ability to Discard Packets

Starting with Junos OS Release 12.1, unicast RPF loose mode has the ability to discard packets with the source address pointing to the discard interface. This feature is supported on MX Series routers and on T Series routers with Type 1 FPCs, Type 2 FPCs, and Type 3 FPCs. Using unicast RPF loose mode, along with Remote Triggered Black Hole (RTBH) filtering, provides an efficient way to discard packets coming from known attack sources. BGP policies in edge routers ensure that packets with untrusted source addresses have their next hop set to a discard route. When a packet arrives at the router with an untrusted source address, unicast RPF performs a route lookup of the source address. Because the source address route points to a discard next hop, the packet is dropped and a counter is incremented. This feature is supported on both IPv4 (inet) and IPv6 (inet6) address families.

To configure unicast RPF loose mode with the ability to discard packets, include the **rpf-loose-mode-discard family inet** statement at the [edit forwarding-options] hierarchy level:

```
rpf-loose-mode-discard {
  family {
    inet;
  }
}
```

Unicast RPF and Default Routes

When the active route cannot be chosen from the routes in a routing table, the router chooses a default route. A default route is equivalent to an IP address of 0.0.0.0/0. If you configure a default route, and you configure unicast RPF on an interface that the default route uses, unicast RPF behaves differently than it does otherwise. For information about configuring default routes, see the *Junos OS Routing Protocols Library for Routing Devices*.

To determine whether the default route uses an interface, enter the **show route** command:

```
user@host> show route address
```

address is the next-hop address of the configured default route. The default route uses the interfaces shown in the output of the **show route** command.

The following sections describe how unicast RPF behaves when a default route uses an interface and when a default route does not use an interface:

- [Unicast RPF Behavior with a Default Route on page 54](#)
- [Unicast RPF Behavior Without a Default Route on page 54](#)

Unicast RPF Behavior with a Default Route

If you configure a default route that uses an interface configured with unicast RPF, unicast RPF behaves as follows:

- Loose mode—All packets are automatically accepted. For this reason, we recommend that you not configure unicast RPF loose mode on interfaces that the default route uses.
- Strict mode—The packet is accepted when either of the following is true:
 - The source address of the packet matches any of the routes (either default or learned) that can be originated from the interface. Note that routes can have multiple destinations associated with them; therefore, if one of the destinations matches the incoming interface of the packet, the packet is accepted.
 - The source address of the packet does not match any of the routes.

The packet is not accepted when either of the following is true:

- The source address of the packet does not match a prefix in the routing table.
- The interface does not expect to receive a packet with this source address prefix.

Unicast RPF Behavior Without a Default Route

If you do not configure a default route, or if the default route does not use an interface configured with unicast RPF, unicast RPF behaves as described in “[Configuring Unicast RPF Strict Mode](#)” on page 51 and “[Configuring Unicast RPF Loose Mode](#)” on page 52. To summarize, unicast RPF without a default route behaves as follows:


- Strict mode—The packet is not accepted when either of the following is true:
 - The packet has a source address that does not match a prefix in the routing table.
 - The interface does not expect to receive a packet with this source address prefix.
- Loose mode—The packet is not accepted when the packet has a source address that does not match a prefix in the routing table.

Unicast RPF with Routing Asymmetry

In general, we recommend that you not enable unicast RPF on interfaces that are internal to the network because internal interfaces are likely to have *routing asymmetry*. Routing asymmetry means that a packet’s outgoing and return paths are different. Routers in the

core of the network are more likely to have asymmetric reverse paths than routers at the customer or provider edge. [Figure 3 on page 55](#) shows unicast RPF in an environment with routing asymmetry.

Figure 3: Unicast RPF with Routing Asymmetry



In [Figure 3 on page 55](#), if you enable unicast RPF on interface `so-0/0/0`, traffic destined for Router A is not rejected. If you enable unicast RPF on interface `so-1/0/1`, traffic from Router A is rejected.

If you need to enable unicast RPF in an asymmetric routing environment, you can use fail filters to allow the router to accept incoming packets that are known to be arriving by specific paths. For an example of a fail filter that accepts packets with a specific source and destination address, see [“Example: Configuring Unicast RPF” on page 56](#).

Configuring Unicast RPF on a VPN

You can configure unicast RPF on a VPN interface by enabling unicast RPF on the interface and including the `interface` statement at the `[edit routing-instances routing-instance-name]` hierarchy level.

You can configure unicast RPF only on the interfaces you specify in the routing instance. This means the following:

- For Layer 3 VPNs, unicast RPF is supported on the CE router interface.
- Unicast RPF is not supported on core-facing interfaces.
- For virtual-router routing instances, unicast RPF is supported on all interfaces you specify in the routing instance.
- If an input filter forwards packets anywhere other than the routing instance the input interface is configured for, the unicast RPF check is not performed.

For more information about VPNs and virtual-router routing instances, see the *Junos OS VPNs Library for Routing Devices*. For more information about FBF, see the *Junos OS Routing Protocols Library for Routing Devices*.

Example: Configuring Unicast RPF on a VPN

Configure unicast RPF on a Layer 3 VPN interface:

```
[edit interfaces]
so-0/0/0 {
  unit 0 {
    family inet {
      rpf-check;
    }
  }
}
```

```
}
[edit routing-instance]
VPN-A {
  interface so-0/0/0.0;
}
```

Example: Configuring Unicast RPF

Configure unicast RPF strict mode, and apply a fail filter that allows the interface to accept BOOTP packets and DHCP packets. The filter accepts all packets with a source address of **0.0.0.0** and a destination address of **255.255.255.255**.

```
[edit firewall]
filter rpf-special-case-dhcp-bootp {
  term allow-dhcp-bootp {
    from {
      source-address {
        0.0.0.0/32;
      }
      address {
        255.255.255.255/32;
      }
    }
    then {
      count rpf-dhcp-bootp-traffic;
      accept;
    }
  }
  term default {
    then {
      log;
      reject;
    }
  }
}
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        rpf-check fail-filter rpf-special-case-dhcp-bootp;
      }
    }
  }
}
```

- Related Documentation**
- [unicast-reverse-path](#)
 - [Example: Configuring Unicast Reverse-Path-Forwarding Check on page 60](#)

Configuring Targeted Broadcast

The following sections explain how to configure targeted broadcast on an egress interface and its options:

- [Configuring Targeted Broadcast and Its Options on page 57](#)
- [Display Targeted Broadcast Configuration Options on page 58](#)

Configuring Targeted Broadcast and Its Options

You can configure targeted broadcast on an egress interface with different options. You can either allow the IP packets destined for a Layer 3 broadcast address to be forwarded on the egress interface and to send a copy of the IP packets to the Routing Engine or you can allow the IP packets to be forwarded on the egress interface only. Note that the packets are broadcast only if the egress interface is a LAN interface.

To configure targeted broadcast and its options:

1. Configure the physical interface.

```
[edit]
user@host# set interfaces interface-name
```

2. Configure the logical unit number at the `[edit interfaces interface-name` hierarchy level.

```
[edit interfaces interface-name]
user@host# set unit logical-unit-number
```

3. Configure the protocol family as `inet` at the `[edit interfaces interface-name unit interface-unit-number` hierarchy level.

```
[edit interfaces interface-name unit interface--unit-number]
user@host# set family inet
```

4. Configure targeted broadcast at the `[edit interfaces interface-name unit interface-unit-number family inet` hierarchy level

```
[edit interfaces interface-name unit interface--unit-number family inet]
user@host# set targeted-broadcast
```

5. Specify one of the following options as per requirement:

- To allow IP packets destined for a Layer 3 broadcast address to be forwarded on the egress interface and to send a copy of the IP packets to the Routing Engine.

```
[edit interfaces interface-name unit interface-unit-number family inet
targeted-broadcast]
user@host# set forward-and-send-to-re
```

- To allow IP packets to be forwarded on the egress interface only.

```
[edit interfaces interface-name unit interface-unit-number family inet
targeted-broadcast]
user@host# set forward-only
```



NOTE: Targeted broadcast does not work when the targeted broadcast option `forward-and-send-to-re` and the traffic sampling option `sampling` are configured on the same egress interface of an M320 router, a T640 router, or an MX960 router. To overcome this scenario, you must either disable one of these options or enable the sampling option with the targeted broadcast option `forward-only` on the egress interface. For information about traffic sampling, see *Configuring Traffic Sampling*.

Display Targeted Broadcast Configuration Options

The following topics display targeted broadcast configuration with its various options:

- [Forward IP Packets On the Egress Interface and To the Routing Engine on page 58](#)
- [Forward IP Packets On the Egress Interface Only on page 58](#)

Forward IP Packets On the Egress Interface and To the Routing Engine

Purpose Display the configuration when targeted broadcast is configured on the egress interface to forward the IP packets on the egress interface and to send a copy of the IP packets to the Routing Engine.

Action To display the configuration run the **show** command at the **[edit interfaces *interface-name* unit *interface-unit-number* family inet]** where the interface name is `ge-2/0/0`, the unit value is set to `0`, the protocol family is set to `inet`.

```
[edit interfaces interface-name unit interface-unit-number family inet]
user@host#show
targeted-broadcast {
    forward-and-send-to-re;
}
```

Forward IP Packets On the Egress Interface Only

Purpose Display the configuration when targeted broadcast is configured on the egress interface to forward the IP packets on the egress interface only.

Action To display the configuration run the **show** command at the **[edit interfaces *interface-name* unit *interface-unit-number* family inet]** where the interface name is `ge-2/0/0`, the unit value is set to `0`, the protocol family is set to `inet`.

```
[edit interfaces interface-name unit interface-unit-number family inet]
user@host#show
targeted-broadcast {
    forward-only;
}
```

Related Documentation

- [targeted-broadcast on page 164](#)
- [Understanding Targeted Broadcast on page 59](#)

Understanding Targeted Broadcast

Targeted broadcast is a process of flooding a target subnet with Layer 3 broadcast IP packets originating from a different subnet. The intent of targeted broadcast is to flood the target subnet with the broadcast packets on a LAN interface without broadcasting to the entire network. Targeted broadcast is configured with various options on the egress interface of the router or switch and the IP packets are broadcast only on the LAN (egress) interface. Targeted broadcast helps you implement remote administration tasks such as backups and wake-on LAN (WOL) on a LAN interface, and supports virtual routing and forwarding (VRF) instances.

Regular Layer 3 broadcast IP packets originating from a subnet are broadcast within the same subnet. When these IP packets reach a different subnet, they are forwarded to the Routing Engine (to be forwarded to other applications). Because of this, remote administration tasks such as backups cannot be performed on a particular subnet through another subnet. As a workaround you can enable targeted broadcast, to forward broadcast packets that originate from a different subnet.

Layer 3 broadcast IP packets have a destination IP address that is a valid broadcast address for the target subnet. These IP packets traverse the network in the same way as unicast IP packets until they reach the destination subnet. In the destination subnet, if the receiving router has targeted broadcast enabled on the egress interface, the IP packets are forwarded to an egress interface and the Routing Engine or to an egress interface only. The IP packets are then translated into broadcast IP packets which flood the target subnet only through the LAN interface (if there is no LAN interface, the packets are discarded), and all hosts on the target subnet receive the IP packets. If targeted broadcast is not enabled on the receiving router, the IP packets are treated as regular Layer 3 broadcast IP packets and are forwarded to the Routing Engine. If targeted broadcast is enabled without any options, the IP packets are discarded.

Targeted broadcast can be configured to forward the IP packets only to an egress interface, which is helpful when the router is flooded with packets to process, or to both an egress interface and the Routing Engine.



NOTE: Targeted broadcast does not work when the targeted broadcast option `forward-and-send-to-re` and the traffic sampling option `sampling` are configured on the same egress interface of an M320 router, a T640 router, or an MX960 router. To overcome this scenario, you must either disable one of these options or enable the `sampling` option with the targeted broadcast option `forward-only` on the egress interface. For information about traffic sampling, see *Configuring Traffic Sampling*.



NOTE: Any firewall filter that is configured on the Routing Engine loopback interface (lo0) cannot be applied to IP packets that are forwarded to the Routing Engine as a result of a targeted broadcast. This is because broadcast packets are forwarded as flood next hop and not as local next hop traffic, and you can only apply a firewall filter to local next hop routes for traffic directed towards the Routing Engine.

**Related
Documentation**

- [Configuring Targeted Broadcast on page 57](#)
- [targeted-broadcast on page 164](#)

Example: Configuring Unicast Reverse-Path-Forwarding Check

- [Understanding Unicast Reverse Path Forwarding on page 60](#)
- [Example: Configuring Unicast Reverse-Path-Forwarding Check on page 60](#)

Understanding Unicast Reverse Path Forwarding

IP spoofing can occur during a denial-of-service (DoS) attack. IP spoofing allows an intruder to pass IP packets to a destination as genuine traffic, when in fact the packets are not actually meant for the destination. This type of spoofing is harmful because it consumes the destination's resources.

A unicast reverse-path-forwarding (RPF) check is a tool to reduce forwarding of IP packets that might be spoofing an address. A unicast RPF check performs a route table lookup on an IP packet's source address, and checks the incoming interface. The router or switch determines whether the packet is arriving from a path that the sender would use to reach the destination. If the packet is from a valid path, the router or switch forwards the packet to the destination address. If it is not from a valid path, the router or switch discards the packet. Unicast RPF is supported for the IPv4 and IPv6 protocol families, as well as for the virtual private network (VPN) address family.



NOTE: Reverse path forwarding is not supported on the interfaces you configure as tunnel sources. This affects only the transit packets exiting the tunnel.

Example: Configuring Unicast Reverse-Path-Forwarding Check

Unicast reverse path forwarding (RPF) helps protect against DoS and DDoS attacks by verifying the unicast source address of each packet that arrives on an ingress interface where unicast RPF is enabled.

This example shows how to help defend ingress interfaces against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by configuring unicast RPF to filter incoming traffic.

- [Requirements on page 61](#)
- [Overview on page 61](#)
- [Configuration on page 62](#)
- [Verification on page 67](#)

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

Large amounts of unauthorized traffic such as attempts to flood a network with fake (bogus) service requests in a DoS attack can consume network resources and deny service to legitimate users. One way to help prevent DoS and DDoS attacks is to verify that incoming traffic originates from legitimate network sources.


Unicast RPF helps ensure that a traffic source is legitimate (authorized) by comparing the source address of each packet that arrives on an interface to the forwarding table entry for its source address. If the device uses the same interface that the packet arrived on to reply to the packet's source, this verifies that the packet originated from an authorized source, and the device forwards the packet. If the device does not use the same interface that the packet arrived on to reply to the packet's source, the packet might have originated from an unauthorized source, and the device discards the packet.

In this example, Device B has unicast RPF configured. Device A is using OSPF to advertise a prefix for the link that connects to Device D. OSPF is enabled on the links between Device B and Device C and the links between Device A and Device C, but not on the links between Device A and Device B. Therefore, Device B learns about the route to Device D through Device C.

This example also includes a fail filter. When a packet fails the unicast RPF check, the fail filter is evaluated to determine if the packet should be accepted anyway. The fail filter in this example allows Device B's interfaces to accept Dynamic Host Configuration Protocol (DHCP) packets. The filter accepts all packets with a source address of 0.0.0.0 and a destination address of 255.255.255.255.

[Figure 4 on page 62](#) shows the sample network.

Figure 4: Unicast RPF Sample Topoolgy



Configuration

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

Device A

```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces fe-0/0/2 unit 5 family inet address 10.0.0.5/30
set interfaces fe-0/0/1 unit 17 family inet address 10.0.0.17/30
set interfaces fe-0/1/1 unit 25 family inet address 10.0.0.25/30
set interfaces fe-1/1/1 unit 29 family inet address 10.0.0.29/30
set protocols ospf export send-direct
set protocols ospf area 0.0.0.0 interface fe-0/1/1.25
set protocols ospf area 0.0.0.0 interface fe-1/1/1.29
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct from route-filter 10.0.0.16/30 exact
set policy-options policy-statement send-direct then accept

```

Device B

```

set interfaces fe-1/2/0 unit 2 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/1/1 unit 6 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-1/1/1 unit 6 family inet address 10.0.0.6/30
set interfaces fe-0/1/1 unit 9 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-0/1/1 unit 9 family inet address 10.0.0.9/30
set interfaces fe-0/1/0 unit 13 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-0/1/0 unit 13 family inet address 10.0.0.13/30
set protocols ospf area 0.0.0.0 interface fe-0/1/1.9
set protocols ospf area 0.0.0.0 interface fe-0/1/0.13
set routing-options forwarding-table unicast-reverse-path active-paths
set firewall filter rpf-special-case-dhcp term allow-dhcp from source-address 0.0.0.0/32
set firewall filter rpf-special-case-dhcp term allow-dhcp then count rpf-dhcp-traffic
set firewall filter rpf-special-case-dhcp term allow-dhcp then accept
set firewall filter rpf-special-case-dhcp term default then log
set firewall filter rpf-special-case-dhcp term default then reject

```

Device C

```

set interfaces fe-1/2/0 unit 10 family inet address 10.0.0.10/30
set interfaces fe-0/0/2 unit 14 family inet address 10.0.0.14/30
set interfaces fe-1/0/2 unit 21 family inet address 10.0.0.21/30
set interfaces fe-1/2/2 unit 26 family inet address 10.0.0.26/30
set interfaces fe-1/2/1 unit 30 family inet address 10.0.0.30/30

```



```

set protocols ospf area 0.0.0.0 interface fe-1/2/0.10
set protocols ospf area 0.0.0.0 interface fe-0/0/2.14
set protocols ospf area 0.0.0.0 interface fe-1/2/2.26
set protocols ospf area 0.0.0.0 interface fe-1/2/1.30

```

Device D set interfaces fe-1/2/0 unit 18 family inet address 10.0.0.18/30

Device E set interfaces fe-1/2/0 unit 22 family inet address 10.0.0.22/30

Configuring Device A

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device A:

1. Configure the interfaces.


```

[edit interfaces]
user@A# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

user@A# set fe-0/0/2 unit 5 family inet address 10.0.0.5/30

user@A# set fe-0/0/1 unit 17 family inet address 10.0.0.17/30

user@A# set fe-0/1/1 unit 25 family inet address 10.0.0.25/30

user@A# set fe-1/1/1 unit 29 family inet address 10.0.0.29/30

```
2. Configure OSPF.


```

[edit protocols ospf]
user@A# set export send-direct
user@A# set area 0.0.0.0 interface fe-0/1/1.25
user@A# set area 0.0.0.0 interface fe-1/1/1.29

```
3. Configure the routing policy.


```

[edit policy-options policy-statement send-direct]
user@A# set from protocol direct
user@A# set from route-filter 10.0.0.16/30 exact
user@A# set then accept

```
4. If you are done configuring Device A, commit the configuration.


```

[edit]
user@A# commit

```

Configuring Device B

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device B:

1. Configure the interfaces.

```
[edit interfaces]
user@B# set fe-1/2/0 unit 2 family inet address 10.0.0.2/30

user@B# set fe-1/1/1 unit 6 family inet address 10.0.0.6/30

user@B# set fe-0/1/1 unit 9 family inet address 10.0.0.9/30

user@B# set fe-0/1/0 unit 13 family inet address 10.0.0.13/30
```
2. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@B# set interface fe-0/1/1.9
user@B# set interface fe-0/1/0.13
```
3. Configure unicast RPF, and apply the optional fail filter.

```
[edit interfaces]
user@B# set fe-1/2/0 unit 2 family inet rpf-check fail-filter rpf-special-case-dhcp

user@B# set fe-1/1/1 unit 6 family inet rpf-check fail-filter rpf-special-case-dhcp

user@B# set fe-0/1/1 unit 9 family inet rpf-check fail-filter rpf-special-case-dhcp

user@B# set fe-0/1/0 unit 13 family inet rpf-check fail-filter rpf-special-case-dhcp
```
4. (Optional) Configure the fail filter that gets evaluated if a packet fails the RPF check.

```
[edit firewall filter rpf-special-case-dhcp]
user@B# set term allow-dhcp from source-address 0.0.0.0/32
user@B# set term allow-dhcp then count rpf-dhcp-traffic
user@B# set term allow-dhcp then accept
user@B# set term default then log
user@B# set term default then reject
```
5. (Optional) Configure only active paths to be considered in the RPF check.
This is the default behavior.

```
[edit routing-options forwarding-table]
user@B# set unicast-reverse-path active-paths
```
6. If you are done configuring Device B, commit the configuration.

```
[edit]
user@B# commit
```

Results

Confirm your configuration by issuing the **show firewall**, **show interfaces**, **show protocols**, **show routing-options**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

Device A user@A# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
fe-0/0/2 {
  unit 5 {
    family inet {
      address 10.0.0.5/30;
    }
  }
}
fe-0/0/1 {
  unit 17 {
    family inet {
      address 10.0.0.17/30;
    }
  }
}
fe-0/1/1 {
  unit 25 {
    family inet {
      address 10.0.0.25/30;
    }
  }
}
fe-1/1/1 {
  unit 29 {
    family inet {
      address 10.0.0.29/30;
    }
  }
}

user@A# show protocols
ospf {
  export send-direct;
  area 0.0.0.0 {
    interface fe-0/1/1.25;
    interface fe-1/1/1.29;
  }
}

user@A# show policy-options
policy-statement send-direct {
  from {

```

```
        protocol direct;
        route-filter 10.0.0.16/30 exact;
    }
    then accept;
}

Device B user@B# show firewall
filter rpf-special-case-dhcp {
    term allow-dhcp {
        from {
            source-address {
                0.0.0.0/32;
            }
        }
        then {
            count rpf-dhcp-traffic;
            accept;
        }
    }
    term default {
        then {
            log;
            reject;
        }
    }
}
user@B# show interfaces
fe-1/2/0 {
    unit 2 {
        family inet {
            rpf-check fail-filter rpf-special-case-dhcp;
            address 10.0.0.2/30;
        }
    }
}
fe-1/1/1 {
    unit 6 {
        family inet {
            rpf-check fail-filter rpf-special-case-dhcp;
            address 10.0.0.6/30;
        }
    }
}
fe-0/1/1 {
    unit 9 {
        family inet {
            rpf-check fail-filter rpf-special-case-dhcp;
            address 10.0.0.9/30;
        }
    }
}
fe-0/1/0 {
    unit 13 {
        family inet {
            rpf-check fail-filter rpf-special-case-dhcp;
            address 10.0.0.13/30;
        }
    }
}
```

```
    }  
  }  
}  
  
user@B# show protocols  
ospf {  
  area 0.0.0.0 {  
    interface fe-0/1/1.9;  
    interface fe-0/1/0.13;  
  }  
}  
  
user@B# show routing-options  
forwarding-table {  
  unicast-reverse-path active-paths;  
}
```

Enter the configurations on Device C, Device D, and Device E, as shown in “[CLI Quick Configuration](#)” on [page 62](#).

Verification

Confirm that the configuration is working properly.

- [Confirm That Unicast RPF Is Enabled on page 67](#)
- [Confirm That the Source Addresses Are Blocked on page 68](#)
- [Confirm That the Source Addresses Are Unblocked on page 68](#)

Confirm That Unicast RPF Is Enabled

Purpose Make sure that the interfaces on Device B have unicast RPF enabled.

Action user@B> show interfaces fe-0/1/0.13 extensive
Logical interface fe-0/1/0.13 (Index 73) (SNMP ifIndex 553) (Generation 208)
Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
Traffic statistics:
Input bytes : 999390
Output bytes : 1230122
Input packets: 12563
Output packets: 12613
Local statistics:
Input bytes : 998994
Output bytes : 1230122
Input packets: 12563
Output packets: 12613
Transit statistics:
Input bytes : 396 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol inet, MTU: 1500, Generation: 289, Route table: 22
Flags: Sendbcst-pkt-to-re, uRPF
RPF Failures: Packets: 0, Bytes: 0
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.0.0.12/30, Local: 10.0.0.13, Broadcast: 10.0.0.15,
Generation: 241

Meaning The uRPF flag confirms that unicast RPF is enabled on this interface.

Confirm That the Source Addresses Are Blocked

Purpose Use the ping command to make sure that Device B blocks traffic from unexpected source addresses.

Action From Device A, ping Device B's interfaces, using 10.0.0.17 as the source address.

```
user@A> ping 10.0.0.6 source 10.0.0.17
PING 10.0.0.6 (10.0.0.6): 56 data bytes
^C
--- 10.0.0.6 ping statistics ---
3 packets transmitted, 0 packets received, 100% packet loss
```

Meaning As expected, the ping operation fails.

Confirm That the Source Addresses Are Unblocked

Purpose Use the ping command to make sure that Device B does not block traffic when the RPF check is deactivated.

Action 1. Deactivate the RPF check on one of the interfaces.
2. Rerun the ping operation.

```
user@B> deactivate interfaces fe-1/1/1.6 family inet rpf-check
user@A> ping 10.0.0.6 source 10.0.0.17
PING 10.0.0.2 (10.0.0.2): 56 data bytes
64 bytes from 10.0.0.2: icmp_seq=0 ttl=63 time=1.316 ms
```

```
64 bytes from 10.0.0.2: icmp_seq=1 ttl=63 time=1.263 ms
^C
--- 10.0.0.2 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.263/1.289/1.316/0.027 ms
```

Meaning As expected, the ping operation succeeds.

Related Documentation

- *Example: Enabling Indirect Next Hops on the Packet Forwarding Engine*

CHAPTER 3

Network Interfaces Configuration Statements and Hierarchy

- [\[edit firewall\] Hierarchy Level on page 71](#)
- [\[edit interfaces\] Hierarchy Level on page 72](#)
- [\[edit logical-systems\] Hierarchy Level on page 89](#)
- [\[edit protocols pppoe\] Hierarchy Level on page 93](#)

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

The following CoS statements can be configured at the **[edit firewall]** hierarchy level. This is not a comprehensive list of statements available at the **[edit firewall]** hierarchy level.

```
[edit firewall]
  atm-policer policer-name {
    cdvt rate;
    logical-interface-policer;
    max-burst-size max-burst-size;
    peak-rate rate;
    policing-action (discard | discard-tag | count);
    sustained-rate rate;
  }
  family family-name {
    filter filter-name {
      term term-name {
        from {
          match-conditions;
        }
        then {
          dscp 0;
          forwarding-class class-name;
          loss-priority (high | low);
          three-color-policer {
            (single-rate | two-rate) policer-name;
          }
        }
      }
    }
  }
  simple-filter filter-name {
```




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

```

interfaces {
  traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable |
      no-world-readable> ;
    flag flag <disable>;
  }
  interface-name {
    account-layer2-overhead (Interface Level) {
      value;
      egress bytes;
      ingress bytes;
    }
    accounting-profile name;
    aggregated-ether-options {
      (flow-control | no-flow-control);
      lacp {
        (active | passive);
        link-protection {
          disable;
          (revertive | non-revertive);
          periodic interval;
          system-priority priority;
        }
        link-protection;
        link-speed speed;
        (loopback | no-loopback);
        mc-ae {
          chassis-id chassis-id;
          mc-ae-id mc-ae-id;
          mode (active-active | active-standby);
          redundancy-group group-id;
          status-control (active | standby);
        }
        minimum-links number;
        source-address-filter {
          mac-address;
        }
        (source-filtering | no-source-filtering);
      }
      shared-scheduler;
      aggregated-sonet-options {
        link-speed speed | mixed;
        minimum-links number;
      }
      atm-options {
        cell-bundle-size cells;
        ilmi;
      }
    }
  }
}

```

```
linear-red-profiles profile-name {
    high-plp-max-threshold percent;
    low-plp-max-threshold percent;
    queue-depth cells high-plp-threshold percent low-plp-threshold percent;
}
mpls {
    pop-all-labels {
        required-depth number;
    }
}
pic-type (atm1 | atm2);
plp-to-clp;
promiscuous-mode {
    vpi vpi-identifier;
}
scheduler-maps map-name {
    forwarding-class class-name {
        epd-threshold cells plp1 cells;
        linear-red-profile profile-name;
        priority (high | low);
        transmit-weight (cells number | percent number);
    }
    vc-cos-mode (alternate | strict);
}
use-null-cw;
vpi vpi-identifier {
    maximum-vcs maximum-vcs;
    oam-liveness {
        down-count cells;
        up-count cells;
    }
    oam-period (seconds | disable);
    shaping {
        (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
        burst length);
        queue-length number;
    }
}
}
clocking clock-source;
data-input (system | interface interface-name);
dce;
serial-options {
    clock-rate rate;
    clocking-mode (dce | internal | loop);
    control-polarity (negative | positive);
    cts-polarity (negative | positive);
    dcd-polarity (negative | positive);
    dce-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
    }
}
```

```

    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
}
dsr-polarity (negative | positive);
dte-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
}
dtr-circuit (balanced | unbalanced);
dtr-polarity (negative | positive);
encoding (nrz | nrzi);
indication-polarity (negative | positive);
line-protocol protocol;
loopback mode;
rts-polarity (negative | positive);
tm-polarity (negative | positive);
transmit-clock invert;
}
description text;
dialer-options {
    pool pool-name <priority priority>;
}
disable;
ds0-options {
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    byte-encoding (nx56 | nx64);
    fcs (16 | 32);
    idle-cycle-flag (flags | ones);
    invert-data;
    loopback payload;
    start-end-flag (filler | shared);
}
e1-options {
    bert-error-rate rate;
    bert-period seconds;
    fcs (16 | 32);
    framing (g704 | g704-no-crc4 | unframed);
    idle-cycle-flag (flags | ones);
    invert-data;
    loopback (local | remote);
    start-end-flag (filler | shared);
    timeslots time-slot-range;
}
e3-options {
    atm-encapsulation (direct | plcp);
    bert-algorithm algorithm;
    bert-error-rate rate;

```

```
bert-period seconds;  
framing feet;  
compatibility-mode (digital-link | kentrox | larscom) <subrate value>;  
fcs (16 | 32);  
framing (g.751 | g.832);  
idle-cycle-flag (filler | shared);  
invert-data;  
loopback (local | remote);  
(payload-scrambler | no-payload-scrambler);  
start-end-flag (filler | shared);  
(unframed | no-unframed);  
}  
encapsulation type;  
es-options {  
    backup-interface es-fpc/pic/port;  
}  
fastether-options {  
    802.3ad aex;  
(flow-control | no-flow-control);  
ignore-l3-incompletes;  
ingress-rate-limit rate;  
(loopback | no-loopback);  
mpls {  
    pop-all-labels {  
        required-depth number;  
    }  
}  
source-address-filter {  
    mac-address;  
}  
(source-filtering | no-source-filtering);  
}  
flexible-vlan-tagging;  
gether-options {  
    802.3ad aex;  
(asynchronous-notification | no-asynchronous-notification);  
(auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |  
    local-interface-offline>;  
auto-reconnect seconds;  
(flow-control | no-flow-control);  
ignore-l3-incompletes;  
(loopback | no-loopback);  
mpls {  
    pop-all-labels {  
        required-depth number;  
    }  
}  
no-auto-mdix;  
source-address-filter {  
    mac-address;  
}  
(source-filtering | no-source-filtering);  
ethernet-switch-profile {  
    (mac-learn-enable | no-mac-learn-enable);  
tag-protocol-id [ tpids ];  
ethernet-policer-profile {
```

```

    input-priority-map {
        ieee802.1p premium [ values ];
    }
    output-priority-map {
        classifier {
            premium {
                forwarding-class class-name {
                    loss-priority (high | low);
                }
            }
        }
    }
    policer cos-policer-name {
        aggregate {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
        premium {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
    }
}

(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time up milliseconds down milliseconds;
ima-group-options {
    differential-delay number;
    frame-length (32 | 64 | 128 | 256);
    frame-synchronization {
        alpha number;
        beta number;
        gamma number;
    }
    minimum-links number;
    symmetry (symmetrical-config-and-operation |
        symmetrical-config-asymmetrical-operation);
    test-procedure {
        ima-test-start;
        ima-test-stop;
        interface name;
        pattern number;
        period number;
    }
    transmit-clock (common | independent);
    version (1.0 | 1.1);
}
ima-link-options group-id group-id;
interface-set interface-set-name {
    interface ethernet-interface-name {
        (unit unit-number | vlan-tags-outer vlan-tag);
    }
    interface interface-name {
        (unit unit-number);
    }
}

```

```
}
isdn-options {
    bchannel-allocation (ascending | descending);
    calling-number number;
    pool pool-name <priority priority>;
    spid1 spid-string;
    spid2 spid-string;
    static-tei-val value;
    switch-type (att5e | etsi | nil | ntdms100 | ntt);
    t310 seconds;
    tei-option (first-call | power-up);
}
keepalives <down-count number> <interval seconds> <up-count number>;
link-mode mode;
lmi {
    lmi-type (ansi | itu | c-lmi);
    n391dte number;
    n392dce number;
    n392dte number;
    n393dce number;
    n393dte number;
    t391dte seconds;
    t392dce seconds;
}
lsq-failure-options {
    no-termination-request;
    [ trigger-link-failure interface-name ];
}
mac mac-address;
mlfr-uni-nni-bundle-options {
    acknowledge-retries number;
    acknowledge-timer milliseconds;
    action-red-differential-delay (disable-tx | remove-link);
    drop-timeout milliseconds;
    fragment-threshold bytes;
    cisco-interoperability send-lip-remove-link-for-link-reject;
    hello-timer milliseconds;
    link-layer-overhead percent;
    lmi-type (ansi | itu | c-lmi);
    minimum-links number;
    mrru bytes;
    n391 number;
    n392 number;
    n393 number;
    red-differential-delay milliseconds;
    t391 seconds;
    t392 seconds;
    yellow-differential-delay milliseconds;
}
modem-options {
    dialin (console | routable);
    init-command-string initialization-command-string;
}
mtu bytes;
multi-chassis-protection {
    peer a.b.c.d {
```



```

    interface interface-name;
  }
}
multiservice-options {
  (core-dump | no-core-dump);
  (syslog | no-syslog);
}
native-vlan-id number;
no-gratuitous-arp-request;
no-keepalives;
no-partition {
  interface-type type;
}
no-vpivci-swapping;
optics-options {
  alarm low-light-alarm {
    (link-down | syslog);
  }
  tx-power dbm;
  warning low-light-warning {
    (link-down | syslog);
  }
  wavelength nm;
}
otn-options {
  bytes transmit-payload-type value;
  fec (efec | gfec | gfec-sdfec | none);
  (is-ma | no-is-ma);
  (laser-enable | no-laser-enable);
  (line-loopback | no-line-loopback);
  (local-loopback | no-local-loopback);
  (odu-ttim-action-enable | no-odu-ttim-action-enable);
  (otu-ttim-action-enable | no-otu-ttim-action-enable);
  odu-delay-management {
    (bypass | no-bypass);
    (monitor-end-point | no-monitor-end-point);
    (number-of-frames | no-number-of-frames);
    (start-measurement | no-start-measurement);
  }
  (prbs | no-prbs);
  preemptive-fast-reroute {
    (backward-frr-enable | no-backward-frr-enable);
    (signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);
  }
  rate {
    (fixed-stuff-bytes | no-fixed-stuff-bytes);
    otu4;
    (pass-through | no-pass-through);
  }
  signal-degrade {
    ber-threshold-clear value;
    ber-threshold-signal-degrade value;
    interval value;
  }
  trigger trigger-identifier;
  tti tti-identifier;
}

