



Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Solution



Published: 2012-02-28

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.

Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Solution
Copyright © 2012, Juniper Networks, Inc.
All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

	About the Documentation	xi
	Documentation and Release Notes	xi
	Using the Examples in This Manual	xi
	Merging a Full Example	xii
	Merging a Snippet	xii
	Documentation Conventions	xiii
	Documentation Feedback	xiv
	Requesting Technical Support	xv
	Self-Help Online Tools and Resources	xv
	Opening a Case with JTAC	xvi
Part 1	Overview	
Chapter 1	Broadband Subscriber Management Basics Overview	3
	Broadband Subscriber Management Overview	3
	Broadband Subscriber Management Platform Support	4
	Broadband Subscriber Management Solutions Terms and Acronyms	5
	Supporting Documentation for Broadband Subscriber Management	6
Chapter 2	Broadband Subscriber Management Solution Hardware Overview	9
	Broadband Subscriber Management Edge Router Overview	9
	Broadband Services Router Overview	9
	High-Speed Internet Access Support	9
	IPTV Support	10
	Video Services Router	10
	Services Router Placement	10
	Single-Edge Placement	10
	Multiedge Placement	11
	Multiservice Access Node Overview	11
	Ethernet MSAN Aggregation Options	13
	Direct Connection	13
	Ethernet Aggregation Switch Connection	13
	Ring Aggregation Connection	14
Chapter 3	Broadband Subscriber Management Solution Software Overview	15
	Broadband Subscriber Management VLAN Architecture Overview	15
	Broadband Subscriber Management VLANs Across an MSAN	16
	Customer VLANs and Ethernet Aggregation	16

	VLANs and Residential Gateways	17
	DHCP and Broadband Subscriber Management Overview	17
	Extended DHCP Local Server and Broadband Subscriber Management Overview	18
	Extended DHCP Relay and Broadband Subscriber Management Overview	18
	AAA Service Framework and Broadband Subscriber Management Overview	18
Chapter 4	Broadband Subscriber Management DHCP Layer 3 Wholesale Overview	21
	Layer 2 and Layer 3 Wholesale Overview	21
	Wholesale Network Configuration Options and Considerations	22
	DHCP Layer 3 Wholesale Configuration Interface Support	23
	Layer 3 Wholesale Configuration DHCP Support	23
	Subscriber to Logical System and Routing Instance Relationship	24
	RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview	24
Part 2	Configuration	
Chapter 5	Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Network Configuration	29
	Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements	29
	DHCPv4 Layer 3 Wholesale Network Topology Overview	30
	Configuring Loopback Interfaces for the DHCPv4 Layer 3 Wholesale Solution . .	31
	Configuring VLANs for the DHCPv4 Layer 3 Wholesale Network Solution	32
	Configuring Static Customer VLANs for the DHCPv4 Layer 3 Wholesale Network Solution	32
	Configuring Dynamic VLANs for the DHCPv4 Layer 3 Wholesale Network Solution	33
	Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution	35
	Configuring RADIUS Server Access	35
	Configuring a DHCP Wholesaler Access Profile	36
	Configuring DHCP Retailer Access Profiles	36
	Configuring Dynamic Profiles for the DHCPv4 Layer 3 Wholesale Network Solution	37
	Configuring a Wholesale Dynamic Profile for use in the DHCPv4 Solution	38
	Configuring a Dynamic Profile for use by a Retailer in the DHCPv4 Solution	39
	Configuring Separate Routing Instances for DHCPv4 Service Retailers	40
	Configure Default Forwarding Options for the DHCPv4 Wholesale Network Solution	42
Chapter 6	Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Network Configuration Examples	45
	Example: Wholesaler Dynamic Profile for a DHCPv4 Wholesale Network	45
	Example: Retailer Dynamic Profile for a DHCPv4 Wholesale Network	46

	Example: Default Forwarding Options Configuration for the DHCPv4 Wholesale Network	46
	Example: Retailer Routing Instances for a DHCPv4 Wholesale Network	47
Part 3	Administration	
Chapter 7	Subscriber Management AAA and Address Assignment Pool CLI Commands	53
	show network-access aaa statistics	54
	show network-access aaa statistics authentication	57
	show network-access aaa subscribers	60
	show network-access address-assignment pool	63
Chapter 8	Subscriber Management DHCPv4 Local Server CLI Commands	65
	show dhcp server binding	66
	show dhcp server statistics	70
	clear dhcp server binding	73
	clear dhcp server statistics	76
Chapter 9	Subscriber Management DHCPv4 Relay CLI Commands	79
	show dhcp relay binding	80
	show dhcp relay statistics	85
	clear dhcp relay binding	88
	clear dhcp relay statistics	90
Chapter 10	Subscriber Management Interface CLI Commands	93
	show interfaces (Loopback)	94
	show interfaces (Aggregated Ethernet)	101
	show interfaces (Fast Ethernet)	110
	show interfaces (Gigabit Ethernet)	126
	show interfaces demux0 (Demux Interfaces)	147
	show interfaces filters	155
	show interfaces routing	157
Chapter 11	Subscriber Management Subscriber CLI Commands	163
	show subscribers	164
Part 4	Index	
	Index	177

List of Figures

Part 1	Overview	
Chapter 2	Broadband Subscriber Management Solution Hardware Overview	9
	Figure 1: Choosing an MSAN Type	12
Part 2	Configuration	
Chapter 5	Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Network Configuration	29
	Figure 2: Basic Subscriber Management Layer 3 Wholesale Solution Topology	30
	Figure 3: DHCPv4 Layer 3 Wholesale Network Reference Topology	31

List of Tables

	About the Documentation	xi
	Table 1: Notice Icons	xiii
	Table 2: Text and Syntax Conventions	xiii
Part 1	Overview	
Chapter 2	Broadband Subscriber Management Solution Hardware Overview	9
	Table 3: Ethernet MSAN Aggregation Methods	13
Chapter 4	Broadband Subscriber Management DHCP Layer 3 Wholesale Overview	21
	Table 4: Required Juniper Networks VSAs for the Broadband Subscriber Management Wholesale Network Solution	24
Part 3	Administration	
Chapter 7	Subscriber Management AAA and Address Assignment Pool CLI Commands	53
	Table 5: show network-access aaa statistics Output Fields	54
	Table 6: show network-access aaa statistics authentication Output Fields	57
	Table 7: show network-access aaa subscribers Output Fields	60
	Table 8: show network-access address-assignment pool Output Fields	63
Chapter 8	Subscriber Management DHCPv4 Local Server CLI Commands	65
	Table 9: show dhcp server binding Output Fields	67
	Table 10: show dhcp server statistics Output Fields	71
Chapter 9	Subscriber Management DHCPv4 Relay CLI Commands	79
	Table 11: show dhcp relay binding Output Fields	81
	Table 12: show dhcp relay statistics Output Fields	86
	Table 13: clear dhcp relay statistics Output Fields	91
Chapter 10	Subscriber Management Interface CLI Commands	93
	Table 14: Loopback show interfaces Output Fields	94
	Table 15: Aggregated Ethernet show interfaces Output Fields	101
	Table 16: show interfaces Fast Ethernet Output Fields	110
	Table 17: show interfaces Gigabit Ethernet Output Fields	127
	Table 18: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type	140
	Table 19: Demux show interfaces Output Fields	147
	Table 20: show interfaces filters Output Fields	155
	Table 21: show interfaces routing Output Fields	157

Chapter 11	Subscriber Management Subscriber CLI Commands	163
	Table 22: show subscribers Output Fields	166