```

```
}
partition partition-number oc-slice oc-slice-range interface-type type;
timeslots time-slot-range;
passive-monitor-mode;
per-unit-scheduler;
ppp-options {
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
    dynamic-profile profile-name;
    no-termination-request;
    pap {
        access-profile name;
        local-name name;
        local-password password;
        compression;
    }
}
psn-vcip psn-vci-identifier;
psn-vpip psn-vpi-identifier;
receive-bucket {
    overflow (discard | tag);
    rate percentage;
    threshold bytes;
}
redundancy-options {
    priority sp-fpc/pic/port;
    secondary sp-fpc/pic/port;
    hot-standby;
}
satop-options {
    payload-size n;
}
schedulers number;
serial-options {
    clock-rate rate;
    clocking-mode (dce | internal | loop);
    control-polarity (negative | positive);
    cts-polarity (negative | positive);
    dcd-polarity (negative | positive);
    dce-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
    }
}
```

```

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

```

```
    hold-time milliseconds;
    lockout;
    neighbor address;
    paired-group group-name;
    preserve-interface;
    protect-circuit group-name;
    request;
    revert-time seconds;
    switching-mode (bidirectional | unidirectional);
    working-circuit group-name;
}
bytes {
    c2 value;
    e1-quiet value;
    f1 value;
    f2 value;
    s1 value;
    z3 value;
    z4 value;
}
fcs (16 | 32);
loopback (local | remote);
mpls {
    pop-all-labels {
        required-depth number;
    }
}
path-trace trace-string;
(payload-scrambler | no-payload-scrambler);
rfc-2615;
trigger {
    defect ignore;
    hold-time up milliseconds down milliseconds;
}
vtmapping (itu-t | klm);
(z0-increment | no-z0-increment);
}
speed (10m | 100m | 1g | oc3 | oc12 | oc48);
stacked-vlan-tagging;
switch-options {
    switch-port port-number {
        (auto-negotiation | no-auto-negotiation);
        speed (10m | 100m | 1g);
        link-mode (full-duplex | half-duplex);
    }
}
}
t1-options {
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout value;
    byte-encoding (nx56 | nx64);
    crc-major-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5);
    crc-minor-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5 | 5e-6 | 1e-6);
    fcs (16 | 32);
    framing (esf | sf);
```

```

idle-cycle-flag (flags | ones);
invert-data;
line-encoding (ami | b8zs);
loopback (local | payload | remote);
remote-loopback-respond;
start-end-flag (filler | shared);
timeslots time-slot-range;
}
t3-options {
  atm-encapsulation (direct | plcp);
  bert-algorithm algorithm;
  bert-error-rate rate;
  bert-period seconds;
  buildout feet;
  (cbit-parity | no-cbit-parity);
  compatibility-mode (adtran | digital-link | kentrox | larscom | verilink) <subrate
    value>;
  fcs (16 | 32);
  (feac-loop-respond | no-feac-loop-respond);
  idle-cycle-flag value;
  (long-buildout | no-long-buildout);
  (loop-timing | no-loop-timing);
  loopback (local | payload | remote);
  (mac | no-mac);
  (payload-scrambler | no-payload-scrambler);
  start-end-flag (filler | shared);
}
traceoptions {
  flag flag <flag-modifier> <disable>;
}
transmit-bucket {
  overflow discard;
  rate percentage;
  threshold bytes;
}
(traps | no-traps);
unidirectional;
vlan-tagging;
vlan-vci-tagging;
unit logical-unit-number {
  accept-source-mac {
    mac-address mac-address {
      policer {
        input cos-policer-name;
        output cos-policer-name;
      }
    }
  }
}
account-layer2-overhead {
  value;
  egress bytes;
  ingress bytes;
}
accounting-profile name;
advisory-options {
  downstream-rate rate;
}

```

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

```

dynamic-call-admission-control {
    activation-priority priority;
    bearer-bandwidth-limit kilobits-per-second;
}
encapsulation type;
epd-threshold cells plp1 cells;
fragment-threshold bytes;
inner-vlan-id-range start start-id end end-id;
input-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
interleave-fragments;
inverse-arp;
layer2-policer {
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
link-layer-overhead percent;
minimum-links number;
mrru bytes;
multicast-dlci dlci-identifier;
multicast-vci vpi-identifier.vci-identifier;
multilink-max-classes number;
multipoint;
oam-liveness {
    down-count cells;
    up-count cells;
}
oam-period (seconds | disable);
output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
}

```

```
pap;
  default-pap-password password;
  local-name name;
  local-password password;
  passive;
}
dynamic-profile profile-name;
lcp-max-conf-req number;
lcp-restart-timer milliseconds;
loopback-clear-timer seconds;
ncp-max-conf-req number;
ncp-restart-timer milliseconds;
}
pppoe-options {
  access-concentrator name;
  auto-reconnect seconds;
  (client | server);
  service-name name;
  underlying-interface interface-name;
}
proxy-arp;
service-domain (inside | outside);
shaping {
  (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
   burst length);
  queue-length number;
}
short-sequence;
transmit-weight number;
(traps | no-traps);
trunk-bandwidth rate;
trunk-id number;
tunnel {
  backup-destination address;
  destination address;
  key number;
  routing-instance {
    destination routing-instance-name;
  }
  source source-address;
  ttl number;
}
vci vpi-identifier.vci-identifier;
vci-range start start-vci end end-vci;
vpi vpi-identifier;
vlan-id number;
vlan-id-list [vlan-id vlan-id-vlan-id];
vlan-id-range number-number;
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
vlan-tags-outer tpid.vlan-id inner-list [vlan-id vlan-id-vlan-id];
family family {
  accounting {
    destination-class-usage;
    source-class-usage {
      direction;
    }
  }
}
```



```

}
access-concentrator name;
address address {
    destination address;
}
bundle ml-fpc/pic/port | ls-fpc/pic/port);
duplicate-protection;
dynamic-profile profile-name;
filter {
    group filter-group-number;
    input filter-name;
    input-list {
        [ filter-names ];
        output filter-name;
    }
    output-list {
        [ filter-names ];
    }
}
ipsec-sa sa-name;
keep-address-and-control;
max-sessions number;
max-sessions-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 dlcid-identifier | vci vcid-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 vpid-identifier.vcid-identifier;
    }
    preferred;
    primary;
    (vrrp-group | vrrp-inet6-group) group-number {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        authentication-type authentication;
        authentication-key key;
        fast-interval milliseconds;
        (preempt | no-preempt) {
            hold-time seconds;
        }
        priority-number number;
        track {
            priority-cost seconds;
            priority-hold-time interface-name {
                bandwidth-threshold bits-per-second {
                    priority;
                }
                interface priority;
            }
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-address [ addresses ];
    }
}

```

```

    }
  }
}

```

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

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

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

```

logical-systems logical-system-name {
  interfaces interface-name {
    unit logical-unit-number {
      accept-source-mac {
        mac-address mac-address {
          policer {
            input cos-policer-name;
            output cos-policer-name;
          }
        }
      }
    }
    allow-any-vci;
    atm-scheduler-map (map-name | default);
    bandwidth rate;
    backup-options {
      interface interface-name;
    }
    cell-bundle-size cells;
    clear-dont-fragment-bit;
    compression {
      rtp {
        f-max-period number;
        port {
          minimum port-number;
          maximum port-number;
        }
      }
      queues [ queue-numbers ];
    }
  }
  compression-device interface-name;
  description text;
  interface {
    l2tp-interface-id name;
    (dedicated | shared);
  }
  dialer-options {
    activation-delay seconds;
    deactivation-delay seconds;
  }
}

```

```
dial-string [ dial-string-numbers ];
idle-timeout seconds;
initial-route-check seconds;
load-threshold number;
pool pool;
remote-name remote-callers;
watch-list {
    [ routes ];
}
}
disable;
dlci dlci-identifier;
drop-timeout milliseconds;
dynamic-call-admission-control {
    activation-priority priority;
    bearer-bandwidth-limit kilobits-per-second;
}
encapsulation type;
epd-threshold cells plp1 cells;
fragment-threshold bytes;
input-vlan-map {
    inner-tag-protocol-id;
    inner-vlan-id;
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    tag-protocol-id tpid;
    vlan-id number;
}
interleave-fragments;
inverse-arp;
layer2-policer {
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
link-layer-overhead percent;
minimum-links number;
mrru bytes;
multicast-dlci dlci-identifier;
multicast-vci vpi-identifier.vci-identifier;
multilink-max-classes number;
multipoint;
oam-liveness {
    up-count cells;
    down-count cells;
}
oam-period (seconds | disable);
output-vlan-map {
    inner-tag-protocol-id;
    inner-vlan-id;
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-swap);
    tag-protocol-id tpid;
    vlan-id number;
}
passive-monitor-mode;
peer-unit unit-number;
```

```

plp-to-clp;
point-to-point;
ppp-options {
  chap {
    access-profile name;
    default-chap-secret name;
    local-name name;
    passive;
  }
  compression {
    acfc;
    pfc;
  }
}
dynamic-profile profile-name;
pap {
  default-pap-password password;
  local-name name;
  local-password password;
  passive;
}
}
proxy-arp;
service-domain (inside | outside);
shaping {
  (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst length);
  queue-length number;
}
short-sequence;
transmit-weight number;
(traps | no-traps);
trunk-bandwidth rate;
trunk-id number;
tunnel {
  backup-destination address;
  destination address;
  key number;
  routing-instance {
    destination routing-instance-name;
  }
  source source-address;
  ttl number;
}
vci vpi-identifier.vci-identifier;
vlan-id number;
vlan-id-list [vlan-id vlan-id-vlan-id]
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
vlan-tags outer tpid.vlan-id inner-list [vlan-id vlan-id-vlan-id]
vpi vpi-identifier;
family family {
  accounting {
    destination-class-usage;
    source-class-usage {
      direction;
    }
  }
}

```

```
}
bundle interface-name;
filter {
  group filter-group-number;
  input filter-name;
  input-list {
    [ filter-names ];
  }
  output filter-name;
  output-list {
    [ filter-names ];
  }
}
ipsec-sa sa-name;
keep-address-and-control;
mtu bytes;
multicast-only;
no-redirects;
policer {
  arp policer-template-name;
  input policer-template-name;
  output policer-template-name;
}
primary;
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check <fail-filter filter-name> {
  <mode loose>;
}
sampling {
  direction;
}
service {
  input {
    service-set service-set-name <service-filter filter-name>;
    post-service-filter filter-name;
  }
  output {
    service-set service-set-name <service-filter filter-name>;
  }
}
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
unnumbered-address interface-name destination address destination-profile
profile-name;
address address {
  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 plp cells;
  }
}
```

```

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

```

Related Documentation

- [Junos OS Hierarchy and RFC Reference](#)
- [Ethernet Interfaces](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)

[edit protocols pppoe] Hierarchy Level

```

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

```

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 4

Statement Summary

address

```

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

```

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

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

Description Configure the interface address.

Options *address*—Address of the interface.

The remaining statements are explained separately.



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

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

Related Documentation

- [Configuring the Protocol Family on page 32](#)
- [negotiate-address on page 144](#)
- [unnumbered-address \(Ethernet\) on page 178](#)
- *Junos OS Administration Library for Routing Devices*
- *family*


aggregate (Hierarchical Policer)

Syntax	<pre>aggregate { if-exceeding { bandwidth-limit <i>bandwidth</i>; burst-size-limit <i>burst</i>; } then { discard; } }</pre>
Hierarchy Level	[edit firewall hierarchical-policer]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	On M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, configure an aggregate hierarchical policer.
Options	Options are described separately.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Applying Policers on page 11• <i>Junos OS Class of Service Library for Routing Devices</i>

accounting

Syntax	<pre>accounting { destination-class-usage; source-class-usage { direction; } }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Enable IP packet counters on an 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">• Enabling Source Class and Destination Class Usage on page 24


arp (Interfaces)

Syntax	<code>arp ip-address (mac multicast-mac) mac-address publish;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inetaddress <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inetaddress <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, configure Address Resolution Protocol (ARP) table entries, mapping IP addresses to MAC addresses.
Options	<p>ip-address—IP address to map to the MAC address. The IP address specified must be part of the subnet defined in the enclosing address statement.</p> <p>mac mac-address—MAC address to map to the IP address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i>. For example, 0011.2233.4455 or 00:11:22:33:44:55.</p> <p>multicast-mac mac-address—Multicast MAC address to map to the IP address. Specify the multicast MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i>. For example, 0011.2233.4455 or 00:11:22:33:44:55.</p> <p>publish—(Optional) Have the router or switch reply to ARP requests for the specified IP address. If you omit this option, the router or switch uses the entry to reach the destination but does not reply to ARP requests.</p>
<div> NOTE: The edit logical-systems hierarchy is not available on QFabric systems.</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">• <i>Configuring Static ARP Table Entries</i>• <i>Configuring Static ARP Entries</i>


bandwidth-limit (Hierarchical Policer)

Syntax	<code>bandwidth-limit <i>bps</i>;</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall hierarchical-policer aggregate if-exceeding], [edit dynamic-profiles <i>profile-name</i> firewall hierarchical-policer premium if-exceeding], [edit firewall hierarchical-policer aggregate if-exceeding], [edit firewall hierarchical-policer premium if-exceeding]
Release Information	Statement introduced in Junos OS Release 9.5. Support at the [edit dynamic-profiles ... if-exceeding] hierarchy level introduced in Junos OS Release 11.4.
Description	For M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, configure the maximum average bandwidth for premium or aggregate traffic in a hierarchical policer.
Options	<i>bps</i> —You can specify the number of bits per second either as a decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). Range: 32,000 through 50,000,000,000
Required Privilege Level	firewall —To view this statement in the configuration. firewall-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Hierarchical Policer Configuration Overview</i> • <i>Policer Bandwidth and Burst-Size Limits</i> • <i>Policer Color-Marking and Actions</i> • <i>Single Token Bucket Algorithm</i> • <i>Determining Proper Burst Size for Traffic Policers</i> • <i>aggregate (Hierarchical Policer)</i> • burst-size-limit (Hierarchical Policer) on page 104 • premium (Hierarchical Policer) on page 151

broadcast

Syntax	<code>broadcast address;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the broadcast address on the network or subnet. On a subnet you cannot specify a host address of 0, nor can you specify a broadcast address.
Default	The default broadcast address has a host portion of all ones.
Options	address —Broadcast address. The address must have a host portion of either all ones or all zeros. You cannot specify the addresses 0.0.0.0 or 255.255.255.255.
<div> NOTE: The edit logical-systems hierarchy is not available on QFabric systems.</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">• Configuring the Interface Address on page 35

bundle

Syntax	<code>bundle (ml-<i>fpc/pic/port</i> ls-<i>fpc/pic/port</i>);</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Associate the multilink interface with the logical interface it is joining. You can include this statement for the mlfr-end-to-end and mlfr-uni-nni protocol families only.
<div>  <p>NOTE:</p> <p>For M Series routers and T Series routers, the following caveats apply:</p> <ul style="list-style-type: none"> • Maximum supported throughput on the bundle interfaces is 45 Mbps. • Bundling of the logical interfaces under a T3 physical interface into the same or different bundles is not supported. </div>	
Options	<p>ml-<i>fpc/pic/port</i>—Name of the multilink interface you are linking.</p> <p>ls-<i>fpc/pic/port</i>—Name of the link services interface you are linking.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Services Interfaces Library for Routing Devices</i>

burst-size-limit (Hierarchical Policer)

Syntax	<code>burst-size-limit bytes;</code>
Hierarchy Level	[edit dynamic-profiles profile-name firewall hierarchical-policer aggregate if-exceeding], [edit dynamic-profiles profile-name firewall hierarchical-policer premium if-exceeding], [edit firewall hierarchical-policer aggregate if-exceeding], [edit firewall hierarchical-policer premium if-exceeding]
Release Information	Statement introduced in Junos OS Release 9.5. Support at the [edit dynamic-profiles ... if exceeding] hierarchy level introduced in Junos OS Release 11.4.
Description	For M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, configure the burst-size limit for premium or aggregate traffic in a hierarchical policer.
Options	bytes —Burst-size limit in bytes. The minimum recommended value is the maximum transmission unit (MTU) of the IP packets being policed. You can specify the value either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). Range: 1500 through 2,147,450,880
Required Privilege Level	firewall —To view this statement in the configuration. firewall-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Hierarchical Policer Configuration Overview</i>• <i>Policer Bandwidth and Burst-Size Limits</i>• <i>Policer Color-Marking and Actions</i>• <i>Single Token Bucket Algorithm</i>• <i>Determining Proper Burst Size for Traffic Policers</i>• <i>Hierarchical Policers</i>• <i>aggregate (Hierarchical Policer)</i>• bandwidth-limit (Hierarchical Policer) on page 101• premium (Hierarchical Policer) on page 151

cbr

Syntax	<code>cbr rate;</code>
Hierarchy Level	<p>[edit interfaces at-<i>fpc/pic/port</i> atm-options vpi <i>vpi-identifier</i> shaping],</p> <p>[edit interfaces at-<i>fpc/pic/port</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping],</p> <p>[edit interfaces at-<i>fpc/pic/ port</i> unit <i>logical-unit-number</i> shaping],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces at-<i>fpc/pic/port</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces at-<i>fpc/pic/port</i> unit <i>logical-unit-number</i> shaping]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For ATM encapsulation only, define a constant bit rate bandwidth utilization in the traffic-shaping profile.
Default	Unspecified bit rate (UBR); that is, bandwidth utilization is unlimited.
Options	<p>rate—Peak rate, in bits per second (bps) or cells per second (cps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation c; values expressed in cells per second are converted to bits per second by means of the formula 1 cps = 384 bps.</p> <p>For ATM1 and ATM2 OC3 interfaces, the maximum available rate is 100 percent of <i>line-rate</i>, or 135,600,000 bps. For ATM1 OC12 interfaces, the maximum available rate is 50 percent of <i>line-rate</i>, or 271,263,396 bps. For ATM2 IQ interfaces, the maximum available rate is 542,526,792 bps.</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"> • Defining the ATM Traffic-Shaping Profile • rtvbr on page 158 • shaping on page 162 • vbr on page 180

demux0 (Dynamic Interface)

Syntax `demux0 {
 unit logical-unit-number {
 demux-options {
 underlying-interface interface-name
 }
 family family {
 access-concentrator name;
 address address;
 demux-source {
 source-prefix;
 }
 duplicate-protection;
 dynamic-profile profile-name;
 filter {
 input filter-name;
 output filter-name;
 }
 mac-validate (loose | strict):
 max-sessions number;
 max-sessions-vsa-ignore;
 rpf-check {
 fail-filter filter-name;
 mode loose;
 }
 service-name-table table-name
 short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
 maximum-seconds>;
 unnumbered-address interface-name <preferred-source-address address>;
 }
 filter {
 input filter-name;
 output filter-name;
 }
 vlan-id number;
 }
 }`

Hierarchy Level [edit [dynamic-profiles](#) *profile-name* [interfaces](#)]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure the logical demultiplexing (demux) interface in a dynamic profile.

Logical IP demux interfaces do not support IPv4 and IPv6 dual stack.

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 Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*
 - For information about static IP demux interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

destination (IPCP)

Syntax	<code>destination address destination-profile profile-name;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For unnumbered interfaces with PPP encapsulation, specify the IP address of the remote interface.
Options	<i>address</i> —IP address of the remote interface. The remaining statement is explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring IPCP Options on page 41 • address on page 96 • negotiate-address on page 144 • <i>Junos OS Administration Library for Routing Devices</i>

destination (Tunnels)

Syntax	<code>destination address;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <i>family</i> inet address <i>address</i>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <i>family</i> inet <i>unnumbered-address</i></code> <code> <i>interface-name</i>],</code> <code>[edit interfaces <i>interface-name</i> <i>unit</i> <i>logical-unit-number</i> tunnel],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code> <i>family</i> inet address <i>address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code> <i>family</i> inet <i>unnumbered-address</i> <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <i>unit</i> <i>logical-unit-number</i></code> <code> tunnel]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for EX Series switches.
Description	For encrypted, PPP-encapsulated, and tunnel interfaces, specify the remote address of the connection.
Options	<i>address</i> —Address of the remote side of the connection.
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 the Interface Address on page 35• <i>Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</i>• <i>Junos OS Services Interfaces Library for Routing Devices</i>• <i>point-to-point</i>

destination-class-usage

Syntax	<code>destination-class-usage;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet accounting], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet accounting]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Enable packet counters on an interface that count packets that arrive from specific customers and are destined for specific prefixes on the provider core router.
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"> • Enabling Source Class and Destination Class Usage on page 24 • accounting on page 99 • source-class-usage on page 163

destination-profile

Syntax	<code>destination-profile <i>name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i> destination <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i> destination <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For interfaces with PPP encapsulation, assign PPP properties to the remote destination end. You define the profile at the [edit access group-profile <i>name</i> ppp] hierarchy level.
Options	<i>name</i> —Profile name defined at the [edit access group-profile <i>name</i> ppp] 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"> • Configuring IPCP Options on page 41 • destination (IPCP) on page 107 • Junos OS Administration Library for Routing Devices

dynamic-profiles

```

Syntax  dynamic-profiles {
        profile-name {
            class-of-service {
                interfaces {
                    interface-name ;
                }
                unit logical-unit-number {
                    classifiers {
                        type (classifier-name | default);
                    }
                    output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
                    rewrite-rules {
                        dscp (rewrite-name | default);
                        dscp-ipv6 (rewrite-name | default);
                        ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                        inet-precedence (rewrite-name | default);
                    }
                }
            }
        }
        scheduler-maps {
            map-name {
                forwarding-class class-name scheduler scheduler-name;
            }
        }
        schedulers {
            (scheduler-name) {
                buffer-size (seconds | percent percentage | remainder | temporal microseconds);
                drop-profile-map loss-priority (any | low | medium-low | medium-high | high)
                    protocol (any | non-tcp | tcp) drop-profile profile-name;
                excess-priority (low | high | $junos-cos-scheduler-excess-priority);
                excess-rate (percent percentage | percent $junos-cos-scheduler-excess-rate);
                overhead-accounting (shaping-mode) <bytes (byte-value)>;
                priority priority-level;
                shaping-rate (rate | predefined-variable);
                transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
            }
        }
        traffic-control-profiles profile-name {
            delay-buffer-rate (percent percentage | rate | $junos-cos-delay-buffer-rate);
            excess-rate (percent percentage | proportion value | percent $junos-cos-excess-rate);
            guaranteed-rate (percent percentage | rate | $junos-cos-guaranteed-rate);
            overhead-accounting (shaping-mode) <bytes (byte-value)>;
            scheduler-map map-name;
            shaping-rate (rate | predefined-variable);
        }
    }
    firewall {
        family family {
            fast-update-filter filter-name {
                interface-specific;
                match-order [match-order];
            }
        }
    }

```



```

term term-name {
  from {
    match-conditions;
  }
  then {
    action;
    action-modifiers;
  }
  only-at-create;
}
}
firewall {
  family family {
    fast-update-filter filter-name {
      interface-specific;
      match-order [match-order];
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
        only-at-create;
      }
    }
    filter filter-name {
      interface-specific;
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      }
    }
  }
  policer policer-name {
    filter-specific;
    if-exceeding {
      (bandwidth-limit bps | bandwidth-percent percentage);
      burst-size-limit bytes;
    }
    logical-bandwidth-policer;
    logical-interface-policer;
    physical-interface-policer;
    then {
      policer-action;
    }
  }
}
hierarchical-policer policer-name {
  aggregate {
    if-exceeding {
      bandwidth-limit-limit bps;
      burst-size-limit bytes;
    }
    then {

```

```
        policer-action;
    }
}
premium {
    if-exceeding {
        bandwidth-limit bps;
        burst-size-limit bytes;
    }
    then {
        policer-action;
    }
}
}
three-color-policer policer-name {
    action {
        loss-priority high then discard;
    }
    logical-interface-policer;
    single-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        excess-burst-size bytes;
    }
    two-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        peak-burst-size bytes;
        peak-information-rate bps;
    }
}
}
}
policy-options {
    prefix-list name {
        ip-addresses;
    }
}
}
}
interfaces interface-name {
    interface-set interface-set-name {
        interface interface-name {
            unit logical unit number {
                advisory-options {
                    downstream-rate rate;
                    upstream-rate rate;
                }
            }
        }
    }
}
}
unit logical-unit-number {
    auto-configure {
        agent-circuit-identifier {
            dynamic-profile profile-name;
        }
    }
}
```

```

    }
}
encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid |
atm-tcc-vc-mux | atm-mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux |
atm-snap | atm-tcc-snap | atm-vc-mux | ether-over-atm-llc |
ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp | ethernet |
frame-relay-ccc | frame-relay-ppp | frame-relay-tcc | frame-relay-ether-type |
frame-relay-ether-type-tcc | multilink-frame-relay-end-to-end | multilink-ppp |
ppp-over-ether | ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc | vlan-vci-ccc
| vlan-tcc | vlan-vpls);
family family {
    address address;
    filter {
        adf {
            counter;
            input-precedence precedence;
            not-mandatory;
            output-precedence precedence;
            rule rule-value;
        }
        input filter-name (
            precedence precedence;
        )
        output filter-name {
            precedence precedence;
        }
    }
}
rpf-check {
    fail-filter filter-name;
    mode loose;
}
service {
    input {
        service-set service-set-name {
            service-filter filter-name;
        }
        post-service-filter filter-name;
    }
    input-vlan-map {
        inner-tag-protocol-id tpid;
        inner-vlan-id number;
        (push | swap);
        tag-protocol-id tpid;
        vlan-id number;
    }
    output {
        service-set service-set-name {
            service-filter filter-name;
        }
    }
    output-vlan-map {
        inner-tag-protocol-id tpid;
        inner-vlan-id number;
        (pop | swap);
        tag-protocol-id tpid;
        vlan-id number;
    }
}

```

```
    }
  }
  unnumbered-address interface-name <preferred-source-address address>;
}
ppp-options {
  chap;
  pap;
}
vlan-id number;
vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
}
}
interfaces {
  demux0 {...}
}
interfaces {
  pp0 {...}
}
protocols {
  igmp {
    interface interface-name {
      accounting;
      disable;
      group-policy;
      immediate-leave;
      no-accounting;
      promiscuous-mode;
      ssm-map ssm-map-name;
      static {
        group group {
          source source;
        }
      }
      version version;
    }
  }
  mld {
    interface interface-name {
      disable;
      (accounting | no-accounting);
      group-policy;
      immediate-leave;
      oif-map;
      passive;
      ssm-map ssm-map-name;
      static {
        group multicast-group-address {
          exclude;
          group-count number;
          group-increment increment;
          source ip-address {
            source-count number;
            source-increment increment;
          }
        }
      }
    }
  }
  version version;
}
```

```

    }
  }
  router-advertisement {
    interface interface-name {
      current-hop-limit number;
      default-lifetime seconds;
      (managed-configuration | no-managed-configuration);
      max-advertisement-interval seconds;
      min-advertisement-interval seconds;
      (other-stateful-configuration | no-other-stateful-configuration);
      prefix prefix;
      reachable-time milliseconds;
      retransmit-timer milliseconds;
    }
  }
}
routing-instances routing-instance-name {
  interface interface-name;
  routing-options {
    access {
      route prefix {
        next-hop next-hop;
        metric route-cost;
        preference route-distance;
        tag route-tag;
      }
    }
    access-internal {
      route subscriber-ip-address {
        qualified-next-hop underlying-interface {
          mac-address address;
        }
      }
    }
    multicast {
      interface interface-name {
        no-qos-adjust;
      }
    }
  }
}
rib routing-table-name {
  access {
    route prefix {
      next-hop next-hop;
      metric route-cost;
      preference route-distance;
      tag route-tag;
    }
  }
  access-internal {
    route subscriber-ip-address {
      qualified-next-hop underlying-interface {
        mac-address address;
      }
    }
  }
}

```

```
    }
  }
}
routing-options {
  access {
    route prefix {
      next-hop next-hop;
      metric route-cost;
      preference route-distance;
      tag route-tag;
    }
  }
  access-internal {
    route subscriber-ip-address {
      qualified-next-hop underlying-interface {
        mac-address address;
      }
    }
  }
}
multicast {
  interface interface-name {
    no-qos-adjust;
  }
}
}
variables {
  variable-name {
    default-value default-value;
    equals expression;
    mandatory;
    radius {
      vendor-id id {
        attribute attribute-number;
        tag tag-number;
      }
    }
    uid;
    uid-reference;
  }
}
}
```

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

Release Information Statement introduced in Junos OS Release 9.2.
Support at the **filter**, **policer**, **hierarchical-policer**, **three-color-policer**, and **policy options** hierarchy levels introduced in Junos OS Release 11.4.

Description Create dynamic profiles for use with DHCP or PPP client access.

Options *profile-name*—Name of the dynamic profile; string of up to 80 alphanumeric characters.
The remaining statements are explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring a Basic Dynamic Profile</i> • <i>Configuring Dynamic VLANs Based on Agent Circuit Identifier Information</i> • <i>Dynamic Profiles Overview</i>

epd-threshold (Logical Interface)

Syntax	<code>epd-threshold cells plp1 cells;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i>], [edit logical-systems <i>logical-system-name</i> 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> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>For ATM2 IQ interfaces only, define the early packet discard (EPD) threshold on a VC. The EPD threshold is a limit on the number of transmit packets that can be queued. Packets that exceed the limit are discarded. For interfaces configured in trunk mode, you can also configure dual EPD thresholds depending on the packet loss priorities (PLPs).</p>
Default	<p>Approximately 1 percent of the available cell buffers. If shaping is enabled, the default EPD threshold is proportional to the shaping rate according to the following formula:</p> $\text{default epd-threshold} = \text{number of buffers} * \text{shaping rate} / \text{line rate}$ <p>The minimum EPD threshold value is 48 cells. If the default EPD threshold formula results in an EPD threshold of less than 48 cells, the result will be ignored, and the minimum value of 48 cells will be used.</p>
Options	<p>cells—Maximum number of cells. Range: For 1-port and 2-port OC12 interfaces, 48 through 425,984 cells</p>
Required Privilege Level	<p>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the ATM2 IQ EPD Threshold</i> • <i>Configuring Two EPD Thresholds per Queue</i>

eui-64

Syntax	eui-64;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>number</i> family inet6 address <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 12.2 for the QFX Series.
Description	For interfaces that carry IP version 6 (IPv6) traffic, automatically generate the host number portion of interface addresses. Not supported on QFX Series switches.
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 the Interface Address on page 35

family (Dynamic Standard Interface)

```

Syntax  family family {
    access-concentrator name;
    address address;
    duplicate-protection;
    dynamic-profile profile-name;
    filter {
        adf {
            counter;
            input-precedence precedence;
            not-mandatory;
            output-precedence precedence;
            rule rule-value;
        }
        input filter-name {
            precedence precedence;
        }
        output filter-name {
            precedence precedence;
        }
    }
    mac-validate (loose | strict);
    max-sessions number;
    max-sessions-vs-a-ignore;
    rpf-check {
        fail-filter filter-name;
        mode loose;
    }
    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;
            }
        }
    }
    service-name-table table-name
    short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
        maximum-seconds>;
    unnumbered-address interface-name <preferred-source-address address>;
}

```

Hierarchy Level [edit [dynamic-profiles](#) *profile-name* [interfaces](#) *interface-name* [unit](#) *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.2.
Option **pppoe** introduced in Junos OS Release 11.2.

Description Configure protocol family information for the logical interface.



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

Options *family*—Protocol family:

- **inet**—IP version 4 suite
- **inet6**—IP version 6 suite
- **pppoe**—(MX Series routers with MPCs only) Point-to-Point Protocol over Ethernet
- **vpls**—Virtual private LAN service

The remaining statements are explained separately.

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

Related Documentation

- For general information about configuring static interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.
- “Configuring the Protocol Family,” in the *Junos OS Network Interfaces Library for Routing Devices*.

family

```

Syntax  family family {
        accounting {
            destination-class-usage;
            source-class-usage {
                (input | output | input output);
            }
        }
        access-concentrator name;
        address address {
            ... the address subhierarchy appears after the main [edit interfaces interface-name unit
                logical-unit-number family family-name] hierarchy ...
        }
        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;
        max-sessions-vsa-ignore;
        mtu bytes;
        multicast-only;
        negotiate-address;
        no-redirects;
        policer {
            arp policer-template-name;
            input policer-template-name;
            output policer-template-name;
        }
        primary;
        protocols [inet iso mpls];
        proxy inet-address address;
        receive-options-packets;
        receive-ttl-exceeded;
        remote (inet-address address | mac-address address);

```

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

```

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

```

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

Release Information Statement introduced before Junos OS Release 7.4.
Option **max-sessions-vs-a-ignore** introduced in Junos OS Release 11.4.

Description Configure protocol family information for the logical interface.



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

Options *family*—Protocol family:

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



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

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


The remaining statements are explained separately.

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

Related Documentation

- [Configuring the Protocol Family on page 32](#)
- *Example: Configuring E-LINE and E-LAN Services for a PBB Network on MX Series Routers*
- *Junos OS Services Interfaces Library for Routing Devices*

fast-aps-switch

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

filter

Syntax	<pre>filter { group <i>filter-group-number</i>; input <i>filter-name</i>; input-list [<i>filter-names</i>]; output <i>filter-name</i>; output-list [<i>filter-names</i>]; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Apply a filter to an interface. You can also use filters for encrypted traffic. When you configure filters, you can configure them under the family ethernet-switching , inet , inet6 , mpls , or vpls only.
Options	<p>group <i>filter-group-number</i>—Define an interface to be part of a filter group. Range: 1 through 255</p> <p>input <i>filter-name</i>—Name of one filter to evaluate when packets are received on the interface.</p> <p>output <i>filter-name</i>—Name of one filter to evaluate when packets are transmitted on the interface.</p> <p>The remaining statements are explained separately.</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">• Applying a Filter to an Interface on page 19• <i>Junos OS Services Interfaces Library for Routing Devices</i>• <i>Routing Policy Feature Guide for Routing Devices</i>• <i>Junos OS Administration Library for Routing Devices</i>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i>• <i>Configuring Firewall Filters (CLI Procedure)</i>• <i>Configuring Firewall Filters and Policers for VPLS</i>• <i>family</i>

forward-and-send-to-re

Syntax	forward-and-send-to-re;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet targeted-broadcast], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet targeted-broadcast]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	Specify that IP packets destined for a Layer 3 broadcast address be forwarded to an egress interface and the Routing Engine. The packets are broadcast only if the egress interface is a LAN interface.
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 Targeted Broadcast on page 57 • targeted-broadcast on page 164 • Understanding Targeted Broadcast on page 59

forward-only

Syntax	forward-only;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet targeted-broadcast], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet targeted-broadcast]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	Specify that IP packets destined for a Layer 3 broadcast address be forwarded to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.
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 Targeted Broadcast on page 57 • targeted-broadcast on page 164 • Understanding Targeted Broadcast on page 59

hierarchical-policer

Syntax hierarchical-policer *name* {
aggregate {
if-exceeding {
bandwidth-limit *bandwidth*;
burst-size-limit *burst*;
}
then {
discard;
}
}
premium {
if-exceeding {
bandwidth-limit *bandwidth*;
burst-size-limit *burst*;
}
then {
discard;
}
}
}

Hierarchy Level [edit firewall]

Release Information Statement introduced in Junos OS Release 9.5.

Description For M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, specify a hierarchical policer.

Options Options are described separately.

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

Related Documentation

- [Applying Policers on page 11](#)
- *Junos OS Class of Service Library for Routing Devices*

if-exceeding

Syntax	if-exceeding { bandwidth-limit <i>bandwidth</i> ; burst-size-limit <i>burst</i> ; }
Hierarchy Level	[edit firewall hierarchical-policer aggregate], [edit firewall hierarchical-policer premium]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	For M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, specify bandwidth and burst limits for an aggregate level of a hierarchical policer.
Options	Options are described separately.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Applying Policers on page 11 • <i>Junos OS Class of Service Library for Routing Devices</i>


input

Syntax	input { service-set <i>service-set-name</i> <service-filter <i>filter-name</i> >; post-service-filter <i>filter-name</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define one or more input service sets and filters, and one postservice filter to be applied to traffic.
Options	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	<ul style="list-style-type: none"> • <i>Junos OS Services Interfaces Library for Routing Devices</i>

input-list

Syntax	<code>input-list [<i>filter-names</i>];</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> filter], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> filter]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	Apply a group of filters to evaluate when packets are received on an interface.
Options	[<i>filter-names</i>] —Name of a filter to evaluate when packets are received on the interface. Up to 16 filters can be included in a filter input list.
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">• Applying a Filter to an Interface on page 19• <i>Routing Policy Feature Guide for Routing Devices</i>• <i>Junos OS Administration Library for Routing Devices</i>• output-list on page 147

interface-mode

Syntax	interface-mode (access trunk);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching] [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge]
Release Information	Statement introduced in Junos OS Release 9.2. Statement introduced in Junos Release 12.3R2 for EX Series switches.
Description	Determines whether the logical interface accepts or discards packets based on VLAN tags. Specify the trunk option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the vlan-id-list statement, then forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the access option to accept packets with no VLAN ID, then forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the vlan-id statement.
<div style="display: flex; align-items: center;">  <div> <p>NOTE: On MX Series routers, if you want IGMP snooping to be functional for a bridge domain, then you should not configure interface-mode and irb for that bridge. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed. For more information, see <i>Configuring a Trunk Interface on a Bridge Network</i>.</p> </div> </div>	
Options	<p>access—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the vlan-id statement.</p> <p>trunk—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the vlan-id-list statement.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring a Logical Interface for Access Mode</i> • <i>Configuring a Logical Interface for Trunk Mode</i>

interfaces

Syntax	interfaces { ... }
Hierarchy Level	[edit]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure interfaces on the router or switch.
Default	The management and internal Ethernet interfaces are automatically configured. You must configure all other 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">• <i>Physical Interface Configuration Statements Overview</i>• <i>Configuring Aggregated Ethernet Link Protection</i>

interfaces (Static and Dynamic Subscribers)

```

Syntax  interfaces {
        interface-name {
            unit logical-unit-number {
                auto-configure {
                    agent-circuit-identifier {
                        dynamic-profile profile-name;
                    }
                }
            }
            family family {
                access-concentrator name;
                address address;
                duplicate-protection;
                dynamic-profile profile-name;
                filter {
                    adf {
                        counter;
                        input-precedence precedence;
                        not-mandatory;
                        output-precedence precedence;
                        rule rule-value;
                    }
                    input filter-name (
                        precedence precedence;
                        shared-name filter-shared-name;
                    )
                    output filter-name {
                        precedence precedence;shared-name filter-shared-name;
                    }
                }
                max-sessions number;
                max-sessions-vsa-ignore;
                rpf-check {
                    mode loose;
                }
                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;
                        }
                    }
                }
                service-name-table table-name
                short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
                    maximum-seconds>;
                unnumbered-address interface-name <preferred-source-address address>;
            }
        }
    }

```

```
filter {
  input filter-name;
  shared-name filter-shared-name;
  output filter-name;
  shared-name filter-shared-name;
}
ppp-options {
  chap;
  pap;
}
proxy-arp;
vlan-id;
vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
}
vlan-tagging;
}
interface-set interface-set-name {
  interface interface-name {
    unit logical unit number {
      advisory-options {
        downstream-rate rate;
        upstream-rate rate;
      }
    }
  }
}
pppoe-underlying-options {
  max-sessions number;
}
}
demux0 {
  unit logical-unit-number {
    demux-options {
      underlying-interface interface-name
    }
    family family {
      access-concentrator name;
      address address;
      duplicate-protection;
      dynamic-profile profile-name;
      demux-source {
        source-prefix;
      }
    }
    filter {
      input filter-name {
        precedence precedence;
        shared-name filter-shared-name;
      }
      output filter-name {
        precedence precedence;
        shared-name filter-shared-name;
      }
    }
  }
  mac-validate (loose | strict);
  max-sessions number;
  max-sessions-vsa-ignore;
  rpf-check {
```



```

    fail-filter filter-name;
    mode loose;
  }
  service-name-table table-name
  short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
    maximum-seconds>;
  unnumbered-address interface-name <preferred-source-address address>;
}
filter {
  input filter-name;
  output filter-name;
}
vlan-id number;
vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
}
}
pp0 {
  unit logical-unit-number {
    keepalives interval seconds;
    no-keepalives;
    pppoe-options {
      underlying-interface interface-name;
      server;
    }
    ppp-options {
      authentication [ authentication-protocols ];
      chap {
        challenge-length minimum minimum-length maximum maximum-length;
      }
    }
    pap;
  }
  family inet {
    unnumbered-address interface-name;
    address address;
    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;
        }
      }
    }
  }
  filter {
    input filter-name {
      precedence precedence;
      shared-name filter-shared-name;
    }
    output filter-name {
      precedence precedence;
      shared-name filter-shared-name;
    }
  }
}

```

```
    }  
  }  
}
```

Hierarchy Level [edit **dynamic-profiles** *profile-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Define interfaces for dynamic profiles.

Options *interface-name*—The interface variable (**\$junos-interface-ifd-name**). The interface variable is dynamically replaced with the interface the DHCP client accesses when connecting to the router.



NOTE: Though we do not recommend it, you can also enter the specific name of the interface you want to assign to the dynamic profile.

The remaining statements are explained separately.

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

Related Documentation

- *Configuring Static Subscriber Interfaces in Dynamic Profiles*
- *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*
- *Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles*
- *Configuring Dynamic VLANs Based on Agent Circuit Identifier Information*
- *Subscriber Interface Overview*
- *Relationship Between Subscribers and Interfaces in an Access Network*
- For general information about configuring static interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*
- For information about static IP demux interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

inverse-arp

Syntax	inverse-arp;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> multipoint-destination <i>destination</i>], [edit logical-systems <i>logical-system-name</i> 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> family inet address <i>address</i> multipoint-destination <i>destination</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For ATM encapsulation, enable responses to receive inverse ATM ARP requests. For Frame Relay encapsulation, enable responses to receive inverse Frame Relay ARP requests.
Default	Inverse ARP is disabled on all ATM and Frame Relay 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"> • <i>Configuring Inverse ATM1 or ATM2 ARP</i> • <i>Configuring Inverse Frame Relay ARP</i>

ipsec-sa

Syntax	ipsec-sa <i>sa-name</i> ;
Hierarchy Level	[edit interfaces <i>es-fpc/pic/port</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>es-fpc/pic/port</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the IP Security (IPsec) security association (SA) name associated with the interface.
Options	sa-name —IPsec security association name.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Services Interfaces Library for Routing Devices</i> • <i>Junos OS Administration Library for Routing Devices</i>

keep-address-and-control

Syntax	keep-address-and-control;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>ccc</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>ccc</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For interfaces with encapsulation type PPP CCC, do not remove the address and control bytes before encapsulating the packet into a tunnel.
Default	If you do not include this statement, address and control bytes are removed before encapsulating the packet into a tunnel.
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">• Disabling the Removal of Address and Control Bytes on page 10

logical-systems

Syntax	logical-systems { <i>logical-system-name</i> { ... <i>logical-system-configuration</i> ... } }
Hierarchy Level	[edit]
Release Information	Statement introduced before Junos OS Release 7.4. Statement name changed from logical-routers in Junos OS Release 9.3.
Description	Configure a logical system.
Options	<i>logical-system-name</i> —Name of the logical system.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Logical Systems Feature Guide for Routing Devices</i>

mode (Dynamic Profiles)

Syntax	mode loose;
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet) rpf-check]
Release Information	Statement introduced in Junos OS Release 9.6.
Description	Check whether the packet has a source address with a corresponding prefix in the routing table. If a corresponding prefix is not found, unicast reverse path forwarding (RPF) loose mode does not accept the packet. Unlike strict mode, loose mode does not check whether the interface expects to receive a packet with a specific source address prefix.
Default	If you do not include this statement, unicast RPF is in strict mode.
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 Unicast RPF Strict Mode on page 51

mode (Interfaces)

Syntax	mode loose;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) rpf-check], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) rpf-check]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Check whether the packet has a source address with a corresponding prefix in the routing table. If a corresponding prefix is not found, unicast reverse path forwarding (RPF) loose mode does not accept the packet. Unlike strict mode, loose mode does not check whether the interface expects to receive a packet with a specific source address prefix.
Default	If you do not include this statement, unicast RPF is in strict mode.
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 Unicast RPF Strict Mode on page 51

mtu

Syntax	<code>mtu bytes;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>],</p> <p>[edit interfaces <i>interface-range name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols l2circuit local-switching interface <i>interface-name</i> backup-neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls],</p> <p>[edit protocols l2circuit local-switching interface <i>interface-name</i> backup-neighbor <i>address</i>],</p> <p>[edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</p> <p>[edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Layer 2 VPNs and VPLS introduced in Junos OS Release 10.4.</p> <p>Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Support at the [set interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>ccc</i>] hierarchy level introduced in Junos OS Release 12.3R3 for MX Series routers.</p>
Description	<p>Specify the maximum transmission unit (MTU) size for the media or protocol. The default MTU size depends on the device type. Changing the media MTU or protocol MTU causes an interface to be deleted and added again.</p> <p>To route jumbo data packets on the routed VLAN interface (RVI) on EX Series switches, you must configure the jumbo MTU size on the member physical interfaces and also on the RVI itself (the vlan interface).</p>



CAUTION: For EX Series switches, setting or deleting the jumbo MTU size on the RVI (the **vlan** interface) while the switch is transmitting packets might cause packets to be dropped.



NOTE: If a packet whose size is larger than the configured MTU size is received on the receiving interface, the packet is eventually dropped. The value

considered for MRU (maximum receive unit) size is also the same as the MTU size configured on that interface.



NOTE: Not all devices allow you to set an MTU value, and some devices have restrictions on the range of allowable MTU values. You cannot configure an MTU for management Ethernet interfaces (fxp0, em0, or me0) or for loopback, multilink, and multicast tunnel devices.

For more information about configuring MTU for specific interfaces and router or switch combinations, see *Configuring the Media MTU*.

Options *bytes*—MTU size.

Range: 256 through 9192 bytes, 256 through 9500 bytes (Junos OS 12.1X48R2 for PTX Series routers)

Default: 1500 bytes (INET, INET6, and ISO families), 1448 bytes (MPLS), 1514 bytes (EX Series switch interfaces)

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

Related Documentation

- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring the Media MTU*
- *Configuring the MTU for Layer 2 Interfaces*
- *Configuring Routed VLAN Interfaces (CLI Procedure)*
- [Setting the Protocol MTU on page 9](#)

multicast-only

Syntax	multicast-only;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the unit and family so that it can transmit and receive multicast traffic only. You can configure this property on the IP family only.
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 the Protocol Family on page 32• <i>Junos OS Services Interfaces Library for Routing Devices</i>• <i>tunnel</i>

multipoint-destination

Syntax	<pre> multipoint-destination address dlcid dlcid-identifier; multipoint-destination address { epd-threshold cells; inverse-arp; oam-liveness { down-count cells; up-count cells; } oam-period (disable seconds); shaping { (cbr rate rtvbr peak rate sustained rate burst length vbr peak rate sustained rate burst length); queue-length number; } vci vpi-identifier.vci-identifier; } </pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For point-to-multipoint Frame Relay or ATM interfaces only, enable the support of multicast on the interface. You can configure multicast support on the interface if the Frame Relay or ATM switch performs multicast replication.
Options	<p>address—Address of the remote side of the point-to-multipoint connection.</p> <p>dlcid-identifier—For Frame Relay interfaces, the data-link connection identifier. Range: 0 through 0xFFFFFFFF (24 bits)</p> <p>vci-identifier—For ATM interfaces, the virtual circuit identifier. Range: 0 through 16,384</p> <p>vpi-identifier—For ATM interfaces, the virtual path identifier. Range: 0 through 255 Default: 0</p> <p>The remaining 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"> • <i>Configuring a Point-to-Point ATM1 or ATM2 IQ Connection</i> • <i>Configuring a Point-to-Multipoint Frame Relay Connection</i> • <i>dlci</i> • <i>encapsulation (Logical Interface)</i>

negotiate-address

Syntax	negotiate-address;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For interfaces with PPP encapsulation, enable the interface to be assigned an IP address by the remote end.
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 IPCP Options on page 41• address on page 96• unnumbered-address (PPP) on page 179• <i>Junos OS Administration Library for Routing Devices</i>

no-redirects

Syntax	no-redirects;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Do not send protocol redirect messages on the interface. To disable the sending of protocol redirect messages for the entire router or switch, include the no-redirects statement at the [edit system] hierarchy level.
Default	Interfaces send protocol redirect messages.
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">• Disabling the Transmission of Redirect Messages on an Interface on page 10• <i>Junos OS Administration Library for Routing Devices</i>

oam-liveness

Syntax	<pre>oam-liveness { down-count <i>cells</i>; up-count <i>cells</i>; }</pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i>], [edit logical-systems <i>logical-system-name</i> 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> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>For ATM encapsulation only, configure Operation, Administration, and Maintenance (OAM) F5 loopback cell count thresholds. Not supported on ATM-over-SHDSL interfaces.</p> <p>For ATM2 IQ PICs only, configure OAM F4 loopback cell count thresholds at the [edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i>] hierarchy level.</p>
Options	<p>down-count <i>cells</i>—Minimum number of consecutive OAM F4 or F5 loopback cells lost before a VC is declared down. Range: 1 through 255 Default: 5 cells</p> <p>up-count <i>cells</i>—Minimum number of consecutive OAM F4 or F5 loopback cells received before a VC is declared up. Range: 1 through 255 Default: 5 cells</p>
Required Privilege Level	<p>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring the ATM OAM F5 Loopback Cell Threshold

oam-period

Syntax	oam-period (disable seconds);
Hierarchy Level	[edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i>], [edit logical-systems <i>logical-system-name</i> 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> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For ATM encapsulation only, configure the OAM F5 loopback cell period. Not supported on ATM-over-SHDSL interfaces. For ATM2 IQ PICs only, configure the OAM F4 loopback cell period at the [edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i>] hierarchy level.
Default	If you omit this statement, OAM F5 loopback cells are not initiated, but the interface still responds if it receives OAM F5 loopback cells.
Options	disable —Disable the OAM loopback cell transmit feature. seconds —OAM loopback cell period. Range: 1 through 900 seconds
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">Defining the ATM OAM F5 Loopback Cell Period

output

Syntax	<code>output { service-set service-set-name <service-filter filter-name>; }</code>
Hierarchy Level	[edit interfaces interface-name unit logical-unit-number family inet service], [edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet service]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define one or more output service sets and filters to be applied to traffic.
Options	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	<ul style="list-style-type: none"> • <i>Junos OS Services Interfaces Library for Routing Devices</i>

output-list

Syntax	<code>output-list [filter-names];</code>
Hierarchy Level	[edit interfaces interface-name unit logical-unit-number family family filter], [edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family filter]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	Apply a group of filters to evaluate when packets are transmitted on an interface.
Options	[<i>filter-names</i>]—Name of a filter to evaluate when packets are transmitted on the interface. Up to 16 filters can be included in a filter input list.
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"> • Applying a Filter to an Interface on page 19 • input-list on page 130 • <i>Routing Policy Feature Guide for Routing Devices</i> • <i>Junos OS Services Interfaces Library for Routing Devices</i> • <i>Junos OS Administration Library for Routing Devices</i>


policer (Interface)

Syntax	<pre>policer { arp <i>policer-template-name</i>; input <i>policer-template-name</i>; output <i>policer-template-name</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Apply a policer to an interface.
Options	<p>arp <i>policer-template-name</i>—For inet family only, name of one policer to evaluate when ARP packets are received on the interface.</p> <p>input <i>policer-template-name</i>—Name of one policer to evaluate when packets are received on the interface.</p> <p>output <i>policer-template-name</i>—Name of one policer to evaluate when packets are transmitted on the interface.</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">• Applying Policers on page 11• <i>Configuring Firewall Filters and Policers for VPLS</i>• <i>Routing Policy Feature Guide for Routing Devices</i>• <i>Junos OS Services Interfaces Library for Routing Devices</i>

post-service-filter

Syntax	<code>post-service-filter <i>filter-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service input], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service input]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define the filter to be applied to traffic after service processing. The filter is applied only if a service set is configured and selected.
Options	<i>filter-name</i> —Identifier for postservice filter.
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"> • Junos OS Services Interfaces Library for Routing Devices

preferred

Syntax	<code>preferred;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure this address to be the preferred address on the interface. If you configure more than one address on the same subnet, the preferred source address is chosen by default as the source address when you initiate frame transfers to destinations on the subnet.
<div style="display: flex; align-items: center;">  <div> <p>NOTE: The <code>edit logical-systems</code> hierarchy is not available on QFabric systems.</p> </div> </div>	
Default	The lowest-numbered address on the subnet is the preferred address.
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 the Interface Address on page 35


preferred-source-address

Syntax	<code>preferred-source-address address;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> unnumbered-address <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> unnumbered-address <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.0.
Description	<p>For unnumbered Ethernet interfaces configured with a loopback interface as the donor interface, specify one of the loopback interface's secondary addresses as the preferred source address for the unnumbered Ethernet interface. Configuring the preferred source address enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet interfaces in your network.</p> <p>Configuration of a preferred source address for unnumbered Ethernet interfaces is supported for the IPv4 and IPv6 address families.</p>
Options	address —Secondary IP address of the donor loopback interface.
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 a Preferred Source Address for Unnumbered Ethernet or Demux Interfaces on page 46• address on page 96• <i>Junos OS Administration Library for Routing Devices</i>

premium (Hierarchical Policer)

Syntax	<pre> premium { if-exceeding { bandwidth-limit <i>bandwidth</i>; burst-size-limit <i>burst</i>; } then { discard; } } </pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall hierarchical-policer], [edit firewall hierarchical-policer]
Release Information	Statement introduced in Junos OS Release 9.5. Support at the [edit dynamic-profiles ... hierarchical-policer <i>name</i>] hierarchy level introduced in Junos OS Release 11.4.
Description	On M40e, M120, and M320 edge routers with FPC input as FFPC and FPC output as SFPC, and on MX Series, T320, T640, and T1600 edge routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, specify a premium level for a hierarchical policer.
Options	Options are described separately.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Applying Policers on page 11 • <i>Junos OS Class of Service Library for Routing Devices</i> • <i>Hierarchical Policer Configuration Overview</i> • <i>Hierarchical Policers</i> • <i>aggregate (Hierarchical Policer)</i> • bandwidth-limit (Hierarchical Policer) on page 101 • burst-size-limit (Hierarchical Policer) on page 104 • <i>hierarchical-policer</i> • <i>if-exceeding (Hierarchical Policer)</i>

primary (Address on Interface)

Syntax	<code>primary;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure this address to be the primary address of the protocol on the interface. If the logical unit has more than one address, the primary address is used by default as the source address when packet transfer originates from the interface and the destination address does not indicate the subnet.
	<div> NOTE: The edit logical-systems hierarchy is not available on QFabric systems.</div>
Default	For unicast traffic, the primary address is the lowest non-127 (in other words, non-loopback) preferred address on the unit.
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 the Interface Address on page 35

protocols

Syntax	<code>protocols [inet iso mpls];</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit logical-unit-number family <i>family</i> tcc]
Release Information	Statement introduced in Junos OS Release 8.3.
Description	For Layer 2.5 VPNs on T Series, MX Series, M120, and M320 routers support, configure IS-IS (ISO traffic) or MPLS traffic to traverse a TCC interface. By default, IPv4 (inet) traffic runs on T Series, MX, Series, M120, and M320 routers and over TCC interfaces. You must configure the same traffic type on both ends of the Layer 2.5 VPN.
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 IS-IS or MPLS Traffic for TCC Interfaces

proxy

Syntax	<code>proxy inet-address <i>address</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family tcc], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family tcc]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For Layer 2.5 VPNs using an Ethernet interface as the TCC router, configure the IP address for which the TCC router is proxying. Ethernet TCC is supported on interfaces that carry IPv4 traffic only. Ethernet TCC encapsulation is supported on 1-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Gigabit Ethernet, and 4-port Fast Ethernet PICs only. Ethernet TCC is not supported on the T640 router.
Options	inet-address —Configure the IP address of the neighbor to the TCC router.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Ethernet TCC</i> • <i>Example: Configuring an Ethernet TCC or Extended VLAN TCC</i> • remote on page 155 • <i>Junos OS VPNs Library for Routing Devices</i>

queue-length

Syntax	queue-length <i>number</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> shaping], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> shaping]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For ATM1 interfaces only, define the maximum queue length in the traffic-shaping profile. For ATM1 PICs, each VC has its own independent shaping parameters.
Default	Buffer usage is unregulated.
Options	<i>number</i> —Maximum number of packets the queue can contain. Range: 1 through 16,383 packets Default: 16,383 packets
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the ATM1 Queue Length</i>

receive-options-packets

Syntax	receive-options-packets;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For a Monitoring Services PIC and an ATM or SONET/SDH PIC installed in an M160, M40e, or T Series router, guarantee conformity with cflowd records structure. This statement is required when you enable passive monitoring.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Enabling Passive Monitoring on ATM Interfaces</i>• <i>Enabling Passive Monitoring on SONET/SDH Interfaces</i>

receive-ttl-exceeded

Syntax	receive-ttl-exceeded;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For Monitoring Services PIC and an ATM or SONET/SDH PIC installed in an M160, M40e, or T Series router, guarantee conformity with cflowd records structure. This statement is required when you enable passive monitoring.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Enabling Passive Monitoring on ATM Interfaces</i> • <i>Enabling Passive Monitoring on SONET/SDH Interfaces</i>

remote

Syntax	remote { (inet-address <i>address</i> mac-address <i>address</i>); }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family tcc], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family tcc]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For Layer 2.5 VPNs using an Ethernet interface as the TCC router, configure the location of the remote router. Ethernet TCC is supported on interfaces that carry IPv4 traffic only. Ethernet TCC encapsulation is supported on 1-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Gigabit Ethernet, and 4-port Fast Ethernet PICs only.
Options	<p>mac-address—Configure the MAC address of the remote site.</p> <p>inet-address—Configure the IP address of the remote site.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Ethernet TCC</i> • <i>Example: Configuring an Ethernet TCC or Extended VLAN TCC</i> • proxy on page 153 • <i>Junos OS VPNs Library for Routing Devices</i>

rpf-check (Dynamic Profiles)

Syntax	<pre>rpf-check { fail-filter <i>filter-name</i>; mode loose; }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced in Junos OS Release 9.6.
Description	<p>Check whether traffic is arriving on an expected path. You can include this statement with the inet protocol family only.</p> <p>The remaining statements are explained separately.</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 Unicast RPF Strict Mode on page 51• <i>Configuring Unicast RPF and Fail Filters in Dynamic Profiles for Subscriber Interfaces</i>

rpf-check (interfaces)

Syntax	<pre>rpf-check { fail-filter <i>filter-name</i>; mode loose; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Check whether traffic is arriving on an expected path. You can include this statement with the inet or inet6 protocol family only.</p> <p>The mode statement is explained separately.</p>
Options	fail-filter —A filter to evaluate when packets are received on the interface. If the RPF check fails, this optional filter is evaluated. If the fail filter is not configured, the default action is to silently discard the packet.
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 Unicast RPF Strict Mode on page 51• Configuring Unicast RPF Loose Mode on page 52• Example: Configuring Unicast Reverse-Path-Forwarding Check on page 60

rtvbr

Syntax	<code>rtvbr peak rate sustained rate burst length;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i> shaping]</code> , <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping]</code> , <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> shaping]</code> , <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> shaping]</code> , <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>For ATM2 IQ PICs only, define the real-time variable bandwidth utilization in the traffic-shaping profile.</p> <p>When you configure the real-time bandwidth utilization, you must specify all three options (burst, peak, and sustained). You can specify the rate in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). You can also specify the rate in cells per second by entering a decimal number followed by the abbreviation c; values expressed in cells per second are converted to bits per second using the formula 1 cps = 384 bps.</p>
Default	If the rtvbr statement is not included, bandwidth utilization is unlimited.
Options	<p>burst length—Burst length, in cells. If you set the length to 1, the peak traffic rate is used. Range: 1 through 4000 cells</p> <p>peak rate—Peak rate, in bits per second or cells per second. Range: For ATM2 IQ OC3 and OC12 interfaces, 33 Kbps through 542,526,792 bps. For ATM2 IQ OC48 interfaces, 33 Kbps through 2,170,107,168 bps. For ATM2 IQ DS3 and E3 interfaces, 33 Kbps through the maximum rate, which depends on the ATM encapsulation and framing you configure..</p> <p>sustained rate—Sustained rate, in bps or cps. Range: For ATM2 IQ OC3 and OC12 interfaces, 33 Kbps through 542,526,792 bps. For ATM2 IQ OC48 interfaces, 33 Kbps through 2,170,107,168 bps. For ATM2 IQ DS3 and E3 interfaces, from 33 Kbps through the maximum rate, which depends on the ATM encapsulation and framing you configure.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring ATM CBR</i>• <i>Configuring ATM2 IQ Real-Time VBR</i>• <i>Applying Scheduler Maps to Logical ATM Interfaces</i>

- [cbr on page 105](#)
- [vbr on page 180](#)

sampling (Interfaces)

Syntax	sampling <i>direction</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the direction of traffic to be sampled.
Options	<p><i>direction</i> can be one of the following:</p> <p>input—Configure at least one expected ingress point.</p> <p>output—Configure at least one expected egress point.</p> <p>input output—On a single interface, configure at least one expected ingress point and one expect egress point.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Services Interfaces Library for Routing Devices</i> • <i>Configuring Flow Monitoring</i>

service (Logical Interfaces)

Syntax	<pre>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>; } }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define one or more service sets and filters, and one postservice filter to be applied to an interface.
Options	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	<ul style="list-style-type: none">• <i>Junos OS Services Interfaces Library for Routing Devices</i>

service-filter (Interfaces)

Syntax	<code>service-filter <i>filter-name</i>;</code>
Hierarchy Level	[edit <code>interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>inet</i> service (<i>input</i> <i>output</i>) service-set <i>service-set-name</i></code>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>inet</i> service (<i>input</i> <i>output</i>) service-set <i>service-set-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define the filter to be applied to traffic before it is accepted for service processing. Configuration of a service filter is optional; if you include the service-set statement without a service-filter definition, Junos OS assumes the match condition is true and selects the service set for processing automatically.
Options	<i>filter-name</i> —Identifies the filter to be applied in service processing.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Applying Filters and Services to Interfaces</i> • <i>Junos OS Services Interfaces Library for Routing Devices</i>

service-set

Syntax	<code>service-set <i>service-set-name</i>;</code>
Hierarchy Level	[edit <code>interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>inet</i> service (<i>input</i> <i>output</i>) service-set <i>service-set-name</i></code>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>inet</i> service (<i>input</i> <i>output</i>) service-set <i>service-set-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define one or more service sets to be applied to an interface. If you define multiple service sets, the Junos OS evaluates the filters in the order in which they appear in the configuration.
Options	<i>service-set-name</i> —Identifies the service set.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Services Interfaces Library for Routing Devices</i>

shaping

Syntax	<pre>shaping { (cbr <i>rate</i> rtvbr <i>peak rate</i> <i>sustained rate</i> <i>burst length</i> vbr <i>peak rate</i> <i>sustained rate</i> <i>burst length</i>); queue-length <i>number</i>; }</pre>
Hierarchy Level	<pre>[edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i>], [edit logical-systems <i>logical-system-name</i> 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> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i>]</pre>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>For ATM encapsulation only, define the traffic-shaping profile.</p> <p>For Circuit Emulation PICs, specify traffic shaping in the ingress and egress directions.</p> <p>For ATM2 IQ interfaces, changing or deleting VP tunnel traffic shaping causes all logical interfaces on a VP to be deleted and then re-added.</p> <p>VP tunnels are not supported on multipoint interfaces.</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">• <i>Defining Virtual Path Tunnels</i>• <i>Defining the ATM Traffic-Shaping Profile</i>• <i>Configuring ATM QoS or Shaping</i>• <i>Applying Scheduler Maps to Logical ATM Interfaces</i>

source-class-usage

Syntax	source-class-usage { <i>direction</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet accounting], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet accounting], [edit routing-instances <i>routing-instance-name</i> vrf-table-label]
Release Information	Statement introduced before Junos OS Release 7.4. Support for the vrf-table-label statement added in Junos OS Release 9.3.
Description	Enable packet counters on an interface that count packets that arrive from specific prefixes on the provider core router and are destined for specific prefixes on the customer edge router.
Options	<i>direction</i> can be one of the following: input —Configure at least one expected ingress point. output —Configure at least one expected egress point. input output —On a single interface, configure at least one expected ingress point and one expect egress point.
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"> • Enabling Source Class and Destination Class Usage on page 24 • accounting on page 99 • destination-class-usage on page 109 • <i>Junos OS Services Interfaces Library for Routing Devices</i> • <i>vrf-table-label</i>

targeted-broadcast

Syntax	targeted-broadcast { forward-and-send-to-re; forward-only; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	<p>Specify the IP packets destined for a Layer 3 broadcast address to be forwarded to both an egress interface and the Routing Engine, or to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.</p> <p>The statements are explained separately.</p>
Default	When this statement is not included, broadcast packets are sent to the Routing Engine only.
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 Targeted Broadcast on page 57• Understanding Targeted Broadcast on page 59

then

Syntax	then { discard; }
Hierarchy Level	[edit firewall hierarchical-policer aggregate], [edit firewall hierarchical-policer premium]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	On M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, discard packets when a specified bandwidth or burst limits for an aggregate level of a hierarchical policer is reached.
Options	discard —Discard packets if condition is met.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Applying Policers on page 11 • <i>Junos OS Class of Service Library for Routing Devices</i>

translate-discard-eligible

Syntax	(translate-discard-eligible no-translate-discard-eligible);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>ccc</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>ccc</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For interfaces with encapsulation type Frame Relay CCC, enable or disable translation of Frame Relay discard eligible (DE) control bits.
Default	DE bit translation is disabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Frame Relay Control Bit Translation</i>

translate-fecn-and-becn

Syntax	(translate-fecn-and-becn no-translate-fecn-and-becn);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>ccc</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>ccc</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For interfaces with encapsulation type Frame Relay CCC, enable or disable translation of Frame Relay forward explicit congestion notification (FECN) control bits and Frame Relay backward explicit congestion notification (BECN) control bits.
Default	FECN and BECN bit translation is disabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Frame Relay Control Bit Translation</i>

unit (Dynamic Profiles Standard Interface)

```

Syntax  unit logical-unit-number {
        auto-configure {
            agent-circuit-identifier {
                dynamic-profile profile-name;
            }
        }
        dial-options {
            ipsec-interface-id name;
            l2tp-interface-id name;
            (shared | dedicated);
        }
        encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-tcc-vc-mux
            | atm-mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap
            | atm-vc-mux | ether-over-atm-llc | ether-vpls-over-atm-llc | ether-vpls-over-fr |
            ether-vpls-over-ppp | ethernet | frame-relay-ccc | frame-relay-ppp | frame-relay-tcc |
            frame-relay-ether-type | frame-relay-ether-type-tcc | multilink-frame-relay-end-to-end
            | multilink-ppp | ppp-over-ether | ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc |
            vlan-vci-ccc | vlan-tcc | vlan-vpls);
        family family {
            access-concentrator name;
            address address;
            duplicate-protection;
            dynamic-profile profile-name;
            filter {
                adf {
                    counter;
                    input-precedence precedence;
                    not-mandatory;
                    output-precedence precedence;
                    rule rule-value;
                }
                input filter-name (
                    precedence precedence;
                )
                output filter-name {
                    precedence precedence;
                }
            }
            max-sessions number;
            max-sessions-vsa-ignore;
            rpf-check {
                fail-filter filter-name;
                mode loose;
            }
            service {
                input {
                    service-set service-set-name {
                        service-filter filter-name;
                    }
                }
                post-service-filter filter-name;
            }
            input-vlan-map {

```

```
    inner-tag-protocol-id tpid;  
    inner-vlan-id number;  
    (push | swap);  
    tag-protocol-id tpid;  
    vlan-id number;  
  }  
  output {  
    service-set service-set-name {  
      service-filter filter-name;  
    }  
  }  
  output-vlan-map {  
    inner-tag-protocol-id tpid;  
    inner-vlan-id number;  
    (pop | swap);  
    tag-protocol-id tpid;  
    vlan-id number;  
  }  
}  
service-name-table table-name  
short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max  
  maximum-seconds>;  
unnumbered-address interface-name <preferred-source-address address>;  
filter {  
  input filter-name;  
  output filter-name;  
}  
keepalives {  
  interval seconds;  
}  
ppp-options {  
  chap;  
  pap;  
}  
vlan-id number;  
vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];  
}
```

Hierarchy Level [edit [dynamic-profiles](#) *profile-name* [interfaces](#) *interface-name*]

Release Information Statement introduced in Junos OS Release 9.2.