About the Documentation

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

Documentation and Release Notes

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

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

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

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.xsl;
    }
  }
}
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.xsl; }
```

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 [Junos OS CLI User Guide](#).

Documentation Conventions

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

Table 1: Notice Icons





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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces 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>

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

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

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at

<https://www.juniper.net/cgi-bin/docbugreport/> . If you are using e-mail, be sure to include the following information with your comments:

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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://www.juniper.net/alerts/>
- Join and participate in the Juniper Networks Community Forum: <http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

- [Broadband Subscriber Management Basics Overview on page 3](#)
- [Broadband Subscriber Management Solution Hardware Overview on page 9](#)
- [Broadband Subscriber Management Solution Software Overview on page 15](#)
- [Broadband Subscriber Management DHCP Layer 3 Wholesale Overview on page 21](#)

CHAPTER 1

Broadband Subscriber Management Basics Overview

- [Broadband Subscriber Management Overview on page 3](#)
- [Broadband Subscriber Management Platform Support on page 4](#)
- [Broadband Subscriber Management Solutions Terms and Acronyms on page 5](#)
- [Supporting Documentation for Broadband Subscriber Management on page 6](#)

Broadband Subscriber Management Overview

Broadband Subscriber Management is a method of dynamically provisioning and managing subscriber access in a multiplay or triple play network environment. This method uses AAA configuration in conjunction with dynamic profiles to provide dynamic, per-subscriber authentication, addressing, access, and configuration for a host of broadband services including Internet access, gaming, IPTV, Video on Demand (VoD), and subscriber wholesaling.



NOTE: The Junos OS broadband subscriber management solution currently supports Dynamic Host Configuration Protocols (DHCP)-based and Point-to-Point Protocol /Point-to-Point Protocol over Ethernet (PPP/PPPoE)-based configuration and RADIUS authentication and authorization.

This guide focuses on the general components necessary for configuring a Juniper Networks MX Series 3D Universal Edge Router to dynamically provision and manage subscribers. However, you can also use a Juniper Networks EX Series Ethernet Switch in a subscriber network.

Managing subscribers in a DHCP-based or PPP/PPPoE-based residential broadband network using an MX Series router requires the following:

- Planning and configuring a virtual LAN (VLAN) architecture for the access network.
- Configuring an authentication, authorization, and accounting (AAA) framework for subscriber authentication and authorization through external servers (for example, RADIUS) as well as accounting and dynamic-request change of authorization (CoA)

and disconnect operations through external servers, and address assignment through a combination of local address-assignment pools and RADIUS.

- Configuring DHCP local server or DHCP relay for subscriber address assignment for DHCP-based networks.
- Configuring address assignment pools for PPPoE-based networks.
- Configuring dynamic profiles to include dynamic IGMP, firewall filter, and class of service (CoS) configuration for subscriber access.
- Configuring multicast access to the core network.

To better understand the subscriber access network, this guide also provides general information about some hardware not from Juniper Networks and suggests methods for choosing different network configuration options. You can configure a subscriber network in many different ways. This guide does not cover all configuration scenarios. It is intended as a starting point for understanding subscriber management and how you can use Juniper Networks hardware and software to plan and build your own subscriber management solution.

**Related
Documentation**

- [Broadband Subscriber Management Platform Support on page 4](#)
- [Broadband Subscriber Management Network Topology Overview](#)
- [Broadband Subscriber Management Solutions Terms and Acronyms on page 5](#)
- [Supporting Documentation for Broadband Subscriber Management on page 6](#)
- [Triple Play and Multiplay Overview](#)
- [Broadband History](#)

Broadband Subscriber Management Platform Support

Juniper Networks currently supports DHCP and PPP/PPPoE broadband subscriber management solutions on MX Series routers and PPP/PPPoE broadband subscriber management solutions on M120 and M320 routers.



NOTE: This guide describes configuration on MX Series routers.

**Related
Documentation**

- [Broadband Subscriber Management Overview on page 3](#)
- [Broadband Subscriber Management Edge Router Overview on page 9](#)

Broadband Subscriber Management Solutions Terms and Acronyms

- **AAA (authentication, authorization, and accounting)**—An IP-based networking system that controls user access to computer resources and manages the activity of users over a network.
- **ASM (Any Source Multicast)**—A method of allowing a multicast receiver to listen to all traffic sent to a multicast group, regardless of its source.
- **BSR (broadband services router)**—A router used for subscriber management and edge routing.
- **CoA (change of authorization)**—RADIUS messages that contain information for dynamically changing session authorizations.
- **CoS (class of service)**—A method of managing network traffic by grouping similar types of traffic together and treating each traffic type as a “class” with a defined service priority.
- **DHCP (Dynamic Host Configuration Protocol)**—A mechanism through which hosts using TCP/IP can obtain protocol configuration parameters automatically from a DHCP server on the network; allocates IP addresses dynamically so that they can be reused when no longer needed.
- **IGMP (Internet Group Membership Protocol)**—A host-to-router signaling protocol for IPv4 used to support IP multicasting.
- **IS-IS (Intermediate System-to-Intermediate System)**—A link-state interior gateway routing protocol (IGRP) for IP networks that uses the shortest-path-first (SPF) algorithm to determine routes.
- **LSP (label-switched path)**—The path traversed by a packet that is routed by MPLS. Some LSPs act as tunnels. LSPs are unidirectional, carrying traffic only in the downstream direction from an ingress node to an egress node.
- **MPLS (Multiprotocol Label Switching)**—A mechanism for engineering network traffic patterns that functions by assigning to network packets short labels that describe how to forward the packets through the network.
- **MSAN (multiservice access node)**—A group of commonly used aggregation devices including digital subscriber line access multiplexers (DSLAMs) used in xDSL networks, optical line termination (OLT) for PON/FTTx networks, and Ethernet switches for Active Ethernet connections.
- **Multiplay**—A networking paradigm that enables the ability to add new and robust networking services that individual subscribers can access.
- **OIF (outgoing interface)**—An interface used by multicast functions within a router to determine which egress ports to use for forwarding multicast groups.
- **OSPF (Open Shortest Path First)**—A link-state interior gateway protocol (IGP) that makes routing decisions based on the shortest-path-first (SPF) algorithm (also referred to as the Dijkstra algorithm).

- **PIM (Protocol Independent Multicast)**—A multicast routing protocol used for delivering multicast messages in a routed environment.
- **PPP (Point-to-Point Protocol)**—A link-layer protocol that provides multiprotocol encapsulation. PPP is used for link-layer and network-layer configuration. Provides a standard method for transporting multiprotocol datagrams over point-to-point links.
- **PPPoE (Point-to-Point Protocol over Ethernet)**—A network protocol that encapsulates PPP frames in Ethernet frames and connects multiple hosts over a simple bridging access device to a remote access concentrator.
- **RADIUS (Remote Authentication Dial-In User Service)**—A networking protocol that provides centralized access, authorization, and accounting management for subscribers to connect and use a network service.
- **Residential gateway**—A firewall, Network Address Translation (NAT) router, or other routing device used as a customer premises equipment (CPE) terminator in the home, office, or local point of presence (POP).
- **SSM (single-source multicast)**—A routing method that allows a multicast receiver to detect only a specifically identified sender within a multicast group.
- **set-top box**—The end host or device used to receive IPTV video streams.
- **Triple play**—A networking paradigm that dedicates bandwidth to data, voice, and video service.
- **VOD (video on demand)**—A unicast streaming video offering by service providers that enables the reception of an isolated video session per user with rewind, pause, and similar VCR-like capabilities.
- **VSR (video services router)**—A router used in a video services network to route video streams between an access network and a metro or core network. The video services router is any M Series Multiservice Edge Router or MX Series router that supports the video routing package provided with Junos OS Release 8.3 or later.

**Related
Documentation**

- [Broadband Subscriber Management Overview on page 3](#)

Supporting Documentation for Broadband Subscriber Management

The *Junos OS Broadband Subscriber Management Solutions Guide* relies heavily on existing configuration documentation. In particular, this guide references configuration material presented in the *Junos OS Subscriber Access Configuration Guide*. We recommend you become familiar with the configuration options presented for subscriber access before reading this guide.

Several guides in the Junos OS documentation set provide detailed configuration information that is not fully covered in this guide. This guide might reference other Junos OS configuration and solutions documents that can provide more detail about a specific feature or configuration option.

For more detailed configuration information, see the following Junos OS documents:

- [Junos OS Subscriber Access Configuration Guide](#)
- [Junos OS Layer 2 Configuration Guide](#)
- [Junos OS Multicast Protocols Configuration Guide](#)
- [Junos OS Network Interfaces Configuration Guide](#)
- [Junos OS Policy Framework Configuration Guide](#)

For other solution examples, see the following Junos OS solutions guides:

- [Junos OS MX Series 3D Universal Edge Routers Solutions Guide](#)
- [Session Border Control Solutions Guide Using BGF and IMSG](#)

In addition to related Junos OS documentation, you can obtain useful information from the JunosE Software documentation. Many features described in the *JunosE Broadband Access Configuration Guide* are similar to those described in both this guide and the [Junos OS Subscriber Access Configuration Guide](#).

**Related
Documentation**

- [Broadband Subscriber Management Overview on page 3](#)

CHAPTER 2

Broadband Subscriber Management Solution Hardware Overview

- [Broadband Subscriber Management Edge Router Overview on page 9](#)
- [Multiservice Access Node Overview on page 11](#)
- [Ethernet MSAN Aggregation Options on page 13](#)

Broadband Subscriber Management Edge Router Overview

The edge router is the demarcation point between the residential broadband access network and the core network. The Juniper Networks MX Series router (along with the Juniper Networks EX Series Ethernet Switch) can play multiple roles as an edge router. The most common include the following:

- **Broadband services router (BSR)**—This router supports high speed Internet access along with several other subscriber-based services including VoIP, IPTV, and gaming.
- **Video services router (VSR)**—The video services router capabilities are a subset of those provided by a broadband services router. In general, using the MX Series router as a video services router provides bi-directional traffic destined for the set-top box (STB). This traffic includes IPTV and video on demand (VoD) streams as well as associated control traffic such as IGMP and electronic program guide (EPG) updates.

You can also use the MX Series router in certain Layer 2 solutions. For information about configuring the MX Series router in Layer 2 scenarios, see the [Junos OS Layer 2 Configuration Guide](#) or the [Junos OS MX Series 3D Universal Edge Routers Solutions Guide](#).

Broadband Services Router Overview

A broadband services router is an edge router that traditionally supports primarily Internet-bound traffic. This router replaces and provides a superset of the functionality provided by a Broadband Remote Access Server (B-RAS). The broadband services router functions can be broken into two key areas—high speed Internet access and IPTV support.

High-Speed Internet Access Support

The broadband services router communicates with the RADIUS server to enforce which services each subscriber can access. For example, one subscriber might have signed up for a smaller Internet access service of 1 Mbps where another subscriber might have

signed up for a higher, 10 Mbps service. The broadband services router manages the traffic to each subscriber, ensuring that each subscriber obtains the level of access service they have purchased, while also ensuring that any VoIP traffic receives priority. The broadband services router also makes traffic forwarding decisions based on aggregate bandwidth detected on any adjacent multiservice access node (MSAN).

IPTV Support

The broadband services router supports IPTV traffic including support for IGMP multicast group start and stop requests from downstream MSANs. The broadband services router manages the bandwidth allocations associated with high-bandwidth IPTV as well as video on demand (VoD) traffic to ensure high quality service delivery.

Video Services Router

When configuring a multiedge network, you can use the MX Series router as a video services router (VSR) to support only video traffic without supporting the high-speed Internet access (HSIA) capabilities.



NOTE: We recommend a single-edge network model but the MX Series router allows for flexibility when defining a multiplay network topology.

Some advantages of using a separate video services router for video traffic include the following:

- Provides the ability to add IPTV service without the need to modify an existing edge router that is performing other functions.
- Reduces network bandwidth by moving the video edge further out to the network edge while still allowing for centralized broadband services router operation.
- Typically requires less capital investment because the video services router does not need to provide per-subscriber management.

Services Router Placement

Depending on the type of network you are creating—single edge or multiedge—you can place a broadband services router or video services router in various locations.

Single-Edge Placement

In a single-edge network, you use only broadband services routers because the single device must perform all of the necessary edge functions—providing subscriber management for high-speed Internet access and IPTV services. You can use the two following topology models when placing the broadband services router:

- **Centralized single edge**—The edge router is centrally located and placed at one location to cover a particular region. A secondary router is sometimes placed in this location to act as a backup. Downstream MSANs are connected to the broadband services router using a ring or mesh topology.

- **Distributed single edge**—The edge router is placed further out into the network, typically in the central office (CO) closest to the subscribers that it services. Downstream MSANs are typically connected directly to the broadband services router (in a true, single edge topology) or through an Ethernet aggregation switch.

In general, the addition of IPTV service favors a more distributed model because it pushes the need for subscriber management farther out into the network.

Multiedge Placement

In a multiedge network, you use both broadband services routers and video services routers. The broadband services router controls any high-speed Internet traffic and the video services router controls video traffic. You can use the two following topology models when placing service routers in a multiedge network topology:

- **Co-located multiedge**—The broadband services router and video services router are housed in the same location and an Ethernet switch directs traffic in the CO to the appropriate edge router.



NOTE: A single MX Series router can serve as both Ethernet switch and video services router. For information about configuring the MX Series router in Layer 2 scenarios, see the [Junos OS Layer 2 Configuration Guide](#) or the [Junos OS MX Series 3D Universal Edge Routers Solutions Guide](#).

- **Split multiedge**—The video services router and broadband services router reside in different locations. In this model, the broadband services router is typically located more centrally and video services routers are distributed.

Related Documentation

- [Multiservice Access Node Overview on page 11](#)
- [Ethernet MSAN Aggregation Options on page 13](#)
- [Broadband Subscriber Management Platform Support on page 4](#)

Multiservice Access Node Overview

A *multiservice access node* is a broader term that refers to a group of commonly used aggregation devices. These devices include digital subscriber line access multiplexers (DSLAMs) used in xDSL networks, optical line termination (OLT) for PON/FTTx networks, and Ethernet switches for Active Ethernet connections. Modern MSANs often support all of these connections, as well as providing connections for additional circuits such as plain old telephone service (referred to as POTS) or Digital Signal 1 (DS1 or T1).

The defining function of a multiservice access node is to aggregate traffic from multiple subscribers. At the physical level, the MSAN also converts traffic from the *last mile technology* (for example, ADSL) to Ethernet for delivery to subscribers.

You can broadly categorize MSANs into three types based on how they forward traffic in the network:


- **Layer–2 MSAN**—This type of MSAN is essentially a Layer 2 switch (though typically not a fully functioning switch) with some relevant enhancements. These MSANs use Ethernet (or ATM) switching to forward traffic. The MSAN forwards all subscriber traffic upstream to an edge router that acts as the centralized control point and prevents direct subscriber-to-subscriber communication. Ethernet Link Aggregation (LAG) provides the resiliency in this type of network.

Layer 2 DSLAMs cannot interpret IGMP, so they cannot selectively replicate IPTV channels.

- **Layer–3 aware MSAN**—This IP-aware MSAN can interpret and respond to IGMP requests by locally replicating a multicast stream and forwarding the stream to any subscriber requesting it. Layer 3 awareness is important when supporting IPTV traffic to perform channel changes (sometimes referred to as *channel zaps*). Static IP-aware MSANs always receive all multicast television channels. They do not have the ability to request that specific channels be forwarded to the DSLAM. Dynamic IP-aware DSLAMs, however, can inform the network to begin (or discontinue) sending individual channels to the DSLAM. Configuring IGMP proxy or IGMP snooping on the DSLAM accomplishes this function.
- **Layer–3 MSAN**—These MSANs use IP routing functionality rather than Layer 2 technologies to forward traffic. The advantage of this forwarding method is the ability to support multiple upstream links going to different upstream routers and improving network resiliency. However, to accomplish this level of resiliency, you must assign a separate IP subnetwork to each MSAN, adding a level of complexity that can be more difficult to maintain or manage.

In choosing a MSAN type, refer to [Figure 1 on page 12](#):

Figure 1: Choosing an MSAN Type



Related Documentation

- [Ethernet MSAN Aggregation Options on page 13](#)

Ethernet MSAN Aggregation Options

Each MSAN can connect directly to an edge router (broadband services router or video services router), or an intermediate device (for example, an Ethernet switch) can aggregate MSAN traffic before being sent to the services router. [Table 3 on page 13](#) lists the possible MSAN aggregation methods and under what conditions they are used.

Table 3: Ethernet MSAN Aggregation Methods

Method	When Used
Direct connection	Each MSAN connects directly to the broadband services router and optional video services router.
Ethernet aggregation switch connection	Each MSAN connects directly to an intermediate Ethernet switch. The switch, in turn, connects to the broadband services router or optional video services router.
Ethernet ring aggregation connection	Each MSAN connects to a ring topology of MSANs. The head-end MSAN (the device closest to the upstream edge router) connects to the broadband services router.

You can use different aggregation methods in different portions of the network. You can also create multiple layers of traffic aggregation within the network. For example, an MSAN can connect to a central office terminal (COT), which, in turn, connects to an Ethernet aggregation switch, or you can create multiple levels of Ethernet aggregation switches prior to connecting to the edge router.

Direct Connection

In the direct connection method, each MSAN has a point-to-point connection to the broadband services router. If an intermediate central office exists, traffic from multiple MSANs can be combined onto a single connection using wave-division multiplexing (WDM). You can also connect the MSAN to a video services router. However, this connection method requires that you use a Layer 3 MSAN that has the ability to determine which link to use when forwarding traffic.

When using the direct connection method, keep the following in mind:

- We recommend this approach when possible to simplify network management.
- Because multiple MSANs are used to connect to the services router, and Layer 3 MSANs generally require a higher equipment cost, this method is rarely used in a multiedge subscriber management model.
- Direct connection is typically used when most MSAN links are utilized less than 33 percent and there is little value in combining traffic from multiple MSANs.

Ethernet Aggregation Switch Connection

An Ethernet aggregation switch aggregates traffic from multiple downstream MSANs into a single connection to the services router (broadband services router or optional video services router).

When using the Ethernet aggregation switch connection method, keep the following in mind:

- Ethernet aggregation is typically used when most MSAN links are utilized over 33 percent or to aggregate traffic from lower speed MSANs (for example, 1 Gbps) to a higher speed connection to the services router (for example, 10 Gbps).
- You can use an MX Series router as an Ethernet aggregation switch. For information about configuring the MX Series router in Layer 2 scenarios, see the [Junos OS Layer 2 Configuration Guide](#) or the [Junos OS MX Series 3D Universal Edge Routers Solutions Guide](#).

Ring Aggregation Connection

In a ring topology, the remote MSAN that connects to subscribers is called the remote terminal (RT). This device can be located in the outside plant (OSP) or in a remote central office (CO). Traffic traverses the ring until it reaches the central office terminal (COT) at the head-end of the ring. The COT then connects directly to the services router (broadband services router or video services router).



NOTE: The RT and COT must support the same ring resiliency protocol.

You can use an MX Series router in an Ethernet ring aggregation topology. For information about configuring the MX Series router in Layer 2 scenarios, see the [Junos OS Layer 2 Configuration Guide](#) or the [Junos OS MX Series 3D Universal Edge Routers Solutions Guide](#).

Related Documentation

- [Multiservice Access Node Overview on page 11](#)

CHAPTER 3

Broadband Subscriber Management Solution Software Overview

- [Broadband Subscriber Management VLAN Architecture Overview on page 15](#)
- [DHCP and Broadband Subscriber Management Overview on page 17](#)
- [AAA Service Framework and Broadband Subscriber Management Overview on page 18](#)

Broadband Subscriber Management VLAN Architecture Overview

The subscriber management logical network architecture is as important as the physical network architecture. You configure the logical portion of the subscriber management network using virtual local area networks (VLANs).

Three VLAN models deliver multiple services to subscribers. These models include the following:

- **Service VLAN**—The service VLAN (S-VLAN) provides many-to-one (N:1) subscriber-to-service connectivity: The service VLAN carries a service (for example, data, video, or voice) to all subscribers instead of having different services share a VLAN. Adding a new service requires adding a new VLAN and allocating bandwidth to the new service. The service VLAN model enables different groups that are using the broadband network (for example, external application providers) to manage a given service. One limitation of service VLANs is the absence of any logical isolation between user sessions at the VLAN level. This lack of isolation requires that the multiservice access node (MSAN) and broadband services router provide the necessary security filtering.
- **Customer VLAN**—The customer VLAN (C-VLAN) provides one-to-one (1:1) subscriber-to-service connectivity: One VLAN carries all traffic to each subscriber on the network. Having a single VLAN per subscriber simplifies operations by providing a 1:1 mapping of technology (VLANs) to subscribers. You can also understand what applications any subscriber is using at any given time. Because you use only one VLAN to carry traffic to each subscriber, this approach is not affected when adding new services. However, using a pure C-VLAN model consumes more bandwidth because a single television channel being viewed by multiple subscribers is carried across the network several times—once on each C-VLAN. This approach requires a more scalable, robust edge router that can support several thousand VLANs.

- **Hybrid C-VLAN**—The hybrid VLAN combines the best of both previous VLANs by using one VLAN per subscriber to carry unicast traffic and one shared multicast VLAN (M-VLAN) for carrying broadcast (multicast) television traffic. You can use both the *pure* and *hybrid* C-VLAN models in different portions of the network, depending upon available bandwidth and MSAN capabilities.



NOTE: The term *C-VLAN*, when used casually, often refers to a *hybrid* C-VLAN implementation.

We recommend using one of the C-VLAN models to simplify configuration and management when expanding services. However, some MSANs are limited to the number of VLANs they can support, limiting the ability to use either C-VLAN model.



NOTE: Most MSANs can support the service VLAN model.

Broadband Subscriber Management VLANs Across an MSAN

You configure VLANs to operate between the MSAN and the edge router (broadband services router or video services router). However, the MSAN might modify VLAN identifiers before forwarding information to the subscriber in the following ways:



NOTE: Not all MSANs support these options.

- The VLAN identifiers can be carried within the ATM VCs or they can be removed. The value of keeping the VLAN header is that it carries the IEEE 802.1p Ethernet priority bits. These priority bits can be added to upstream traffic by the residential gateway, allowing the DSLAM to easily identify and prioritize more important traffic (for example, control and VoIP traffic). Typically, a VLAN identifier of zero (0) is used for this purpose.
- In a C-VLAN model, the MSAN might modify the VLAN identifier so that the same VLAN is sent to each subscriber. This enables the use of the same digital subscriber line (DSL) modem and residential gateway configuration for all subscribers without the need to define a different VLAN for each device.

Customer VLANs and Ethernet Aggregation

The 12-bit VLAN identifier (VLAN ID) can support up to 4095 subscribers. When using an aggregation switch with a C-VLAN topology, and fewer than 4095 subscribers are connected to a single edge router port, the aggregation switch can transparently pass all VLANs. However, if the VLAN can exceed 4095 subscribers per broadband services router port, you must use VLAN stacking (IEEE 802.1ad, also known as Q-in-Q). VLAN stacking includes two VLAN tags—an outer tag to identify the destination MSAN and an inner tag to identify the subscriber. For downstream traffic (that is, from the broadband services router or Ethernet switch to the MSAN), the outer tag determines which port to forward traffic. The forwarding device then uses the VLAN pop function on this tag before forwarding the traffic with a single tag. The reverse process occurs for upstream traffic.

VLAN stacking is not necessary for S-VLANs or M-VLANs. However, for the hybrid (C-VLAN and M-VLAN) model, the Ethernet switch or services router must be able to pop or push tags onto C-VLAN traffic while not modifying M-VLAN packets.

VLANs and Residential Gateways

One function provided by a residential gateway is to enable each subscriber to have a private (in-home) network, unseen by other broadband subscribers, while enabling the subscriber to have multiple devices connected to the broadband network. This private network is made possible by using Network Address Translation (NAT).

Most conditional access systems (for example, video on demand) require detecting the real IP address of the set-top box (STB). This security measure means that traffic to and from the STB must be bridged, not routed, across all network elements including aggregation switches, MSANs, and residential gateways. NAT cannot be used at the residential gateway for traffic to and from the STB. In addition, some residential gateways associate VLANs (or ATM virtual circuits) with ports. Traffic on a given VLAN is always forwarded to specific downstream port. Use caution when mapping VLANs on an MSAN.

Related Documentation

- Static Subscriber Interfaces and VLAN Overview in the [Junos OS Subscriber Access Configuration Guide](#).

DHCP and Broadband Subscriber Management Overview

You use DHCP in broadband networks to provide IP address configuration and service provisioning. DHCP, historically a popular protocol in LANs, works well with Ethernet connectivity and is becoming increasingly popular in broadband networks as a simple, scalable solution for assigning IP addresses to subscriber home PCs, set-top boxes (STBs), and other devices.

The Junos OS broadband subscriber management solution currently supports the following DHCP allocation models:

- DHCP Local Server
- DHCP Relay

DHCP uses address assignment pools from which to allocate subscriber addresses. Address-assignment pools support both dynamic and static address assignment:

- Dynamic address assignment—A subscriber is automatically assigned an address from the address-assignment pool.
- Static address assignment—Addresses are reserved and always used by a particular subscriber.



NOTE: Addresses that are reserved for static assignment are removed from the dynamic address pool and cannot be assigned to other clients.

Extended DHCP Local Server and Broadband Subscriber Management Overview

You can enable the services router to function as an extended DHCP local server. As an extended DHCP local server the services router, and not an external DHCP server, provides an IP address and other configuration information in response to a client request. The extended DHCP local server supports the use of external AAA authentication services, such as RADIUS, to authenticate DHCP clients.

Extended DHCP Relay and Broadband Subscriber Management Overview

You can configure extended DHCP relay options on the router and enable the router to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server. You can use DHCP relay in carrier edge applications such as video and IPTV to obtain configuration parameters, including an IP address, for your subscribers. The extended DHCP relay agent supports the use of external AAA authentication services, such as RADIUS, to authenticate DHCP clients.

Related Documentation

- Extended DHCP Local Server Overview in the [Junos OS Subscriber Access Configuration Guide](#).
- Extended DHCP Relay Agent Overview in the [Junos OS Subscriber Access Configuration Guide](#).
- Address-Assignment Pools Overview in the [Junos OS Subscriber Access Configuration Guide](#).

AAA Service Framework and Broadband Subscriber Management Overview

You use AAA Service Framework for all authentication, authorization, accounting, address assignment, and dynamic request services that the services router uses for network access. The framework supports authentication and authorization through external servers, such as RADIUS. The framework also supports accounting and dynamic-request CoA and disconnect operations through external servers, and address assignment through a combination of local address-assignment pools and RADIUS.



NOTE: The broadband subscriber management solution currently supports the use of only RADIUS servers.

The broadband services router interacts with external servers to determine how individual subscribers access the broadband network. The router also obtains information from external servers for the following:

- Methods used for authentication and accounting.
- How accounting statistics are collected and used.
- How dynamic requests are handled.

**Related
Documentation**

- AAA Service Framework Overview in the [Junos OS Subscriber Access Configuration Guide](#).
- RADIUS-Initiated Change of Authorization (CoA) Overview in the [Junos OS Subscriber Access Configuration Guide](#).
- RADIUS-Initiated Disconnect Overview in the [Junos OS Subscriber Access Configuration Guide](#).

CHAPTER 4

Broadband Subscriber Management DHCP Layer 3 Wholesale Overview

- [Layer 2 and Layer 3 Wholesale Overview on page 21](#)
- [Wholesale Network Configuration Options and Considerations on page 22](#)
- [DHCP Layer 3 Wholesale Configuration Interface Support on page 23](#)
- [Layer 3 Wholesale Configuration DHCP Support on page 23](#)
- [Subscriber to Logical System and Routing Instance Relationship on page 24](#)
- [RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview on page 24](#)

Layer 2 and Layer 3 Wholesale Overview

In general, wholesaling broadband services allows service providers to resell broadband services and allows other providers to deploy their own services over the incumbent network. There are different methods to partitioning an access network for resale. The two most common approaches are based on either Layer 2 or Layer 3 information. Wholesale access is the process by which the access network provider (the *wholesaler*) partitions the access network into separately manageable and accountable subscriber segments for resale to other network providers (or *retailers*).

In a Layer 3 wholesale configuration, you partition the wholesaler access network at the network layer or the subscriber IP component by associating the IP component with a distinct Layer 3 domain. In a Layer 2 wholesale configuration, you partition the access network at the subscriber circuit or customer VLAN (C-VLAN) by backhauling the connection through the service provider backbone network to the subscribing retailer network where the access traffic can be managed at higher layers.

In a Junos OS Dynamic Host Configuration Protocol (DHCP) or Point-to-Point Protocol over Ethernet (PPPoE) subscriber access configuration, wholesale partitioning is accomplished through the use of logical systems and routing instances within the router. Logical systems offer a stricter partitioning of routing resources than routing instances. The purpose behind the use of logical systems is to distinctly partition the physical router into separate administrative domains. This partitioning enables multiple providers to administer the router simultaneously, with each provider having access only to the portions of the configuration relevant to their logical system. Junos OS supports up to 15 named

logical systems in addition to the default logical system (that is, **inet.0**). Unless otherwise specified in configuration, all interfaces belong to the default logical system.



NOTE: This Junos OS release supports the use of only the default logical system. Partitioning currently occurs through the use of separate routing instances.

A logical system can have one or more routing instances. Typically used in Layer 3 VPN scenarios, a routing instance does not have the same level of administrative separation as a logical system because it does not offer administrative isolation. However, the routing instance defines a distinct routing table, set of routing policies, and set of interfaces.

Related Documentation

- [Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements on page 29](#)
- [Broadband Subscriber Management PPPoE Layer 3 Wholesale Topology and Configuration Elements](#)
- [Broadband Subscriber Management Layer 2 Wholesale Topology and Configuration Elements](#)

Wholesale Network Configuration Options and Considerations

You can configure a wholesale network any number of ways using Juniper Hardware and JUNOS software. For information about subscriber management hardware support, see [Subscriber Access Support Considerations](#) in the *Junos OS Subscriber Access Configuration Guide*. The general configuration options, and considerations for each, are provided below:

Wholesale Configuration Options	Considerations
Fully Static (all interfaces, VLANs, and routing instances are configured statically)	Providing more control over retailer space and access, this option is more labor intensive and can require more detailed planning of the network, address allocation, and so on.
Static VLANs and Dynamic Demux Interfaces	Service VLANs are created statically and must be managed. Demux interfaces are dynamically created over the service VLANs. This option uses more logical interfaces; one for each VLAN and one for each dynamic demux interface that runs over each VLAN.
Dynamic VLANs Only (dedicated customer VLANs for each subscriber)	Dynamic (auto-sensed) VLANs are authenticated and installed in the correct non-default routing instance before DHCP is instantiated. This method helps to conserve logical interfaces by avoiding the need for additional logical interfaces being created for each demux interface. NOTE: In a customer VLAN model, each VLAN functions on a 1:1 basis for each customer (in this case, per household).

Wholesale Configuration Options	Considerations
Dynamic VLANs and Dynamic Demux Interfaces	Allows for the greatest ease of use and flexibility in configuring subscribers, by enabling access over a service VLAN and targetting more service levels over individual, dynamically-created demux interfaces over the service VLAN. This option uses more logical interfaces; one for each VLAN and one for each demux interface that runs over each VLAN.

DHCP Layer 3 Wholesale Configuration Interface Support

DHCP Layer 3 wholesale currently supports only the use of IP demux interfaces.

For general additional information about configuring IP demux interfaces, see the [Junos OS Network Interfaces Configuration Guide](#).

Related Documentation

- [Junos OS Network Interfaces Configuration Guide](#)
- Subscriber Interfaces and Demultiplexing Overview in the [Junos OS Subscriber Access Configuration Guide](#).
- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles in the [Junos OS Subscriber Access Configuration Guide](#).
- Configuring a Subscriber Interface Using a Set of Static IP Demux Interfaces in the [Junos OS Subscriber Access Configuration Guide](#).

Layer 3 Wholesale Configuration DHCP Support

DHCP Layer 3 wholesale supports the following DHCP configuration options:

- DHCP Relay
- DHCP Relay Proxy
- DHCP Local Server



NOTE: All routing instances within the same wholesale network must use the same DHCP configuration option.

For additional information about any of these DHCP options, see the AAA Service Framework Overview in the [Junos OS Subscriber Access Configuration Guide](#).

Related Documentation

- Extended DHCP Relay Agent Overview in the [Junos OS Subscriber Access Configuration Guide](#).
- DHCP Relay Proxy Overview in the [Junos OS Subscriber Access Configuration Guide](#).
- Extended DHCP Local Server Overview in the [Junos OS Subscriber Access Configuration Guide](#).

Subscriber to Logical System and Routing Instance Relationship

As subscriber sessions are established, subscriber to logical system/routing instance memberships are established by the AAA framework configured for the default logical system. When configuring Layer 3 wholesaling, you typically configure global (wholesale) information within the default (master) logical system and default routing instance. Incoming subscribers must then be authenticated, but this authentication can be handled in one of two ways:

- Single (wholesaler only) authentication—Incoming subscribers are authenticated by the wholesaler RADIUS server. After authentication, the subscribers are assigned values specified by dynamic profiles (routing instances, interfaces, and any configuration values) specific to a particular retailer.
- Dual (wholesaler and retailer) authentication—Sometimes referred to as *double-dip authentication*. Incoming subscribers are initially authenticated by RADIUS using the wholesale configuration. Authenticated subscribers are then redirected to other routing instances associated with individual retailer network space. When you redirect subscribers, and those subscribers are to be authenticated by AAA servers owned by individual retailers, the subscribers must be authenticated again by the AAA servers before they are provided an address and any dynamic profile values are assigned. After reauthentication, however, the subscribers are managed normally using any values specific to the retailer routing instance to which they are assigned.

Related Documentation • See Routing Instances Overview in the [Junos OS Routing Protocols Configuration Guide](#).

RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview

You can use RADIUS to assign various values through the use of dynamic variables within dynamic profiles. However, the configuration of at least one of the two VSAs described in [Table 4 on page 24](#) is required for a wholesale network to function.

Table 4: Required Juniper Networks VSAs for the Broadband Subscriber Management Wholesale Network Solution

Attribute Number	Attribute Name	Description	Value
26-1	LSRI-Name	Client logical system/routing instance membership name. Allowed only from RADIUS server for "default" logical system/routing instance membership.	string: logical system:routing instance

Table 4: Required Juniper Networks VSAs for the Broadband Subscriber Management Wholesale Network Solution (*continued*)

Attribute Number	Attribute Name	Description	Value
26-25	Redirect-LSRI-Name	Client logical system/routing instance membership name indicating to which logical system/routing instance membership the request is redirected for user authentication.	string: logical system:routing instance

Specifying the **\$junos-routing-instance** dynamic variable in a dynamic profile triggers a RADIUS access-accept response of either the LSRI-Name VSA or the Redirect-LSRI-Name VSA. Returning an LSRI-Name attribute in the access-accept response provides the logical system and routing instance in which the logical interface is to be created and the router updates the session database with the specified routing instance value. Returning a Redirect-LSRI-Name attribute in the access-accept response results in the router immediately sending a second access-request message (sometimes referred to as a *double-dip*) to the RADIUS server specified by the logical system:routing instance attribute specified by the Redirect-LSRI-Name VSA.



NOTE: Attributes returned as a result of a second access-request message to the logical system/routing instance membership specified by the Redirect-LSRI-Name VSA override any prior attributes returned by initial access-accept responses to the default logical system/routing instance membership.

Related Documentation

- Juniper Networks VSAs Supported by the AAA Service Framework in the [Junos OS Subscriber Access Configuration Guide](#).

PART 2

Configuration

- [Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Network Configuration on page 29](#)
- [Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Network Configuration Examples on page 45](#)

CHAPTER 5

Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Network Configuration

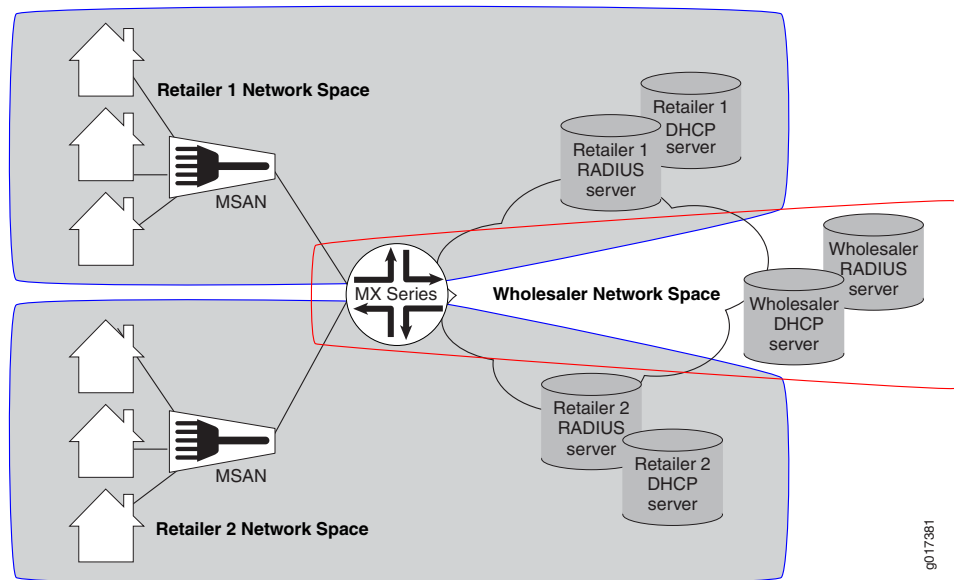
- [Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements on page 29](#)
- [DHCPv4 Layer 3 Wholesale Network Topology Overview on page 30](#)
- [Configuring Loopback Interfaces for the DHCPv4 Layer 3 Wholesale Solution on page 31](#)
- [Configuring VLANs for the DHCPv4 Layer 3 Wholesale Network Solution on page 32](#)
- [Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution on page 35](#)
- [Configuring Dynamic Profiles for the DHCPv4 Layer 3 Wholesale Network Solution on page 37](#)
- [Configuring Separate Routing Instances for DHCPv4 Service Retailers on page 40](#)
- [Configure Default Forwarding Options for the DHCPv4 Wholesale Network Solution on page 42](#)

Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements

The network topology for the subscriber management DHCPv4 Layer 3 wholesale solution includes configuring separate routing instances for individual retailers that use a portion of the router. This solution uses a DHCPv4 relay configuration. However, you can also implement DHCPv4 Relay Proxy or DHCPv4 Local Server configuration.

To explain the concept, but to limit complexity, this solution provides a configuration with one wholesaler and only two retailers. [Figure 2 on page 30](#) illustrates a basic Layer 3 wholesale topology model from which you can expand.

Figure 2: Basic Subscriber Management Layer 3 Wholesale Solution Topology



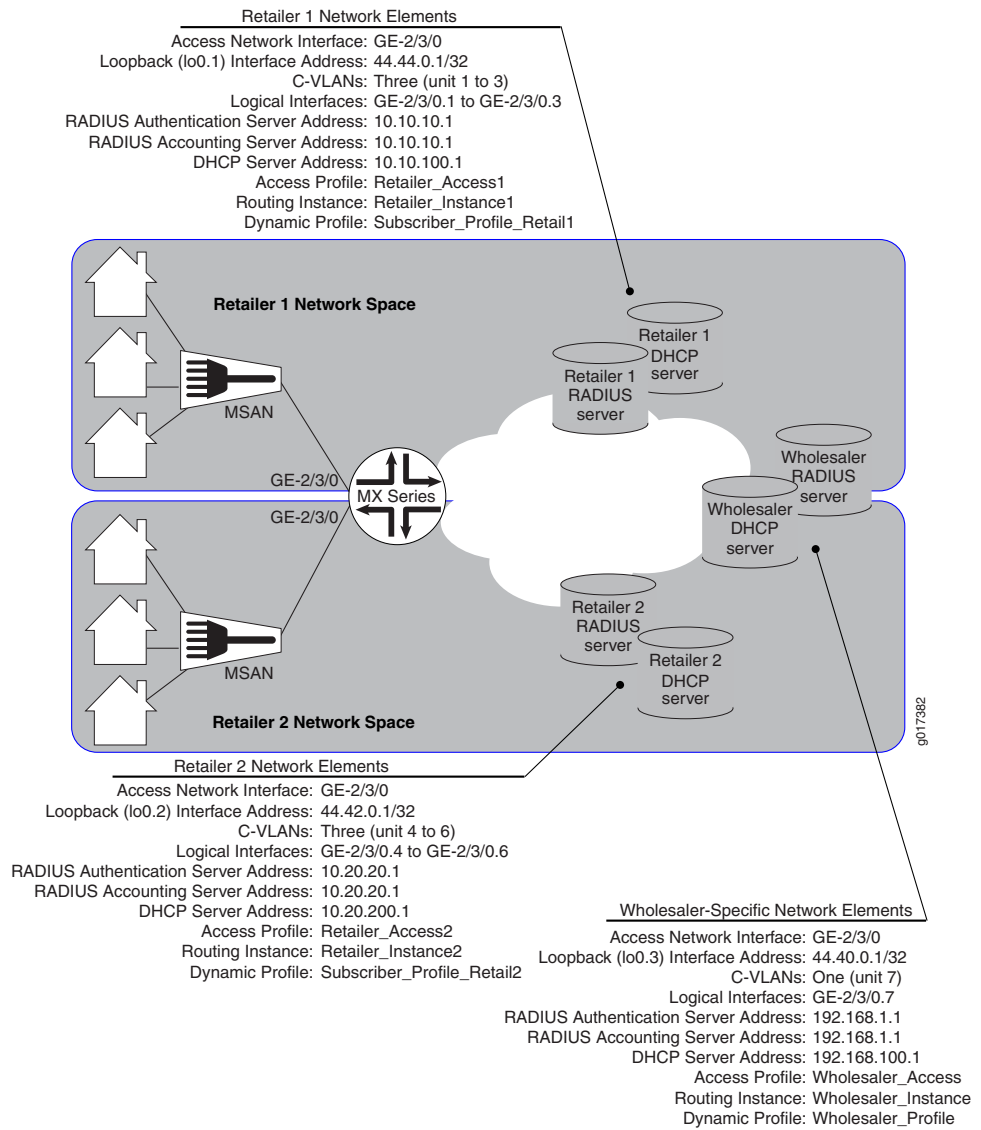
A DHCP Layer 3 wholesale network solution can use various combinations of the following configuration elements:

- Subscriber network VLAN configuration
- DHCPv4 configuration (DHCPv4 Relay, DHCPv4 Relay Proxy, or DHCPv4 Local Server)
- Addressing server or addressing server access configuration (if not using DHCPv4 Local Server)
- RADIUS server access configuration
- Dynamic profile configuration for default (wholesaler) access
- Dynamic profile configuration for retailer access (following subscriber redirection, if applicable)
- Routing instance configuration for individual retailers
- Group configuration and forwarding options for the network
- Core network configuration

DHCPv4 Layer 3 Wholesale Network Topology Overview

This configuration explains how to configure a simple DHCPv4 Layer 3 wholesale subscriber access network. This solution incorporates two retailers sharing resources on a wholesaler router. [Figure 3 on page 31](#) provides the reference topology for this configuration example.

Figure 3: DHCPv4 Layer 3 Wholesale Network Reference Topology



- Related Documentation**
- [Layer 2 and Layer 3 Wholesale Overview on page 21](#)
 - [Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements on page 29](#)

Configuring Loopback Interfaces for the DHCPv4 Layer 3 Wholesale Solution

You must configure loopback interfaces for use in the subscriber management access network. The loopback interfaces are automatically used for unnumbered interfaces.

To configure loopback interfaces:

1. Edit the loopback interface.