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*—The specific unit number of the interface you want to assign to the dynamic profile, or one of the following Junos OS predefined variables:

- **\$junos-underlying-interface-unit**—For static VLANs, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP when it accesses the subscriber network.
- **\$junos-interface-unit**—The unit number variable on a dynamic underlying VLAN interface for which you want to enable the creation of dynamic VLAN subscriber interfaces based on agent circuit identifier information.

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 Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information*
- *Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information*
- *Agent Circuit Identifier-Based Dynamic VLANs Components Overview*

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 <plp 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 <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i>], [edit interfaces interface-set <i>interface-set-name</i> interface <i>interface-name</i>]
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	<p><i>logical-unit-number</i>—Number of the logical unit.</p> <p>Range: 0 through 1,073,741,823 for demux and PPPoE static interfaces only. 0 through 16,385 for all other static interface types.</p> <p>The remaining 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**
- *Configuring Logical Interface Properties*
 - *Example: Configuring E-LINE and E-LAN Services for a PBB Network on MX Series Routers*
 - *Junos OS Services Interfaces Library for Routing Devices*

unnumbered-address (Demux)

- Syntax** `unnumbered-address interface-name <preferred-source-address address>;`
- Hierarchy Level** [edit [interfaces interface-name unit logical-unit-number family inet](#)],
[edit [logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet](#)]
- Release Information** Statement introduced in Junos OS Release 8.2.
preferred-source-address option introduced in Junos OS Release 9.0.
IP demultiplexing interfaces supported in Junos OS Release 9.2.
- Description** For IP demultiplexing interfaces, enable the local address to be derived from the specified interface. Configuring an unnumbered interface enables IP processing on the interface without assigning an explicit IP address to the interface.
- Options** **interface-name**—Name of the interface from which the local address is derived. The specified interface must have a logical unit number and a configured IP address, and must not be an unnumbered interface.
- The **preferred-source-address** statement is explained separately.
- Required Privilege Level** interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.
- Related Documentation**
- [Configuring an Unnumbered Interface on page 43](#)
 - [address on page 96](#)
 - *Junos System Basics Configuration Guide*

unnumbered-address (Dynamic Profiles)

Syntax	<code>unnumbered-address interface-name <preferred-source-address address>;</code>
Hierarchy Level	<p>[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family],</p> <p>[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>\$junos-preferred-source-address variable support added in Junos OS Release 9.6.</p> <p>Support for the \$junos-loopback-interface predefined variable introduced in Junos OS Release 9.6.</p>
Description	<p>For Ethernet interfaces, enable the local address to be derived from the specified interface. Configuring unnumbered Ethernet interfaces enables IP processing on the interface without assigning an explicit IP address to the interface. To configure unnumbered address dynamically, include the \$junos-loopback-interface-address predefined variable.</p> <p>You can configure unnumbered address support on Ethernet interfaces for IPv4 and IPv6 address families.</p>
Options	<p>interface-name—Name of the interface from which the local address is derived. Use the \$junos-loopback-interface dynamic variable to dynamically apply a loopback interface. The loopback interface used is based on the routing instance of the subscriber. The specified interface must have a logical unit number and a configured IP address, and must not be an unnumbered interface.</p> <p>preferred-source-address address—(Optional) Secondary IP address of the donor loopback interface. Use the \$junos-preferred-source-address dynamic variable to dynamically apply a preferred source address to the unnumbered Ethernet interface. When you use the dynamic variable, the address that is selected resides in the same network as the IP address of the subscriber. Configuring the preferred source address enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet interfaces in your network</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 an Unnumbered Interface on page 43 in <i>Junos OS Network Interfaces Library for Routing Devices</i>. • <i>Junos OS Network Interfaces Library for Routing Devices</i>

unnumbered-address (Ethernet)

Syntax	<code>unnumbered-address interface-name <preferred-source-address address>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced in Junos OS Release 8.2. preferred-source-address option introduced in Junos OS Release 9.0.
Description	For Ethernet interfaces, enable the local address to be derived from the specified interface. Configuring an unnumbered Ethernet interface enables IP processing on the interface without assigning an explicit IP address to the interface.
Options	interface-name —Name of the interface from which the local address is derived. The specified interface must have a logical unit number and a configured IP address, and must not be an unnumbered interface. The preferred-source-address statement is explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring an Unnumbered Interface on page 43• address on page 96• <i>Junos System Basics Configuration Guide</i>

unnumbered-address (PPP)

Syntax	<code>unnumbered-address interface-name destination address destination-profile profile-name;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For interfaces with PPP encapsulation, enable the local address to be derived from the specified interface.
Options	<p><i>interface-name</i>—Interface from which the local address is derived. The interface name must include a logical unit number and must have a configured address.</p> <p>The remaining 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"> • Configuring IPCP Options on page 41 • address on page 96 • negotiate-address on page 144 • <i>Junos OS Administration Library for Routing Devices</i>

vbr

Syntax	<code>vbr peak rate sustained rate burst length;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i> shaping],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i></code> <code> multipoint-destination <i>address</i> shaping],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> shaping],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code> shaping]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>For ATM encapsulation only, define the variable bandwidth utilization in the traffic-shaping profile.</p> <p>When you configure the variable bandwidth utilization, you must specify all three options (burst, peak, and sustained). You can specify the rate in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). You can also specify the rate in cells per second by entering a decimal number followed by the abbreviation c; values expressed in cells per second are converted to bits per second by means of the formula 1 cps = 384 bps.</p>
Default	If the vbr statement is not specified, bandwidth utilization is unlimited.
Options	<p>burst length—Burst length, in cells. If you set the length to 1, the peak traffic rate is used. Range: 1 through 4000 cells</p> <p>peak rate—Peak rate, in bits per second or cells per second. Range: For ATM1 interfaces, 33 Kbps through 135.6 Mbps (ATM OC3); 33 Kbps through 276 Mbps (ATM OC12). For ATM2 IQ OC3 and OC12 interfaces, 33 Kbps through 542,526,792 bps. For ATM2 IQ OC48 interfaces, 33 Kbps through 2,170,107,168 bps. For ATM2 IQ DS3 and E3 interfaces, from 33 Kbps through the maximum rate, which depends on the ATM encapsulation and framing you configure.</p> <p>sustained rate—Sustained rate, in bits per second or cells per second. Range: For ATM1 interfaces, 33 Kbps through 135.6 Mbps (ATM OC3); 33 Kbps through 276 Mbps (ATM OC12). For ATM2 IQ OC3 and OC12 interfaces, 33 Kbps through 542,526,792 bps. For ATM2 IQ OC48 interfaces, 33 Kbps through 2,170,107,168 bps. For ATM2 IQ DS3 and E3 interfaces, from 33 Kbps through the maximum rate, which depends on the ATM encapsulation and framing you configure.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring ATM CBR</i>• <i>Applying Scheduler Maps to Logical ATM Interfaces</i>

- [cbr on page 105](#)
- [rtvbr on page 158](#)
- [shaping on page 162](#)

vci

Syntax	<code>vci vpi-identifier.vci-identifier;</code>
Hierarchy Level	<p>[edit interfaces at-<i>fpc/pic/port</i> unit <i>logical-unit-number</i>],</p> <p>[edit interfaces at-<i>fpc/pic/port</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces at-<i>fpc/pic/port</i> unit <i>logical-unit-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces at-<i>fpc/pic/port</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.</p>
Description	<p>For ATM point-to-point logical interfaces only, configure the virtual circuit identifier (VCI) and virtual path identifier (VPI).</p> <p>To configure a VPI for a point-to-multipoint interface, specify the VPI in the multipoint-destination statement.</p> <p>VCIs 0 through 31 are reserved for specific ATM values designated by the ATM Forum.</p>
Options	<p>vci-identifier—ATM virtual circuit identifier. Unless you configure the interface to use promiscuous mode, this value cannot exceed the highest-numbered VC configured for the interface with the maximum-vcs option of the vpi statement.</p> <p>Range: 0 through 4089 or 0 through 65,535 with promiscuous mode, with VCIs 0 through 31 reserved.</p> <p>vpi-identifier—ATM virtual path identifier.</p> <p>Range: 0 through 255</p> <p>Default: 0</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 a Point-to-Point ATM1 or ATM2 IQ Connection • Applying Scheduler Maps to Logical ATM Interfaces • multipoint-destination on page 143 • promiscuous-mode • vpi (ATM CCC Cell-Relay Promiscuous Mode)

vlan-id (Logical Port in Bridge Domain)

Syntax	<code>vlan-id <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	The VLAN ID configured on the logical port. Received packets with no VLAN tags are forwarded within the bridge domain with the matching VLAN ID.
Options	number —The VLAN ID. Range: 1 through 4095
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring a Logical Interface for Access Mode</i>

vlan-id-list (Interface in Bridge Domain)

Syntax	<code>vlan-id-list [<i>number number-number</i>];</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Configure a logical interface to forward packets and learn MAC addresses within each bridge domain configured with a VLAN ID that matches a VLAN ID specified in the list. VLAN IDs can be entered individually using a space to separate each ID, entered as an inclusive list separating the starting VLAN ID and ending VLAN ID with a hyphen, or a combination of both.
Options	<i>number number</i> —Individual VLAN IDs separated by a space. <i>number-number</i> —Starting VLAN ID and ending VLAN ID in an inclusive range. Range: 1 through 4095
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring a Logical Interface for Trunk Mode</i>• <i>Configuring the VLAN ID List for a Trunk Interface</i>

PART 3

Administration

- [Monitoring Commands on page 185](#)
- [Command Summaries on page 323](#)

CHAPTER 5

Monitoring Commands

clear firewall

Syntax	clear firewall (all counter <i>counter-name</i> filter <i>filter-name</i> log (all <i>logical-system-name</i>) logical-system <i>logical-system-name</i>)
Syntax (EX Series Switches)	clear firewall (all counter <i>counter-name</i> filter <i>filter-name</i> log (all <i>logical-system-name</i>) policer counter (all counter-id <i>counter-index</i>))
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. logical-system option introduced in Junos OS Release 9.3. log option introduced before Junos OS Release 11.4.
Description	Clear statistics about configured firewall filters. When you clear the counters of a filter, this impacts not only the counters shown by the CLI, but also the ones tracked by SNMP2.



NOTE: The clear firewall command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for graceful Routing Engine switchover (GRES).

If you clear statistics for firewall filters that are applied to Trio-based DPCs and that also use the **prefix-action** action on matched packets, wait at least 5 seconds before you enter the **show firewall prefix-action-stats** command. A 5-second pause between issuing the **clear firewall** and **show firewall prefix-action-stats** commands avoids a possible timeout of the **show firewall prefix-action-stats** command.

- Options**
- all**—Clear the packet and byte counts for all filters. On EX Series switches, this option also clears the packet counts for all policer counters.
 - counter *counter-name***—Clear the packet and byte counts for a filter counter that has been configured with the counter firewall filter action.
 - filter *filter-name***—Clear the packet and byte counts for the specified firewall filter.
 - log (all | *logical-system-name*)**—Clear log entries for IPv4 firewall filters that have **then log** as an action. Use **log all** to clear all log entries or **log *logical-system-name*** to clear log entries for the specified logical system.
 - logical-system *logical-system-name***—Clear the packet and byte counts for the specified logical system.
 - policer counter (all | counter-id *counter-index*)**—(EX8200 switches only) Clear all policer counters using the **policer counter all** command, or clear a specific policer counter using the **policer counter counter-id *counter-index*** command. The value of *counter-index* can be 0, 1, or 2.

Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show firewall on page 202
List of Sample Output	clear firewall all on page 187 clear firewall (counter counter-name) on page 187 clear firewall (filter filter-name) on page 187 clear firewall (policer counter all) (EX8200 Switch) on page 187 clear firewall (policer counter counter-id counter-index) (EX8200 Switch) on page 187

Sample Output

[clear firewall all](#)

```
user@host> clear firewall all
```

[clear firewall \(counter counter-name\)](#)

```
user@host> clear firewall counter port-filter-counter
```

[clear firewall \(filter filter-name\)](#)

```
user@host> clear firewall filter ingress-port-filter
```

[clear firewall \(policer counter all\) \(EX8200 Switch\)](#)

```
user@switch> clear firewall policer counter all
```

[clear firewall \(policer counter counter-id counter-index\) \(EX8200 Switch\)](#)

```
user@switch> clear firewall policer counter counter-id 0
```

clear vpls mac-address

Syntax	<code>clear vpls mac-address</code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><mac-address></code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	(T Series and M Series routers, except for the M160 router) Clear media access control (MAC) address entries from the virtual private LAN service (VPLS) table.
Options	<p>none—Clear all MAC address entries from the VPLS table for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Clear all MAC address entries for a VPLS instance from the VPLS table.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>mac-address—(Optional) Clear a specific MAC address in a VPLS instance from the VPLS table.</p>
Required Privilege Level	maintenance
List of Sample Output	clear vpls mac-address on page 188
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear vpls mac-address

```
user@host> clear vpls mac-address
```

clear vpls mac-table

Syntax	clear vpls mac-table <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> <mac-address> <vlan-id>
Release Information	Command introduced before Junos OS Release 9.5.
Description	(MX Series routers) Clear media access control (MAC) addresses from the virtual private LAN service (VPLS) MAC table.
Options	<p>none—Clear all MAC addresses from the VPLS table for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Clear all MAC addresses for a VPLS instance from the VPLS table.</p> <p>interface <i>interface-name</i>—(Optional) Clear all MAC addresses for a VPLS interface from the VPLS table.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>mac-address—(Optional) Clear a specific MAC address in a VPLS instance from the VPLS table.</p> <p>vlan-id—(Optional) Clear MAC addresses on a specified VLAN (0 through 4095).</p>
Required Privilege Level	maintenance
List of Sample Output	clear vpls mac-table on page 189
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear vpls mac-table

```
user@host> clear vpls mac-table
```

request l2circuit-switchover

Syntax	<code>request l2circuit-switchover</code> <code><logical-system (all logical-system-name)></code> <code><neighbor <i>address</i>></code> <code><virtual-circuit-id <i>identifier</i>></code>
Release Information	Command introduced in Junos OS Release 9.2.
Description	Manually trigger a switch from the active pseudowire to the redundant pseudowire. This command can be useful when performing network maintenance.
Options	<p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor <i>address</i>—(Optional) Trigger a switch of all of the active pseudowire connections with the specified neighbor to their respective redundant pseudowires.</p> <p>virtual-circuit-id <i>identifier</i>—(Optional) Trigger a switch from the active pseudowire connection of the specified Layer 2 circuit to its redundant pseudowire.</p>
Required Privilege Level	maintenance
List of Sample Output	request l2circuit-switchover virtual-circuit-id on page 190
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request l2circuit-switchover virtual-circuit-id

```
user@host>request l2circuit-switchover virtual-circuit-id 12
```


show as-path

Syntax	show as-path <brief detail> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show as-path <brief detail>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Display the distribution of autonomous system (AS) paths that the local routing device is using (usually through the routing table). Use this command to debug problems for AS paths and to understand how AS paths have been manipulated through a policy (through the as-path-prepend action) or through aggregation.</p> <p>AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.</p>
Options	<p>none—Display basic information about AS paths that the local routing device is using (same as brief).</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show as-path summary on page 197
List of Sample Output	show as-path on page 192 show as-path detail on page 193
Output Fields	Table 3 on page 191 lists the output fields for the show as-path command. Output fields are listed in the approximate order in which they appear.

Table 3: show as-path Output Fields

Field Name	Field Description	Level of Output
Total AS paths	Total number of AS paths.	brief none
Bucket	Bucket number.	All levels

Table 3: show as-path Output Fields (*continued*)

Field Name	Field Description	Level of Output
Count	Number of AS path entries in this bucket.	All levels
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. • Atomic—Route is an aggregate of several route prefixes. • Aggregator—Routing device has summarized a range of prefixes. 	All levels
domain	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.	detail
neighbor as	AS peer address.	detail
length	Length of the AS path.	detail
segments	Length of the AS segment descriptor.	detail
references	Path reference count.	detail

Sample Output

show as-path

```

user@host> show as-path
Total AS paths: 30382
Bucket 0      Count: 36
  I
  14203 2914 174 31752 I
  14203 2914 701 21512 I
  14203 2914 1239 26632 I
  14203 2914 1239 29704 I
  14203 2914 4323 10248 I
  14203 2914 4766 23560 I
  14203 2914 6395 32776 I
  14203 2914 7911 11272 I
  14203 2914 12180 18440 I
  14203 2914 17408 17416 I
  14203 2914 701 702 24586 I
  14203 2914 1239 4657 9226 I
  14203 2914 1239 7132 16394 I
  14203 2914 1299 8308 34826 I
  14203 2914 3320 5603 28682 I
  14203 2914 3491 1680 33802 I
  14203 2914 3549 7908 27658 I
  14203 2914 3549 20804 30730 I
  14203 2914 7018 2687 9226 I
  14203 2914 174 9318 9318 23564 I

```

```

14203 2914 701 3786 3786 23564 I
14203 2914 701 4761 4795 9228 I
14203 2914 1239 7132 5673 18444 I
14203 2914 3491 20485 24588 24588 I
14203 2914 5511 2200 1945 2060 I
14203 2914 7911 14325 14325 14348 I
14203 2914 701 4637 9230 9230 9230 I
14203 2914 6395 14 14 14 14 I
14203 2914 9299 6163 6163 6163 6163 9232 I
14203 2914 3356 3356 3356 3356 3356 11955 21522 I
14203 2914 9837 9837 9219 I Aggregator: 9219 202.27.91.253
14203 2914 174 30209 30222 30222 30222 ?
14203 2914 1299 5377 I (Atomic) Aggregator: 5377 193.219.192.22
14203 2914 4323 36097 I (Atomic) Aggregator: 36097 216.69.252.254
14203 2914 209 2516 17676 23813 I (Atomic) Aggregator: 23813 219.127.233.66
Bucket 1    Count: 28
14203 2914 35847 I
14203 2914 174 19465 I
14203 2914 174 35849 I
14203 2914 2828 32777 I
14203 2914 4323 14345 I
14203 2914 4323 29705 I
14203 2914 6395 32777 I
...

```

show as-path detail

```

user@host> show as-path detail
Total AS paths: 30410
Bucket 0    Count: 36
AS path: I
  domain 0, length 0, segments 0, references 54
AS path: 14203 2914 174 31752 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 701 21512 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 1239 26632 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 1239 29704 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 4323 10248 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 4766 23560 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 6395 32776 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 7911 11272 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 12180 18440 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 17408 17416 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 701 702 24586 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 3
AS path: 14203 2914 1239 4657 9226 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 7
AS path: 14203 2914 1239 7132 16394 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 1299 8308 34826 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 2

```

```
AS path: 14203 2914 3320 5603 28682 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3491 1680 33802 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3549 7908 27658 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3549 20804 30730 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 7018 2687 9226 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 3
AS path: 14203 2914 174 9318 9318 23564 I
  domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 3786 3786 23564 I
  domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 4761 4795 9228 I
  domain 1, neighbor as: 14203, length 6, segments 1, references 14
AS path: 14203 2914 1239 7132 5673 18444 I
  domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 3491 20485 24588 24588 I
  domain 1, neighbor as: 14203, length 6, segments 1, references 4
AS path: 14203 2914 5511 2200 1945 2060 I
  domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 7911 14325 14325 14348 I
  domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 4637 9230 9230 9230 I
  domain 1, neighbor as: 14203, length 7, segments 1, references 3
AS path: 14203 2914 6395 14 14 14 14 I
  domain 1, neighbor as: 14203, length 7, segments 1, references 10
...
```

show as-path domain

Syntax	show as-path domain <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show as-path domain
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display autonomous system (AS) path domain information.
Options	none —(Optional) Display AS path domain information for all routing instances. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show as-path domain on page 196
Output Fields	Table 4 on page 195 lists the output fields for the show as-path domain command. Output fields are listed in the approximate order in which they appear

Table 4: show as-path domain Output Fields

Field Name	Field Description
Domain	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.
Primary	Primary AS number.
References	Path reference count.
Number Paths	Number of known AS paths.
Flags	Information about the AS path: <ul style="list-style-type: none"> • ASLoop—Path contains an AS loop. • Atomic—Path includes the ATOMIC_AGGREGATE path attribute. • Local—Path was created by local aggregation. • Master—Path was created by the master routing instance.
Local AS	AS number of the local routing device.
Loops	How many times this AS number can appear in an AS path.

Sample Output

show as-path domain

```
user@host> show as-path domain
Domain: 1          Primary: 10458
References:        3 Paths:      30383
Flags: Master
Local AS: 10458   Loops: 1
```

show as-path summary

Syntax	show as-path summary <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show as-path summary
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display autonomous system (AS) path summary information. AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.
Options	none —(Optional) Display AS path summary information for all routing instances. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show as-path on page 191
List of Sample Output	show as-path summary on page 198
Output Fields	Table 5 on page 197 lists the output fields for the show as-path summary command. Output fields are listed in the approximate order in which they appear.

Table 5: show as-path summary Output Fields

Field Name	Field Description
AS Paths	Number of AS paths.
Buckets	Number of hash buckets in use.
Max	Maximum number of AS path entries per bucket.
Min	Minimum number of AS path entries per bucket.
Avg	Average number of AS path entries per bucket.
Std deviation	Standard deviation of AS path entries per bucket.

Sample Output

show as-path summary

```
user@host> show as-path summary
AS Paths Buckets Max Min Avg Std deviation
30425    1024    95  12  29  6.481419
```


show dynamic-tunnels database

Syntax	show dynamic-tunnels database <destination> <logical-system (all <i>logical-system-name</i>) > <table <i>routing-table-name</i> >
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display dynamic tunnel database information.
Options	<p>none—Display dynamic tunnel database information for all destinations and routing tables.</p> <p>destination—(Optional) Display database entries for the specified IP address (with optional destination prefix length) only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>table <i>routing-table-name</i>—(Optional) Display database entries for the specified table only.</p>
Required Privilege Level	view
List of Sample Output	show dynamic-tunnels database (Tunnel Is Up) on page 200 show dynamic-tunnels database (No Tunnel PIC) on page 200 show dynamic-tunnels database (Tunnel Is Expiring) on page 200 show dynamic-tunnels database (Destination Specified) on page 200
Output Fields	Table 6 on page 199 lists the output fields for the show dynamic-tunnels database command. Output fields are listed in the approximate order in which they appear.

Table 6: show dynamic-tunnels database Output Fields

Field Name	Field Description
Table	Name of the routing table (for example, inet.0).
Destination-network	Destination IP address and subnet.
Tunnel to	Destination IP address and prefix of the tunnel.
State	State of the tunnel: Up , Up (expires in <i>nn:nn:nn</i>seconds) , or Dn (down).
Reference count	Number of routes across the dynamic tunnel that are currently being resolved.
Next-hop type	Type of tunnel: GRE.
Source address	Source IP address of the tunnel.

Table 6: show dynamic-tunnels database Output Fields (*continued*)

Field Name	Field Description
Next-hop	IP address of the destination interface.
State	State of the destination interface: Up, Dn, or Dn (no tunnel pic).

Sample Output

show dynamic-tunnels database (Tunnel Is Up)

```

user@host> show dynamic-tunnels database
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Up
Reference count: 2
Next-hop type: gre
Source address: 10.255.120.92
Next hop: gr-4/3/0.32769
State: Up

```

show dynamic-tunnels database (No Tunnel PIC)

```

user@host> show dynamic-tunnels database
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Dn
Reference count: 2
Next-hop type: gre
Source address: 10.255.120.92
State: Dn (no tunnel pic)

```

show dynamic-tunnels database (Tunnel Is Expiring)

```

user@host> show dynamic-tunnels database
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Up (expires in 00:14:56 seconds)
Reference count: 0
Next-hop type: gre
Source address: 10.255.120.92
Next hop: gr-4/3/0.32769
State: Up

```

show dynamic-tunnels database (Destination Specified)

```

user@host> show dynamic-tunnels database 10.255.120.94
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Up
Reference count: 2
Next-hop type: gre
Source address: 10.255.120.92

```

Next hop: gr-4/3/0.32769
State: Up

show firewall

Syntax	<pre>show firewall <counter <i>counter-name</i>> <filter <i>filter-name</i>> <log> <logical-system (all <i>logical-system-name</i>)> <terse></pre>
Syntax (EX Series Switches)	<pre>show firewall <counter <i>counter-name</i>> <detail> <filter <i>filter-name</i>> <log <(detail interface <i>interface-name</i>)>> <policer counters <(detail counter-id <i>counter-index</i> <detail>)>> <terse></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> <p>logical-system option introduced in Junos OS Release 9.3.</p> <p>terse option introduced in Junos OS Release 9.4.</p> <p>policer counters option introduced in Junos OS Release 12.2 for EX Series switches.</p> <p>detail option introduced in Junos OS Release 12.3.</p>
Description	Display statistics about configured firewall filters.
Options	<p>none—(Optional) Display statistics about all configured firewall filters and counters. For EX Series switches, this command also displays statistics about all configured policers.</p> <p>counter <i>counter-name</i>—(Optional) Name of a filter counter.</p> <p>detail—(EX Series switches only) (Optional) Display firewall filter statistics with enhanced policer.</p> <p>filter <i>filter-name</i>—(Optional) Name of a configured filter.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>log—(Optional) Display log entries for firewall filters.</p> <p>log <(detail interface <i>interface-name</i>)>—(EX Series switches only) (Optional) Display detailed log entries of firewall activity or log information about a specific interface.</p> <p>policer counters <(detail counter-id <i>counter-index</i> <detail>)>—(EX8200 switches only) (Optional) Display policer counter statistics in brief or in detail.</p> <p>terse—(Optional) Display firewall filter names only.</p>
Required Privilege Level	view

- Related Documentation**
- [clear firewall on page 186](#)
 - [show firewall log on page 209](#)
 - *Verifying That Firewall Filters Are Operational*
 - *Verifying That Policers Are Operational*

- List of Sample Output**
- [show firewall filter \(MX Series Router and EX Series Switch\) on page 205](#)
 - [show firewall filter \(non MX Series Router and EX Series Switch\) on page 205](#)
 - [show firewall filter \(Hierarchical Policier, MX Series with MPC\) on page 205](#)
 - [show firewall filter \(Dynamic Input Filter\) on page 205](#)
 - [show firewall \(Logical Systems\) on page 205](#)
 - [show firewall \(counter counter-name\) on page 206](#)
 - [show firewall log on page 206](#)
 - [show firewall policer counters \(EX8200 Switch\) on page 206](#)
 - [show firewall policer counters \(detail\) \(EX8200 Switch\) on page 207](#)
 - [show firewall policer counters \(counter-id counter-index\) \(EX8200 Switch\) on page 207](#)
 - [show firewall policer counters \(counter-id counter-index detail\) \(EX8200 Switch\) on page 207](#)
 - [show firewall detail on page 208](#)

- Output Fields** [Table 7 on page 203](#) lists the output fields for the **show firewall** command. Output fields are listed in the approximate order in which they appear.

Table 7: show firewall Output Fields

Field Name	Field Description
Filter	<p>Name of a filter that has been configured with the filter statement at the [edit firewall] hierarchy level.</p> <p>Except on EX Series switches:</p> <ul style="list-style-type: none"> • When an interface-specific filter is displayed, the name of the filter is followed by the full interface name and by either -i for an input filter or -o for an output filter. • When dynamic filters are displayed, the name of the filter is followed by the full interface name and by either -in for an input filter or -out for an output filter. When a logical system-specific filter is displayed, the name of the filter is prefixed with two underscore (__) characters and the name of the logical system (for example, __ls1/filter1). • When a service filter is displayed that uses a service set, the separator between the service-set name and the service-filter name is a semicolon (:).
Counters	<p>Display filter counter information:</p> <ul style="list-style-type: none"> • Name—Name of a filter counter that has been configured with the counter firewall filter action. • Bytes—Number of bytes that match the filter term under which the counter action is specified. • Packets—Number of packets that matched the filter term under which the counter action is specified. <p>NOTE: On M and T series routers, firewall filters cannot count ip-options packets on a per option type and per interface basis. A limited work around is to use the show pfe statistics ip options command to see ip-options statistics on a per Packet Forwarding Engine (PFE) basis. See <i>show pfe statistics ip</i> for sample output.</p>

Table 7: show firewall Output Fields (*continued*)

Field Name	Field Description
Policers	<p>Display policer information:</p> <ul style="list-style-type: none"> • Name—Name of policer. • Bytes—(For two-color policers on MX Series routers and EX Series switches, and for hierarchical policers on MS-DPC, MIC, and MPC interfaces on MX Series routers) Number of bytes that match the filter term under which the policer action is specified. This is only the number out-of-specification (out-of-spec) byte counts, not all the bytes in all packets policed by the policer. For other combinations of policer type, device, and line card type, this field is blank. • Packets—Number of packets that matched the filter term under which the policer action is specified. This is only the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer.
Policer Counter Index	(EX8200 switch only) Global management counter ID. The counter ID value (<i>counter-index</i>) can be 0, 1, or 2.
Green	(EX8200 switch only) Number of packets within the limits. The number of packets is smaller than the committed information rate (CIR).
Yellow	(EX8200 switch only) Number of packets partially within the limits. The number of packets is greater than the CIR, but the burst size is within the excess burst size (EBS) limit.
Discard	(EX8200 switch only) Number of discarded packets.
Bytes	(EX8200 switch only) Number of green, yellow, red, or discarded packets in bytes.
Packets	(EX8200 switch only) Number of green, yellow, red, or discarded packets.
Filter name	(EX8200 switch only) Name of the filter with a term associated to a policer.
Term name	(EX8200 switch only) Name of the term associated with a policer.
Policer name	(EX8200 switch only) Name of the policer that is associated with a global management counter.

Sample Output

show firewall filter (MX Series Router and EX Series Switch)

```
user@host> show firewall filter test
Filter: test
Counters:
Name                               Bytes          Packets
Counter-1                          0              0
Counter-2                          0              0
Policers:
Name                               Bytes          Packets
Policer-1                         2770           70
```

show firewall filter (non MX Series Router and EX Series Switch)

```
user@host> show firewall filter test
Filter: test
Counters:
Name                               Bytes          Packets
Counter-1                          0              0
Counter-2                          0              0
Policers:
Name                               Bytes          Packets
Policer-1                         70
```

show firewall filter (Hierarchical Policier, MX Series with MPC)

```
user@host> show firewall filter
FL_V4_PHY-HP-EF-AWARE-Gold=400k-MCAST=200k-Total=1M-ds-10/0/0:2:1-i

Filter: FL_V4_PHY-HP-EF-AWARE-Gold=400k-MCAST=200k-Total=1M-ds-10/0/0:2:1-i
Counters:
Name                               Bytes          Packets
AF1x_counter-ds-10/0/0:2:1-i      0              0
AF2x_counter-ds-10/0/0:2:1-i      25529445976    24500428
AF3x_counter-ds-10/0/0:2:1-i      2182022        39482
AF4x_counter-ds-10/0/0:2:1-i      0              0
BE_counter-ds-10/0/0:2:1-i        0              0
EF_counter-ds-10/0/0:2:1-i        14817044120    12265765
STD_counter-ds-10/0/0:2:1-i       0              0
Policers:
Name                               Bytes          Packets
POL_CE-PE_M=200k-filter-ds-10/0/0:2:1-i  5948099658    5708349
POL_CE-PE_G=400K_R=1M-filter-ds-10/0/0:2:1-i  ??????????    3572794
?????????????                    ??????????    ????????
```

show firewall filter (Dynamic Input Filter)

```
user@host> show firewall filter dfwd-ge-5/0/0.1-in
Filter: dfwd-ge-5/0/0.1-in
Counters:
Name                               Bytes          Packets
c1-ge-5/0/0.1-in                  0              0
```

show firewall (Logical Systems)

```
user@host> show firewall
```

```

Filter: __lr1/test
Counters:
Name                               Bytes      Packets
icmp                               420        5
Filter: __default_bpdu_filter__
Filter: __lr1/inet_filter1
Counters:
Name                               Bytes      Packets
inet_tcp_count                     0          0
inet_udp_count                     0          0
Filter: __lr1/inet_filter2
Counters:
Name                               Bytes      Packets
inet_icmp_count                    0          0
inet_pim_count                     0          0
Filter: __lr2/inet_filter1
Counters:
Name                               Bytes      Packets
inet_tcp_count                     0          0
inet_udp_count                     0          0

```

show firewall (counter counter-name)

```

user@host> show firewall counter icmp-counter
Filter: ingress-port-voip-class-filter
Counters:
Name                               Bytes      Packets
icmp-counter                       0          0

```

show firewall log

```

user@host> show firewall log
Log :

Time      Filter  Action Interface  Protocol  Src Addr
      Dest Addr
08:00:53 pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:52 pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:51 pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:50 pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:49 pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:48 pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:47 pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4

```

show firewall policer counters (EX8200 Switch)

```

user@switch> show firewall policer counters
Policer Counter Index 0:
Green:           73      15914
Yellow:          9      1962
Discard:        119     25942

```



```

Policer Counter Index 1:
      Bytes      Packets
Green:         0         0
Yellow:        0         0
Discard:       0         0

Policer Counter Index 2:
      Bytes      Packets
Green:         0         0
Yellow:        0         0
Discard:       0         0

```

show firewall policer counters (detail) (EX8200 Switch)

```

user@switch> show firewall policer counters detail
Policer Counter Index 0:
      Bytes      Packets
Green:         73      15914
Yellow:         9      1962
Discard:       119     25942

Filter name      Term name      Policer name
myfilter         polcr-term-1   myfilter-polcr-1
inet-filter-ae   ae-snmp        policer-1
inet-filter-ae   ae-ssh         policer-2

Policer Counter Index 1:
      Bytes      Packets
Green:         0         0
Yellow:        0         0
Discard:       0         0

Filter name      Term name      Policer name

Policer Counter Index 2:
      Bytes      Packets
Green:         0         0
Yellow:        0         0
Discard:       0         0

Filter name      Term name      Policer name

```

show firewall policer counters (counter-id counter-index) (EX8200 Switch)

```

user@switch> show firewall policer counters counter-id 0
Policer Counter Index 0:
      Bytes      Packets
Green:         73      15914
Yellow:         9      1962
Discard:       119     25942

```

show firewall policer counters (counter-id counter-index detail) (EX8200 Switch)

```

user@switch> show firewall policer counters counter-id 0 detail
Policer Counter Index 0:
      Bytes      Packets
Green:         73      15914
Yellow:         9      1962
Discard:       119     25942

Filter name      Term name      Policer name

```

myfilter	polcr-term-1	myfilter-polcr-1
inet-filter-ae	ae-snmp	policer-1
inet-filter-ae	ae-ssh	policer-2

show firewall detail

```
user@host> show firewall detail
Filter: __default_bpdu_filter__
```

Filter: foo

Counters:

Name

c1

Bytes	Packets
17652140	160474

Policers:

Name

P1-t1

Bytes	Packets
-------	---------

OOS

0	18286
---	-------

Offered

0	18446744073709376546
---	----------------------

Transmitted

0	18446744073709358260
---	----------------------

show firewall log

Syntax	show firewall log <detail> <interface <i>interface-name</i> > <logical-system (<i>logical-system-name</i> all)>
Syntax (EX Series Switches)	show firewall log <detail> <interface <i>interface-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. logical-system option introduced in Junos OS Release 9.3.
Description	Display log information about firewall filters.
Options	none —Display log information about firewall filters. detail —(Optional) Display detailed information. interface <i>interface-name</i> —(Optional) Display log information about a specific interface. logical-system (<i>logical-system-name</i> all) —(Optional) Perform this operation on all logical systems or on a particular system.
Required Privilege Level	view
List of Sample Output	show firewall log on page 210 show firewall log detail on page 210
Output Fields	Table 8 on page 209 lists the output fields for the show firewall log command. Output fields are listed in the approximate order in which they appear.