```
[edit]
user@host# edit interfaces lo0
```

2. Edit the unit for the wholesale loopback interface.

```
[edit interfaces lo0]
user@host# edit unit 3
```

3. Edit the loopback interface family that belongs to the wholesaler.

```
[edit interfaces lo0 unit 3]
user@host# edit family inet
```

4. Specify the loopback interface address that belongs to the wholesaler.

```
[edit interfaces lo0 unit 3]
user@host# set address 44.40.0.1/32
```

5. Edit the unit for a retail loopback interface to be assigned to the retailer.

```
[edit interfaces lo0]
user@host# edit unit 1
```

6. Edit the loopback interface family that will be assigned to the retailer.

```
[edit interfaces lo0 unit 1]
user@host# edit family inet
```

7. Specify the loopback interface address that will be assigned to the retailer.

```
[edit interfaces lo0 unit 1]
user@host# set address 44.42.0.1/32
```

8. Repeat steps 5 through 7 for additional retailers, making sure to use unique unit and address values for each retailer loopback interface.

Related Documentation

- [Configuring Top-Level Broadband Subscriber Management Elements](#)
- [Junos OS Network Interfaces Configuration Guide](#)

Configuring VLANs for the DHCPv4 Layer 3 Wholesale Network Solution

You can configure either static or dynamic customer VLANs for use in the DHCPv4 wholesale network solution.

- [Configuring Static Customer VLANs for the DHCPv4 Layer 3 Wholesale Network Solution on page 32](#)
- [Configuring Dynamic VLANs for the DHCPv4 Layer 3 Wholesale Network Solution on page 33](#)

Configuring Static Customer VLANs for the DHCPv4 Layer 3 Wholesale Network Solution

In this example configuration, the access interface (**ge-2/3/0**) connects to a device (that is, a DSLAM) on the access side of the network. You can define static VLANs for use by the access network subscribers.

To configure the static VLANs:

1. Edit the access side interface.

```
[edit]
user@host# edit interfaces ge-2/3/0
```

2. Specify the use of stacked VLAN tagging.

```
[edit interfaces ge-2/3/0]
user@host# set stacked-vlan-tagging
```

3. Edit the interface unit for the first VLAN.

```
[edit interfaces ge-2/3/0]
user@host# edit unit 1
```

4. Define the VLAN tags for the first VLAN.

```
[edit interfaces ge-2/3/0 unit 1]
user@host# set vlan-tags outer 3 inner 1
```

5. Specify that you want to create IPv4 demux interfaces.

```
[edit interfaces ge-2/3/0 unit 1]
user@host# set demux-source inet
```

6. Edit the family for the first VLAN.

```
[edit interfaces ge-2/3/0 unit 1]
user@host# edit family inet
```

7. (Optional) Define the unnumbered address and the preferred source address for the first VLAN.

```
[edit interfaces ge-2/3/0 unit 1 family inet]
user@host# set unnumbered-address lo0.1 preferred-source-address 44.44.0.1
```

8. Repeat steps 2 through 7 for additional VLAN interface units.

Configuring Dynamic VLANs for the DHCPv4 Layer 3 Wholesale Network Solution

To configure dynamic VLANs for the solution:

1. Configure a dynamic profile for dynamic VLAN creation.

- a. Name the profile.

```
[edit]
user@host# edit dynamic-profiles VLAN-PROF
```

- b. Define the **interfaces** statement with the internal **\$junos-interface-ifd-name** variable used by the router to match the interface name of the receiving interface.

```
[edit dynamic-profiles VLAN-PROF]
user@host# edit interfaces $junos-interface-ifd-name
```

- c. Define the **unit** statement with the predefined **\$junos-interface-unit** variable:

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name"]
user@host# edit unit $junos-interface-unit
```

- d. (Optional) To configure the router to respond to any ARP request, specify the **proxy-arp** statement.

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit"]
user@host# set proxy-arp
```

- e. Specify that you want to create IPv4 demux interfaces.

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit"]
user@host# set demux-source inet
```

- f. Specify the VLAN ID variable.

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit"]
user@host# set vlan-tags outer $junos-stacked-vlan-id
```

The variable is dynamically replaced with an outer VLAN ID within the VLAN range specified at the **[interfaces]** hierarchy level.

- g. Specify the inner VLAN ID variable.

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit"]
user@host# set vlan-tags inner $junos-vlan-id
```

The variable is dynamically replaced with an inner VLAN ID within the VLAN range specified at the **[interfaces]** hierarchy level.

- h. Access the family type.

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit"]
user@host# edit family inet
```

- i. (Optional) Enable IP and MAC address validation for dynamic IP demux interfaces in a dynamic profile.

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit" family inet]
user@host# set mac-validate strict
```

- j. (Optional) Specify the unnumbered address and preferred source address.

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit" family inet]
user@host# set unnumbered-address lo.0 preferred-source-address 33.33.0.1
```

2. Associate the dynamic profile with the interface on which the dynamic VLANs will be created.

- a. Access the interface that you want to use for creating VLANs.

```
[edit interfaces]
user@host# edit interfaces ge-2/3/0
```

- b. Specify the use of stacked VLAN tagging.

```
[edit interfaces ge-2/3/0]
user@host# set stacked-vlan-tagging
```

- c. Specify that you want to automatically configure VLAN interfaces.

```
[edit interfaces ge-2/3/0]
user@host# edit auto-configure
```

- d. Specify that you want to configure stacked VLANs.

```
[edit interfaces ge-2/3/0 auto-configure]
user@host# edit stacked-vlan-ranges
```

- e. Specify the dynamic VLAN profile that you want the interface to use.

```
[edit interfaces ge-2/3/0 auto-configure stacked-vlan-ranges]
user@host# set dynamic-profile VLAN-PROF
```

- f. Repeat steps a through e for any other interfaces that you want to use for creating VLANs.

3. Specify the Ethernet packet type that the VLAN dynamic profile can accept.

```
[edit interfaces ge-2/3/0 auto-configure stacked-vlan-ranges dynamic-profile
VLAN-PROF]
user@host# set accept inet
```

4. Define VLAN ranges for use by the dynamic profile when dynamically creating VLAN IDs. For this solution, specify the outer and inner stacked VLAN ranges that you want the dynamic profile to use. The following example specifies an outer stacked VLAN ID range of 3–3 (enabling only the outer range of 3) and an inner stacked VLAN ID range of 1–3 (enabling a range from 1 through 3 for the inner stacked VLAN ID).

```
[edit interfaces ge-0/0/0 auto-configure stacked-vlan-ranges dynamic-profile
VLAN-PROF]
user@host# set stacked-vlan-ranges 3–3,1–3
```

Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution

When configuring a wholesale network, you must configure several components globally. This configuration provides access to RADIUS servers that you want the wholesaler and any configured retailers to use globally. The access configuration includes the following general steps:

- [Configuring RADIUS Server Access on page 35](#)
- [Configuring a DHCP Wholesaler Access Profile on page 36](#)
- [Configuring DHCP Retailer Access Profiles on page 36](#)

Configuring RADIUS Server Access

You can globally define any RADIUS servers in your network that either the wholesaler access profile or retailer access profile can use. After you define the global RADIUS servers, you can specify specific RADIUS servers within individual access profiles.

To define RADIUS servers for profile access:

1. Access the **[edit access radius-server]** hierarchy level.

```
[edit ]
user@host# edit access radius-server
```

2. Specify the address and secret for any RADIUS servers in the network.

```
[edit access radius-server]
user@host# set 192.168.10.1 secret $9$CzBxBBfleWx-wM8xgaU.m345B02EcyKXL
user@host# set 10.10.10.1 secret $7$OsCsBAf1fXx-wY3xgaU.m123A02ZtyNMT
```

Configuring a DHCP Wholesaler Access Profile

You must define the network and interface over which you want subscribers to initially access the network with a wholesale access profile. When a subscriber attempts to access the network, the access profile provides initial access information including authentication and accounting values that the router uses for the accessing subscriber.

To define a wholesale access profile:

1. Create the wholesale access profile.

```
[edit]
user@host# edit access-profile Wholesaler_Access
```

2. Specify the authentication methods for the profile and the order in which they are used.

```
[edit access profile Wholesaler1]
user@host# set authentication-order radius password
```

3. Specify that you want to configure RADIUS support.

```
[edit access profile Wholesaler1]
user@host# edit radius
```

4. Specify the IP address of the RADIUS server used for authentication.

```
[edit access profile Wholesaler1 radius]
user@host# set authentication-server 192.168.10.1
```

5. Specify the IP address of the RADIUS server used for accounting.

```
[edit access profile Wholesaler1 radius]
user@host# set accounting-server 192.168.10.1
```

6. Configure any desired options for the RADIUS server.

See [Configuring RADIUS Server Options for Subscriber Access](#).

7. Configure subscriber accounting (RADIUS accounting).

See [Configuring Per-Subscriber Session Accounting](#).

Configuring DHCP Retailer Access Profiles

In this solution, subscribers are redirected to a networking space used by a specific retailer and defined by a unique routing instance. This method requires that you define the network

and interface over which you want subscribers to access the network after being redirected by the wholesale access profile.

To define a retailer access profile:

1. Create the retailer access profile.

```
[edit]
user@host# edit access-profile Retailer_Access1
```

2. Specify the authentication methods for the profile and the order in which they are used.

```
[edit access profile Retailer1]
user@host# set authentication-order radius password
```

3. Specify that you want to configure RADIUS support.

```
[edit access profile Retailer1]
user@host# edit radius
```

4. Specify the IP address of the RADIUS server used for authentication.

```
[edit access profile Retailer1 radius]
user@host# set authentication-server 10.10.10.1
```

5. Specify the IP address of the RADIUS server used for accounting.

```
[edit access profile Retailer1 radius]
user@host# set accounting-server 10.10.10.1
```

6. Configure any desired options for the RADIUS server.

See Configuring RADIUS Server Options for Subscriber Access.

7. Configure subscriber accounting (RADIUS accounting).

See Configuring Per-Subscriber Session Accounting.

Configuring Dynamic Profiles for the DHCPv4 Layer 3 Wholesale Network Solution

A dynamic profile is a set of characteristics, defined in a type of template, that you can use to provide services for broadband applications. These services are assigned dynamically to interfaces as they access the network. When configuring dynamic profiles for the DHCPv4 Layer 3 wholesale network, you can choose to configure one dynamic profile to address all incoming subscribers or you can configure individual dynamic profiles for use by the different network management groups (that is, the wholesaler and any retailers). In fact, you can create multiple dynamic profiles that you can use to roll out different services and selectively apply those dynamic profiles to different subscriber groups as necessary.

In this solution example, one dynamic profile is created for use by the wholesaler when subscribers initially access the network. Other dynamic profiles are created for the

subscribers for each individual retailer to use after they are redirected to that retailer network space.

- [Configuring a Wholesale Dynamic Profile for use in the DHCPv4 Solution on page 38](#)
- [Configuring a Dynamic Profile for use by a Retailer in the DHCPv4 Solution on page 39](#)

Configuring a Wholesale Dynamic Profile for use in the DHCPv4 Solution

You can configure a basic access profile to initially manage subscribers that access the network.

To configure a dynamic profile for use by the wholesaler:

1. Create a wholesale dynamic profile.

```
[edit]
user@host# edit dynamic-profiles Wholesaler_Profile
```

2. Specify that you want to configure the **demux0** interface in the dynamic profile.

```
[edit dynamic-profiles Subscriber_Profile_Retail1]
user@host# edit interfaces demux0
```

3. Configure the unit for the **demux0** interface.

- a. Configure the variable for the unit number of the **demux0** interface.

The variable is dynamically replaced with the unit number that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 demux0]
user@host# edit unit $junos-interface-unit
```

- b. Configure the variable for the underlying interface of the demux interfaces and specify the **\$junos-underlying-interface** variable.

The variable is dynamically replaced with the underlying interface that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces demux0 unit
"$junos-interface-unit"]
user@host# set demux-options underlying-interface $junos-underlying-interface
```

4. Configure the family for the demux interfaces.

- a. Specify that you want to configure the family.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces demux0 unit
"$junos-interface-unit"]
user@host# edit family inet
```

- b. Configure the unnumbered address for the family.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 demux0 unit "$junos-interface-unit"
family inet6]
user@host# set unnumbered-address lo0.0
```

- c. Configure the variable for the IPv4 address of the demux interface.

The variable is dynamically replaced with the IPv4 address that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles business-profile interfaces demux0 unit "$junos-interface-unit"]
user@host# set demux-source $junos-subscriber-ip-address
```

Configuring a Dynamic Profile for use by a Retailer in the DHCPv4 Solution

To configure a dynamic profile for use with retailer access:

1. Create a retail dynamic profile.

```
[edit]
user@host# edit dynamic-profiles Subscriber_Profile_Retail1
```

2. Define the dynamic routing instance variable in the dynamic profile.

```
[edit dynamic-profiles Subscriber_Profile_Retail1]
user@host# edit routing-instances $junos-routing-instance
```

3. Set the dynamic interface variable for the dynamic routing instance.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 routing-instances
"$junos-routing-instance"]
user@host# set interface $junos-interface-name
```

4. Specify that you want to configure the **demux0** interface in the dynamic profile.

```
[edit dynamic-profiles Subscriber_Profile_Retail1]
user@host# edit interfaces demux0
```

5. Configure the unit for the **demux0** interface.

- a. Configure the variable for the unit number of the **demux0** interface.

The variable is dynamically replaced with the unit number that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 demux0]
user@host# edit unit $junos-interface-unit
```

- b. Configure the variable for the underlying interface of the demux interfaces and specify the **\$junos-underlying-interface** variable.

The variable is dynamically replaced with the underlying interface that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces demux0 unit
"$junos-interface-unit"]
user@host# set demux-options underlying-interface $junos-underlying-interface
```

6. Configure the family for the demux interfaces.

- a. Specify that you want to configure the family.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces demux0 unit
"$junos-interface-unit"]
user@host# edit family inet
```

- b. Configure the unnumbered address for the family.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 demux0 unit "$junos-interface-unit"
family inet6]
user@host# set unnumbered-address lo0.0
```

- c. Configure the variable for the IPv6 address of the demux interface.

The variable is dynamically replaced with the IPv6 address that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles business-profile interfaces demu0 unit "$junos-interface-unit"]
user@host# set demux-source $junos-subscriber-ip-address
```

Configuring Separate Routing Instances for DHCPv4 Service Retailers

As the owner of the system, the wholesaler typically uses the default routing instance. You must create separate routing instances for each individual retailer to keep routing information for individual retailers separate and to define any servers and forwarding options specific to each retailer.

To define a retailer routing instance:

1. Create the retailer routing instance.

```
[edit]
user@host# edit routing-instances RetailerInstance1
```

2. Specify the routing instance type for the retailer.

```
[edit routing-instances "RetailerInstance1"]
user@host# set instance-type vrf
```

3. Specify the access profile that you want the routing instance to use.

```
[edit routing-instances "RetailerInstance1"]
user@host# set access-profile Retailer1
```

4. Specify the interface that faces the Retailer1 RADIUS server.

```
[edit routing-instances "RetailerInstance1"]
user@host# set interface ge-11/1/9.10
```

5. Specify the interface that faces the Retailer1 DHCP server.

```
[edit routing-instances "RetailerInstance1"]
user@host# set interface ge-11/1/10.100
```

6. Specify the loopback interface unit for this routing instance.

```
[edit routing-instances "RetailerInstance1"]
user@host# set interface lo0.1
```



NOTE: Loopback interfaces must be unique for each routing instance.

7. Access the DHCP Relay forwarding options hierarchy for the routing instance.

```
[edit routing-instances "RetailerInstance1"]
user@host# edit forwarding-options dhcp-relay
```




NOTE: The configuration for this wholesale solution uses DHCP Relay. However, you can also configure DHCP Proxy Relay or DHCP Local Server for the DHCP Layer 3 wholesale network.

8. Specify that you want to configure authentication options and use external AAA authentication services.

```
[edit routing-instances "RetailerInstance1" forwarding-options dhcp-relay]
user@host# edit authentication
```

9. (Optional) Configure a password that authenticates the username to the external authentication service.

See Configuring Passwords for Usernames.

10. (Optional) Configure optional features to create a unique username.

See Creating Unique Usernames for DHCP Clients.

11. Specify the default dynamic profile that you want to attach to DHCP subscriber for this retailer.

```
[edit routing-instances "RetailerInstance1" forwarding-options dhcp-relay]
user@host# set dynamic-profile Subscriber_Profile_Retail1
```

12. Specify any overrides for the default DHCP Relay configuration.

See Overriding the Default DHCP Relay Configuration Settings.

13. Configure a named server group for the retailer.

```
[edit routing-instances "RetailerInstance1" forwarding-options dhcp-relay]
user@host# edit server-group Retailer1_Group
```

14. Specify the DHCP server address for the retailer group.

```
[edit routing-instances "RetailerInstance1" forwarding-options dhcp-relay server-group
"Retailer1_Group"]
user@host# set 10.10.100.1
```

15. Specify the retailer group as the active server group for this routing instance.

```
[edit routing-instances "RetailerInstance1" forwarding-options dhcp-relay]
user@host# set active-server-group Retailer1_Group
```

16. Configure a group you can use to define the retailer dynamic profile and DHCP access interface.

```
[edit routing-instances "RetailerInstance1" forwarding-options dhcp-relay]
user@host# edit group Retailer1_Group
```

17. Specify the dynamic profile that the retailer DHCP subscribers use.

```
[edit routing-instances "RetailerInstance1" forwarding-options dhcp-relay group
"Retailer1_Group"]
user@host# set dynamic-profile Subscriber_Profile_Retailer1
```

18. Specify the retailer interface that the retailer DHCP subscribers use.

```
[edit routing-instances "RetailerInstance1" forwarding-options dhcp-relay group  
"Retailer1_Group"]
```

```
user@host# set interface ge-2/3/0.2
```

19. (Optional) Configure any passwords that authenticate the username to the external authentication service for the retailer groups that you created.

See Configuring Passwords for Usernames.

20. (Optional) Configure any unique username values for the retailer groups that you created.

See Creating Unique Usernames for DHCP Clients.

21. (Optional) Specify any overrides for any of the DHCP Relay group configurations that you created.

See Overriding the Default DHCP Relay Configuration Settings.

22. Repeat this procedure for other retailers.

Related Documentation

- [Configuring Routing Instances](#)

Configure Default Forwarding Options for the DHCPv4 Wholesale Network Solution

You can use DHCP Relay, DHCP Relay Proxy, or DHCP Local Server configuration in a DHCP wholesale network. DHCP configuration is defined at the **[edit forwarding-options]** hierarchy level.



NOTE: The configuration for this wholesale solution uses DHCP Relay.

To configure DHCPv4 Relay forwarding options:

1. Access the **[edit forwarding-options dhcp-relay]** hierarchy.

```
[edit]  
user@host# edit forwarding-options dhcp-relay
```

2. Specify that you want to configure authentication options and use external AAA authentication services.

```
[edit forwarding-options dhcp-relay]  
user@host# edit authentication
```

3. (Optional) Configure a password that authenticates the username to the external authentication service.

See Configuring Passwords for Usernames.

4. (Optional) Configure optional features to create a unique username.

See Creating Unique Usernames for DHCP Clients.

5. Specify the default dynamic profile that you want to attach to all DHCP subscriber that access the router.

```
[edit forwarding-options dhcp-relay]
user@host# set dynamic-profile Wholesaler_Profile
```

6. Specify any overrides for the default DHCP Relay configuration.

See [Overriding the Default DHCP Relay Configuration Settings](#).

7. Configure a named server group for default (wholesaler) DHCP server access.

```
[edit forwarding-options dhcp-relay]
user@host# edit server-group Wholesaler_Group
```

8. Specify the DHCP server address for the default (wholesale) group.

```
[edit forwarding-options dhcp-relay server-group "Wholesaler_Group"]
user@host# set 192.168.100.1
```

9. Specify the default (wholesale) group as the active server group.

```
[edit forwarding-options dhcp-relay]
user@host# set active-server-group Wholesaler_Group
```

10. Configure a group you can use to define the wholesale DHCP access interface.

```
[edit forwarding-options dhcp-relay]
user@host# edit group Wholesaler_Group
```

11. Specify the default (wholesale) interface that all DHCP subscribers use when first accessing the router.

```
[edit forwarding-options dhcp-relay group "Wholesaler_Group"]
user@host# set interface ge-2/3/0.1
```

12. Configure a group you can use to define a retail DHCP interface.

```
[edit forwarding-options dhcp-relay]
user@host# edit group Retailer1_Group
```

13. Specify the logical interface the DHCP subscribers use once redirected.

```
[edit forwarding-options dhcp-relay group "Retailer1_Group"]
user@host# set interface ge-2/3/0.2
```

14. Repeat steps 12 and 13 for other retailer groups.

In this solution example, you configure another group name of "Retailer2_Group" and specify **ge-2/3/0.3** for the logical interface.

15. (Optional) Configure any passwords that authenticate the username to the external authentication service for any of the groups that you created.

See [Configuring Passwords for Usernames](#).

16. (Optional) Configure optional features to create a unique username for any of the groups that you created.

See [Creating Unique Usernames for DHCP Clients](#).

17. (Optional) Specify any overrides for any of the DHCP Relay group configurations that you created.

See [Overriding the Default DHCP Relay Configuration Settings](#).

**Related
Documentation**

- [Extended DHCP Relay Agent Overview](#)
- [DHCP Relay Proxy Overview](#)
- [Configuring Passwords for Usernames](#)
- [Creating Unique Usernames for DHCP Clients](#)
- [Overriding the Default DHCP Relay Configuration Settings](#)

CHAPTER 6

Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Network Configuration Examples

- [Example: Wholesaler Dynamic Profile for a DHCPv4 Wholesale Network on page 45](#)
- [Example: Retailer Dynamic Profile for a DHCPv4 Wholesale Network on page 46](#)
- [Example: Default Forwarding Options Configuration for the DHCPv4 Wholesale Network on page 46](#)
- [Example: Retailer Routing Instances for a DHCPv4 Wholesale Network on page 47](#)

Example: Wholesaler Dynamic Profile for a DHCPv4 Wholesale Network

This example specifies a dynamic profile name of *Wholesaler_Profile*, uses dynamic IP demux interfaces, and references the predefined input firewall filter.

```
dynamic-profiles {
  Wholesaler_Profile {
    interfaces {
      demux0 {
        unit "$junos-interface-unit" {
          demux-options {
            underlying-interface "$junos-underlying-interface";
          }
          family inet {
            demux-source {
              $junos-subscriber-ip-address;
            }
            filter {
              input "$junos-input-filter";
            }
            unnumbered-address "$junos-loopback-interface" preferred-source-address
              $junos-preferred-source-address;
          }
        }
      }
    }
  }
}
```

Example: Retailer Dynamic Profile for a DHCPv4 Wholesale Network

```
dynamic-profiles {
  Subscriber_Profile_Retailer1 {
    routing-instances {
      "$junos-routing-instance" {
        interface "$junos-interface-name";
      }
    }
    interfaces {
      demux0 {
        unit "$junos-interface-unit" {
          demux-options {
            underlying-interface "$junos-underlying-interface";
          }
          family inet {
            demux-source {
              "$junos-subscriber-ip-address";
            }
            unnumbered-address "$junos-loopback-interface" preferred-source-address
              "$junos-preferred-source-address";
          }
        }
      }
    }
  }
}
```

Example: Default Forwarding Options Configuration for the DHCPv4 Wholesale Network

```
forwarding-options {
  dhcp-relay {
    traceoptions {
      file size 1g;
      inactive: flag all;
    }
    authentication {
      password psswd;
      username-include {
        user-prefix WholesaleNetwork;
      }
    }
    dynamic-profile Wholesaler_Profile;
    overrides {
      always-write-giaddr;
      always-write-option-82;
      layer2-unicast-replies;
      trust-option-82;
      client-discover-match;
    }
    server-group {
      Wholesaler-Server-Group {
        192.168.100.1;
      }
    }
  }
}
```

```
active-server-group Wholesaler-Server Group;
group Wholesaler-Group {
  authentication {
    password psswd;
    username-include {
      user-prefix WholesaleNetwork;
    }
  }
  interface ge-2/3/0.1;
}
group Retailer1-Group {
  authentication {
    password psswd1;
    username-include {
      user-prefix WholesaleNetwork_Retailer1;
    }
  }
  interface ge-2/3/0.2;
}
group Retailer2-Group {
  authentication {
    password psswd2;
    username-include {
      user-prefix WholesaleNetwork_Retailer1;
    }
  }
  interface ge-2/3/0.3;
}
}
```

Example: Retailer Routing Instances for a DHCPv4 Wholesale Network

```
routing-instances {
  Retailer_Instance1 {
    instance-type vrf;
    access-profile Retailer_Access1;
    interface ge-11/1/9.10;
    interface ge-11/1/10.100;
    interface lo0.1;
    route-distinguisher 1:1;
    forwarding-options {
      dhcp-relay {
        authentication {
          password psswd1;
          username-include {
            user-prefix WholesaleNetwork_Retailer1;
          }
        }
      }
    }
    dynamic-profile Subscriber_Profile_Retailer1;
    overrides {
      always-write-giaddr;
      always-write-option-82;
      layer2-unicast-replies;
      trust-option-82;
    }
  }
}
```

```
        client-discover-match;
    }
    server-group {
        Retailer1-Server-Group {
            10.10.100.1;
        }
    }
    active-server-group Retailer1-Server-Group;
    group Retailer1-Group {
        authentication {
            password psswd1;
            username-include {
                user-prefix WholesaleNetwork_Retailer1;
            }
        }
        dynamic-profile Subscriber_Profile_Retailer1;
        overrides {
            always-write-giaddr;
            trust-option-82;
            client-discover-match;
        }
        interface ge-2/3/0.2;
    }
}
}
}
Retailer_Instance2 {
    instance-type vrf;
    access-profile Retailer_Access2;
    interface ge-7/1/9.10;
    interface ge-7/1/9.100;
    interface lo0.2;
    route-distinguisher 2:2;
    forwarding-options {
        dhcp-relay {
            authentication {
                password psswd2;
                username-include {
                    user-prefix WholesaleNetwork_Retailer2;
                }
            }
        }
        dynamic-profile Subscriber_Profile_Retailer2;
        overrides {
            always-write-giaddr;
            trust-option-82;
            client-discover-match;
        }
        server-group {
            Retailer2-Group {
                10.20.200.1;
            }
        }
        active-server-group Retailer2-Group;
        group Retailer2-Group {
            authentication {
                password psswd2;
            }
        }
    }
}
```



```
        username-include {  
            user-prefix psswd2;  
        }  
    }  
    dynamic-profile Subscriber_Profile_Retailer2;  
    overrides {  
        always-write-giaddr;  
        trust-option-82;  
        client-discover-match;  
    }  
    interface ge-2/3/0.3;  
} }  
}
```


PART 3

Administration

- [Subscriber Management AAA and Address Assignment Pool CLI Commands on page 53](#)
- [Subscriber Management DHCPv4 Local Server CLI Commands on page 65](#)
- [Subscriber Management DHCPv4 Relay CLI Commands on page 79](#)
- [Subscriber Management Interface CLI Commands on page 93](#)
- [Subscriber Management Subscriber CLI Commands on page 163](#)

CHAPTER 7

Subscriber Management AAA and Address Assignment Pool CLI Commands

show network-access aaa statistics

Syntax	<pre>show network-access aaa statistics <accounting> <address-assignment (client pool <i>pool-name</i>)> <dynamic-requests> <radius></pre>
Release Information	<p>Command introduced in Junos OS Release 9.1.</p> <p>Option address-assignment introduced in Junos OS Release 10.0.</p> <p>Option radius introduced in Junos OS Release 11.4.</p>
Description	Display AAA accounting, address-assignment, dynamic request statistics, and RADIUS settings and statistics.
Options	<p>accounting—(Optional) Display AAA accounting statistics.</p> <p>address-assignment (client pool <i>pool-name</i>)—(Optional) Display AAA address-assignment client and pool statistics.</p> <p>dynamic-requests—(Optional) Display AAA dynamic requests.</p> <p>radius— (Optional) Display RADIUS settings and statistics.</p>
Required Privilege Level	view
List of Sample Output	<p>show network-access aaa statistics accounting on page 56</p> <p>show network-access aaa statistics address-assignment client on page 56</p> <p>show network-access aaa statistics address-assignment pool on page 56</p> <p>show network-access aaa statistics dynamic-requests on page 56</p> <p>show network-access aaa statistics radius on page 56</p>
Output Fields	Table 5 on page 54 lists the output fields for the show network-access aaa statistics command. Output fields are listed in the approximate order in which they appear.

Table 5: show network-access aaa statistics Output Fields

Field Name	Field Description
Requests received	<ul style="list-style-type: none"> Number of accounting requests generated by the AAA framework. Number of dynamic requests received from the external server.
Accounting Response failures	Number of accounting requests not acknowledged (NAK) by the accounting server.
Accounting Response Success	Number of accounting requests acknowledged by the accounting server.
Requests timedout	Number of accounting requests to the accounting server that timed out.
Client	Client type; for example, DHCP, Mobile IP, PPP.

Table 5: show network-access aaa statistics Output Fields (*continued*)

Field Name	Field Description
Out of Memory	Number of times an address was not given to the client due to memory issues.
No Matches	Number of times there were no network matches for the pool.
Pool Name	Name of the address-assignment pool for this client.
Out of Addresses	Number of times there were no available addresses in the pool.
Address total	Number of addresses in the pool.
Addresses in use	Number of addresses in use.
Address Usage (percent)	Percentage of total addresses in use.
processed successfully	Number of dynamic requests processed successfully by the AAA framework.
errors during processing	Number of dynamic requests that resulted in processing errors by the AAA framework.
Link Name	Name of the secondary address-assignment pool to which the primary pool is linked.
Pool Usage	Percentage of allocated addresses in the specified address pool.
silently dropped	Number of dynamic requests dropped by the AAA framework due to multiple back-to-back or duplicate requests.
RADIUS Server	IP address of the RADIUS server to which the router is sending requests.
Profile	Name of the RADIUS profile associated with the RADIUS server. A RADIUS server can be associated with more than one RADIUS profile.
Configured	Configured maximum number of outstanding requests from the router to the RADIUS server for a specific profile. An outstanding request is a request to which the RADIUS server has not yet responded. The range of values is 0 through 2000 outstanding requests. The default value is 1000.
Current	Current number of outstanding requests from the router to the RADIUS server for a specific profile. An outstanding request is a request to which the RADIUS server has not yet responded.
Peak	Highest number of outstanding requests from the router to the RADIUS server for a specific profile at any point in time since the router was started or since the counter was last cleared. NOTE: If the value of this field is equal to the value of the Configured field, you may want to increase the value of the Configured field.
Exceeded	Number of times that the router attempted to send requests to the RADIUS server in excess of the configured maximum value for a specific profile. NOTE: If the value of this field is nonzero, you may want to increase the value of the Configured field.