Table 8: show firewall log Output Fields

Field Name	Field Description
Time of Log	Time that the event occurred.
Filter	<ul style="list-style-type: none"> Displays the name of a configured firewall filter or service filter only if the packet hit the filter's log action in a kernel filter (in the control plane). For any traffic that reaches the Routing Engine, the packets hit the log action in the kernel. For all other logged packets (packet hit the filter's log action in the Packet Forwarding Engine), this field displays pfe instead of a configured filter name.

Table 8: show firewall log Output Fields (*continued*)

Field Name	Field Description
Filter Action	Filter action: <ul style="list-style-type: none"> • A—Accept • D—Discard • R—Reject
Name of Interface	<ul style="list-style-type: none"> • Displays a physical interface name if the packet arrived at a port on a line card. • Displays local if the packet was generated by the device's internal Ethernet interface, em1 or fxp1, which connects the Routing Engine with the router's packet-forwarding components.
Name of protocol	Packet's protocol name: egp , gre , icmp , ipip , ospf , pim , rsvp , tcp , or udp .
Packet length	Length of the packet.
Source address	Packet's source address.
Destination address	Packet's destination address and port.

Sample Output

show firewall log

```

user@host>show firewall log
Time      Filter  Action Interface    Protocol  Src Addr    Dest Addr
13:10:12  pfe      D      rlsq0.902     ICMP      180.1.177.2 180.1.177.1
13:10:11  pfe      D      rlsq0.902     ICMP      180.1.177.2 180.1.177.1

```

show firewall log detail

```

user@host> show firewall log detail
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0Name of protocol: TCP, Packet Length: 50824, Source address:
172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 1020, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0

```

```
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,  
Destination address: 192.168.70.66:513  
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of  
interface: fxp0.0  
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,  
Destination address: 192.168.70.66:513  
....
```

show firewall prefix-action-stats

Syntax	show firewall prefix-action-stats filter <i>filter-name</i> prefix-action <i>prefix-action-name</i> <from <i>number</i> to <i>number</i> > <logical-system (<i>logical-system-name</i> all)>
Release Information	Command introduced before Junos OS Release 7.4. logical-system option introduced in Junos OS Release 9.3.
Description	Display prefix action statistics about configured firewall filters. If you clear statistics for firewall filters that are applied to Trio-based DPCs and that also use the prefix-action action on matched packets, wait at least 5 seconds before you enter the show firewall prefix-action-stats command. A 5-second pause between issuing the clear firewall and show firewall prefix-action-stats commands avoids a possible timeout of the show firewall prefix-action-stats command.
Options	filter <i>filter-name</i> —Name of a filter. prefix-action <i>prefix-action-name</i> —Name of a prefix action. from <i>number</i> to <i>number</i> —(Optional) Starting and ending counter or policer. logical-system (<i>logical-system-name</i> all) —(Optional) Perform this operation on all logical systems or on a particular system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear firewall on page 186
List of Sample Output	show firewall prefix-action-stats on page 212
Output Fields	Table 9 on page 212 lists the output fields for the show firewall prefix-action-stats command. Output fields are listed in the approximate order in which they appear.

Table 9: show firewall prefix-action-stats Output Fields

Field Name	Field Description
Filter	Filter name. Filters configured for logical systems include the name of the filter prefixed with the two underscore characters (__) and the name of the logical system (for example, __ls1/filter1).

Sample Output

show firewall prefix-action-stats

```
user@host> show firewall prefix-action-stats filter test prefix-action act1
```

Filter: __ls2/test

show hfrf profiles

Syntax	show hfrf profiles <brief extensive>
Release Information	Command introduced in Junos OS Release 12.2.
Description	<p>Display Host fast reroute (HFRR) profile information.</p> <p>HFRR adds a precomputed protection path into the Packet Forwarding Engine (PFE), such that if a link between a provider edge device and a server farm becomes unusable for forwarding, the PFE can use another path without having to wait for the router or the protocols to provide updated forwarding information.</p>
Options	<p>none—Display information about HFRR profiles.</p> <p>brief extensive—(Optional) Display the specified level of output.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Host Fast Reroute</i>
List of Sample Output	show hfrf profiles on page 215
Output Fields	<p>Table 10 on page 214 describes the output fields for the show hfrf profiles command. Output fields are listed in the approximate order in which they appear.</p>

Table 10: show hfrf profiles Output Fields

Field Name	Field Description
HFRR pointer	
HFRR current state	Status of the HFRR profile: HFRR_ACTIVE , HFRR_INACTIVE , HFRR_IFLH-NOT-CONF , and so on.
HFRR Prefix limit blackout timer expiry (in secs)	Time interval between a HFRR profile becoming inactive on exceeding the ARP prefix limit, and the profile starting the SYNC process.
HFRR prefix limit hit count	Number of times that a HFRR profile becomes inactive on exceeding the ARP prefix limit.
HFRR protected IFL name	Interface configured for the HFRR feature.
HFRR protected IFL handle	
HFRR routing instance name	The routing instance in which the HFRR interface is configured.
HFRR routing instance handle	

Table 10: show hfrf profiles Output Fields (*continued*)

Field Name	Field Description
HFRR sync BG scheduled	
HFRR RTS filter on	
HFRR delete BG scheduled	
HFRR ARP prefix limit	Configured ARP prefix limit.
HFRR ARP supplementary blackout timeout (in mins)	Supplementary time-out value configured for profile to be inactive when it hits ARP prefix limit
HFRR number of ARP routes learned	Number of ARP routes learned on the configured interface.
HFRR number of FRR routes created	Number of ARP routes learned on the configured interface.

Sample Output

show hfrf profiles

```

user@host> show hfrf profiles
HFRR pointer: 0x9254000
HFRR current state: HFRR_ACTIVE
HFRR Prefix limit blackout timer expiry (in secs): 0
HFRR prefix limit hit count: 0
HFRR protected IFL name: ge-4/1/0.0
HFRR protected IFL handle: 0x9248738
HFRR routing instance name: test
HFRR routing instance handle: 0x9145740
HFRR sync BG scheduled: NO
HFRR RTS filter on: YES
HFRR delete BG scheduled: NO
HFRR ARP prefix limit: 0
HFRR ARP supplementary blackout timeout (in mins): 1
HFRR number of ARP routes learned: 4
HFRR number of FRR routes created: 2

```

show interfaces (10-Gigabit Ethernet)

Syntax	<code>show interfaces <i>xe-fpc/pic/port</i></code> <code><brief detail extensive terse></code> <code><descriptions></code> <code><media></code> <code><snmp-index <i>snmp-index</i>></code> <code><statistics></code>
Release Information	Command introduced in Junos OS Release 8.0.
Description	(M320, M120, MX Series, and T Series routers and EX Series switches only) Display status information about the specified 10-Gigabit Ethernet interface.
Options	<p><code><i>xe-fpc/pic/port</i></code>—Display standard information about the specified 10-Gigabit Ethernet interface.</p> <p><code>brief detail extensive terse</code>—(Optional) Display the specified level of output.</p> <p><code>descriptions</code>—(Optional) Display interface description strings.</p> <p><code>media</code>—(Optional) Display media-specific information about network interfaces.</p> <p><code>snmp-index <i>snmp-index</i></code>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><code>statistics</code>—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
List of Sample Output	<p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2) on page 231</p> <p>show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode) on page 234</p> <p>show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC) on page 236</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode) on page 238</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only) on page 238</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only) on page 239</p>
Output Fields	See Table 11 on page 217 for the output fields for the show interfaces (10-Gigabit Ethernet) command.

Table 11: show interfaces Gigabit Ethernet Output Fields

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

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

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

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

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Output bytes—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 11 on page 217.</p>	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the <code>ignore-l3-incompletes</code> statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

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

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

Table 11: 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 11: 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 12 on page 231 • 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 11: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the routing device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local routing device (which the routing device is rejecting). • Input SA rejects—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • Output packet count—Number of packets that the filter has given to the MAC hardware. • Output packet pad count—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • Output packet error count—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • CAM destination filters, CAM source filters—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0. 	extensive
PMA PHY	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • PHY Lock—Phase-locked loop • PHY Light—Loss of optical signal 	extensive

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

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner status—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the attached Ethernet device, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is None. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Local resolution—Information from the link partner: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive). 	extensive
Received path trace, Transmitted path trace	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the routing device at the other end of the fiber. The transmitted path trace value is the message that this routing device transmits.</p>	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive

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

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

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

Field Name	Field Description	Level of Output
VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. • pop-swap—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none
Demux:	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> • Source Family Inet • Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the routing device.	extensive

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

Field Name	Field Description	Level of Output
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the Output bytes and Output packets interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	detail extensive none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	detail extensive
Output Filters	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	detail extensive
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about address flag (possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i>).	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interlace.	detail extensive none

Table 11: 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 12 on page 231](#) 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 12 on page 231](#), 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 12: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound physical interface	show interfaces ge-0/3/0 extensive	Traffic statistics: Input bytes: 496 bytes per packet, representing the Layer 2 packet MAC statistics: Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes	The additional 4 bytes are for the CRC.
Inbound logical interface	show interfaces ge-0/3/0.50 extensive	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	
Outbound physical interface	show interfaces ge-0/0/0 extensive	Traffic statistics: Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes MAC statistics: Received octets: 478 bytes per packet, representing the Layer 3 packet	For input bytes, the additional 12 bytes includes 6 bytes for the destination MAC address + 4 bytes for VLAN + 2 bytes for the Ethernet type.
Outbound logical interface	show interfaces ge-0/0/0.50 extensive	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	

Sample Output

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2)

```

user@host> show interfaces xe-5/0/0 extensive
Physical interface: xe-5/0/0, Enabled, Physical link is Up
  Interface index: 177, SNMP ifIndex: 99, Generation: 178
  Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Speed: 10Gbps, Loopback:

```

```

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

0 best-effort          81049992          81049992              0

1 expedited-fo              0              0              0

2 assured-forw           0              0              0

3 network-cont           0              0              0

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

0 best-effort          0              0              0

1 expedited-fo          0              0              0

2 assured-forw          0              0              0

3 network-cont          0              0              0

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

```

```

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

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

Packet Forwarding Engine configuration:
Destination slot: 5

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

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

Logical interface xe-5/0/0.0 (Index 71) (SNMP ifIndex 95) (Generation 195)
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ] Encapsulation: ENET2
Egress accounting overhead: 100
Ingress accounting overhead: 90

Traffic statistics:
Input bytes : 0
Output bytes : 46
Input packets: 0
Output packets: 1

IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Local statistics:
Input bytes : 0
Output bytes : 46
Input packets: 0
Output packets: 1

Transit statistics:
Input bytes : 0
Output bytes : 0

```

```

Input packets:                0                0 pps
Output packets:               0                0 pps
IPv6 transit statistics:
  Input bytes :                0
  Output bytes :               0
  Input packets:              0
  Output packets:             0
Protocol inet, MTU: 1500, Generation: 253, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 192.1.1/24, Local: 192.1.1.1, Broadcast: 192.1.1.255,
Generation: 265
Protocol multiservice, MTU: Unlimited, Generation: 254, Route table: 0
  Flags: None
  Policer: Input: __default_arp_policer__

```

show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode)

```

user@host> show interfaces xe-1/0/0 extensive
Physical interface: xe-1/0/0, Enabled, Physical link is Up
Interface index: 141, SNMP ifIndex: 34, Generation: 47
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Loopback: Disabled
WAN-PHY mode
Source filtering: Disabled, Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps 16384
Link flags : None
CoS queues : 4 supported
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:05:85:a2:10:9d, Hardware address: 00:05:85:a2:10:9d
Last flapped : 2005-07-07 11:22:34 PDT (3d 12:28 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes :                0                0 bps
  Output bytes :               0                0 bps
  Input packets:              0                0 pps
  Output packets:             0                0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  HS Link CRC errors: 0, HS Link FIFO overflows: 0,
  Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0,
  Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
  Resource errors: 0
Queue counters:
  Queued packets  Transmitted packets  Dropped packets
0 best-effort    0                0                0
1 expedited-fo   0                0                0
2 assured-forw   0                0                0
3 network-cont   0                0                0
Active alarms : LOL, LOS, LBL
Active defects: LOL, LOS, LBL, SEF, AIS-L, AIS-P
PCS statistics
  Seconds  Count
Bit errors 0        0
Errored blocks 0      0
MAC statistics:
  Receive  Transmit
Total octets 0        0
Total packets 0        0
Unicast packets 0      0
Broadcast packets 0    0
Multicast packets 0    0

```

```

CRC/Align errors                0          0
FIFO errors                     0          0
MAC control frames              0          0
MAC pause frames               0          0
Oversized frames               0
Jabber frames                  0
Fragment frames                0
VLAN tagged frames             0
Code violations                 0
Filter statistics:
  Input packet count            0
  Input packet rejects          0
  Input DA rejects              0
  Input SA rejects              0
  Output packet count           0
  Output packet pad count       0
  Output packet error count     0
CAM destination filters: 0, CAM source filters: 0
PMA PHY:
  Seconds      Count  State
  PLL lock     0      0 OK
  PHY light    63159  1 Light Missing
WIS section:
  BIP-B1        0      0
  SEF           434430  434438 Defect Active
  LOS           434430  1 Defect Active
  LOF           434430  1 Defect Active
  ES-S          434430
  SES-S         434430
  SEFS-S        434430
WIS line:
  BIP-B2        0      0
  REI-L         0      0
  RDI-L         0      0 OK
  AIS-L         434430  1 Defect Active
  BERR-SF       0      0 OK
  BERR-SD       0      0 OK
  ES-L          434430
  SES-L         434430
  UAS-L         434420
  ES-LFE        0
  SES-LFE       0
  UAS-LFE       0
WIS path:
  BIP-B3        0      0
  REI-P         0      0
  LOP-P         0      0 OK
  AIS-P         434430  1 Defect Active
  RDI-P         0      0 OK
  UNEQ-P        0      0 OK
  PLM-P         0      0 OK
  ES-P          434430
  SES-P         434430
  UAS-P         434420
  ES-PFE        0
  SES-PFE       0
  UAS-PFE       0
Received path trace:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted path trace: orissa so-1/0/0
6f 72 69 73 73 61 20 73 6f 2d 31 2f 30 2f 30 00 orissa so-1/0/0.
Packet Forwarding Engine configuration:

```

```

Destination slot: 1
CoS information:
  CoS transmit queue      Bandwidth      Buffer      Priority  Limit
                           %      bps      %      bytes
0 best-effort             95      950000000  95        0      low      none
3 network-control         5       50000000   5         0      low      none

```

show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC)

```

user@host> show interfaces ge-7/0/0 extensive
Physical interface: ge-7/0/0, Enabled, Physical link is Down
Interface index: 143, SNMP ifIndex: 508, Generation: 208
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled
Device flags   : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
Link flags     : None
Wavelength     : 1550.12 nm, Frequency: 193.40 THz
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:05:85:70:2b:72, Hardware address: 00:05:85:70:2b:72
Last flapped   : 2011-04-20 15:48:54 PDT (18:39:49 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   : 0          0 bps
Output bytes  : 0          0 bps
Input packets: 0          0 pps
Output packets: 0         0 pps
IPv6 transit statistics:
Input bytes   : 0
Output bytes  : 0
Input packets: 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 2, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

0 best-effort        0          0          0

1 expedited-fo       0          0          0

2 assured-forw       0          0          0

3 network-cont
Queue number:        Mapped forwarding classes
0                    best-effort
1                    expedited-forwarding
2                    assured-forwarding
3                    network-control
Active alarms  : LINK
Active defects : LINK
MAC statistics:
Total octets      Receive      Transmit
Total packets     0           0

```

```

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

```

user@host> show interfaces xe-7/0/0-tx extensive
Physical interface: xe-7/0/0-tx, Enabled, Physical link is Up
  Interface index: 176, SNMP ifIndex: 137, Generation: 177
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
  Unidirectional: Tx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:73:e4:83, Hardware address: 00:05:85:73:e4:83
  Last flapped   : 2007-06-01 09:08:19 PDT (3d 02:31 ago)
  Statistics last cleared: Never
Traffic statistics:
  Input bytes :                0                0 bps

```



```

Output bytes :      322891152287160      9627472888 bps
Input packets:              0              0 pps
Output packets:    328809727380      1225492 pps

...

Filter statistics:
  Output packet count      328810554250
  Output packet pad count      0
  Output packet error count    0
...

Logical interface xe-7/0/0-tx.0 (Index 73) (SNMP ifIndex 138) (Generation 139)

Flags: SNMP-Traps Encapsulation: ENET2
Egress accounting overhead: 100
Ingress accounting overhead: 90
Traffic statistics:
  Input bytes :      0
  Output bytes :    322891152287160
  Input packets:      0
  Output packets:    328809727380
IPv6 transit statistics:
  Input bytes :      0
  Output bytes :      0
  Input packets:      0
  Output packets:      0
Local statistics:
  Input bytes :      0
  Output bytes :      0
  Input packets:      0
  Output packets:      0
Transit statistics:
  Input bytes :      0      0 bps
  Output bytes :    322891152287160      9627472888 bps
  Input packets:      0      0 pps
  Output packets:    328809727380      1225492 pps
IPv6 transit statistics:
  Input bytes :      0
  Output bytes :      0
  Input packets:      0
  Output packets:      0
Protocol inet, MTU: 1500, Generation: 147, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.11.12/24, Local: 10.11.12.13, Broadcast: 10.11.12.255,
Generation: 141
  Protocol multiservice, MTU: Unlimited, Generation: 148, Route table: 0
  Flags: None
  Policer: Input: __default_arp_policer__

```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only)

```

user@host> show interfaces xe-7/0/0-rx extensive
Physical interface: xe-7/0/0-rx, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 118, Generation: 175
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Rx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues

```

```
Hold-times      : Up 0 ms, Down 0 ms
Current address: 00:05:85:73:e4:83, Hardware address: 00:05:85:73:e4:83
Last flapped   : 2007-06-01 09:08:22 PDT (3d 02:31 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes :      322857456303482      9627496104 bps
  Output bytes :              0          0 bps
  Input packets:      328775413751      1225495 pps
  Output packets:              0          0 pps

...

Filter statistics:
  Input packet count      328775015056
  Input packet rejects    1
  Input DA rejects        0

...

Logical interface xe-7/0/0-rx.0 (Index 72) (SNMP ifIndex 120) (Generation 138)

Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
  Input bytes :      322857456303482
  Output bytes :              0
  Input packets:      328775413751
  Output packets:              0
IPv6 transit statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:              0
Local statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:              0
Transit statistics:
  Input bytes :      322857456303482      9627496104 bps
  Output bytes :              0          0 bps
  Input packets:      328775413751      1225495 pps
  Output packets:              0          0 pps
IPv6 transit statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:              0
Protocol inet, MTU: 1500, Generation: 145, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 192.1.1/24, Local: 192.1.1.1, Broadcast: 192.1.1.255,
Generation: 139
  Protocol multiservice, MTU: Unlimited, Generation: 146, Route table: 0
  Flags: None
  Policer: Input: __default_arp_policer__
```

show interfaces (Fast Ethernet)

Syntax	show interfaces <i>interface-type</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i> > <statistics>
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	show interfaces (Fast Ethernet) on page 254 show interfaces brief (Fast Ethernet) on page 255 show interfaces detail (Fast Ethernet) on page 255 show interfaces extensive (Fast Ethernet) on page 255
Output Fields	<p>Table 13 on page 241 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 13: 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 <i>Common Output Fields Description</i> .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none

Table 13: 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 <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> .	All levels
Link flags	Information about the link. Possible values are described in the "Links Flags" section under <i>Common Output Fields Description</i> .	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels

Table 13: 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 13: 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 13: 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 13: 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 routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
OTN 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 13: 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 13: 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 routing device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local routing device (which the routing device is rejecting). • Input SA rejects—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • Output packet count—Number of packets that the filter has given to the MAC hardware. • Output packet pad count—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • Output packet error count—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • CAM destination filters, CAM source filters—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0. 	extensive
PMA PHY	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • PHY Lock—Phase-locked loop • PHY Light—Loss of optical signal 	extensive

Table 13: 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 13: 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 13: 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 routing device manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the routing device at the other end of the fiber. The transmitted path trace value is the message that this routing device transmits.</p>	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive

Table 13: 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 <i>Common Output Fields Description</i> .	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 13: 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 <i>Common Output Fields Description</i> .	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	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 routing device.	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 <i>Common Output Fields Description</i> .	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none

Table 13: 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 <i>Common Output Fields Description</i> .	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about address flag (possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	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 (Gigabit Ethernet)

Syntax	<pre>show interfaces <i>ge-fpc/pic/port</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	(M Series, T Series, and MX Series routers and EX Series switches only) Display status information about the specified Gigabit Ethernet interface.
Options	<p><i>ge-fpc/pic/port</i>—Display standard information about the specified 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>
Additional Information	In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> <i>Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration</i>
List of Sample Output	<p>show interfaces (Gigabit Ethernet) on page 272</p> <p>show interfaces (Gigabit Ethernet on MX Series Routers) on page 272</p> <p>show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration) on page 273</p> <p>show interfaces brief (Gigabit Ethernet) on page 273</p> <p>show interfaces detail (Gigabit Ethernet) on page 273</p> <p>show interfaces extensive (Gigabit Ethernet IQ2) on page 275</p> <p>show interfaces (Gigabit Ethernet Unnumbered Interface) on page 278</p> <p>show interfaces (ACI Interface Set Configured) on page 278</p>
Output Fields	<p>Table 14 on page 258 describes the output fields for the show interfaces (Gigabit Ethernet) command. Output fields are listed in the approximate order in which they appear. For Gigabit Ethernet IQ and IQE PICs, the traffic and MAC statistics vary by interface type. For more information, see Table 15 on page 271.</p>

Table 14: show interfaces Gigabit Ethernet Output Fields

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

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

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

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

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Output bytes—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the show interfaces (10-Gigabit Ethernet) command.</p>	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that 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 14: 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. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the Drops field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> <ul style="list-style-type: none"> • 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. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the Dropped packets field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>	detail extensive
Ingress queues	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	extensive

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

Field Name	Field Description	Level of Output
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. <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
Active alarms and Active defects	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 . <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
Interface transmit statistics	(On MX Series devices) Status of the interface-transmit-statistics configuration: Enabled or Disabled. <ul style="list-style-type: none"> • Enabled—When the interface-transmit-statistics statement is included in the configuration. If this is configured, the interface statistics show the actual transmitted load on the interface. • Disabled—When the interface-transmit-statistics statement is not included in the configuration. If this is not configured, the interface statistics show the offered load on the interface. 	detail extensive
OTN FEC statistics	The forward error correction (FEC) counters provide the following statistics: <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
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 the PCS receiver is operating in normal mode. 	detail extensive

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

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • 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—There are two possible conditions regarding the number of oversized frames: <ul style="list-style-type: none"> • Packet length exceeds 1518 octets, or • Packet length exceeds MRU • 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. <p>NOTE: The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the VLAN tagged frames field displays 0 when the show interfaces command is executed on a 20-port Gigabit Ethernet MIC. In other words, the number of VLAN tagged frames cannot be determined for the 20-port Gigabit Ethernet MIC.</p> <ul style="list-style-type: none"> • 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 14: 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 14: 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. Any 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 14: 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 (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 14: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • 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—Information from the remote Ethernet device: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the link partner, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the link partner. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), Symmetric/Asymmetric (link partner supports PAUSE on receive and transmit or only PAUSE on transmit), and None (link partner does not support flow control). • 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 local Ethernet device: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the local device. For Gigabit Ethernet interfaces, advertised capabilities are Symmetric/Asymmetric (local device supports PAUSE on receive and transmit or only PAUSE on receive) and None (local device does not support flow control). Depending on the result of the negotiation with the link partner, local resolution flow control type will display Symmetric (local device supports PAUSE on receive and transmit), Asymmetric (local device supports PAUSE on receive), and None (local device does not support flow control). • 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	(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.	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 14: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 14: 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
ACI VLAN: Dynamic Profile	Name of the dynamic profile that defines the agent circuit identifier (ACI) interface set. If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information.	brief detail extensive none
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Dynamic Profile	(MX Series routers with Trio MPCs only) Name of the dynamic profile that was used to create this interface configured with a Point-to-Point Protocol over Ethernet (PPPoE) family.	detail extensive none
Service Name Table	(MX Series routers with Trio MPCs only) Name of the service name table for the interface configured with a PPPoE family.	detail extensive none
Max Sessions	(MX Series routers with Trio MPCs only) Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	detail extensive none

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

Field Name	Field Description	Level of Output
Duplicate Protection	(MX Series routers with Trio MPCs only) State of PPPoE duplicate protection: On or Off . When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the router.	extensive
Transit statistics	<p>Number and rate of bytes and packets transiting the switch.</p> <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 <i>Common Output Fields Description</i> .	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	detail extensive none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	detail extensive

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

Field Name	Field Description	Level of Output
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 parentheses next to all interfaces.	detail extensive
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about the address flag. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Table 15: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound physical interface	show interfaces ge-0/3/0 extensive	<p>Traffic statistics:</p> <p>Input bytes: 496 bytes per packet, representing the Layer 2 packet</p> <p>MAC statistics:</p> <p>Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes</p>	The additional 4 bytes are for the CRC.
Inbound logical interface	show interfaces ge-0/3/0.50 extensive	<p>Traffic statistics:</p> <p>Input bytes: 478 bytes per packet, representing the Layer 3 packet</p>	
Outbound physical interface	show interfaces ge-0/0/0 extensive	<p>Traffic statistics:</p> <p>Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes</p> <p>MAC statistics:</p> <p>Received octets: 478 bytes per packet, representing the Layer 3 packet</p>	For input bytes, the additional 12 bytes include 6 bytes for the destination MAC address plus 4 bytes for VLAN plus 2 bytes for the Ethernet type.

Table 15: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type (*continued*)

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
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 (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues    : 4 supported, 4 maximum usable queues
  Current address: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
  Last flapped  : 2006-08-10 17:25:10 PDT (00:01:08 ago)
  Input rate    : 0 bps (0 pps)
  Output rate   : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine      : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms : None
  Active defects: None

Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69)
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  Egress accounting overhead: 100
  Ingress accounting overhead: 90
  Input packets : 0
  Output packets: 0
  Protocol ccc, MTU: 1522
  Flags: Is-Primary

```

show interfaces (Gigabit Ethernet on MX Series Routers)

```

user@host> show interfaces ge-2/2/2
Physical interface: ge-2/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 188
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, MAC-REWRITE Error: None,
  Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues    : 8 supported, 4 maximum usable queues
  Schedulers    : 0
  Current address: 00:1f:12:b7:d7:c0, Hardware address: 00:1f:12:b7:d6:76
  Last flapped  : 2008-09-05 16:44:30 PDT (3d 01:04 ago)
  Input rate    : 0 bps (0 pps)

```

```

Output rate      : 0 bps (0 pps)
Active alarms   : None
Active defects   : None
Logical interface ge-2/2/2.0 (Index 82) (SNMP ifIndex 219)
  Flags: SNMP-Traps 0x20000000 Encapsulation: Ethernet-Bridge
  Egress accounting overhead: 100
  Ingress accounting overhead: 90
  Input packets : 0
  Output packets: 0
  Protocol aenet, AE bundle: ae0.0    Link Index: 4

```

show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration)

```

user@host> show interfaces ge-2/1/2 extensive | match "output|interface"
Physical interface: ge-2/1/2, Enabled, Physical link is Up
Interface index: 151, SNMP ifIndex: 530, Generation: 154
Interface flags: SNMP-Traps Internal: 0x4000
Output bytes      :          240614363944          772721536 bps
Output packets:    3538446506             1420444 pps
Direction : Output
Interface transmit statistics: Enabled

Logical interface ge-2/1/2.0 (Index 331) (SNMP ifIndex 955) (Generation 146)
Output bytes      :          195560312716          522726272 bps
Output packets:    4251311146             1420451 pps

```

show interfaces brief (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2 brief
Physical interface: ge-3/0/2, Enabled, Physical link is Up
Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags      : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags        : None

Logical interface ge-3/0/2.0
Flags: SNMP-Traps 0x4000
VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
0x8100.512 0x8100.513)
Encapsulation: VLAN-CCC
ccc

Logical interface ge-3/0/2.32767
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2

```

show interfaces detail (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2 detail
Physical interface: ge-3/0/2, Enabled, Physical link is Up
Interface index: 167, SNMP ifIndex: 35, Generation: 177
Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags      : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags        : None
CoS queues        : 4 supported, 4 maximum usable queues
Hold-times        : Up 0 ms, Down 0 ms

```

```

Current address: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
Last flapped : 2006-08-09 17:17:00 PDT (01:31:33 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Ingress traffic statistics at Packet Forwarding Engine:
Input bytes : 0 0 bps
Input packets: 0 0 pps
Drop bytes : 0 0 bps
Drop packets: 0 0 pps
Ingress queues: 4 supported, 4 in use
Queue counters:

```

	Queued packets	Transmitted packets	Dropped packets
0 best-effort	0	0	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	0	0	0

```

Egress queues: 4 supported, 4 in use
Queue counters:

```

	Queued packets	Transmitted packets	Dropped packets
0 best-effort	0	0	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	0	0	0

```

Active alarms : None
Active defects : None

Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69) (Generation 140)
Flags: SNMP-Traps 0x4000
VLAN-Tag [0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530)
Out(swap-push 0x8100.512 0x8100.513)
Encapsulation: VLAN-CCC
Egress accounting overhead: 100
Ingress accounting overhead: 90
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol ccc, MTU: 1522, Generation: 149, Route table: 0

```

Flags: Is-Primary

Logical interface ge-3/0/2.32767 (Index 71) (SNMP ifIndex 70)
(Generation 139)
Flags: SNMP-Traps 0x4000 VLAN-Tag [0x0000.0] Encapsulation: ENET2
Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps

show interfaces extensive (Gigabit Ethernet IQ2)

```
user@host> show interfaces ge-7/1/3 extensive
Physical interface: ge-7/1/3, Enabled, Physical link is Up
Interface index: 170, SNMP ifIndex: 70, Generation: 171
Link-level type: Ethernet, MTU: 1514, Speed: 1000Mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4004000
Link flags : None
CoS queues : 8 supported, 4 maximum usable queues
Schedulers : 256
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:14:f6:30:5e:74, Hardware address: 00:14:f6:30:5e:74
Last flapped : 2007-11-07 21:31:41 PST (02:03:33 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes : 38910844056 7952 bps
  Output bytes : 7174605 8464 bps
  Input packets: 418398473 11 pps
  Output packets: 78903 12 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Ingress traffic statistics at Packet Forwarding Engine:
  Input bytes : 38910799145 7952 bps
  Input packets: 418397956 11 pps
  Drop bytes : 0 0 bps
  Drop packets: 0 0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
```

```

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

  0 best-effort      418390823                418390823                0
  1 expedited-fo          0                      0                      0
  2 assured-forw          0                      0                      0
  3 network-cont       7133                7133                0

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

  0 best-effort      1031                1031                0
  1 expedited-fo          0                      0                      0
  2 assured-forw          0                      0                      0
  3 network-cont       77872                77872                0

Active alarms : None
Active defects : None
MAC statistics:
  Receive      Transmit
  Total octets  38910844056    7174605
  Total packets 418398473    78903
  Unicast packets 408021893366 1026
  Broadcast packets 10      12
  Multicast packets 418398217 77865
  CRC/Align errors 0      0
  FIFO errors 0      0
  MAC control frames 0      0
  MAC pause frames 0      0
  Oversized frames 0
  Jabber frames 0
  Fragment frames 0
  VLAN tagged frames 0
  Code violations 0 OTN Received Overhead Bytes:
  APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58
  Payload Type: 0x08
OTN Transmitted Overhead Bytes:
  APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
  Payload Type: 0x08
Filter statistics:
  Input packet count      418398473
  Input packet rejects    479
  Input DA rejects        479
  Input SA rejects        0
  Output packet count      78903
  Output packet pad count  0
  Output packet error count 0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
    Link mode: Full-duplex, Flow control: Symmetric/Asymmetric,
    Remote fault: OK
  Local resolution:
    Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:

```

```

    Destination slot: 7
    CoS information:
    Direction : Output
    CoS transmit queue      Bandwidth      Buffer      Priority      Limit
                             %      bps      %      usec
    0 best-effort           95      950000000  95      0
low  none
    3 network-control       5      500000000   5      0
low  none
    Direction : Input
    CoS transmit queue      Bandwidth      Buffer      Priority      Limit
                             %      bps      %      usec
    0 best-effort           95      950000000  95      0
low  none
    3 network-control       5      500000000   5      0
low  none

Logical interface ge-7/1/3.0 (Index 70) (SNMP ifIndex 85) (Generation 150)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
Input bytes :      812400
Output bytes :    1349206
Input packets:      9429
Output packets:    9449
IPv6 transit statistics:
Input bytes :      0
Output bytes :      0
Input packets:      0
Output packets:      0
Local statistics:
Input bytes :      812400
Output bytes :    1349206
Input packets:      9429
Output packets:    9449
Transit statistics:
Input bytes :      0      7440 bps
Output bytes :      0      7888 bps
Input packets:      0      10 pps
Output packets:      0      11 pps
IPv6 transit statistics:
Input bytes :      0
Output bytes :      0
Input packets:      0
Output packets:      0
Protocol inet, MTU: 1500, Generation: 169, Route table: 0
Flags: Is-Primary, Mac-Validate-Strict
Mac-Validate Failures: Packets: 0, Bytes: 0
Addresses, Flags: Is-Preferred Is-Primary
Input Filters: F1-ge-3/0/1.0-in, F3-ge-3/0/1.0-in
Output Filters: F2-ge-3/0/1.0-out (53)
Destination: 10.74.2/24, Local: 10.74.2.2, Broadcast: 10.74.2.255,
Generation: 196
Protocol multiservice, MTU: Unlimited, Generation: 170, Route table: 0
Flags: Is-Primary
Policer: Input: __default_arp_policer__

```

NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics displayed in the **show interfaces** command output might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the interface counters. For detailed

information, see the description of the logical interface **Transit statistics** fields in [Table 14 on page 258](#).

show interfaces (Gigabit Ethernet Unnumbered Interface)

```
user@host> show interfaces ge-3/2/0
Physical interface: ge-3/2/0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 50
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Current address: 00:14:f6:11:26:f8, Hardware address: 00:14:f6:11:26:f8
  Last flapped   : 2006-10-27 04:42:23 PDT (08:01:52 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 624 bps (1 pps)
  Active alarms  : None
  Active defects : None

Logical interface ge-3/2/0.0 (Index 67) (SNMP ifIndex 85)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 0
  Output packets: 6
  Protocol inet, MTU: 1500
    Flags: Unnumbered
    Donor interface: lo0.0 (Index 64)
    Preferred source address: 22.22.22.22
```

show interfaces (ACI Interface Set Configured)

```
user@host> show interfaces ge-1/0/0.4001
Logical interface ge-1/0/0.4001 (Index 340) (SNMP ifIndex 548)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4001 ] Encapsulation: PPP-over-

Ethernet
ACI VLAN:
  Dynamic Profile: aci-vlan-set-profile
  PPPoE:
    Dynamic Profile: aci-vlan-pppoe-profile,
    Service Name Table: None,
    Max Sessions: 32000, Max Sessions VSA Ignore: Off,
    Duplicate Protection: On, Short Cycle Protection: Off,
    AC Name: nbc
  Input packets : 9
  Output packets: 8
  Protocol multiservice, MTU: Unlimited
```


show l2circuit connections

Syntax	<pre>show l2circuit connections <brief extensive summary> <down up up-down> <history> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor <i>neighbor</i>> <status></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Display enhancements in Junos OS Release 9.6.</p> <p>Display enhancements in Junos OS Release 10.2.</p> <p>Display enhancements in Junos OS Release 12.1.</p> <p>Display enhancements in Junos OS Release 13.2.</p>
Description	Display status information about Layer 2 virtual circuits from the local provider edge (PE) router to its neighbors.
Options	<p>none—Display standard information about Layer 2 virtual circuits on all interfaces for all neighbors.</p> <p>brief extensive summary—(Optional) Display the specified level of output. Use history to display information about connection history. Use status to display information about the connection and interface status.</p> <p>down up up-down—(Optional) Display nonoperational, operational, or both kinds of connections.</p> <p>history—(Optional) Display information about connection history.</p> <p>interface <i>interface-name</i>—(Optional) Show all Layer 2 virtual circuits on an interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor <i>neighbor</i>—(Optional) IP address of a specific neighbor.</p> <p>status—(Optional) Display information about the connection and interface status.</p>
Required Privilege Level	view
List of Sample Output	<p>show l2circuit connections on page 282</p> <p>show l2circuit connections interface on page 283</p> <p>show l2circuit connections extensive on page 283</p> <p>show l2circuit connections extensive (Pseudowire Redundancy with Hot Standby) on page 284</p>
Output Fields	<p>Table 16 on page 280 lists the output fields for the show l2circuit connections command. Output fields are listed in the approximate order in which they appear.</p>

Table 16: show l2circuit connections Output Fields

Field Name	Field Description
Layer-2 Circuit Connections	Displays the legends for connection and interface status.
Neighbor	Remote PE neighbor.
Interface	Logical PE-to-CE interface on which the virtual circuit is configured.
Type	VC type: rmt (remote) or loc (local).
Legend for connection status (St)	<p>Status of the virtual circuit connection:</p> <ul style="list-style-type: none"> • EI—The local virtual circuit interface is configured with an encapsulation that is not supported. • MM—The two routers do not agree on an MTU value, which causes an MTU mismatch. • EM—The encapsulation type received on this virtual circuit from the neighbor does not match the local virtual circuit interface encapsulation type. • CM—The two routers do not agree on a control word, which causes a control word mismatch. • VM—The remote and local VLAN IDs do not match across the Layer 2 circuit. • OL—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit. • NC—The interface is not configured as a CCC or TCC interface. • BK—The virtual circuit has switched to a backup connection. • CB—The remote PE router is advertising a different cell bundle from that configured on the local PE router. • LD—The connection to the local site is signaled down, because the CE-facing interface to the local site is down. • RD—The remote neighbor is down. It has signaled a problem using the pseudowire status code. • NP—The router detects that interface hardware is not present. The hardware may be offline, a PIC may not be of the desired type, or the interface may be configured in a different routing instance. • Dn—The virtual circuit is down. • VC-Dn—The virtual circuit is down because there is no tunnel LSP from the local PE router to the neighbor. • UP—The virtual circuit is operational. • CF—The router cannot find enough bandwidth to the remote router to satisfy the Layer 2 circuit bandwidth requirement. • IB—The bit rate is incompatible for Time Division Multiplexing (TDM). • TDM—TDM is not configured correctly. • ST—The virtual circuit has been switched to a standby connection. • SP—The virtual circuit connection is using a static pseudowire. • RS—The remote site is in a standby state. • XX—The virtual circuit is down for an unknown reason. This is a programming error.
Time last up	Date and time the virtual circuit was last operational.

Table 16: show l2circuit connections Output Fields (*continued*)

Field Name	Field Description
# Up trans	Number of times the virtual circuit came up.
<i>local-interface-name</i>	Name of the local PE-to-CE interface.
Status	Status of the local interface.
Up	Interface is operational.
Dn	Interface is not operational.
NP	Not present. Interface does not exist.
DS	Disabled. Interface has been administratively disabled.
WE	Wrong encapsulation. The interface is not configured as CCC.
UN	Interface status is initialized.
Encapsulation	Encapsulation of the local interface.
Remote PE	Prefix of the remote PE router.
Negotiated control-word	Whether the use of the control word has been negotiated for this virtual circuit: Yes (Null) or No .
Incoming label	Label used by the remote side of the virtual circuit to send packets destined to the local side. This label is routed to the local virtual circuit interface.
Outgoing label	Label used by the local side of the virtual circuit to send packets to the remote side of the virtual circuit. Packets originated on the local virtual circuit interface are encapsulated with this label before being placed on the tunnel LSP to the neighbor for this virtual circuit. This label is allocated by the neighbor and is used in demultiplexing incoming packets destined for this virtual circuit.
Negotiated PW status TLV	Displays the pseudowire status type, length, and value (TLV). TLVs are a method of encoding variable-length or optional information. If the pseudowire status TLV is used, the corresponding local or neighbor PE router status code is also displayed.
local PW status code	If the pseudowire status TLV is used, displays the local PE router status code.
Neighbor PW status code	If the pseudowire status TLV is used, displays the neighbor PE router status code.
Local interface	Name of the local interface used for the Layer 2 circuit connection.
Status	Status of the local interface (Up or Down).

Table 16: show l2circuit connections Output Fields (*continued*)

Field Name	Field Description
Encapsulation	Encapsulation configured for the local interface.
APS-active	Indicates that the interface belongs to the working circuit.
APS-inactive	Indicates that the interface belongs to the protect circuit.
Connection protection	Whether or not connection protection is configured for the Layer 2 circuit to the neighbor: Yes or No .
VC bandwidth	Bandwidth requirement of the Layer 2 circuit.
Time	Time at which the event occurred.
Connection History	<p>Event types logged in history.</p> <ul style="list-style-type: none"> • loc intf up—Local virtual circuit interface went up. • loc intf down—Local virtual circuit interface went down. • In lbl Update—Incoming label has been updated. • Out lbl Update—Outgoing label has been updated. • PE route changed—Route to PE router has been updated. • PE route down—Route to PE router is down. • rmt side marked—Remote side is marked. • VC Dn—Remote side indicated that its end of the virtual circuit is down (if the tunnel LSP from the remote side to the local side is down). • status update timer—Status update timer processing. It computes the state of the virtual circuit, and determines whether it should be advertised to or withdrawn from the remote side.

Sample Output

show l2circuit connections

```
user@host> show l2circuit connections
```

```
Layer-2 Circuit Connections:
```

```
Legend for connection status (St)
```

```

EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch              Dn -- down
EM -- encapsulation mismatch    VC-Dn -- Virtual circuit Down
CM -- control-word mismatch     Up -- operational
VM -- vlan id mismatch          CF -- Call admission control failure
OL -- no outgoing label         IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC   TM -- TDM misconfiguration
BK -- Backup Connection         ST -- Standby Connection
CB -- rcvd cell-bundle size bad SP -- Static Pseudowire
LD -- local site signaled down  RS -- remote site standby
RD -- remote site signaled down HS -- hot standby
XX -- unknown

```

```
Legend for interface status
```

```
Up -- operational
```

```

Dn -- down
Neighbor: 10.255.245.51
  Interface                Type St   Time last up      # Up trans
ge-2/0/2.600(vc 5)        rmt  Up    Dec 7 18:11:18 2009      1
  Remote PE: 10.255.245.51, Negotiated control-word: No
  Incoming label: 299856, Outgoing label: 299808
  Negotiated PW status TLV: No
  Local interface: ge-2/0/2.600, Status: Up, Encapsulation: VLAN

```

Sample Output

show l2circuit connections interface

```

user@host> show l2circuit connections interface t1-2/0/0:1:1.0
Layer-2 Circuit Connections:

Legend for connection status (St)
EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch              Dn -- down
EM -- encapsulation mismatch     VC-Dn -- Virtual circuit Down
CM -- control-word mismatch      Up -- operational
VM -- vlan id mismatch          CF -- Call admission control failure
OL -- no outgoing label         IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC    TM -- TDM misconfiguration
BK -- Backup Connection          ST -- Standby Connection
CB -- rcvd cell-bundle size bad  SP -- Static Pseudowire
LD -- local site signaled down   RS -- remote site standby
RD -- remote site signaled down  HS -- hot standby
XX -- unknown

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 10.1.1.1
  Interface                Type St   Time last up      # Up trans
t1-2/0/0:1:1.0(vc 1)(SP)  rmt  Up    Apr 27 04:21:02 2011      1
  Remote PE: 10.1.1.1, Negotiated control-word: Yes (Non-null)
  Incoming label: 1010001, Outgoing label: 1000001
  Negotiated PW status TLV: No
  Local interface: t1-1/0/0:1:1.0, Status: Up, Encapsulation: SATOP-T1,
APS-active
  Local interface: t1-2/0/0:1:1.0, Status: Up, Encapsulation: SATOP-T1,
APS-inactive

```

Sample Output

show l2circuit connections extensive

```

user@host> show l2circuit connections extensive
Layer-2 Circuit Connections:

Legend for connection status (St)
EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch              Dn -- down
EM -- encapsulation mismatch     VC-Dn -- Virtual circuit Down
CM -- control-word mismatch      Up -- operational
VM -- vlan id mismatch          CF -- Call admission control failure
OL -- no outgoing label         IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC    TM -- TDM misconfiguration
BK -- Backup Connection          ST -- Standby Connection
CB -- rcvd cell-bundle size bad  SP -- Static Pseudowire

```

LD -- local site signaled down RS -- remote site standby
RD -- remote site signaled down HS -- hot standby
XX -- unknown

Legend for interface status

Up -- operational
Dn -- down

Neighbor: 10.255.49.149

Interface	Type	St	Time last up	# Up trans
ae0.0(vc 100)	rmt	Up	Aug 31 09:36:12 2009	1

Remote PE: 10.255.49.149, Negotiated control-word: Yes (Null)
Incoming label: 299824, Outgoing label: 299776
Negotiated PW status TLV: Yes
Local PW status code: 0x00000000, Neighbor PW status code: 0x00000000
Local interface: ae0.0, Status: Up, Encapsulation: ETHERNET
Connection protection: Yes

Connection History:

Aug 31 09:36:12 2009	status update timer	
Aug 31 09:36:12 2009	PE route changed	
Aug 31 09:36:12 2009	Out lbl Update	299776
Aug 31 09:36:12 2009	In lbl Update	299824
Aug 31 09:36:12 2009	loc intf up	ae0.0

Sample Output

show l2circuit connections extensive (Pseudowire Redundancy with Hot Standby)

user@host>show l2circuit connections extensive

Layer-2 Circuit Connections:

Legend for connection status (St)

EI -- encapsulation invalid	NP -- interface h/w not present
MM -- mtu mismatch	Dn -- down
EM -- encapsulation mismatch	VC-Dn -- Virtual circuit Down
CM -- control-word mismatch	Up -- operational
VM -- vlan id mismatch	CF -- Call admission control failure
OL -- no outgoing label	IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC	TM -- TDM misconfiguration
BK -- Backup Connection	ST -- Standby Connection
CB -- rcvd cell-bundle size bad	SP -- Static Pseudowire
LD -- local site signaled down	RS -- remote site standby
RD -- remote site signaled down	HS -- Hot-standby Connection
XX -- unknown	

Legend for interface status

Up -- operational
Dn -- down

Neighbor: 88.0.0.101

Interface	Type	St	Time last up	# Up trans
ge-1/3/2.600(vc 1)	rmt	Up	Jan 24 11:00:26 2013	1

Remote PE: 88.0.0.101, Negotiated control-word: Yes (Null)
Incoming label: 299776, Outgoing label: 299776
Negotiated PW status TLV: Yes
Local PW status code: 0x00000000, Neighbor PW status code: 0x00000000
Local interface: ge-1/3/2.600, Status: Up, Encapsulation: VLAN

Connection History:

Jan 24 11:00:26 2013	status update timer	
Jan 24 11:00:26 2013	PE route changed	
Jan 24 11:00:26 2013	Out lbl Update	299776
Jan 24 11:00:26 2013	In lbl Update	299776
Jan 24 11:00:26 2013	loc intf up	ge-1/3/2.600

```
Neighbor: 88.0.0.102
Interface          Type  St    Time last up      # Up trans
ge-1/3/2.600(vc 2) rmt   HS    -----          ----
Remote PE: 88.0.0.102, Negotiated control-word: Yes (Null)
Incoming label: 299792, Outgoing label: 299776
Negotiated PW status TLV: Yes
local PW status code: 0x00000020, Neighbor PW status code: 0x00000000
Local interface: ge-1/3/2.600, Status: Up, Encapsulation: VLAN
```

show l2vpn connections

Syntax	<code>show l2vpn connections</code> <code><brief extensive></code> <code><down up up-down></code> <code><history></code> <code><instance <i>instance</i>></code> <code><instance-history></code> <code><local-site <i>local-site</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><remote-site <i>remote-site</i>></code> <code><status></code> <code><summary></code>
Release Information	Command introduced before Junos OS Release 7.4. instance-history option introduced in Junos OS Release 12.3R2.
Description	Display Layer 2 virtual private network (VPN) connections.
Options	none —Display all Layer 2 VPN connections for all routing instances. brief extensive —(Optional) Display the specified level of output. down up up-down —(Optional) Display nonoperational, operational, or both kinds of connections. history —(Optional) Display information about connection history. instance <i>instance</i> —(Optional) Display connections for the specified routing instance only. instance-history —(Optional) Display information about connection history for a particular instance. local-site <i>local-site</i> —(Optional) Display connections for the specified Layer 2 VPN local site name or ID only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. remote-site <i>remote-site</i> —(Optional) Display connection for the specified Layer 2 VPN remote site ID only. status —(Optional) Display information about the connection and interface status. summary —(Optional) Display summary of all Layer 2 VPN connections information.
Required Privilege Level	view
List of Sample Output	show l2vpn connections on page 289 show l2vpn connections extensive on page 290 show l2vpn connections extensive (VPWS) on page 290

Output Fields Table 17 on page 287 lists the output fields for the **show l2vpn connections** command. Output fields are listed in the approximate order in which they appear.

Table 17: show l2vpn connections Output Fields

Field Name	Field Description
Instance	Name of Layer 2 VPN instance.
L2vpn-id	For BGP autodiscovery, a globally unique Layer 2 VPN community identifier for the instance.
Local-ID	BGP local-address assigned to the local routing device.
Local site	Name of local site.
Local source-attachment-id	For FEC 129, the VPWS source attachment identifier. The point-to-point nature of VPWS requires that you specify the source access individual identifier (SAII) and the target access individual identifier (TAII). This SAII-TAII pair defines a unique pseudowire between two PE devices.
Target-attachment-id	For FEC 129, the VPWS target attachment identifier. If the configured target identifier matches a source identifier advertised by a remote PE device by way of a BGP auto-discovery message, the pseudowire between that source-target pair is signaled. If there is no match between an advertised source identifier and the configured target identifier, the pseudowire is not established.
Interface name	Name of interface.
Remote Site ID	Remote site ID.
Label Offset	Numbers within the label block that are skipped to find the next label base.
Label-base	Advertises the first label in a block of labels. A remote PE router uses this first label when sending traffic toward the advertising PE router.
Range	Advertises the label block size.
status-vector	Bit vector advertising the state of local PE-CE circuits to remote PE routers. A bit value of 0 indicates that the local circuit and LSP tunnel to the remote PE router are up, whereas a value of 1 indicates either one or both are down.
connection-site	Name of the connection site.
Type	Type of connection: loc (local) or rmt (remote).
St	Status of the connection. (For a list of possible values, see the Legend for connection status (St) field.)
Time last up	Time that the connection was last in the Up condition.
# Up trans	Number of transitions from Down to Up condition.

Table 17: show l2vpn connections Output Fields (*continued*)

Field Name	Field Description
Local circuit	Address and status of local circuit.
Remote circuit	Address and status of remote circuit.
St	<p>Status of the Layer 2 VPN connection (corresponds with Legend for Connection Status):</p> <ul style="list-style-type: none"> • EI—The local Layer 2 VPN interface is configured with an encapsulation that is not supported. • EM—The encapsulation type received on this Layer 2 VPN connection from the neighbor does not match the local Layer 2 VPN connection interface encapsulation type. • VC-Dn—The virtual circuit is currently down. • CM—The two routers do not agree on a control word, which causes a control word mismatch. • CN—The virtual circuit is not provisioned properly. • OR—The label associated with the virtual circuit is out of range. • OL—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit. • LD—All of the CE-facing interfaces to the local site are down. Therefore, the connection to the local site is signaled as down to the other PE routers. No pseudowires can be established. • RD—All the interfaces to the remote neighbor are down. Therefore, the remote site has been signaled as down to the other PE routers. No pseudowires can be established. • LN—The local site has lost path selection to the remote site and therefore no pseudowires can be established from this local site. • RN—The remote site has lost path selection to a local site or other remote site and therefore no pseudowires are established to this remote site. • XX—The Layer 2 VPN connection is down for an unknown reason. This is a programming error. • NC—The interface encapsulation is not configured as an appropriate CCC, TCC, or Layer 2 VPN encapsulation. • WE—The encapsulation configured for the interface does not match the encapsulation configured for the associated connection within the Layer 2 VPN routing instance. • NP—The router detects that interface hardware is not present. The hardware might be offline, a PIC might not be of the desired type, or the interface might be configured in a different routing instance. • ->—Only the outbound connection is up. • <-—Only the inbound connection is up. • Up—The Layer 2 VPN connection is operational. • Dn—The Layer 2 VPN connection is down. • CF—The router cannot find enough bandwidth to the remote router to satisfy the Layer 2 VPN connection bandwidth requirement. • SC—The local site identifier matches the remote site identifier. No pseudowire can be established between these two sites. You should configure different values for the local and remote site identifiers.

Table 17: show l2vpn connections Output Fields (*continued*)

Field Name	Field Description
	<ul style="list-style-type: none"> LM—The local site identifier is not the minimum designated, meaning it is not the lowest. There is another local site with a lower site identifier. Pseudowires are not being established to this local site, and the associated local site identifier is not being used to distribute Layer 2 VPN label blocks. However, this is not an error state. Traffic continues to be forwarded to the PE router interfaces connected to the local sites when the local sites are in this state. RM—The remote site identifier is not the minimum designated, meaning it is not the lowest. There is another remote site connected to the same PE router which has lower site identifier. The PE router cannot establish a pseudowire to this remote site and the associated remote site identifier cannot be used to distribute VPLS label blocks. However, this is not an error state. Traffic can continue to be forwarded to the PE router interface connected to this remote site when the remote site is in this state. IL—The incoming packets for the Layer 2 VPN connection have no MPLS label.
Remote PE	Address of the remote provider edge router.
Incoming label	Name of the incoming label.
Outgoing label	Name of the outgoing label.
Time	Date and time of Layer 2 VPN connection event.
Event	Type of event.
Interface/Lbl/PE	Interface, label, or PE router.

Sample Output

show l2vpn connections

```

user@host> show l2vpn connections
L2VPN Connections :
Instance : vpna
Local site: 2 (ce-2)
offset: 1, range: 3, label-base: 32768
  connection-site      Type St  Time last up      # Up trans
  3 (3)                loc  Up   Jul 18 20:45:46 2001      1
    Local circuit: fe-0/0/0.1, Status: Up
    Remote circuit: fe-0/0/3.0, Status: Up
  1                    rmt  Up   Jul 18 21:47:25 2001      1
    Local circuit: fe-0/0/0.0, Status: Up
    Remote PE: 192.168.16.1
    Incoming label: 32768, Outgoing label: 32769
Local site: 3 (ce-3)
offset: 1, range: 2, label-base: 33792
  connection-site      Type St  Time last up      # Up trans
  2 (ce-b)             loc  Up   Jul 18 20:45:46 2001      1
    Local circuit: fe-0/0/0.1, Status: Up
    Remote circuit: fe-0/0/3.0, Status: Up
  1                    rmt  Up   Jul 18 21:47:25 2001      1

```

```

Local circuit: fe-0/0/3.1, Status: Up
Remote PE: 192.168.16.1
Incoming label: 33792, Outgoing label: 32770

```

show l2vpn connections extensive

```

user@host> show l2vpn connections extensive
L2VPN Connections:

```

Legend for connection status (St)

```

EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down   NP -- interface hardware not present
CM -- control-word mismatch     -> -- only outbound connection is up
CN -- circuit not provisioned   <- -- only inbound connection is up
OR -- out of range             Up -- operational
OL -- no outgoing label        Dn -- down
LD -- local site signaled down  CF -- call admission control failure
RD -- remote site signaled down SC -- local and remote site ID collision
LN -- local site not designated LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status IL -- no incoming label

```

```
Instance: vpn-a
```

```
Local site: ce-a (1)
```

Interface name	Remote Site ID
fe-0/0/0.0	2

Label Offset	Offset	Range
32768	1	2

connection-site	Type	St	Time last up	# Up trans
2	rmt	Up	Aug 3 00:08:14 2001	1

```

Local circuit: fe-0/0/0.0, Status: Up
Remote PE: 192.168.24.1
Incoming label: 32769, Outgoing label: 32768

```

Time	Event	Interface/Lbl/PE
Aug 3 00:08:14 2001	PE route up	
Aug 3 00:08:14 2001	Out lbl Update	32768
Aug 3 00:08:14 2001	In lbl Update	32769
Aug 3 00:08:14 2001	ckt0 up	fe-0/0/0.0

show l2vpn connections extensive (VPWS)

```

user@host> show l2vpn connections
Layer-2 VPN connections:

```

Legend for connection status (St)

```

EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down   NP -- interface hardware not present
CM -- control-word mismatch     -> -- only outbound connection is up
CN -- circuit not provisioned   <- -- only inbound connection is up
OR -- out of range             Up -- operational
OL -- no outgoing label        Dn -- down
LD -- local site signaled down  CF -- call admission control failure
RD -- remote site signaled down SC -- local and remote site ID collision
LN -- local site not designated LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status IL -- no incoming label
MM -- MTU mismatch             MI -- Mesh-Group ID not available
BK -- Backup connection        ST -- Standby connection
PF -- Profile parse failure     PB -- Profile busy

```

RS -- remote site standby SN -- Static Neighbor
 LB -- Local site not best-site RB -- Remote site not best-site
 VM -- VLAN ID mismatch

Legend for interface status

Up -- operational
 Dn -- down

Instance: FEC129-VPWS

L2vpn-id: 100:100

Number of local interfaces: 1

Number of local interfaces up: 1

ge-2/0/5.0

Local source-attachment-id: 1 (ONE)

Target-attachment-id	Type	St	Time last up	# Up trans
2	rmt	Up	Nov 28 16:16:14 2012	1

Remote PE: 2.2.2.2, Negotiated control-word: No

Incoming label: 299792, Outgoing label: 299792

Local interface: ge-2/0/5.0, Status: Up, Encapsulation: ETHERNET

Connection History:

Nov 28 16:16:14 2012	status update timer	
Nov 28 16:16:14 2012	PE route changed	
Nov 28 16:16:14 2012	Out lbl Update	299792
Nov 28 16:16:14 2012	In lbl Update	299792
Nov 28 16:16:14 2012	loc intf up	ge-2/0/5.0

show policer

Syntax	show policer <detail> <policer-name>
Release Information	Command introduced before Junos OS Release 7.4. Option detail introduced in Junos OS Release 12.3.
Description	Display the number of policed packets for a given policer or an aggregate policer. An aggregate policer is an aggregate of different policers on the same logical interface.
Options	none —Display the number of policed packets for all configured policers. detail —(Optional) Display enhanced statistics for policers. policer-name —(Optional) Display the number of policed packets for the specified policer.
Required Privilege Level	view
List of Sample Output	show policer (MX Series) on page 292 show policer (non MX Series Router) on page 293 show policer (Aggregate Policer, non MX Series Router) on page 293 show policer detail on page 293
Output Fields	Table 18 on page 292 lists the output fields for the show policer command. Output fields are listed in the approximate order in which they appear.

Table 18: show policer Output Fields

Field Name	Field Description
Name	Name of the policer.
Bytes	(For two-color policers on MX Series routers, and for hierarchical policers on MS-DPC, MIC, and MPC interfaces on MX Series routers) Total number of bytes policed by the specified policer. For other combinations of policer type, device, and line card type, this field is blank.
Packets	Total number of packets policed by the specified policer.

Sample Output

show policer (MX Series)

```

user@host> show policer
Policers:
Name                               Bytes          Packets
__default_arp_policer__            314520         5242
pol-2M-ge-1/2/0.1-inet-i          10372300       103723

```

pol-2M-ge-1/2/0.1-inet6-i	7727800	77278
pol-2M-ge-1/2/0.1-mp1s-i	7070336	67984
pol-2M-ge-1/2/0.1001-vp1s-i	65153700	651537
pol-2M-ge-1/2/0.2001-vp1s-i	65180900	651809
pol-2M-ge-1/2/0.3001-ccc-i	62202144	647939

show policer (non MX Series Router)

```
user@host> show policer
Policers:
Name                               Bytes          Packets
__default_arp_policer__
pol-2M-ge-1/2/0.1-inet-i          103723
pol-2M-ge-1/2/0.1-inet6-i        77278
pol-2M-ge-1/2/0.1-mp1s-i         67984
pol-2M-ge-1/2/0.1001-vp1s-i      651537
pol-2M-ge-1/2/0.2001-vp1s-i      651809
pol-2M-ge-1/2/0.3001-ccc-i       647939
```

show policer (Aggregate Policer, non MX Series Router)

```
user@host> show policer
Policers:
Name                               Bytes          Packets
__default_arp_policer__
P1-ae0.0-log_int-o                0
P2-ge-7/0/2.0-inet-o              0
P2-ge-7/0/2.0-inet6-o             0
__policer_tmpl__-term              0
__policer_tmpl__-fc0               0
__policer_tmpl__-fc0               0
__policer_tmpl__-fc1               0
__policer_tmpl__-fc0               0
__policer_tmpl__-fc1               0
__policer_tmpl__-fc2               0
__policer_tmpl__-fc0               0
__policer_tmpl__-fc1               0
__policer_tmpl__-fc2               0
__policer_tmpl__-fc3               0
```

show policer detail

```
user@host> show policer detail
Policers:
Name                               Bytes          Packets
__default_arp_policer__
  OOS                               0              0
  Offered                           0             496
  Transmitted                        0             496
P1-xe-1/0/0.0-inet-i
  OOS                               0             11329
  Offered                           0             111188
  Transmitted                        0             99859
```

show route export vrf-target

Syntax	show route export vrf-target <brief detail> <community <i>community--regular-expression</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display the VPN routing and forwarding (VRF) target communities for which policy-based route export is currently distributing routes. This command is relevant when there are overlapping virtual private networks (VPNs).
Options	<p>none—Display standard information about all target communities.</p> <p>brief detail—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>community <i>community-regular-expression</i>—(Optional) Display information about the specified community.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route export vrf-target on page 295 show route export vrf-target community on page 295 show route export vrf-target detail on page 295
Output Fields	Table 19 on page 294 lists the output fields for the show route export vrf-target command. Output fields are listed in the approximate order in which they appear.

Table 19: show route export vrf-target Output Fields

Field Name	Field Description	Level of Output
Route target	Target communities for which auto-export is currently distributing routes.	brief none
Family	Routing table entries for the specified family.	brief none
<i>type-of-routing-table(s)</i>	Type of routing tables the feature handles: <ul style="list-style-type: none"> unicast—Indicates <i>instance.inet.0</i>. multicast—Indicates <i>instance.inet.2</i>. unicast multicast—Indicates <i>instance.inet.0</i> and <i>instance.inet.2</i>. 	brief none
Import	Number of routing tables that are currently importing routes with this target community. Omitted for tables that are not importing routes.	brief none

Table 19: show route export vrf-target Output Fields (*continued*)

Field Name	Field Description	Level of Output
Export	Number of routing tables that are currently exporting routes with this target community. Omitted for tables that are not exporting routes.	brief none
Target	Target communities, family, and options for which auto-export is currently distributing routes.	detail
Import table(s)	Name of the routing tables that are importing a particular route target.	detail
Export table(s)	Name of the routing tables that are exporting a particular route target.	detail

Sample Output

show route export vrf-target

```

user@host> show route export vrf-target
Route Target          Family      Import      Export
69:1                  inet        unicast      2          2
69:2                  inet        unicast      2          2

```

show route export vrf-target community

```

user@host> show route export vrf-target community target:69:1
Route Target          Family      Import      Export
69:1                  inet        unicast      2          2

```

show route export vrf-target detail

```

user@host> show route export vrf-target detail
Target: 1:12          inet        unicast
  Import table(s): vrf-11.inet.0 vrf-12.inet.0
  Export table(s): vrf-12.inet.0
Target: 1:13          inet        unicast
  Import table(s): vrf-12.inet.0 vrf-13.inet.0
  Export table(s): vrf-13.inet.0

```

show route forwarding-table interface-name

Syntax	show route forwarding-table interface-name <i>interface-name</i> <detail extensive> <all>
Release Information	Command introduced in Junos OS Release 9.6.
Description	Display the interfaces in the Routing Engine's forwarding table.
Options	<p>none—Display information for the specified interface.</p> <p>detail extensive—(Optional) Display the specified level of output.</p> <p>all—(Optional) Display all interfaces in the routing table.</p>
Required Privilege Level	view
List of Sample Output	show route forwarding-table interface-name fe-0/1/1 on page 297 show route forwarding-table interface-name all on page 297 show route forwarding-table interface-name all detail on page 298
Output Fields	Table 20 on page 296 lists the output fields for the show route forwarding-table interface-name command. Output fields are listed in the approximate order in which they appear.