Sample Output

```
show network-access user@host> show network-access aaa statistics accounting
aaa statistics      Accounting module statistics
                    Requests received: 0
                    Accounting Response failures: 0
                    Accounting Response Success: 0
                    Requests timedout: 0

show network-access user@host> show network-access aaa statistics address-assignment client
aaa statistics      Address-assignment statistics
address-assignment Client: jdhcpd
client              Out of Memory: 0
                    No Matches: 2

show network-access user@host> show network-access aaa statistics address-assignment pool isp_1
aaa statistics      Address-assignment statistics
address-assignment Pool Name: isp_1
pool                Pool Name: (all pools in chain)
                    Out of Memory: 0
                    Out of Addresses: 9
                    Address total: 47
                    Addresses in use: 47
                    Address Usage (percent): 100

show network-access user@host> show network-access aaa statistics dynamic-requests
aaa statistics      requests received: 0
dynamic-requests    processed successfully: 0
                    errors during processing: 0
                    silently dropped: 0

show network-access user@host> show network-access aaa statistics radius
aaa statistics      Outstanding Requests
RADIUS Server      Profile      Configured   Current   Peak   Exceeded
172.28.32.239      prof1        1000         0         1000   14
                   prof2        500          17         432    0
171.27.82.211      myprof       200          0         200    27
12.1.11.254        pppoe-auth   111          0         1       0
```


show network-access aaa statistics authentication

Syntax	show network-access aaa statistics authentication <detail>
Release Information	Command introduced in Junos OS Release 9.1. Option detail introduced in Junos OS Release 12.1.
Description	Display AAA authentication statistics.
Options	detail —(Optional) Displays detailed information about authentication.
Required Privilege Level	view
List of Sample Output	show network-access aaa statistics authentication on page 59 show network-access aaa statistics authentication detail on page 59
Output Fields	Table 6 on page 57 lists the output fields for the show network-access aaa statistics authentication command. Output fields are listed in the approximate order in which they appear.

Table 6: show network-access aaa statistics authentication Output Fields

Field Name	Field Description	Level of Output
Requests received	Number of authentication requests received from clients.	All levels
Multistack requests	Number of authentication requests for dual-stack subscribers.	All levels
Accepts	Number of authentication requests accepted by the authentication server.	All levels
Rejects	Number of authentication requests rejected by the authentication server.	All levels
Challenges	Number of authentication requests challenged by the authentication server.	All levels
Requests timed out	Number of authentication requests that timed out.	All levels
RADIUS authentication failures	Number of RADIUS authentication requests that have failed.	Detail
Queue request deleted	Number of queue requests that have been deleted.	Detail
Malformed reply	Number of malformed replies received from the RADIUS authentication server.	Detail

Table 6: show network-access aaa statistics authentication Output Fields (*continued*)

Field Name	Field Description	Level of Output
No server configured	Number of authentication requests that failed because no authentication server is configured.	Detail
Access Profile configuration not found	Number of authentication requests that failed because no access profile is configured.	Detail
Unable to create client record	Number of times that the router is unable to create the client record for the authentication request.	Detail
Unable to create client request	Number of times that the router is unable to create the client request for the authentication request.	Detail
Unable to build authentication request	Number of times that the router is unable to build the authentication request.	Detail
No server found	Number of requests to the authentication server that have timed out; the server is then considered to be down.	Detail
Unable to create handle	Number of authentication requests that have failed because of an internal allocation failure.	Detail
Unable to queue request	Number of times the router was unable to queue the request to the authentication server.	Detail
Invalid credentials	Number of times the router did not have proper authorization to access the authentication server.	Detail
Malformed request	Number of times the router request to the authentication server is malformed.	Detail
License unavailable	Number of times the router did not have a license to access the authentication server.	Detail
Redirect requested	Number of authentication requests that have been redirected based on routing instance.	Detail
Internal failure	Number of internal failures.	Detail
Local authentication failures	Number of times local authentication failed.	Detail
LDAP lookup failures	Number of times the LDAP lookup operation failed.	Detail

Sample Output

```
show network-access user@host> show network-access aaa statistics authentication
aaa statistics Authentication module statistics
authentication Requests received: 2118
Multistack requests: 0
Accepts: 261
Rejects: 975
Challenges: 0
Requests timed out: 882

show network-access user@host> show network-access aaa statistics authentication detail
aaa statistics Authentication module statistics
authentication detail Requests received: 2118
Multistack requests: 0
Accepts: 261
Rejects: 975
RADIUS authentication failures: 975
Queue request deleted: 0
Malformed reply: 0
No server configured: 0
Access Profile configuration not found: 0
Unable to create client record: 0
Unable to create client request: 0
Unable to build authentication request: 0
No server found: 975
Unable to create handle: 0
Unable to queue request: 0
Invalid credentials: 0
Malformed request: 0
License unavailable: 0
Redirect requested: 0
Internal failure: 0
Local authentication failures: 0
LDAP lookup failures: 0
Challenges: 0
Requests timed out: 882
```

show network-access aaa subscribers

Syntax	show network-access aaa subscribers <logical-system <i>logical-system-name</i>> <routing-instance <i>routing-instance-name</i>> <statistics> <username>
Release Information	Command introduced in Junos OS Release 9.1.
Description	Display subscriber-specific AAA statistics.
Options	<p>logical-system <i>logical-system-name</i>—(Optional) List subscribers in the specific logical system.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) List subscribers for the specific routing instance. If you do not specify a routing instance name, the default routing instance is assumed.</p> <p>statistics—(Optional) Display statistics for the subscriber events.</p> <p>username—(Optional) Display information for the specified subscriber.</p>
Required Privilege Level	view
List of Sample Output	show network-access aaa subscribers logical-system on page 61 show network-access aaa subscribers logical-system routing-instance on page 61 show network-access aaa subscribers statistics username on page 61 show network-access aaa subscribers username on page 62
Output Fields	Table 7 on page 60 lists the output fields for the show network-access aaa subscribers command. Output fields are listed in the approximate order in which they appear.

Table 7: show network-access aaa subscribers Output Fields

Field Name	Field Description
Challenge requests	Number of authentication requests challenged by the authentication server for this subscriber.
Challenge responses	Number of challenge responses sent by the subscriber to the authentication server.
START sent successfully	Number of accounting start requests generated by the AAA framework for this subscriber.
START send failures	Number of accounting start requests that failed to make it to the accounting server for this subscriber.
START ack received	Number of accounting start requests acknowledged by the accounting server for this subscriber.
INTERIM sent successfully	Number of accounting interim requests generated by the AAA framework for this subscriber.

Table 7: show network-access aaa subscribers Output Fields (*continued*)

Field Name	Field Description
INTERIM send failures	Number of accounting interim requests that failed to make it to the accounting server for this subscriber.
INTERIM ack received	Number of accounting interim requests acknowledged by the accounting server for this subscriber.
Requests received	Number of reauthentication requests received by the authentication server.
Successful responses	Number of successful reauthentication requests granted by the authentication server.
Aborts handled	Number of reauthentication requests aborted by the authentication server.
Service name	Name of the subscriber service.
Creation requests	Number of requests to create the service.
Deletion requests	Number of requests to delete the service.
Request timeouts	Number of times the service request was timed out.
Client type	Type of client; for example, DHCP, Mobile IP, PPP.
Session-ID	ID of the subscriber session.
Session uptime	How long the session has been up, in <i>HH:MM:SS</i> .
Accounting	Status of accounting, and type of accounting if accounting is on.

Sample Output

```

show network-access aaa subscribers logical-system
user@host> show network-access aaa subscribers logical-system
Username      Client type  Logical system/Routing instance
cbenson@addr.net  ppp         default
00010e020304.1231 dhcp         isp-bos-metro-12:isp-cmborg-12
conley@isp3.com  dhcp         default:isp-gtown-r3-00
0020df980102.2334 dhcp         isp-bos-metro-16:isp-cmborg-12

show network-access aaa subscribers logical-system routing-instance isp-cmborg-12-32
user@host> show network-access aaa subscribers logical-system routing-instance isp-cmborg-12-32
Username      Client type  Logical system/Routing instance
00010e020304.1231 dhcp         isp-bos-metro-12:isp-cmborg-12
conley@isp3.com  dhcp         default:isp-gtown-r3-00
0020df980102.2334 dhcp         isp-bos-metro-16:isp-cmborg-12

show network-access aaa subscribers statistics username 00010e020304.1231
user@host> show network-access aaa subscribers statistics username 00010e020304.1231
Authentication statistics
  Challenge requests: 0
  Challenge responses: 0
Accounting statistics
  START sent successfully: 1
  START send failures: 0

```

```
START ack received: 1
INTERIM sent successfully: 0
INTERIM send failures: 0
INTERIM ack received: 0
Re-authentication statistics
Requests received: 0
Sucessfull responses: 0
Aborts handled: 0
Service statistics
Service name: filter-serv
Creation requests: 1
Deletion requests: 0
Request timeouts: 0
Service name: filter-serv2
Creation requests: 144
Deletion requests: 0
Request timeouts: 144
```

```
show network-access aaa subscribers username fred@isp5.net
aaa subscribers
username
Logical system/Routing instance  Client type  Session-ID  Session uptime
Accounting
isp-bos-metro-16:isp-cmbrg-12    dhcp        7           01:12:56
on/volume
Service name      Service type  Quota      Accounting
I-Cast            volume       1200 Mbps  on/volume+time
Voip              on/volume
GamingBurst       time         6000 secs  on/volume
```

show network-access address-assignment pool

Syntax	show network-access address-assignment pool <i>pool-name</i> <logical-system <i>logical-system-name</i> > <routing-instance <i>routing-instance-name</i> >
Release Information	Command introduced in Junos OS Release 9.0.
Description	Display state information for each address-assignment pool.
Options	<p>none—Display information about clients that have obtained addresses from the address-assignment pool.</p> <p>pool <i>pool-name</i>—Display information about the specified address-assignment pool.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Perform this operation on the specified logical system.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Perform this operation on the specified routing instance.</p>
Required Privilege Level	view and system
List of Sample Output	show network-access address-assignment pool on page 63
Output Fields	Table 8 on page 63 lists the output fields for the show address-assignment pool command. Output fields are listed in the approximate order in which they appear.

Table 8: show network-access address-assignment pool Output Fields

Field Name	Field Description
IP address	IP address of the client.
Hardware address	MAC address of the client.
Type	Type of client.

Sample Output

```

user@host> show network-access address-assignment pool sunnywest logical-system ls1
routing-instance routinst2
IP address      Hardware address  Type
192.168.2.1     00:05:1b:00:b9:01  DHCP
192.168.2.2     00:05:1b:00:b9:02  DHCP
192.168.2.3     00:05:1b:00:b9:03  DHCP
192.168.2.4     00:05:1b:00:b9:04  DHCP

```


CHAPTER 8

Subscriber Management DHCPv4 Local Server CLI Commands

show dhcp server binding

Syntax	<pre>show dhcp server binding <address> <brief detail summary> <interface interface-name> <interfaces-vlan> <interfaces-wildcard> <logical-system logical-system-name> <routing-instance routing-instance-name></pre>
Release Information	Command introduced in Junos OS Release 9.0. Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.
Description	Display the address bindings in the client table on the extended Dynamic Host Configuration Protocol (DHCP) local server.
Options	<p>address—(Optional) Display DHCP binding information for a specific client identified by one of the following entries:</p> <ul style="list-style-type: none">• <i>ip-address</i>—The specified IP address.• <i>mac-address</i>—The specified MAC address.• <i>session-id</i>—The specified session ID. <p>brief detail summary—(Optional) Display the specified level of output about active client bindings. The default is brief, which produces the same output as show dhcp server binding.</p> <p>interface interface-name—(Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.</p> <p>interfaces-vlan—(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.</p> <p>interfaces-wildcard—(Optional) The set of interfaces on which to show the binding state information. This option supports the use of the wildcard character (*).</p> <p>logical-system logical-system-name—(Optional) Display information about active client bindings for DHCP clients on the specified logical system.</p> <p>routing-instance routing-instance-name—(Optional) Display information about active client bindings for DHCP clients on the specified routing instance.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• Clearing DHCP Bindings for Subscriber Access• clear dhcp server binding on page 73

List of Sample Output

- [show dhcp server binding on page 68](#)
- [show dhcp server binding detail on page 68](#)
- [show dhcp server binding interface <vlan-id> on page 69](#)
- [show dhcp server binding interface <svlan-id> on page 69](#)
- [show dhcp server binding <ip-address> on page 69](#)
- [show dhcp server binding <session-id> on page 69](#)
- [show dhcp server binding summary on page 69](#)
- [show dhcp server binding <interfaces-vlan> on page 69](#)
- [show dhcp server binding <interfaces-wildcard> on page 69](#)

Output Fields Table 9 on page 67 lists the output fields for the **show dhcp server binding** command. Output fields are listed in the approximate order in which they appear.

Table 9: show dhcp server binding Output Fields

Field Name	Field Description	Level of Output
<i>number</i> clients, (<i>number</i> init, <i>number</i> bound, <i>number</i> selecting, <i>number</i> requesting, <i>number</i> renewing, <i>number</i> releasing)	Summary counts of the total number of DHCP clients and the number of DHCP clients in each state.	summary
IP address	IP address of the DHCP client.	brief detail
Session Id	Session ID of the subscriber session.	brief detail
Hardware address	Hardware address of the DHCP client.	brief detail
Expires	Number of seconds in which lease expires.	brief detail
State	State of the address binding table on the extended DHCP local server: <ul style="list-style-type: none"> • BOUND—Client has active IP address lease. • FORCERENEW—Client has received forcerenew message from server. • INIT—Initial state. • RELEASE—Client is releasing IP address lease. • RENEWING—Client sending request to renew IP address lease. • REQUESTING—Client requesting a DHCP server. • SELECTING—Client receiving offers from DHCP servers. 	brief detail
Interface	Interface on which the request was received.	brief
Lease Expires	Date and time at which the client's IP address lease expires.	detail
Lease Expires in	Number of seconds in which lease expires.	detail

Table 9: show dhcp server binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Lease Start	Date and time at which the client's IP address lease started.	detail
Incoming Client Interface	Client's incoming interface.	detail
Server IP Address	IP address of DHCP server.	detail
Server Interface	Interface of DHCP server.	detail
Client Pool Name	Name of address pool used to assign client IP address lease.	detail

Sample Output

```

show dhcp server binding user@host> show dhcp server binding
binding                  IP address    Session Id    Hardware address    Expires    State    Interface
100.20.20.15             6             00:10:94:00:00:01 86180             BOUND      ge-1/0/0.0
100.20.20.16             7             00:10:94:00:00:02 86180             BOUND      ge-1/0/0.0
100.20.20.17             8             00:10:94:00:00:03 86180             BOUND      ge-1/0/0.0
100.20.20.18             9             00:10:94:00:00:04 86180             BOUND      ge-1/0/0.0
100.20.20.19             10            00:10:94:00:00:05 86180             BOUND      ge-1/0/0.0

show dhcp server binding detail user@host> show dhcp server binding detail
binding detail           Client IP Address: 100.20.20.15
                        Hardware Address:      00:10:94:00:00:01
                        State:                   BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)

                        Lease Expires:           2009-07-21 10:10:25 PDT
                        Lease Expires in:        86151 seconds
                        Lease Start:             2009-07-20 10:10:25 PDT
                        Incoming Client Interface: ge-1/0/0.0
                        Server Ip Address:        100.20.20.9
                        Server Interface:         none
                        Session Id:               6
                        Client Pool Name:         6
Client IP Address: 100.20.20.16
                        Hardware Address:        00:10:94:00:00:02
                        State:                   BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)

                        Lease Expires:           2009-07-21 10:10:25 PDT
                        Lease Expires in:        86151 seconds
                        Lease Start:             2009-07-20 10:10:25 PDT
                        Incoming Client Interface: ge-1/0/0.0
                        Server Ip Address:        100.20.20.9
                        Server Interface:         none
                        Session Id:               7
                        Client Pool Name:         7

```

```

show dhcp server binding interface ge-1/1/0:100
<vlan-id>
user@host> show dhcp server binding interface ge-1/1/0:100
IP address      Session Id  Hardware address  Expires  State  Interface
200.20.20.15    6          00:10:94:00:00:01 86124    BOUND  ge-1/1/0:100

show dhcp server binding interface ge-1/1/0:10-100
<svlan-id>
user@host> show dhcp server binding interface ge-1/1/0:10-100
IP address      Session Id  Hardware address  Expires  State  Interface
200.20.20.16    7          00:10:94:00:00:02 86124    BOUND  ge-1/1/0:10-100

show dhcp server binding <ip-address> 100.20.20.19
user@host> show dhcp server binding 100.20.20.19
IP address      Session Id  Hardware address  Expires  State  Interface
100.20.20.19    10         00:10:94:00:00:05 86081    BOUND  ge-1/0/0.0

show dhcp server binding <session-id> 200.20.20.15
user@host> show dhcp server binding 6
IP address      Session Id  Hardware address  Expires  State  Interface
200.20.20.15    6          00:10:94:00:00:01 86124    BOUND  ge-1/0/0.0

show dhcp server binding summary
user@host> show dhcp server binding summary
3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)

show dhcp server binding ge-1/0/0:100-200
<interfaces-vlan>
user@host> show dhcp server binding ge-1/0/0:100-200
IP address      Session Id  Hardware address  Expires  State  Interface
192.168.0.17    42         00:10:94:00:00:02 86346    BOUND  ge-1/0/0.1073741827
192.168.0.16    41         00:10:94:00:00:01 86346    BOUND  ge-1/0/0.1073741827

show dhcp server binding ge-1/3/*
<interfaces-wildcard>
user@host> show dhcp server binding ge-1/3/*
IP address      Session Id  Hardware address  Expires  State  Interface
192.168.0.9     24         00:10:94:00:00:04 86361    BOUND  ge-1/3/0.110
192.168.0.8     23         00:10:94:00:00:03 86361    BOUND  ge-1/3/0.110
192.168.0.7     22         00:10:94:00:00:02 86361    BOUND  ge-1/3/0.110

```

show dhcp server statistics

Syntax	show dhcp server statistics <logical-system <i>logical-system-name</i>> <routing-instance <i>routing-instance-name</i>>
Release Information	Command introduced in Junos OS Release 9.0.
Description	Display extended Dynamic Host Configuration Protocol (DHCP) local server statistics.
Options	logical-system <i>logical-system-name</i> —(Optional) Display information about extended DHCP local server statistics on the specified logical system. If you do not specify a logical system, statistics are displayed for the default logical system. routing-instance <i>routing-instance-name</i> —(Optional) Display information about extended DHCP local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear dhcp server statistics on page 76
List of Sample Output	show dhcp server statistics on page 71
Output Fields	Table 10 on page 71 lists the output fields for the show dhcp server statistics command. Output fields are listed in the approximate order in which they appear.

Table 10: show dhcp server statistics Output Fields

Field Name	Field Description
Packets dropped	<p>Number of packets discarded by the extended DHCP local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.</p> <ul style="list-style-type: none"> • Total—Total number of packets discarded by the extended DHCP local server • Authentication—Number of packets discarded because they could not be authenticated • Bad hardware address—Number of packets discarded because an invalid hardware address was specified • Bad opcode—Number of packets discarded because an invalid operation code was specified • Bad options—Number of packets discarded because invalid options were specified • Dynamic profile—Number of packets discarded due to dynamic profile information • Invalid server address—Number of packets discarded because an invalid server address was specified • No available addresses—Number of packets discarded because there were no addresses available for assignment • No interface match—Number of packets discarded because they did not belong to a configured interface • No routing instance match—Number of packets discarded because they did not belong to a configured routing instance • No valid local address—Number of packets discarded because there was no valid local address • Packet too short—Number of packets discarded because they were too short • Read error—Number of packets discarded because of a system read error • Send error—Number of packets that the extended DHCP local server could not send
Messages received	<p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> • BOOTREQUEST—Number of BOOTP protocol data units (PDUs) received • DHCPDECLINE—Number of DHCP PDUs of type DECLINE received • DHCPDISCOVER—Number of DHCP PDUs of type DISCOVER received • DHCPINFORM—Number of DHCP PDUs of type INFORM received • DHCPRELEASE—Number of DHCP PDUs of type RELEASE received • DHCPREQUEST—Number of DHCP PDUs of type REQUEST received
Messages sent	<p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> • BOOTREPLY—Number of BOOTP PDUs transmitted • DHCPOFFER—Number of DHCP OFFER PDUs transmitted • DHCPACK—Number of DHCP ACK PDUs transmitted • DHCPNACK—Number of DHCP NACK PDUs transmitted • DHCPFORCERENEW—Number of DHCP FORCERENEW PDUs transmitted

Sample Output

```

show dhcp server statistics user@host> show dhcp server statistics
Packets dropped:
    Total                                0

Messages received:
    BOOTREQUEST                        25

```

DHCPDECLINE	0
DHCPDISCOVER	10
DHCPINFORM	0
DHCPRELEASE	4
DHCPREQUEST	10
Messages sent:	
BOOTREPLY	20
DHCPOFFER	10
DHCPACK	10
DHCPNAK	0
DHCPFORCERENEW	0

clear dhcp server binding

Syntax clear dhcp server binding
 <address>
 <all>
 <interface *interface-name*>
 <interfaces-vlan>
 <interfaces-wildcard>
 <logical-system *logical-system-name*>
 <routing-instance *routing-instance-name*>

Release Information Command introduced in Junos OS Release 9.0.
 Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

Description Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table on the extended DHCP local server.

Options *address*—(Optional) Clear the binding state for the DHCP client, using one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

all—(Optional) Clear the binding state for all DHCP clients.

interface interface-name—(Optional) Clear the binding state for DHCP clients on the specified interface.



NOTE: This option clears all bindings whose initial login requests were received over the specified interface. Dynamic demux login requests are not received over the dynamic demux interface, but rather the underlying interface of the dynamic demux interface. To clear a specific dynamic demux interface, use the *ip-address* or *mac-address* options.

interfaces-vlan—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.

interfaces-wildcard—(Optional) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (*).

logical-system logical-system-name—(Optional) Clear the binding state for DHCP clients on the specified logical system.

routing-instance routing-instance-name—(Optional) Clear the binding state for DHCP clients on the specified routing instance.

Required Privilege Level view

Related Documentation

- Clearing DHCP Bindings for Subscriber Access
- [show dhcp server binding on page 66](#)

List of Sample Output

- [clear dhcp server binding <ip-address> on page 74](#)
- [clear dhcp server binding all on page 74](#)
- [clear dhcp server binding interface on page 74](#)
- [clear dhcp server binding <interfaces-vlan> on page 74](#)
- [clear dhcp server binding <interfaces-wildcard> on page 75](#)

Output Fields See [show dhcp server binding](#) for an explanation of output fields.

Sample Output

clear dhcp server binding <ip-address> The following sample output displays the address bindings in the DHCP client table on the extended DHCP local server before and after the **clear dhcp server binding** command is issued.

```
user@host> show dhcp server binding
```

```
2 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)
```

IP address	Hardware address	Type	Lease expires at
100.20.32.1	90:00:00:01:00:01	active	2007-01-17 11:38:47 PST
100.20.32.3	90:00:00:02:00:01	active	2007-01-17 11:38:41 PST

```
user@host> clear dhcp server binding 10.20.32.1
```

```
user@host> show dhcp server binding
```

```
1 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)
```

IP address	Hardware address	Type	Lease expires at
100.20.32.3	90:00:00:02:00:01	active	2007-01-17 11:38:41 PST

clear dhcp server binding all The following command clears all DHCP local server bindings:

```
user@host> clear dhcp server binding all
```

clear dhcp server binding interface The following command clears DHCP local server bindings on a specific interface:

```
user@host> clear dhcp server binding interface fe-0/0/2
```

clear dhcp server binding <interfaces-vlan> The following command uses the *interfaces-vlan* option to clear all DHCP local server bindings on top of the underlying interface **ae0**, which clears DHCP bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcp server binding ae0
```

clear dhcp server binding The following command uses the *interfaces-wildcard* option to clear all DHCP local server bindings over a specific interface:
<interfaces-wildcard> user@host> **clear dhcp server binding ge-1/0/0.***

clear dhcp server statistics

Syntax	clear dhcp server statistics <interface <i>interface-name</i>> <logical-system <i>logical-system-name</i>> <routing-instance <i>routing-instance-name</i>>
Release Information	Command introduced in Junos OS Release 9.0.
Description	Clear all extended Dynamic Host Configuration Protocol (DHCP) local server statistics.
Options	logical-system <i>logical-system-name</i> —(Optional) Clear the statistics for DHCP clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system. routing-instance <i>routing-instance-name</i> —(Optional) Clear the statistics for DHCP clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.
Required Privilege Level	view
List of Sample Output	clear dhcp server statistics on page 76
Output Fields	See show dhcp server statistics for an explanation of output fields.

Sample Output

clear dhcp server statistics The following sample output displays the extended DHCP local server statistics before and after the **clear dhcp server statistics** command is issued.

```
user@host> show dhcp server statistics
Packets dropped:
  Total                  0

Messages received:
  BOOTREQUEST           89163
  DHCPDECLINE            0
  DHCPDISCOVER           8110
  DHCPINFORM             0
  DHCPRELEASE            0
  DHCPREQUEST           81053

Messages sent:
  BOOTREPLY              32420
  DHCPOFFER              8110
  DHCPACK                8110
  DHCPNAK                8100

user@host> clear dhcp server statistics
user@host> show dhcp server statistics
Packets dropped:
  Total                  0
```

```
Messages received:
  BOOTREQUEST      0
  DHCPDECLINE      0
  DHCPDISCOVER     0
  DHCPINFORM       0
  DHCPRELEASE      0
  DHCPREQUEST      0

Messages sent:
  BOOTREPLY        0
  DHCPOFFER        0
  DHCPACK          0
  DHCPNAK          0
```


CHAPTER 9

Subscriber Management DHCPv4 Relay CLI Commands

show dhcp relay binding

Syntax	show dhcp relay binding <address> <brief> <detail> <interface <i>interface-name</i>> <interfaces-vlan> <interfaces-wildcard> <ip-address mac-address> <logical-system <i>logical-system-name</i>> <routing-instance <i>routing-instance-name</i>> <summary>
Release Information	Command introduced in Junos OS Release 8.3. Options interface and mac-address added in Junos OS Release 8.4. Options interfaces-vlan and interfaces-wildcard added in Junos OS Release 12.1.
Description	Display the address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.
Options	<p>address—(Optional) Display DHCP binding information for a specific client identified by one of the following entries:</p> <ul style="list-style-type: none">• ip-address—The specified IP address.• mac-address—The specified MAC address.• session-id—The specified session ID. <p>brief—(Optional) Display brief information about the active client bindings. This is the default, and produces the same output as show dhcp relay binding.</p> <p>detail—(Optional) Display detailed client binding information.</p> <p>interface <i>interface-name</i>—(Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.</p> <p>interfaces-vlan—(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.</p> <p>interfaces-wildcard—(Optional) The set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).</p> <p>logical-system <i>logical-system-name</i>—(Optional) Perform this operation on the specified logical system.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Perform this operation on the specified routing instance.</p> <p>summary—(Optional) Display a summary of DHCP client information.</p>

Required Privilege Level view

Related Documentation

- Clearing DHCP Bindings for Subscriber Access
- [clear dhcp relay binding on page 88](#)

List of Sample Output

- [show dhcp relay binding on page 82](#)
- [show dhcp relay binding detail on page 82](#)
- [show dhcp relay binding interface on page 83](#)
- [show dhcp relay binding interface vlan-id on page 83](#)
- [show dhcp relay binding interface svlan-id on page 83](#)
- [show dhcp relay binding ip-address on page 83](#)
- [show dhcp relay binding mac-address on page 83](#)
- [show dhcp relay binding session-id on page 83](#)
- [show dhcp relay binding <interfaces-vlan> on page 83](#)
- [show dhcp relay binding <interfaces-wildcard> on page 84](#)
- [show dhcp relay binding summary on page 84](#)

Output Fields Table 11 on page 81 lists the output fields for the **show dhcp relay binding** command. Output fields are listed in the approximate order in which they appear.

Table 11: show dhcp relay binding Output Fields

Field Name	Field Description	Level of Output
<i>number</i> clients, (<i>number</i> init, <i>number</i> bound, <i>number</i> selecting, <i>number</i> requesting, <i>number</i> renewing, <i>number</i> rebinding, <i>number</i> releasing)	Summary counts of the total number of DHCP clients and the number of DHCP clients in each state.	summary
IP address	IP address of the DHCP client.	briefdetail
Session Id	Session ID of the subscriber session.	briefdetail
Hardware address	Hardware address of the DHCP client.	briefdetail
Expires	Number of seconds in which the lease expires.	briefdetail
State	State of the DHCP relay address binding table on the DHCP client: <ul style="list-style-type: none"> • BOUND—Client has an active IP address lease. • INIT—Initial state. • REBINDING—Client is broadcasting a request to renew the IP address lease. • RELEASE—Client is releasing the IP address lease. • RENEWING—Client is sending a request to renew the IP address lease. • REQUESTING—Client is requesting a DHCP server. • SELECTING—Client is receiving offers from DHCP servers. 	briefdetail

Table 11: show dhcp relay binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Interface	Incoming client interface.	brief
Lease Expires	Date and time at which the client's IP address lease expires.	detail
Lease Expires in	Number of seconds in which the lease expires.	detail
Lease Start	Date and time at which the client's IP address lease started.	detail
Incoming Client Interface	Client's incoming interface.	detail
Server IP Address	IP address of the DHCP server.	detail
Server Interface	Interface of the DHCP server.	detail
Bootp Relay Address	IP address of BOOTP relay.	detail
Type	Type of DHCP packet processing performed on the router: <ul style="list-style-type: none"> active—Router actively processes and relays DHCP packets. passive—Router passively snoops DHCP packets passing through the router. 	All levels
Lease expires at	Date and time at which the client's IP address lease expires.	All levels

Sample Output

```

show dhcp relay binding user@host> show dhcp relay binding
IP address      Session Id  Hardware address  Expires  State  Interface
100.20.32.11    41         00:10:94:00:00:01 86371    BOUND  ge-1/0/0.0
100.20.32.12    42         00:10:94:00:00:02 86371    BOUND  ge-1/0/0.0
100.20.32.13    43         00:10:94:00:00:03 86371    BOUND  ge-1/0/0.0
100.20.32.14    44         00:10:94:00:00:04 86371    BOUND  ge-1/0/0.0
100.20.32.15    45         00:10:94:00:00:05 86371    BOUND  ge-1/0/0.0

show dhcp relay binding detail user@host> show dhcp relay binding detail
Client IP Address: 100.20.32.11
Hardware Address:   00:10:94:00:00:01
State:              BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
Lease Expires:      2009-07-21 11:00:06 PDT
Lease Expires in:   86361 seconds
Lease Start:        2009-07-20 11:00:06 PDT
Last Packet Received: 2009-07-20 11:00:06 PDT
Incoming Client Interface: ge-1/0/0.0

```

```

Server Ip Address:      100.20.22.2
Server Interface:      none
Bootp Relay Address:   100.20.32.2
Session Id:            41

```

```

Client IP Address: 100.20.32.12
Hardware Address:   00:10:94:00:00:02
State:              BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
Lease Expires:      2009-07-21 11:00:06 PDT
Lease Expires in:   86361 seconds
Lease Start:        2009-07-20 11:00:06 PDT
Last Packet Received: 2009-07-20 11:00:06 PDT
Incoming Client Interface: ge-1/0/0.0
Server Ip Address:   100.20.22.2
Server Interface:    none
Bootp Relay Address: 100.20.32.2
Session Id:          42

```

```

show dhcp relay binding interface user@host> show dhcp relay binding interface fe-0/0/2

```

```

IP address      Hardware address  Type      Lease expires at
100.20.32.1     90:00:00:01:00:01  active    2007-03-27 15:06:20 EDT

```

```

show dhcp relay binding interface user@host> show dhcp relay binding interface ge-1/1/0:100
vlan-id

```

```

IP address      Session Id  Hardware address  Expires  State  Interface
200.20.20.15    6          00:10:94:00:00:01  86124    BOUND  ge-1/1/0:100

```

```

show dhcp relay binding interface user@host> show dhcp relay binding interface ge-1/1/0:10-100
svlan-id

```

```

IP address      Session Id  Hardware address  Expires  State  Interface
200.20.20.16    7          00:10:94:00:00:02  86124    BOUND  ge-1/1/0:10-100

```

```

show dhcp relay binding ip-address user@host> show dhcp relay binding 100.20.32.13

```

```

IP address      Session Id  Hardware address  Expires  State  Interface
100.20.32.13    43         00:10:94:00:00:03  86293    BOUND  ge-1/0/0.0

```

```

show dhcp relay binding mac-address user@host> show dhcp relay binding 00:10:94:00:00:05

```

```

IP address      Session Id  Hardware address  Expires  State  Interface
100.20.32.15    45         00:10:94:00:00:05  86279    BOUND  ge-1/0/0.0

```

```

show dhcp relay binding session-id user@host> show dhcp relay binding 41

```

```

IP address      Session Id  Hardware address  Expires  State  Interface
100.20.32.11    41         00:10:94:00:00:01  86305    BOUND  ge-1/0/0.0

```

```

show dhcp relay binding user@host> show dhcp relay binding ge-1/0/0:100-200
<interfaces-vlan>

```

```

IP address      Session Id  Hardware address  Expires  State  Interface
192.168.0.17    42         00:10:94:00:00:02  86346    BOUND  ge-1/0/0.1073741827
192.168.0.16    41         00:10:94:00:00:01  86346    BOUND  ge-1/0/0.1073741827

```