Table 20: show route forwarding-table interface-name Output Fields

Field Name	Field Description	Level of Output
Name	Name of the interface (for example fe-0/1/1 , lo0 , ae0 , and so on).	All levels
MTU	Interface's maximum transmission unit (MTU).	All levels
Afam	Configured address family (for example inet , tnp , inet6 , and so on).	detail extensive
Network	Network information: <ul style="list-style-type: none"> • <Link>—Physical interface, not a logical interface. • <PtoP>—Point-to-point network. • ipaddress—Network address. 	All levels
Address	Address of the interface. The address can be a MAC address, IPv4 address, IPv6 address, and so on.	All levels
IPkts	Number of packets received on the interface.	All levels
Ierr	Number of packets received on the interface with errors.	All levels
Opkts	Number of packets transmitted or sent from the interface.	All levels

Table 20: show route forwarding-table interface-name Output Fields (*continued*)

Field Name	Field Description	Level of Output
Oerr	Number of packets transmitted or sent from the interface with errors.	All levels
Coll	Number of packets that experienced collisions on the interface.	All levels

Sample Output

show route forwarding-table interface-name fe-0/1/1

```

user@host> show route forwarding-table interface-name fe-0/1/1
Name      Mtu Network      Address      IpKts Ierr      Opkts Oerr  Coll
fe-0/1/1  1514 <Link>      00.05.85.88.cc.20      0    0        0    0    0

```

show route forwarding-table interface-name all

```

user@host> show route forwarding-table interface-name all
Name      Mtu Network      Address      IpKts Ierr      Opkts Oerr  Coll
fxp0      1514 <Link>      00.a0.a5.56.03.83      180965    0      39907    0    0

  unit 0
  fxp1   1514 <Link>      02.00.00.00.00.04      33010497    0      30110800    0    0

  unit 0
    1500 192.168.187.0/ 192.168.187.10
  fxp1   1514 <Link>      02.00.00.00.00.04      33010497    0      30110800    0    0

  unit 0
    1500 10.0.0.0/8      10.0.0.1
    10.0.0.0/8      10.0.0.4
    128.0.0.0/2      128.0.0.1
    128.0.0.0/2      128.0.0.4
    1500 fe80::/64      fe80::200:ff:fe0
    fec0::/64      fec0::a:0:0:4
    1500      4
  lsi     1496 <Link>
  dsc     max <Link>
  lo0     max <Link>
                                8980    0      8980    0    0

  unit 0
    max 127.0.0.1/8      127.0.0.1
    192.168.0.1/8      192.168.0.1
  unit 16384 max 127.0.0.1/8      127.0.0.1
  unit 16385 max
  gre     max <Link>
  ipip    max <Link>
  tap     max <Link>
  pime    max <Link>
  pimd    max <Link>
  mtun    max <Link>
  so-0/0/0 4474 <Link>
                                1679900    0      1068611    0    0

  unit 0
  0       4470 <PtoP>      10.0.60.2
                                0    0        0    0
  so-0/0/1 4474 <Link>
                                0    0        0    0    0

  unit 0
  0       4470 <PtoP>      10.0.80.2
                                0    0        0    0
  so-0/0/2 4474 <Link>
                                0    0        0    0    0
  so-0/0/3 4474 <Link>
                                0    0        0    0    0
  fe-0/1/0 1514 <Link>      00.05.85.88.cc.1f      523120    0      623044    0    0

  unit 0
    1500 10.0.90.12/30 10.0.90.14
                                0    0        0    0

```

```

0
fe-0/1/1 1514 <Link> 00.05.85.88.cc.20 0 0 0 0 0
fe-0/1/2 1514 <Link> 00.05.85.88.cc.21 0 0 0 0 0

```

```
...
```

show route forwarding-table interface-name all detail

```

user@host> show route forwarding-table interface-name all detail
Name      Mtu AFam  Network      Address      Ipkts Ierr  Opkts
Oerr Coll
fxp0      1514      <Link> 00.a0.a5.56.03.83 181005 0 39948
0 0
unit 0     1500 inet  192.168.187.0/ 192.168.187.10
fxp1      1514      <Link> 02.00.00.00.00.04 33012676 0 30112468
0 0
unit 0     1500 inet  10.0.0.0/8     10.0.0.1
                        10.0.0.0/8     10.0.0.4
                        128.0.0.0/2    128.0.0.1
                        128.0.0.0/2    128.0.0.4
                        1500 inet6  fe80::/64      fe80::200:ff:fe0
                        fec0::/64      fec0::a:0:0:4
                        1500 tnp      4
lsi        1496      <Link>
dsc        max      <Link> 0 0 0
0 0
lo0        max      <Link> 8980 0 8980
0 0
unit 0     max inet  127.0.0.1/8    127.0.0.1
                        192.168.0.1/8  192.168.0.1
unit 16384 max inet  127.0.0.1/8    127.0.0.1
unit 16385 max inet
gre        max      <Link>
ipip       max      <Link>
tap        max      <Link>
pime       max      <Link>
pimd       max      <Link>
mtun       max      <Link>
so-0/0/0   4474      <Link> 1679980 0 1068661
0 0
unit 0     4470 inet  <PtoP> 10.0.60.2 0 0 0
0 0
...

```

show route localization

Syntax	show route localization
Release Information	Command introduced in Junos OS Release 11.4.
Description	(T320, T640, and T1600 routers only) Display route localization details.
Options	detail —Display detailed output.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> <i>Example: Configuring Packet Forwarding Engine FIB Localization</i>
Output Fields	Table 21 on page 299 lists the output fields for the show route localization command. Output fields are listed in the approximate order in which they appear.

Table 21: show route localization Output Fields

Field Name	Field Description
FIB-local	FPCs configured as FIB-local.
FIB-remote	FPCs configured as FIB-remote.
Normal	FPCs neither configured as FIB-local or FIB-remote .
Protocols	IPv4 (inet) or IPv6 (inet6) traffic configured for route localization.

Sample Output

```

user@R0> show route localization
FIB localization ready FPCs (and FIB-local Forwarding Engine addresses)
  FIB-local:  FPC2(4,5)
  FIB-remote: FPC0, FPC1
  Normal:     FPC3, FPC4, FPC5, FPC6, FPC7

user@R0> show route localization detail
FIB localization ready FPCs (and FIB-local Forwarding Engine addresses)
  FIB-local:  FPC2(4,5)
  FIB-remote: FPC0, FPC1
  Normal:     FPC3, FPC4, FPC5, FPC6, FPC7
FIB localization configuration
  Protocols:  inet, inet6
  FIB-local:  FPC2
  FIB-remote: FPC0, FPC1
Forwarding Engine addresses
  FPC0: 1
  FPC1: 2
  FPC2: 4, 5
  FPC3: 6

```

FPC4: 8
FPC5: 11
FPC6: 13
FPC7: 15

show route match-prefix

Syntax	<code>show route match-prefix <i>match-prefix</i>;</code>
Release Information	Command introduced in Junos OS Release 11.4.
Description	Allows you to search for routes using regular expressions based on the extended (modern) regular expressions as defined in POSIX 1003.2.
Options	<i>match-prefix</i> —Regular expression to match formatted prefix.
Additional Information	
Required Privilege Level	view
Related Documentation	<i>Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands</i>
List of Sample Output	show route match-prefix *:10.255.2.200:6:* (Show all routes matching route distributor 10.255.2.200:6) on page 301 show route match-prefix 7* (Show all mvpn type-7 routes) on page 301 show route match-prefix *:224.* (Show all routes matching group 224/4) on page 301
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

[show route match-prefix *:10.255.2.200:6:* \(Show all routes matching route distributor 10.255.2.200:6\)](#)

```
user@host> show route match-prefix *:10.255.2.200:6:*
```

[show route match-prefix 7* \(Show all mvpn type-7 routes\)](#)

```
user@host> show route table blue.mvpn.0 match-prefix 7*
Paste
router command output here
```

[show route match-prefix *:224.* \(Show all routes matching group 224/4\)](#)

```
user@host> show route match-prefix *:224.*
```

show route range

Syntax	<code>show route range</code> <code><brief detail extensive terse></code> <code><destination-prefix></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route range</code> <code><brief detail extensive terse></code> <code><destination-prefix></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display routing table entries using a prefix range.
Options	<p>none—Display standard information about all routing table entries using a prefix range.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>destination-prefix—(Optional) Destination and prefix mask for the range.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route range on page 302 show route range destination-prefix on page 303 show route range detail on page 303 show route range extensive on page 304 show route range terse on page 305
Output Fields	For information about output fields, see the output field tables for the <i>show route</i> command, the <i>show route detail</i> command, the <i>show route extensive</i> command, or the <i>show route terse</i> command.

Sample Output

show route range

```
user@host> show route range

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
10.255.71.14/32  *[Direct/0] 00:30:01
                  > via lo0.0
```



```

172.16.0.0/12      *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
192.168.0.0/16    *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
192.168.64.0/21   *[Direct/0] 00:30:01
                  > via fxp0.0
192.168.71.14/32  *[Local/0] 00:30:01
                  Local via fxp0.0
192.168.102.0/23  *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
...

```

show route range destination-prefix

```

user@host> show route range 192.168.0.0

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16    *[Static/5] 00:31:14
                  > to 192.168.71.254 via fxp0.0
192.168.64.0/21   *[Direct/0] 00:31:14
                  > via fxp0.0
192.168.71.14/32  *[Local/0] 00:31:14
                  Local via fxp0.0
192.168.102.0/23  *[Static/5] 00:31:14
                  > to 192.168.71.254 via fxp0.0

```

show route range detail

```

user@host> show route range detail

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.255.71.14/32 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Age: 30:05
    Task: IF

```

```
AS path: I
172.16.0.0/12 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I
```

...

show route range extensive

```
user@host> show route range extensive

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.255.71.14/32 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Age: 30:17
    Task: IF
    AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17
```

Task: RT
 Announcement bits (1): 0-KRT
 AS path: I

...

show route range terse

user@host> show route range terse

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
 + = Active Route, - = Last Active, * = Both

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
* 10.10.0.0/16	S 5			>192.168.71.254	
* 10.209.0.0/16	S 5			>192.168.71.254	
* 10.255.71.14/32	D 0			>lo0.0	
* 172.16.0.0/12	S 5			>192.168.71.254	
* 192.168.0.0/16	S 5			>192.168.71.254	
* 192.168.64.0/21	D 0			>fxp0.0	
* 192.168.71.14/32	L 0			Local	
* 192.168.102.0/23	S 5			>192.168.71.254	
* 207.17.136.0/24	S 5			>192.168.71.254	
* 207.17.136.192/32	S 5			>192.168.71.254	

__juniper_private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
* 10.0.0.0/8	D 0			>fxp2.0	
	D 0			>fxp1.0	
* 10.0.0.4/32	L 0			Local	

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
47.0005.80ff.f800.0000.0108.0001.0102.5507.1014/152					
*	D 0			>lo0.0	

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
abcd::10:255:71:14/128					
*	D 0			>lo0.0	
fe80::280:42ff:fe11:226f/128					
*	D 0			>lo0.0	

__juniper_private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
fe80::280:42ff:fe11:226f/128					
*	D 0			>lo0.16385	

show route snooping

Syntax	<code>show route snooping</code> <code><brief detail extensive terse></code> <code><all></code> <code><best address/prefix></code> <code><exact address></code> <code><range prefix-range></code> <code><summary></code> <code><table table-name></code>
Release Information	Command introduced in Junos OS Release 8.5. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the entries in the routing table that were learned from snooping.
Options	<p>none—Display the entries in the routing table that were learned from snooping.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>all—(Optional) Display all entries, including hidden entries.</p> <p>best address/prefix—(Optional) Display the longest match for the provided address and optional prefix.</p> <p>exact address/prefix—(Optional) Display exact matches for the provided address and optional prefix.</p> <p>range prefix-range—(Optional) Display information for the provided address range.</p> <p>summary—(Optional) Display route snooping summary statistics.</p> <p>table table-name—(Optional) Display information for the named table.</p>
Required Privilege Level	view
List of Sample Output	show route snooping detail on page 306
Output Fields	For information about output fields, see the output field tables for the <i>show route</i> command, the <i>show route detail</i> command, the <i>show route extensive</i> command, or the <i>show route terse</i> command.

Sample Output

show route snooping detail

```
user@host> show route snooping detail
__+domainAll__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
 224.0.0.2/32 (1 entry, 1 announced)
    *IGMP      Preference: 0
               Next hop type: MultiRecv
               Next-hop reference count: 4
               State: <Active NoReadvrt Int>
```

```

Age: 2:24
Task: IGMP
Announcement bits (1): 0-KRT
AS path: I

224.0.0.22/32 (1 entry, 1 announced)
  *IGMP Preference: 0
    Next hop type: MultiRecv
    Next-hop reference count: 4
    State: <Active NoReadvrt Int>
    Age: 2:24
    Task: IGMP
    Announcement bits (1): 0-KRT
    AS path: I

__+domainAll__.inet.1: 36 destinations, 36 routes (36 active, 0 holddown, 0 hidden)

224.0.0.0.0.0.0.0.0/24 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4), Next hop index: 1048584
    Next-hop reference count: 4
    State: <Active Int>
    Age: 2:24
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.2.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.3.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.4.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:17
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.5.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113

```

```
State: <Active Int>
Age: 1:58
Task: MC
Announcement bits (1): 0-KRT
AS path: I

225.0.0.6.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:14
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.7.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.9.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.10.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.1.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.2.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
```

```
Age: 8
Task: MC
Announcement bits (1): 0-KRT
AS path: I

226.0.0.4.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.8.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.10.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:56
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.1.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.2.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.3.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:16
```

```
Task: MC
Announcement bits (1): 0-KRT
AS path: I

227.0.0.4.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.5.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:57
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.7.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:57
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.8.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.10.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.1.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
```



```
Announcement bits (1): 0-KRT
AS path: I

228.0.0.2.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:18
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.7.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:11
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.8.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:17
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.9.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 8
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.10.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.3.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
    Announcement bits (1): 0-KRT
```

AS path: I

229.0.0.4.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
*Multicast Preference: 180
Next hop type: Multicast (IPv4)
Next-hop reference count: 113
State: <Active Int>
Age: 2:12
Task: MC
Announcement bits (1): 0-KRT
AS path: I

229.0.0.5.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
*Multicast Preference: 180
Next hop type: Multicast (IPv4)
Next-hop reference count: 113
State: <Active Int>
Age: 9
Task: MC
Announcement bits (1): 0-KRT
AS path: I

229.0.0.6.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
*Multicast Preference: 180
Next hop type: Multicast (IPv4)
Next-hop reference count: 113
State: <Active Int>
Age: 2:15
Task: MC
Announcement bits (1): 0-KRT
AS path: I

229.0.0.7.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
*Multicast Preference: 180
Next hop type: Multicast (IPv4)
Next-hop reference count: 113
State: <Active Int>
Age: 2:15
Task: MC
Announcement bits (1): 0-KRT
AS path: I

229.0.0.8.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
*Multicast Preference: 180
Next hop type: Multicast (IPv4)
Next-hop reference count: 113
State: <Active Int>
Age: 2:15
Task: MC
Announcement bits (1): 0-KRT
AS path: I

229.0.0.9.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
*Multicast Preference: 180
Next hop type: Multicast (IPv4)
Next-hop reference count: 113
State: <Active Int>
Age: 2:14
Task: MC
Announcement bits (1): 0-KRT
AS path: I

```
229.0.0.10.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I
```

show route snooping

Syntax	<code>show route snooping</code> <code><brief detail extensive terse></code> <code><all></code> <code><best address/prefix></code> <code><exact address></code> <code><range prefix-range></code> <code><summary></code> <code><table table-name></code>
Release Information	Command introduced in Junos OS Release 8.5. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the entries in the routing table that were learned from snooping.
Options	<p>none—Display the entries in the routing table that were learned from snooping.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>all—(Optional) Display all entries, including hidden entries.</p> <p>best address/prefix—(Optional) Display the longest match for the provided address and optional prefix.</p> <p>exact address/prefix—(Optional) Display exact matches for the provided address and optional prefix.</p> <p>range prefix-range—(Optional) Display information for the provided address range.</p> <p>summary—(Optional) Display route snooping summary statistics.</p> <p>table table-name—(Optional) Display information for the named table.</p>
Required Privilege Level	view
List of Sample Output	show route snooping detail on page 314
Output Fields	For information about output fields, see the output field tables for the <i>show route</i> command, the <i>show route detail</i> command, the <i>show route extensive</i> command, or the <i>show route terse</i> command.

Sample Output

show route snooping detail

```
user@host> show route snooping detail
__+domainAll__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
 224.0.0.2/32 (1 entry, 1 announced)
    *IGMP      Preference: 0
               Next hop type: MultiRecv
               Next-hop reference count: 4
               State: <Active NoReadvrt Int>
```

```

Age: 2:24
Task: IGMP
Announcement bits (1): 0-KRT
AS path: I

224.0.0.22/32 (1 entry, 1 announced)
  *IGMP Preference: 0
    Next hop type: MultiRecv
    Next-hop reference count: 4
    State: <Active NoReadvrt Int>
    Age: 2:24
    Task: IGMP
    Announcement bits (1): 0-KRT
    AS path: I

__+domainAll__.inet.1: 36 destinations, 36 routes (36 active, 0 holddown, 0 hidden)

224.0.0.0.0.0.0.0.0/24 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4), Next hop index: 1048584
    Next-hop reference count: 4
    State: <Active Int>
    Age: 2:24
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.2.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.3.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.4.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:17
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.5.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113

```

```
State: <Active Int>
Age: 1:58
Task: MC
Announcement bits (1): 0-KRT
AS path: I

225.0.0.6.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:14
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.7.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.9.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.10.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.1.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.2.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
```

```
Age: 8
Task: MC
Announcement bits (1): 0-KRT
AS path: I

226.0.0.4.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.8.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.10.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:56
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.1.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.2.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.3.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:16
```

```
Task: MC
Announcement bits (1): 0-KRT
AS path: I

227.0.0.4.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.5.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:57
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.7.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:57
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.8.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.10.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.1.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
```



```
Announcement bits (1): 0-KRT
AS path: I

228.0.0.2.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:18
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.7.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:11
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.8.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:17
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.9.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 8
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.10.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.3.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
    Announcement bits (1): 0-KRT
```

AS path: I

229.0.0.4.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
*Multicast Preference: 180
Next hop type: Multicast (IPv4)
Next-hop reference count: 113
State: <Active Int>
Age: 2:12
Task: MC
Announcement bits (1): 0-KRT
AS path: I

229.0.0.5.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
*Multicast Preference: 180
Next hop type: Multicast (IPv4)
Next-hop reference count: 113
State: <Active Int>
Age: 9
Task: MC
Announcement bits (1): 0-KRT
AS path: I

229.0.0.6.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
*Multicast Preference: 180
Next hop type: Multicast (IPv4)
Next-hop reference count: 113
State: <Active Int>
Age: 2:15
Task: MC
Announcement bits (1): 0-KRT
AS path: I

229.0.0.7.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
*Multicast Preference: 180
Next hop type: Multicast (IPv4)
Next-hop reference count: 113
State: <Active Int>
Age: 2:15
Task: MC
Announcement bits (1): 0-KRT
AS path: I

229.0.0.8.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
*Multicast Preference: 180
Next hop type: Multicast (IPv4)
Next-hop reference count: 113
State: <Active Int>
Age: 2:15
Task: MC
Announcement bits (1): 0-KRT
AS path: I

229.0.0.9.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
*Multicast Preference: 180
Next hop type: Multicast (IPv4)
Next-hop reference count: 113
State: <Active Int>
Age: 2:14
Task: MC
Announcement bits (1): 0-KRT
AS path: I

```
229.0.0.10.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I
```


CHAPTER 6

Command Summaries

- [ANCP Operational Mode Commands on page 323](#)
- [BFD Operational Mode Commands on page 324](#)
- [BGP Operational Mode Commands on page 325](#)
- [ES-IS Operational Mode Commands on page 326](#)
- [IP Multicast Operational Mode Commands on page 326](#)
- [IPv6 Operational Mode Commands on page 330](#)
- [IS-IS Operational Mode Commands on page 331](#)
- [LLDP Operational Mode Commands on page 332](#)
- [MVRP Operational Mode Commands on page 333](#)
- [OSPF Operational Mode Commands on page 333](#)
- [Protocol-Independent Routing Operational Mode Commands on page 335](#)
- [RIP Operational Mode Commands on page 337](#)
- [RIPng Operational Mode Commands on page 338](#)
- [Firewall Filter Operational Mode Commands on page 339](#)
- [Layer 2 Bridging and Switching Operational Mode Commands on page 339](#)
- [VPN Operational Mode Commands on page 340](#)

ANCP Operational Mode Commands

[Table 22 on page 323](#) summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot Access Node Control Protocol (ANCP) operations. Commands are listed in alphabetical order.

Table 22: ANCP Operational Mode Commands

Task	Command
Clear ANCP neighbors.	<i>clear ancp neighbor</i>
Clear ANCP subscriber connections.	<i>clear ancp subscriber</i>
Trigger the access node to run a loopback test on the local loop specified by an ANCP interface or interface set.	<i>request ancp oam interface</i>

Table 22: ANCP Operational Mode Commands (*continued*)

Task	Command
Trigger the access node to run a loopback test on the local loop specified by an ANCP neighbor.	<i>request ancp oam neighbor</i>
Display ANCP class-of-service information.	<i>show ancp cos</i>
Display ANCP neighbor information.	<i>show ancp neighbor</i>
Display ANCP subscriber information.	<i>show ancp subscriber</i>
Display a summary of the counts and states for all ANCP neighbors and subscribers.	<i>show ancp summary</i>
Display a summary of the counts and states for all ANCP neighbors and of the neighbor's subscribers when you specify a particular neighbor.	<i>show ancp summary neighbor</i>
Display a summary of the counts and states for all ANCP subscribers.	<i>show ancp summary subscriber</i>



NOTE: For information about how to configure ANCP, see the *Junos Subscriber Access Configuration Guide*.

BFD Operational Mode Commands

Table 23 on page 324 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot Bidirectional Forwarding Detection (BFD) sessions. Commands are listed in alphabetical order.

Table 23: BFD Operational Mode Commands

Task	Command
Clear BFD parameters.	<i>clear bfd adaptation</i>
Clear BFD sessions.	<i>clear bfd session</i>
Display BFD session statistics.	<i>show bfd session</i>



NOTE: The protocol client for which the BFD session is active can be either IS-IS or OSPF.



NOTE: For information about how to configure BFD, see the *Junos Routing Protocols Configuration Guide*.

BGP Operational Mode Commands

Table 24 on page 325 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the Border Gateway Protocol (BGP). Commands are listed in alphabetical order.

Table 24: BGP Operational Mode Commands

Task	Command
Remove damping information.	<code>clear bgp damping</code>
Remove entries from the neighbor database.	<code>clear bgp neighbor</code>
Request BGP to refresh routes.	<code>clear bgp table</code>
Display information about the BGP Monitoring Protocol.	<code>show bgp bmp</code>
Display entries in the BGP group database.	<code>show bgp group</code>
Display traffic statistics for BGP groups.	<code>show bgp group traffic-statistics</code>
Display entries in the BGP neighbor database.	<code>show bgp neighbor</code>
Display the BGP state replication status for nonstop active routing-enabled devices.	<code>show bgp replication</code>
Display BGP summary information.	<code>show bgp summary</code>
Display BGP damping parameters.	<code>show policy damping</code>



NOTE: For more BGP-related commands, such as `show route protocol`, `show route instance`, and `show route table`, see “[Protocol-Independent Routing Operational Mode Commands](#)” on page 335.



NOTE: For information about how to configure BGP, see the *Junos Routing Protocols Configuration Guide*.

ES-IS Operational Mode Commands

Table 25 on page 326 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the End System-to-Intermediate System (ES-IS) protocol. Commands are listed in alphabetical order.

Table 25: ES-IS Operational Mode Commands

Task	Command
Clear ES-IS adjacencies.	<i>clear esis adjacency</i>
Clear ES-IS statistics for packets sent or received.	<i>clear esis statistics</i>
Display ES-IS adjacencies.	<i>show esis adjacency</i>
Display ES-IS interfaces.	<i>show esis interface</i>
Display ES-IS statistics for packets sent or received.	<i>show esis statistics</i>



NOTE: ES-IS is supported only on J Series routers. For information about how to configure ES-IS, see the *J Series Services Router Basic LAN and WAN Access Configuration Guide* or the *Junos OS Routing Protocols Library for Routing Devices*.

IP Multicast Operational Mode Commands

Table 26 on page 326 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot IP multicast. In the table, the commands are listed in alphabetical order.

Table 26: IP Multicast Operational Mode Commands

Task	Command
Clear Automatic Multicast Tunneling (AMT) protocol statistics.	<i>clear amt statistics</i>
Clear Automatic Multicast Tunneling (AMT) protocol state.	<i>clear amt tunnel</i>
Clear Internet Group Management Protocol (IGMP) group members.	<i>clear igmp membership</i>
Clear IGMP snooping membership information.	<i>clear igmp snooping membership</i>
Clear IGMP snooping statistics.	<i>clear igmp snooping statistics</i>

Table 26: IP Multicast Operational Mode Commands (*continued*)

Task	Command
Clear IGMP statistics.	<i>clear igmp statistics</i>
Clear Multicast Listener Discovery (MLD) group members.	<i>clear mld membership</i>
Clear MLD statistics.	<i>clear mld statistics</i>
Clear Multicast Source Discovery Protocol (MSDP) source active cache.	<i>clear msdp cache</i>
Clear MSDP statistics.	<i>clear msdp statistics</i>
Clear multicast bandwidth admissions.	<i>clear multicast bandwidth-admission</i>
Clear IP multicast forwarding cache entries.	<i>clear multicast forwarding-cache</i>
Clear multicast scope.	<i>clear multicast scope</i>
Clear multicast sessions.	<i>clear multicast sessions</i>
Clear multicast snooping statistics.	<i>clear multicast snooping statistics</i>
Clear multicast statistics.	<i>clear multicast statistics</i>
Clear Pragmatic General Multicast (PGM) negative acknowledgments (NAKs).	<i>clear pgm negative-acknowledgments</i>
Clear PGM source-path messages.	<i>clear pgm source-path-messages</i>
Clear PGM statistics.	<i>clear pgm statistics</i>
Clear the Protocol Independent Multicast (PIM) join and prune states.	<i>clear pim join</i>
Redistribute PIM joins among available links.	<i>clear pim join-distribution</i>
Clear PIM register message counters.	<i>clear pim register</i>
Clear PIM snooping joins.	<i>clear pim snooping join</i>
Clear Pim snooping statistics.	<i>clear pim snooping statistics</i>
Clear PIM statistics.	<i>clear pim statistics</i>
Rebalance multicast tunnel (MT) interfaces.	<i>request pim multicast-tunnel rebalance</i>
Display Automatic Multicast Tunneling (AMT) protocol tunnel statistics.	<i>show amt statistics</i>

Table 26: IP Multicast Operational Mode Commands (*continued*)

Task	Command
Display summary information about the Automatic Multicast Tunneling (AMT) protocol.	<i>show amt summary</i>
Display information about the Automatic Multicast Tunneling (AMT) dynamic tunnels.	<i>show amt tunnel</i>
Display the status of interfaces on which Distance Vector Multicast Routing Protocol (DVMRP) is configured.	<i>show dvmrp interfaces</i>
Display DVMRP neighbors.	<i>show dvmrp neighbors</i>
Display DVMRP prefixes.	<i>show dvmrp prefix</i>
Display DVMRP prunes.	<i>show dvmrp prunes</i>
Display members of IGMP groups.	<i>show igmp group</i>
Display members of IGMP groups by interface.	<i>show igmp interface</i>
Display IGMP snooping interface information.	<i>show igmp snooping interface</i>
Display IGMP snooping membership information.	<i>show igmp snooping membership</i>
Display IGMP snooping statistics.	<i>show igmp snooping statistics</i>
Display IGMP statistics.	<i>show igmp statistics</i>
Display members of MLD groups.	<i>show mld group</i>
Display members of MLD groups by interface.	<i>show mld interface</i>
Display MLD statistics.	<i>show mld statistics</i>
Display MSDP peers.	<i>show msdp</i>
Display multicast sources learned from MSDP.	<i>show msdp source</i>
Display the MSDP source-active cache.	<i>show msdp source-active</i>
Display MSDP statistics.	<i>show msdp statistics</i>
Display backup PE router group information when ingress PE redundancy is configured.	<i>show multicast backup-pe-groups</i>
Display configuration information about IP multicast flow maps.	<i>show multicast flow-map</i>

Table 26: IP Multicast Operational Mode Commands (*continued*)

Task	Command
Display IP multicast forwarding cache statistics.	<i>show multicast forwarding-cache statistics</i>
Display multicast interface bandwidth information.	<i>show multicast interface</i>
Display multicast network configuration.	<i>show multicast minfo</i>
Display entries in the multicast next-hop table.	<i>show multicast next-hops</i>
Display configuration information about PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy.	<i>show multicast pim-to-igmp-proxy</i>
Display configuration information about PIM-to-MLD message translation, also known as PIM-to-MLD proxy.	<i>show multicast pim-to-mld-proxy</i>
Display entries in the multicast forwarding cache.	<i>show multicast route</i>
Display multicast reverse-path-forwarding calculations.	<i>show multicast rpf</i>
Display administratively scoped addresses.	<i>show multicast scope</i>
Display multicast snooping next-hops	<i>show multicast snooping next-hops</i>
Display announced multicast sessions.	<i>show multicast sessions</i>
Display multicast snooping route.	<i>show multicast snooping route</i>
Display multicast snooping statistics.	<i>show multicast snooping statistics</i>
Display multicast statistics.	<i>show multicast statistics</i>
Display most active multicast groups.	<i>show multicast usage</i>
Display sent or received NAKs.	<i>show pgm negative-acknowledgments</i>
Display PGM source-path messages.	<i>show pgm source-path-messages</i>
Display PGM statistics.	<i>show pgm statistics</i>
Display bootstrap routers.	<i>show pim bootstrap</i>
Display the status of interfaces on which PIM is configured.	<i>show pim interfaces</i>
Display PIM (*,RP) join and prune states.	<i>show pim join</i>
Display PIM data-driven multicast distribution trees (MDTs).	<i>show pim mdt</i>

Table 26: IP Multicast Operational Mode Commands (*continued*)

Task	Command
Display the information cached from multicast distribution tree (MDT) join TLV packets received by all PE routers in a PIM-enabled VPN routing and forwarding (VRF)-instance.	<i>show pim mdt data-mdt-joins</i>
Display the maximum number configured and the currently active data multicast distribution trees (MDTs) for a specific VPN routing and forwarding (VRF) instance.	<i>show pim mdt data-mdt-limit</i>
Display PIM neighbors.	<i>show pim neighbors</i>
Display rendezvous points.	<i>show pim rps</i>
Display information about PIM snooping interfaces.	<i>show pim snooping interfaces</i>
Display PIM snooping joins.	<i>show pim snooping join</i>
Display information about PIM snooping neighbors.	<i>show pim snooping neighbors</i>
Display PIM snooping statistics.	<i>show pim snooping statistics</i>
Display PIM source RPF state.	<i>show pim source</i>
Display PIM statistics.	<i>show pim statistics</i>
Display Session Announcement Protocol (SAP) addresses.	<i>show sap listen</i>
Test MSDP peers.	<i>test msdp</i>



NOTE: For information about the *mtrace* commands used to monitor IP multicast traffic in real time, see the *Junos System Basics and Services Command Reference*. For information about how to configure IP multicast, see the *Junos Multicast Protocols Configuration Guide*.

IPv6 Operational Mode Commands

Table 27 on page 331 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot Internet Protocol version 6 (IPv6). Commands are listed in alphabetical order.

Table 27: IPv6 Operational Mode Commands

Task	Command
Clear IPv6 neighbor cache information.	<i>clear ipv6 neighbors</i>
Clear IPv6 router advertisement counters.	<i>clear ipv6 router-advertisement</i>
Display neighbor discovery information.	<i>show ipv6 neighbors</i>
Display router advertisement information.	<i>show ipv6 router-advertisement</i>



NOTE: For information about how to configure IPv6 parameters, see the *Junos OS Routing Protocols Library for Routing Devices*.

IS-IS Operational Mode Commands

Table 28 on page 331 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the Intermediate System-to-Intermediate System (IS-IS) protocol. Commands are listed in alphabetical order.

Table 28: IS-IS Operational Mode Commands

Task	Command
Remove adjacencies.	<i>clear isis adjacency</i>
Remove database entries.	<i>clear isis database</i>
Reset IS-IS dynamic overload bit.	<i>clear isis overload</i>
Set IS-IS traffic statistics to zero.	<i>clear isis statistics</i>
Display adjacent routers.	<i>show isis adjacency</i>
Display authentication statistics.	<i>show isis authentication</i>
Display information about the level of backup coverage available for protected routes.	<i>show isis backup coverage</i>
Display information about MPLS LSPs designated as backup paths.	<i>show isis backup label-switched-path</i>
Display SPF calculations for backup paths.	<i>show isis backup spf results</i>
Display IS-IS context identifier information.	<i>show isis context-identifier</i>
Display database entries.	<i>show isis database</i>

Table 28: IS-IS Operational Mode Commands (*continued*)

Task	Command
Display hostname mapping.	<i>show isis hostname</i>
Display the status of interfaces on which IS-IS is running.	<i>show isis interface</i>
Display IS-IS overview information.	<i>show isis overview</i>
Display IS-IS routing table entries.	<i>show isis route</i>
Display SPF calculations.	<i>show isis spf</i>
Display IS-IS traffic statistics.	<i>show isis statistics</i>



NOTE: For more IS-IS-related commands, such as `show route protocol`, `show route instance`, and `show route table`, see “[Protocol-Independent Routing Operational Mode Commands](#)” on page 335. For information about monitoring Bidirectional Forwarding Detection (BFD) sessions for IS-IS clients, see “[BFD Operational Mode Commands](#)” on page 324. For information about how to configure IS-IS, see the *Junos Routing Protocols Configuration Guide*.



NOTE: In IS-IS command output, the CLI displays the system ID numerically by default. To display the hostname instead, include the `static-host-mapping` statement at the `[edit system]` hierarchy level of the configuration.

LLDP Operational Mode Commands

Table 29 on page 332 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the Link Layer Discovery Protocol (LLDP) protocol. Commands are listed in alphabetical order.

Table 29: LLDP Operational Mode Commands

Task	Command
Clear LLDP neighbor information.	<i>clear lldp neighbor</i>
Clear LLDP statistics.	<i>clear lldp statistics</i>
Display basic LLDP information.	<i>show lldp</i>
Display LLDP local information.	<i>show lldp local-information</i>
Display LLDP neighbor information.	<i>show lldp neighbors</i>

Table 29: LLDP Operational Mode Commands (*continued*)

Task	Command
Display LLDP remote global statistics.	<i>show lldp remote-global-statistics</i>
Display LLDP statistics.	<i>show lldp statistics</i>

**Related
Documentation**

- *LLDP Overview*
- *Configuring LLDP*
- *Tracing LLDP Operations*
- *Example: Configuring LLDP*

MVRP Operational Mode Commands

Table 30 on page 333 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the Multiple VLAN Registration Protocol (MVRP). Commands are listed in alphabetical order.

Table 30: MVRP Operational Mode Commands

Task	Command
Display Multiple VLAN Registration Protocol (MVRP) configuration information.	<i>show mvrp</i>
Display Multiple VLAN Registration Protocol (MVRP) applicant state information.	<i>show mvrp applicant-state</i>
Display all Virtual LANs (VLANs) that have been created dynamically using Multiple VLAN Registration Protocol (MVRP) on the router.	<i>show mvrp dynamic-vlan-memberships</i>
Display Multiple VLAN Registration Protocol (MVRP) interface-specific information.	<i>show mvrp interface</i>
Display Multiple VLAN Registration Protocol (MVRP) registration state information.	<i>show mvrp registration-state</i>
Display Multiple VLAN Registration Protocol (MVRP) statistics in the form of Multiple Registration Protocol data unit (MRPDU) messages.	<i>show mvrp statistics</i>

OSPF Operational Mode Commands

Table 31 on page 334 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the Open Shortest Path First (OSPF) protocol. Commands are listed in alphabetical order.

Table 31: OSPF Operational Mode Commands

Task	Command
Clear link-state database entries.	<i>clear (ospf ospf3) database</i>
Clear OSPF input and output statistics.	<i>clear (ospf ospf3) io-statistics</i>
Tear down neighbor connections.	<i>clear (ospf ospf3) neighbor</i>
Clear the OSPF overload bit.	<i>clear (ospf ospf3) overload</i>
Clear OSPF statistics.	<i>clear (ospf ospf3) statistics</i>
Display information about the level of backup coverage available for OSPF nodes and prefixes.	<i>show (ospf ospf3) backup coverage</i>
Display information about MPLS label-switched-paths (LSPs) designated as backup routes for OSPF routes.	<i>show (ospf ospf3) backup lsp</i>
Display the neighbors through which direct next hops for the backup paths are available.	<i>show (ospf ospf3) backup neighbor</i>
Display information about OSPF shortest-path-first calculations for backup paths.	<i>show (ospf ospf3) backup spf</i>
Display context identifier information processed and advertised by OSPF for egress protection.	<i>show ospf context-identifier</i>
Display link-state database entries for OSPFv2.	<i>show ospf database</i>
Display link-state database entries for OSPFv3.	<i>show ospf3 database</i>
Display OSPF interface status.	<i>show (ospf ospf3) interface</i>
Display OSPF input and output statistics.	<i>show (ospf ospf3) io-statistics</i>
Display the SPF log.	<i>show (ospf ospf3) log</i>
Display adjacent routers.	<i>show (ospf ospf3) neighbor</i>
Display overview statistics.	<i>show (ospf ospf3) overview</i>
Display OSPF routing table entries.	<i>show (ospf ospf3) route</i>
Display OSPF statistics.	<i>show (ospf ospf3) statistics</i>



NOTE: For more OSPF-related commands, such as `show route protocol`, `show route instance`, and `show route table`, see “[Protocol-Independent Routing Operational Mode Commands](#)” on page 335. For information about monitoring Bidirectional Forwarding Detection (BFD) sessions for OSPF clients, see “[BFD Operational Mode Commands](#)” on page 324. For information about how to configure OSPF, see the *Junos Routing Protocols Configuration Guide*.

Protocol-Independent Routing Operational Mode Commands

Table 32 on page 335 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot protocol-independent routing properties. Commands are listed in alphabetical order.



NOTE: The `show route` command has a lengthy set of options. Therefore, this chapter describes each option as a separate command. You can, however, combine several options and issue them as single `show route` command. For example, `show route ccc exact`.

The exceptions to this convention are the `show as-path`, `show route damping`, `show route export`, `show route export-vrf-target`, `show route forwarding-table`, `show route instance`, and `show route martians` commands, which cannot be used with any other options (other than level of output options, such as `detail` and `extensive`).

The `show route flow validation` command can only be used with the `table` option.

Table 32: Protocol-Independent Routing Operational Mode Commands

Task	Command
Display known autonomous system (AS) paths.	<code>show as-path</code>
Display AS path domain information.	<code>show as-path domain</code>
Display AS path summary information.	<code>show as-path summary</code>
Display information about the entries in the routing tables.	<code>show route</code>
Display routes that are currently active.	<code>show route active-path</code>
Display routes transmitted by a particular routing protocol.	<code>show route advertising-protocol</code>
Display all information about all routes.	<code>show route all</code>
Display routes containing a specified AS path.	<code>show route aspath-regex</code>

Table 32: Protocol-Independent Routing Operational Mode Commands (*continued*)

Task	Command
Display the best route to the specified address or range of addresses.	<i>show route best</i>
Display brief information about the entries in the routing table.	<i>show route brief</i>
Display circuit cross-connect (CCC) entries in the Multiprotocol Link Switching (MPLS) routing table.	<i>show route ccc</i>
Display routes containing members of a specified BGP community.	<i>show route community</i>
Display routes containing members of a specified BGP community based on a particular community name.	<i>show route community-name</i>
Display routes that have been damped.	<i>show route damping</i>
Display detailed information about the entries in the routing table.	<i>show route detail</i>
Display routes that exactly match the specified address or range of addresses.	<i>show route exact</i>
Display list of instances or routing tables that are importers or exporters of routes.	<i>show route export</i>
Display target communities for which autoexport is currently distributing routes.	<i>show route export vrf-target</i>
Display extensive information about the entries in the routing table.	<i>show route extensive</i>
Display the best route to an address.	<i>show route flow validation</i>
Display the Junos OS forwarding table.	<i>show route forwarding-table</i>
Display information about the interfaces in the Junos OS forwarding table.	<i>show route forwarding-table interface-name</i>
Display hidden routes only.	<i>show route hidden</i>
Display routes that are not preferred.	<i>show route inactive-path</i>
Display routes that are currently inactive.	<i>show route inactive-prefix</i>
Display routing instance information.	<i>show route instance</i>
Display routes corresponding to a specified label value.	<i>show route label</i>

Table 32: Protocol-Independent Routing Operational Mode Commands (*continued*)

Task	Command
Display routes that form a label-switched path.	<i>show route label-switched-path</i>
Display route localization information.	<i>show route localization</i>
Display information about martian addresses.	<i>show route martians</i>
Display routes filtered by a regular expression.	<i>show route match-prefix</i>
Display routes that contain the specified next hop.	<i>show route next-hop</i>
Display routes not associated with any BGP community.	<i>show route no-community</i>
Display routes exiting the router through the specified interface.	<i>show route output</i>
Display routes learned by the specified protocol.	<i>show route protocol</i>
Display routes in a range of destination prefixes.	<i>show route range</i>
Display routes received by a particular routing protocol.	<i>show route receive-protocol</i>
Display entries in the next-hop resolution database.	<i>show route resolution</i>
Display routes learned from snooping.	<i>show route snooping</i>
Display routes learned from the specified source.	<i>show route source-gateway</i>
Display statistics about the routes in all routing tables.	<i>show route summary</i>
Display routes in a particular routing table.	<i>show route table</i>
Display high-level summary of routing table information.	<i>show route terse</i>



NOTE: For information about how to configure protocol-independent features, see the *Junos Routing Protocols Configuration Guide* and the *Junos Policy Framework Configuration Guide*.

RIP Operational Mode Commands

Table 33 on page 338 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the Routing Information Protocol (RIP). Commands are listed in alphabetical order.

Table 33: RIP Operational Mode Commands

Task	Command
Clear RIP general statistics.	<i>clear rip general-statistics</i>
Clear RIP statistics.	<i>clear rip statistics</i>
Display brief RIP statistics.	<i>show rip general-statistics</i>
Display information about RIP neighbors.	<i>show rip neighbor</i>
Display RIP statistics about messages sent and received on an interface, as well as information received through advertisements from other routers.	<i>show rip statistics</i>



NOTE: For more RIP-related commands, such as **show route protocol**, **show route instance**, and **show route table**, see [“Protocol-Independent Routing Operational Mode Commands” on page 335](#).

For information about how to configure RIP, see the *Junos Routing Protocols Configuration Guide*.

RIPng Operational Mode Commands

[Table 34 on page 338](#) summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the Routing Information Protocol next generation (RIPng). Commands are listed in alphabetical order.

Table 34: RIPng Operational Mode Commands

Task	Command
Clear general statistics.	<i>clear ripng general-statistics</i>
Clear statistics.	<i>clear ripng statistics</i>
Display general statistics.	<i>show ripng general-statistics</i>
Display RIPng neighbors.	<i>show ripng neighbor</i>
Display statistics.	<i>show ripng statistics</i>



NOTE: For more RIPng-related commands, such as `show route protocol`, `show route instance`, and `show route table`, see “[Protocol-Independent Routing Operational Mode Commands](#)” on page 335.

For information about how to configure RIPng, see the *Junos Routing Protocols Configuration Guide*.

Firewall Filter Operational Mode Commands

Table 35 on page 339 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot firewall filters. Commands are listed in alphabetical order.

Table 35: Firewall Filter Operational Mode Commands

Task	Command
Clear firewall filter counters.	<code>clear firewall</code>
Operational statistics for firewall filters.	<code>show firewall</code>
Version number of installed firewall filters.	<code>show firewall filter version</code>
Firewall filter log information.	<code>show firewall log</code>
Prefix-action statistics for firewall filters.	<code>show firewall prefix-action-stats</code>
Names of configured filter templates in use by dynamic subscribers and number of times each template is referenced.	<code>show firewall templates-in-use</code>
Counters for policers.	<code>show policer</code>



NOTE: For information about how to configure firewall filters, see the *Junos Policy Framework Configuration Guide*.

For information about the related operational mode commands, `show interfaces filters` and `show interfaces policers`, see the *Junos Interfaces Command Reference*.

Layer 2 Bridging and Switching Operational Mode Commands

Table 36 on page 340 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot Layer 2 bridging and switching. Commands are listed in alphabetical order.

Table 36: Layer 2 Bridging and Switching Operational Mode Commands

Task	Command
Clear learned Layer 2 address information from the media access control (MAC) address table.	<i>clear bridge mac-table</i>
Clear bridge protocol data unit (BPDU) error on interface due to possible bridge spanning tree protocol (STP) loop.	<i>clear error bpdv</i>
Clear a MAC rewrite error condition for Layer 2 protocol tunneling.	<i>clear error mac-rewrite</i>
Display bridge domain information.	<i>show bridge domain</i>
Display bridging flooding information.	<i>show bridge flood</i>
Display learned Layer 2 MAC address information.	<i>show bridge mac-table</i>
Display bridge statistics.	<i>show bridge statistics</i>
Display Layer 2 learning process-related information.	<i>show l2-learning global-information</i>
(MX Series routers only) Display the total number of dynamic and static MAC addresses learned for the entire router.	<i>show l2-learning global-mac-count</i>
Display configured Layer 2 routing instances.	<i>show l2-learning instance</i>
Display configured Layer 2 interfaces.	<i>show l2-learning interface</i>
Display Layer 2 interfaces.	<i>show mac-rewrite interface</i>

VPN Operational Mode Commands

[Table 37 on page 340](#) summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot Layer 2 circuits, Layer 2 virtual private networks (VPNs), virtual private LAN service (VPLS), and Layer 3 VPNs. Commands are listed in alphabetical order.

Table 37: Layer 2 Circuit, Layer 2 VPN, and VPLS Operational Mode Commands

Task	Command
Clear MAC address entries from the VPLS table.	<i>clear vpls mac-address</i>
Clear MAC addresses from the VPLS table.	<i>clear vpls mac-table</i>
Manually trigger a switch from the active pseudowire to the redundant pseudowire.	<i>request l2circuit-switchover</i>

Table 37: Layer 2 Circuit, Layer 2 VPN, and VPLS Operational Mode Commands (*continued*)

Task	Command
Display Layer 3 dynamic tunnel database information.	<code>show dynamic-tunnels database</code>
Display Host fast reroute (HFRR) profile information.	<code>show hfrr profiles</code>
Display ingress replication provider tunnel information.	<code>show ingress-replication mvpn</code>
Display Layer 2 circuit information.	<code>show l2circuit connections</code>
Display Layer 2 VPN information.	<code>show l2vpn connections</code>
Display multicast VPN c-multicast route information.	<code>show mvpn c-multicast</code>
Display multicast VPN instance information.	<code>show mvpn instance</code>
Display multicast VPN neighbor information.	<code>show mvpn neighbor</code>
Display virtual private LAN service (VPLS) information.	<code>show vpls connections</code>
Display the pending events in the level 2 address learning process (l2ald) routing socket code (rtsock) update queue.	<code>show vpls flood event-queue</code>
Display VPLS information related to the level 2 address learning process for the specified routing instance.	<code>show vpls flood instance</code>
Display VPLS route information related to the level 2 address learning process.	<code>show vpls flood route</code>
Display learned VPLS MAC address information.	<code>show vpls mac-table</code>
Display VPLS statistics.	<code>show vpls statistics</code>



NOTE: For information about how to configure Layer 2 circuits, Layer 2 VPNs, VPLS, and Layer 3 VPNs, see the *Junos VPNs Configuration Guide*.

PART 4

Troubleshooting

- [Interface Diagnostics on page 345](#)

CHAPTER 7

Interface Diagnostics

- [Interface Diagnostics on page 345](#)

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 345](#)
- [Interface Diagnostics on page 347](#)

Configuring Loopback Testing

Loopback testing allows you to verify the connectivity of a circuit. You can configure any of the following interfaces to execute a loopback test: Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, E1, E3, NxDS0, serial, SONET/SDH, T1, and T3.

The physical path of a network data circuit usually consists of segments interconnected by devices that repeat and regenerate the transmission signal. The transmit path on one device connects to the receive path on the next device. If a circuit fault occurs in the form of a line break or a signal corruption, you can isolate the problem by using a loopback test. Loopback tests allow you to isolate segments of the circuit and test them separately.

To do this, configure a *line loopback* on one of the routers. Instead of transmitting the signal toward the far-end device, the line loopback sends the signal back to the originating router. If the originating router receives back its own data link layer packets, you have verified that the problem is beyond the originating router. Next, configure a line loopback farther away from the local router. If this originating router does not receive its own data link layer packets, you can assume the problem is on one of the segments between the local router and the remote router's interface card. In this case, the next troubleshooting step is to configure a line loopback closer to the local router to find the source of the problem.

There are several types of loopback testing supported by the Junos OS, as follows:

- DCE local—Loops packets back on the local DCE.
- DCE remote—Loops packets back on the remote DCE.

- **Local**—Useful for troubleshooting physical PIC errors. Configuring local loopback on an interface allows transmission of packets to the channel service unit (CSU) and then to the circuit toward the far-end device. The interface receives its own transmission, which includes data and timing information, on the local router's PIC. The data received from the CSU is ignored. To test a local loopback, issue the **show interfaces *interface-name*** command. If PPP keepalives transmitted on the interface are received by the PIC, the **Device Flags** field contains the output **Loop-Detected**.
- **Payload**—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A payload loopback loops data only (without clocking information) on the remote router's PIC. With payload loopback, overhead is recalculated.
- **Remote**—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A remote loopback loops packets, including both data and timing information, back on the remote router's interface card. A router at one end of the circuit initiates a remote loopback toward its remote partner. When you configure a remote loopback, the packets received from the physical circuit and CSU are received by the interface. Those packets are then retransmitted by the PIC back toward the CSU and the circuit. This loopback tests all the intermediate transmission segments.

Table 38 on page 346 shows the loopback modes supported on the various interface types.

Table 38: Loopback Modes by Interface Type

Interface	Loopback Modes	Usage Guidelines
Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet	Local	<i>Configuring Ethernet Loopback Capability</i>
Circuit Emulation E1	Local and remote	<i>Configuring E1 Loopback Capability</i>
Circuit Emulation T1	Local and remote	<i>Configuring T1 Loopback Capability</i>
E1 and E3	Local and remote	<i>Configuring E1 Loopback Capability and Configuring E3 Loopback Capability</i>
NxDSO	Payload	<i>Configuring Channelized E1 IQ and IQE Interfaces, Configuring T1 and NxDSO Interfaces, Configuring Channelized OC12/STM4 IQ and IQE Interfaces (SONET Mode), Configuring Channelized STM1 IQ and IQE Interfaces, and Configuring Channelized T3 IQ Interfaces</i>
Serial (V.35 and X.21)	Local and remote	<i>Configuring Serial Loopback Capability</i>
Serial (EIA-530)	DCE local, DCE remote, local, and remote	<i>Configuring Serial Loopback Capability</i>
SONET/SDH	Local and remote	<i>Configuring SONET/SDH Loopback Capability</i>

Table 38: Loopback Modes by Interface Type (*continued*)

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

To configure loopback testing, include the **loopback** statement:

loopback mode;

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* aggregated-ether-options]
- [edit interfaces *interface-name* ds0-options]
- [edit interfaces *interface-name* e1-options]
- [edit interfaces *interface-name* e3-options]
- [edit interfaces *interface-name* fastether-options]
- [edit interfaces *interface-name* gigether-options]
- [edit interfaces *interface-name* serial-options]
- [edit interfaces *interface-name* sonet-options]
- [edit interfaces *interface-name* t1-options]
- [edit interfaces *interface-name* t3-options]

Interface Diagnostics

BERT allows you to troubleshoot problems by checking the quality of links. You can configure any of the following interfaces to execute a BERT when the interface receives a request to run this test: E1, E3, T1, T3; the channelized DS3, OC3, OC12, and STM1 interfaces; and the channelized DS3 IQ, E1 IQ, and OC12 IQ interfaces.

A BERT test requires a line loop to be in place on either the transmission devices or the far-end router. The local router generates a known bit pattern and sends it out the transmit path. The received pattern is then verified against the sent pattern. The higher the bit error rate of the received pattern, the worse the noise is on the physical circuit. As you move the position of the line loop increasingly downstream toward the far-end router, you can isolate the troubled portion of the link.

To configure BERT, you must configure the duration of the test, the bit pattern to send on the transmit path, and the error rate to monitor when the inbound pattern is received.

To configure the duration of the test, the pattern to send in the bit stream, and the error rate to include in the bit stream, include the **bert-period**, **bert-algorithm**, and **bert-error-rate** statements, respectively, at the [edit interfaces *interface-name* *interface-type*-options] hierarchy level:

```
[edit interfaces interface-name interface-type-options]
bert-algorithm algorithm;
bert-error-rate rate;
bert-period seconds;
```

By default, the BERT period is 10 seconds. You can configure the BERT period to last from 1 through 239 seconds on some PICs and from 1 through 240 seconds on other PICs.

rate is the bit error rate. This can be an integer from 0 through 7, which corresponds to a bit error rate from 10^{-0} (1 error per bit) to 10^{-7} (1 error per 10 million bits).

algorithm is the pattern to send in the bit stream. For a list of supported algorithms, enter a ? after the **bert-algorithm** statement; for example:

```
[edit interfaces t1-0/0/0 t1-options]
user@host# set bert-algorithm ?
Possible completions:
pseudo-2e11-o152      Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151      Pattern is 2^15 - 1 (per 0.152 standard)
pseudo-2e20-o151      Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e20-o153      Pattern is 2^20 - 1 (per 0.153 standard)
...
```

For specific hierarchy information, see the individual interface types.



NOTE: The 4-port E1 PIC supports only the following algorithms:

pseudo-2e11-o152	Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151	Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151	Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e23-o151	Pattern is 2^23 (per 0.151 standard)

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.



NOTE: The 12-port T1/E1 Circuit Emulation (CE) PIC supports only the following algorithms:

```
all-ones-repeating    Repeating one bits
all-zeros-repeating   Repeating zero bits
alternating-double-ones-zeros Alternating pairs of ones and zeros
alternating-ones-zeros Alternating ones and zeros
pseudo-2e11-o152      Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151      Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151      Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e7            Pattern is 2^7 - 1
pseudo-2e9-o153       Pattern is 2^9 - 1 (per 0.153 standard)
repeating-1-in-4       1 bit in 4 is set
repeating-1-in-8       1 bit in 8 is set
repeating-3-in-24      3 bits in 24 are set
```

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.



NOTE: The IQE PICs support only the following algorithms:

```
all-ones-repeating    Repeating one bits
all-zeros-repeating   Repeating zero bits
alternating-double-ones-zeros Alternating pairs of ones and zeros
alternating-ones-zeros Alternating ones and zeros
pseudo-2e9-o153       Pattern is 2^9 - 1 (per 0.153 (511 type) standard)
pseudo-2e11-o152      Pattern is 2^11 - 1 (per 0.152 and 0.153 (2047 type)
standards)
pseudo-2e15-o151      Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151      Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e20-o153      Pattern is 2^20 - 1 (per 0.153 standard)
pseudo-2e23-o151      Pattern is 2^23 - 1 (per 0.151 standard)
repeating-1-in-4       1 bit in 4 is set
repeating-1-in-8       1 bit in 8 is set
repeating-3-in-24      3 bits in 24 are set
```

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.



NOTE: BERT is supported on the PDH interfaces of the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP and the DS3/E3 MIC. The following BERT algorithms are supported:

all-ones-repeating	Repeating one bits
all-zeros-repeating	Repeating zero bits
alternating-double-ones-zeros	Alternating pairs of ones and zeros
alternating-ones-zeros	Alternating ones and zeros
repeating-1-in-4	1 bit in 4 is set
repeating-1-in-8	1 bit in 8 is set
repeating-3-in-24	3 bits in 24 are set
pseudo-2e9-o153	Pattern is $2^9 - 1$ (per 0.153 standard)
pseudo-2e11-o152	Pattern is $2^{11} - 1$ (per 0.152 standard)
pseudo-2e15-o151	Pattern is $2^{15} - 1$ (per 0.151 standard)
pseudo-2e20-o151	Pattern is $2^{20} - 1$ (per 0.151 standard)
pseudo-2e20-o153	Pattern is $2^{20} - 1$ (per 0.153 standard)
pseudo-2e23-o151	Pattern is $2^{23} - 1$ (per 0.151 standard)

Table 39 on page 350 shows the BERT capabilities for various interface types.

Table 39: 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 355](#)

Index

Symbols

#, comments in configuration statements.....	xvi
(), in syntax descriptions.....	xvi
128-bit IPv6 address.....	38
32-bit IPv4 address.....	36
< >, in syntax descriptions.....	xvi
[], in configuration statements.....	xvi
{ }, in configuration statements.....	xvi
(pipe), in syntax descriptions.....	xvi

A

accounting statement.....	99
usage guidelines.....	24
address statement.....	96
usage guidelines.....	35, 43
aggregate statement	
hierarchical policer.....	98
arp option	
policers.....	11
arp statement.....	100
AS paths	
distribution of, displaying.....	191
domain information, displaying.....	195
summary of, displaying.....	197
Automatic Multicast Tunneling See AMT	
autonomous system paths See AS paths	

B

bandwidth-limit statement	
hierarchical policer.....	101
BERT	
configuring interface diagnostics.....	347
bert-algorithm statement	
usage guidelines.....	347
bert-error-rate statement	
usage guidelines.....	347
bert-period statement	
usage guidelines.....	347
bit error rate test See BERT	
BOOTP	
accepting packets.....	56

borrower interface	
unnumbered Ethernet or demux.....	44
braces, in configuration statements.....	xvi
brackets	
angle, in syntax descriptions.....	xvi
square, in configuration statements.....	xvi
broadcast statement.....	102
usage guidelines.....	35
bundle statement.....	103
burst-size-limit statement	
hierarchical policer.....	104

C

cbr statement.....	105
clear firewall command.....	186
clear vlpls mac-address command.....	188
clear vlpls mac-table command.....	189
comments, in configuration statements.....	xvi
conventions	
text and syntax.....	xv
curly braces, in configuration statements.....	xvi
customer support.....	xvii
contacting JTAC.....	xvii

D

default router addresses.....	49
demux interfaces	
unnumbered.....	44
demux0 statement	
dynamic IP demux interface.....	106
destination statement	
tunnels.....	108
usage guidelines.....	35, 42
destination-class usage	
example configuration.....	28
destination-class-usage statement.....	109
usage guidelines.....	24
destination-profile	
usage guidelines.....	43
destination-profile statement.....	109
usage guidelines.....	35
DHCP	
accepting.....	56
Distance Vector Multicast Routing Protocol See DVMRP	
documentation	
comments on.....	xvii
donor interface	
unnumbered Ethernet or demux.....	44

Dynamic Host Configuration Protocol <i>See</i> DHCP	
dynamic profiles statements	
dynamic-profiles.....	110
interfaces.....	133
unnumbered-address.....	177
dynamic subscribers	
interfaces statement.....	133
dynamic tunnel database, displaying.....	199
dynamic-profiles	
interfaces statement.....	133
dynamic IP demux.....	133
dynamic-profiles statement.....	110
E	
End System-to-Intermediate System <i>See</i> ES-IS	
epd-threshold statement	
logical interface.....	117
Ethernet interfaces	
status information, displaying	
Fast Ethernet.....	241
Gigabit Ethernet.....	216, 257
unnumbered.....	44
preferred source address.....	150
eui-64 statement.....	118
usage guidelines.....	35
F	
fail-filter statement.....	157
usage guidelines.....	51
family bridge	
VLAN ID list.....	182
VLAN IDs.....	182
family statement.....	121
dynamic profiles.....	119
usage guidelines.....	32
Fast Ethernet interfaces	
status information, displaying.....	241
fast-aps-switch statement.....	125
filter statement.....	126
usage guidelines.....	19
firewall	
hierarchical-policer.....	128
statistics	
displaying.....	202
firewall filters	
applying.....	19
example configuration.....	22
log information, displaying.....	209
logical interfaces.....	19
policed packets, displaying.....	292
statistics	
clearing.....	186
displaying.....	212
font conventions.....	xv
forward-and-send-to-re statement.....	127
forward-only statement.....	127
forwarding table	
interfaces, displaying.....	296
G	
Gigabit Ethernet interfaces	
status information, displaying.....	216, 257
Gigabit Ethernet IQ PIC	
traffic and MAC statistics.....	216
group option	
firewall filters.....	19
H	
hierarchical-policer statement.....	128
I	
if-exceeding statement	
hierarchical policer.....	129
inet protocol family	
interface addresses.....	36
inet6 protocol family	
interface addresses.....	38
input option	
firewall filters.....	19
policers.....	11
input statement.....	129
input-list statement.....	130
usage guidelines.....	19
interface addresses	
logical interfaces.....	35
preferred interface addresses.....	49, 50
primary interface addresses.....	49
interface groups.....	21
interface-mode statement.....	131
interfaces	
configuration statements.....	72
firewall filters.....	19
unit statement.....	167
interfaces statement.....	132
dynamic profiles.....	133
interfaces, displaying	
in the forwarding table.....	296

-
- Intermediate System-to-Intermediate System *See*
 - IS-IS
 - Internet Group Management Protocol *See* IGMP
 - Internet Protocol Control Protocol *See* IPCP
 - inverse-arp statement.....137
 - IP addresses
 - 128-bit.....38
 - 32-bit.....36
 - IPCP.....41
 - unnumbered interfaces.....43
 - IP multicast.....326
 - See also* DVMRP, MDT, MLD, MSDP, PGM, PIM, AMT
 - IPCP.....41
 - assigning PPP properties.....43
 - configuring IP address.....41
 - negotiating IP addresses.....42
 - unnumbered interfaces.....42
 - ipsec-sa statement.....137
 - IPv4 Protocol family
 - interface addresses.....35
 - Same IP Address on Multiple Interfaces.....37
 - on logical interfaces.....32
 - IPv6.....34
 - standards documents.....34
 - transition.....34
 - IPv6 Protocol family
 - on logical interfaces.....32
 - ISO Protocol family.....32, 35
- K**
- keep-address-and-control statement.....138
 - usage guidelines.....10
- L**
- Layer 2 bridging.....339
 - Layer 2 circuits
 - connections, displaying.....279
 - redundant pseudowires.....190
 - Layer 2 switching.....339
 - Layer 2 VPNs
 - connections, displaying.....286
 - Layer 3 VPNs
 - dynamic tunnel database, displaying.....199
 - logical interface statements
 - family.....119
 - logical interfaces
 - default router addresses.....49
 - firewall filters.....19
 - interface addresses.....35, 49
 - policers.....11
 - preferred interface addresses.....49, 50
 - primary interface addresses.....49, 50
 - primary router addresses.....49
 - primary router interfaces.....49
 - protocol families.....32
 - protocol MTU.....9
 - protocol redirect messages.....10
 - unnumbered interfaces.....43
 - logical systems
 - configuration statements.....89
 - logical-systems statement.....138
 - loopback testing.....345
- M**
- manuals
 - comments on.....xvii
 - mode statement.....139
 - usage guidelines.....52
 - MPLS
 - protocol family.....32, 35
 - mtu statement.....140
 - logical interfaces
 - usage guidelines.....9
 - MTUs
 - logical interfaces.....9
 - protocol MTUs.....9
 - multicast *See* IP multicast
 - Multicast Listener Discovery *See* MLD
 - Multicast Source Discovery Protocol *See* MSDP
 - multicast-only statement.....142
 - usage guidelines.....33
 - multipoint-destination statement.....143
- N**
- negotiate-address statement.....144
 - usage guidelines.....42
 - negotiating IP addresses
 - IPCP.....42
 - no-redirects statement.....144
 - usage guidelines.....10
 - no-translate-discard-eligible statement.....165
 - no-translate-fecn-and-becn statement.....166

O

oam-liveness statement.....	145
oam-period statement.....	146
output option	
firewall filters.....	19
policers.....	11
output statement.....	147
output-list statement.....	147
usage guidelines.....	19

P

parentheses, in syntax descriptions.....	xvi
point-to-point connections	
unnumbered Ethernet interfaces.....	43
policer	
interface.....	148
policer statement	
usage guidelines.....	11
policers	
applying.....	11
arp option.....	11
burst-size-limit	
statement.....	11
input option.....	11
logical interfaces.....	11
output option.....	11
policers, displaying.....	292
post-service-filter statement.....	149
PPP properties, assigning	
IPCP.....	43
Pragmatic General Multicast See PGM	
preferred interface addresses.....	49, 50
preferred statement.....	149
usage guidelines.....	36, 50
preferred-source-address statement.....	150
example.....	48
usage guidelines.....	46
premium statement	
hierarchical policer.....	151
primary interface addresses.....	49, 50
primary router addresses.....	49
primary router interfaces.....	49
primary statement	
address for interface	
usage guidelines.....	50
address on interface.....	152
interface for router	
usage guidelines.....	49

protocol families

logical interfaces.....	32
protocol MTUs.....	9
unnumbered interfaces.....	43

Protocol Independent Multicast See PIM

protocol MTUs

logical interfaces.....	9
-------------------------	---

protocol redirect messages.....10

protocols statement.....152

proxy statement.....153

pseudowires

redundant switchover.....	190
---------------------------	-----

Q

queue-length statement.....	154
-----------------------------	-----

R

receive-options-packets statement.....	154
--	-----

receive-ttl-exceeded statement.....	155
-------------------------------------	-----

redirect messages.....	10
------------------------	----

redundant pseudowires, switchover.....	190
--	-----

regular expressions

target communities, displaying.....	294
-------------------------------------	-----

remote statement.....	155
-----------------------	-----

request l2circuit-switchover command.....	190
---	-----

routers

default addresses.....	49
------------------------	----

primary addresses.....	49
------------------------	----

primary interfaces.....	49
-------------------------	----

routes, displaying

AS paths

distribution of.....	191
----------------------	-----

domain information.....	195
-------------------------	-----

summary of.....	197
-----------------	-----

in a prefix range.....	302
------------------------	-----

learned from snooping.....	306, 314
----------------------------	----------

Routing Information Protocol See RIP

Routing Information Protocol next generation See

RIPng

rpf-check statement.....	156, 157
--------------------------	----------

usage guidelines.....	51
-----------------------	----

rtvbr statement.....	158
----------------------	-----

S

sampling statement.....	159
-------------------------	-----

service statement

logical interfaces.....	160
-------------------------	-----

service-filter statement.....	161
-------------------------------	-----

service-set statement.....	161
----------------------------	-----

-
- shaping statement.....162
 - show as-path command.....191
 - show as-path domain command.....195
 - show as-path summary command.....197
 - show dynamic-tunnels database command.....199
 - show firewall command.....202
 - show firewall log command.....209
 - show firewall prefix-action-stats command.....212
 - show interfaces (10-Gigabit Ethernet)
 - command.....216
 - show interfaces (Fast Ethernet) command.....241
 - show interfaces (Gigabit Ethernet) command.....257
 - show l2circuit connections command.....279
 - show l2vpn connections command.....286
 - show policer command.....292
 - show route export vrf-target command.....294
 - show route forwarding-table interface-name
 - command.....296
 - show route range command.....302
 - show route snooping command.....306, 314
 - snooping routes, displaying.....306, 314
 - source-class usage
 - example configuration.....28
 - source-class-usage statement.....163
 - usage guidelines.....24
 - static routes
 - unnumbered Ethernet interfaces.....47
 - static subscribers
 - interfaces statement.....133
 - subscriber interface statements
 - demux0.....106
 - family.....119
 - interfaces.....133
 - mode.....139
 - rpf-check.....156
 - unit.....167
 - unnumbered-address.....177
 - support, technical See technical support
 - syntax conventions.....xv
- T**
- targeted-broadcast statement.....164
 - technical support
 - contacting JTAC.....xvii
 - then statement
 - hierarchical policer.....165
 - translate-discard-eligible statement.....165
 - translate-fecn-and-becn statement.....166
- Trivial Network Protocol family**
- interface addresses.....35
 - on logical interfaces.....32
- trunk interface.....131
- trunk port.....131
 - VLAN ID list.....182
- U**
- unicast RPF.....51
 - example configuration.....56, 60
 - fail filters.....54, 56, 60
 - loose mode.....52
 - routing asymmetry.....54
 - strict mode.....51
 - VPNs.....55
 - example configuration.....55
 - unicast-reverse-path statement
 - usage guidelines.....60
 - unit statement.....170
 - interfaces.....167
 - unnumbered interfaces
 - demux.....44
 - Ethernet.....44
 - preferred source address.....150
 - IPCP.....42
 - point-to-point.....43
 - unnumbered-address statement
 - demux interface.....176
 - dynamic profiles.....177
 - Ethernet.....178
 - PPP.....179
 - preferred source address
 - usage guidelines.....46
 - usage guidelines.....42, 44
 - unnumbered-interface statement
 - usage guidelines.....42
- V**
- vbr statement.....180
 - vci statement.....181
 - verification
 - static route.....67
 - vlan-id statement
 - interface in bridge domain.....182
 - vlan-id-list statement
 - bridge domain.....182
 - vpls protocol family
 - interface addresses.....35
 - on logical interfaces.....32

VPN routing and forwarding	See VRF
VPNs	See Layer 2 VPNs, Layer 3 VPNs
unicast RPF	55
VRF	
route export, displaying	294