```
show dhcp relay binding <interfaces-wildcard>
user@host> show dhcp relay binding ge-1/3/*
IP address      Session Id  Hardware address  Expires  State  Interface
192.168.0.9     24         00:10:94:00:00:04  86361   BOUND  ge-1/3/0.110
192.168.0.8     23         00:10:94:00:00:03  86361   BOUND  ge-1/3/0.110
192.168.0.7     22         00:10:94:00:00:02  86361   BOUND  ge-1/3/0.110

show dhcp relay binding summary
user@host> show dhcp relay binding summary
3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 rebinding,
0 releasing)
```

show dhcp relay statistics

Syntax	<pre>show dhcp relay statistics <logical-system <i>logical-system-name</i>> <routing-instance <i>routing-instance-name</i>></pre>
Syntax	<p>Syntax for EX Series switches:</p> <pre>show dhcp relay statistics <routing-instance <i>routing-instance-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	Display Dynamic Host Configuration Protocol (DHCP) relay statistics.
Options	<p>logical-system <i>logical-system-name</i>—(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear dhcp relay statistics on page 90
List of Sample Output	show dhcp relay statistics on page 86
Output Fields	<p>Table 12 on page 86 lists the output fields for the show dhcp relay statistics command. Output fields are listed in the approximate order in which they appear.</p>

Table 12: show dhcp relay statistics Output Fields

Field Name	Field Description
Packets dropped	<p>Number of packets discarded by the extended DHCP relay agent application due to errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.</p> <ul style="list-style-type: none"> • Total—Total number of packets discarded by the extended DHCP relay agent application. • Bad hardware address—Number of packets discarded because an invalid hardware address was specified. • Bad opcode—Number of packets discarded because an invalid operation code was specified. • Bad options—Number of packets discarded because invalid options were specified. • Invalid server address—Number of packets discarded because an invalid server address was specified. • No available addresses—Number of packets discarded because there were no addresses available for assignment. • No interface match—Number of packets discarded because they did not belong to a configured interface. • No routing instance match—Number of packets discarded because they did not belong to a configured routing instance. • No valid local address—Number of packets discarded because there was no valid local address. • Packet too short—Number of packets discarded because they were too short. • Read error—Number of packets discarded because of a system read error. • Send error—Number of packets that the extended DHCP relay application could not send. • Option 60—Number of packets discarded containing DHCP option 60 vendor-specific information. • Option 82—Number of packets discarded because DHCP option 82 information could not be added.
Messages received	<p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> • BOOTREQUEST—Number of BOOTP protocol data units (PDUs) received • DHCPDECLINE—Number of DHCP PDUs of type DECLINE received • DHCPDISCOVER—Number of DHCP PDUs of type DISCOVER received • DHCPINFORM—Number of DHCP PDUs of type INFORM received • DHCPRELEASE—Number of DHCP PDUs of type RELEASE received • DHCPREQUEST—Number of DHCP PDUs of type REQUEST received
Messages sent	<p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> • BOOTREPLY—Number of BOOTP PDUs transmitted • DHCPOFFER—Number of DHCP OFFER PDUs transmitted • DHCPACK—Number of DHCP ACK PDUs transmitted • DHCPNACK—Number of DHCP NACK PDUs transmitted • DHCPFORCERENEW—Number of DHCP FORCERENEW PDUs transmitted

Sample Output

```

show dhcp relay statistics user@host> show dhcp relay statistics
Packets dropped:
    Total                  30
    Bad hardware address   1
    Bad opcode             1
    Bad options            3

```

Invalid server address	5
No available addresses	1
No interface match	2
No routing instance match	9
No valid local address	4
Packet too short	2
Read error	1
Send error	1
Option 60	1
Option 82	2

Messages received:

BOOTREQUEST	116
DHCPDECLINE	0
DHCPDISCOVER	11
DHCPINFORM	0
DHCPRELEASE	0
DHCPREQUEST	105

Messages sent:

BOOTREPLY	0
DHCPOFFER	2
DHCPACK	1
DHCPNAK	0
DHCPFORCERENEW	0

clear dhcp relay binding

Syntax	<pre>clear dhcp relay binding <address> <all> <interface interface-name> <interfaces-vlan> <interfaces-wildcard> <logical-system logical-system-name> <routing-instance routing-instance-name></pre>
Release Information	Command introduced in Junos OS Release 8.3. Options all and interface added in Junos OS Release 8.4. Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.
Description	Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table.
Options	<p>address—(Optional) Clear the binding state for the DHCP client, using one of the following entries:</p> <ul style="list-style-type: none">• <i>ip-address</i>—The specified IP address.• <i>mac-address</i>—The specified MAC address.• <i>session-id</i>—The specified session ID. <p>all—(Optional) Clear the binding state for all DHCP clients.</p> <p>interface interface-name—(Optional) Clear the binding state for DHCP clients on the specified interface.</p> <p>interfaces-vlan—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.</p> <p>interfaces-wildcard—(Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).</p> <p>logical-system logical-system-name—(Optional) Clear the binding state for DHCP clients on the specified logical system.</p> <p>routing-instance routing-instance-name—(Optional) Clear the binding state for DHCP clients on the specified routing instance.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• Clearing DHCP Bindings for Subscriber Access• show dhcp relay binding on page 80
List of Sample Output	clear dhcp relay binding on page 89

[clear dhcp relay binding all on page 89](#)
[clear dhcp relay binding interface on page 89](#)
[clear dhcp relay binding <interfaces-vlan> on page 89](#)
[clear dhcp relay binding <interfaces-wildcard> on page 89](#)

Output Fields See [show dhcp relay binding](#) for an explanation of output fields.

Sample Output

clear dhcp relay binding The following sample output displays the address bindings in the DHCP client table before and after the **clear dhcp relay binding** command is issued.

```
user@host> show dhcp relay binding
IP address      Hardware address  Type    Lease expires at
100.20.32.1     90:00:00:01:00:01 active    2007-02-08 16:41:17 EST
192.168.14.8    90:00:01:01:02:01 active    2007-02-10 10:01:06 EST
```

```
user@host> clear dhcp relay binding 100.20.32.1
```

```
user@host> show dhcp relay binding
IP address      Hardware address  Type    Lease expires at
192.168.14.8    90:00:01:01:02:01 active    2007-02-10 10:01:06 EST
```

clear dhcp relay binding all The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcp relay binding all
```

clear dhcp relay binding interface The following command clears DHCP relay agent bindings on a specific interface:

```
user@host> clear dhcp relay binding interface fe-0/0/3
```

clear dhcp relay binding <interfaces-vlan> The following command uses the *interfaces-vlan* option to clear all DHCP relay agent bindings on top of the underlying interface **ae0**, which clears DHCP bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcp relay binding interface ae0
```

clear dhcp relay binding <interfaces-wildcard> The following command uses the *interfaces-wildcard* option to clear all DHCP relay agent bindings over a specific interface:

```
user@host> clear dhcp relay binding ge-1/0/0.*
```

clear dhcp relay statistics

Syntax	<code>clear dhcp relay statistics</code> <code><logical-system <i>logical-system-name</i>></code> <code><routing-instance <i>routing-instance-name</i>></code>
Syntax	Syntax for EX Series switches: <code>show dhcp relay statistics</code> <code><routing-instance <i>routing-instance-name</i>></code>
Release Information	Command introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.1 for EX Series switches.
Description	Clear all Dynamic Host Configuration Protocol (DHCP) relay statistics.
Options	<code>logical-system <i>logical-system-name</i></code> —(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system. <code>routing-instance <i>routing-instance-name</i></code> —(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show dhcp relay statistics on page 85
List of Sample Output	clear dhcp relay statistics on page 91
Output Fields	Table 13 on page 91 lists the output fields for the <code>clear dhcp relay statistics</code> command.

Table 13: clear dhcp relay statistics Output Fields

Field Name	Field Description
Packets dropped	<p>Number of packets discarded by the extended DHCP relay agent application due to errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.</p> <ul style="list-style-type: none"> • Total—Total number of packets discarded by the extended DHCP relay agent application. • Bad hardware address—Number of packets discarded because an invalid hardware address was specified. • Bad opcode—Number of packets discarded because an invalid operation code was specified. • Bad options—Number of packets discarded because invalid options were specified. • Invalid server address—Number of packets discarded because an invalid server address was specified. • No available addresses—Number of packets discarded because there were no addresses available for assignment. • No interface match—Number of packets discarded because they did not belong to a configured interface. • No routing instance match—Number of packets discarded because they did not belong to a configured routing instance. • No valid local address—Number of packets discarded because there was no valid local address. • Packet too short—Number of packets discarded because they were too short. • Read error—Number of packets discarded because of a system read error. • Send error—Number of packets that the extended DHCP relay application could not send. • Option 60—Number of packets discarded containing DHCP option 60 vendor-specific information. • Option 82—Number of packets discarded because DHCP option 82 information could not be added.
Messages received	<p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> • BOOTREQUEST—Number of BOOTP protocol data units (PDUs) received • DHCPDECLINE—Number of DHCP PDUs of type DECLINE received • DHCPDISCOVER—Number of DHCP PDUs of type DISCOVER received • DHCPINFORM—Number of DHCP PDUs of type INFORM received • DHCPRELEASE—Number of DHCP PDUs of type RELEASE received • DHCPREQUEST—Number of DHCP PDUs of type REQUEST received
Messages sent	<p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> • BOOTREPLY—Number of BOOTP PDUs transmitted • DHCPOFFER—Number of DHCP OFFER PDUs transmitted • DHCPACK—Number of DHCP ACK PDUs transmitted • DHCPNACK—Number of DHCP NACK PDUs transmitted

Sample Output

clear dhcp relay statistics The following sample output displays the DHCP relay statistics before and after the **clear dhcp relay statistics** command is issued.

```
user@host> show dhcp relay statistics
Packets dropped:
    Total                  0
```

```
Messages received:
  BOOTREQUEST      116
  DHCPDECLINE      0
  DHCPDISCOVER     11
  DHCPINFORM       0
  DHCPRELEASE      0
  DHCPREQUEST     105
```

```
Messages sent:
  BOOTREPLY        44
  DHCPOFFER        11
  DHCPACK          11
  DHCPNAK          11
```

```
user@host> clear dhcp relay statistics
```

```
user@host> show dhcp relay statistics
```

```
Packets dropped:
  Total            0
```

```
Messages received:
  BOOTREQUEST      0
  DHCPDECLINE      0
  DHCPDISCOVER     0
  DHCPINFORM       0
  DHCPRELEASE      0
  DHCPREQUEST      0
```

```
Messages sent:
  BOOTREPLY        0
  DHCPOFFER        0
  DHCPACK          0
  DHCPNAK          0
```

CHAPTER 10

Subscriber Management Interface CLI Commands

show interfaces (Loopback)

Syntax `show interfaces lo0`
`<brief | detail | extensive | terse>`
`<descriptions>`
`<media>`
`<snmp-index snmp-index>`
`<statistics>`

Release Information Command introduced before Junos OS Release 7.4.

Description Display status information about the local loopback interface.



NOTE: Logical interface lo0.16385 is the loopback interface for the internal routing instance. Created by the internal routing service process, this interface facilitates internal traffic. It prevents any filter created on loopback lo0.0 from blocking internal traffic.

Options `lo0`—Display standard status information about the local loopback interface.

`brief | detail | extensive | terse`—(Optional) Display the specified level of output.

`descriptions`—(Optional) Display interface description strings.

`media`—(Optional) Display media-specific information.

`snmp-index snmp-index`—(Optional) Display information for the specified SNMP index of the interface.

`statistics`—(Optional) Display static interface statistics.

Required Privilege Level view

List of Sample Output [show interfaces \(Loopback\) on page 97](#)
[show interfaces brief \(Loopback\) on page 98](#)
[show interfaces detail \(Loopback\) on page 98](#)
[show interfaces extensive \(Loopback\) on page 99](#)

Output Fields [Table 14 on page 94](#) lists the output fields for the **show interfaces** (loopback) command. Output fields are listed in the approximate order in which they appear.

Table 14: Loopback show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical Interface	Name of the physical interface.	All levels

Table 14: Loopback show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under Common Output Fields Description.	All levels
Interface index	Physical interface index number, 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
Type	Type of interface.	All levels
Link-level type	Encapsulation type used on the physical interface.	All levels
MTU	Size of the largest packet to be transmitted.	All levels
Clocking	Reference clock source of the interface.	All levels
Speed	Network speed on the interface.	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under Common Output Fields Description.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under Common Output Fields Description.	All levels
Link type	Data transmission type.	detail extensive
Link flags	Information about the link. Possible values are described in the “Link Flags” section under Common Output Fields Description.	detail extensive none
Physical info	Information about the physical interface.	detail extensive
Hold-times	Current interface hold-time up and hold-time down. Value is in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive
Hardware address	Media access control (MAC) address of the interface.	detail extensive
Alternate link address	Backup link address.	detail extensive
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: <i>year-month-day hour:minute:second timezone (hour:minute:second ago)</i> . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive

Table 14: Loopback show interfaces 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, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. 	detail extensive
Input errors	<ul style="list-style-type: none"> • Errors—Input errors on the interface. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Frames received smaller than the runt threshold. • Giants—Frames received larger than the giant threshold. • Policed Discards—Frames that the incoming packet match code discarded because the frames were not recognized or were not of interest. Usually, this field reports protocols that Junos does not support. • Resource errors—Sum of transmit drops. 	extensive
Output errors	<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, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning. • Errors—Sum of 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 dropped by the ASIC RED mechanism. • MTU errors—Number of packets larger than the MTU threshold. • Resource errors—Sum of transmit drops. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	detail extensive
SNMP ifIndex	Logical interface SNMP interface index number.	detail extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface; values are described in the “Logical Interface Flags” section under Common Output Fields Description.	brief detail extensive
Encapsulation	Encapsulation on the logical interface.	brief detail extensive
Input packets	Number of packets received on the logical interface.	None specified

Table 14: Loopback show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output packets	Number of packets transmitted on the logical interface.	None specified
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	detail extensive
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	detail extensive
Protocol	Protocol family configured on the logical interface (such as iso or inet6).	detail extensive none
MTU	MTU size on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which this address exists; for example, Route table:0 refers to inet.0.	detail extensive
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under Common Output Fields Description.	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under Common Output Fields Description.	detail extensive
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 on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

```

show interfaces user@host> show interfaces lo0
  (Loopback) Physical interface: lo0, Enabled, Physical link is Up
              Interface index: 6, SNMP ifIndex: 6
              Type: Loopback, MTU: Unlimited
              Device flags   : Present Running Loopback
              Interface flags: SNMP-Traps
              Link flags     : None
              Last flapped   : Never
              Input packets  : 0
              Output packets : 0

              Logical interface lo0.0 (Index 64) (SNMP ifIndex 16)

```

```
Flags: SNMP-Traps Encapsulation: Unspecified
Input packets : 0
Output packets: 0
Protocol inet, MTU: Unlimited
  Flags: None
  Addresses, Flags: Is-Default Is-Primary
    Local: 10.0.0.1
  Addresses
    Local: 127.0.0.1
Protocol iso, MTU: Unlimited
  Flags: None
  Addresses, Flags: Is-Default Is-Primary
    Local: 49.0004.1000.0000.0001
```

```
Logical interface lo0.16385 (Index 65) (SNMP ifIndex 76)
Flags: SNMP-Traps Encapsulation: Unspecified
Input packets : 0
Output packets: 0
Protocol inet, MTU: Unlimited
  Flags: None
```

**show interfaces brief
(Loopback)**

```
user@host> show interfaces lo0 brief
Physical interface: lo0, Enabled, Physical link is Up
Type: Loopback, Link-level type: Unspecified, MTU: Unlimited,
Clocking: Unspecified, Speed: Unspecified
Device flags   : Present Running Loopback
Interface flags: SNMP-Traps
```

```
Logical interface lo0.0
Flags: SNMP-Traps Encapsulation: Unspecified
inet  10.0.0.1      --> 0/0
      127.0.0.1    --> 0/0
iso   49.0004.1000.0000.0001
```

```
Logical interface lo0.16385
Flags: SNMP-Traps Encapsulation: Unspecified
inet
```

**show interfaces detail
(Loopback)**

```
user@host> show interfaces lo0 detail
Physical interface: lo0, Enabled, Physical link is Up
Interface index: 6, SNMP ifIndex: 6, Generation: 4
Type: Loopback, Link-level type: Unspecified, MTU: Unlimited,
Clocking: Unspecified, Speed: Unspecified
Device flags   : Present Running Loopback
Interface flags: SNMP-Traps
Link type      : Unspecified
Link flags     : None
Physical info  : Unspecified
Hold-times    : Up 0 ms, Down 0 ms
Current address: Unspecified, Hardware address: Unspecified
Alternate link address: Unspecified
Last flapped  : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes   : 0
Output bytes  : 0
Input packets : 0
Output packets: 0
Logical interface lo0.0 (Index 64) (SNMP ifIndex 16) (Generation 3)
Flags: SNMP-Traps Encapsulation: Unspecified
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

Protocol inet, MTU: Unlimited, Generation: 10, Route table: 0
Flags: None
Addresses, Flags: Is-Default Is-Primary
Destination: Unspecified, Local: 10.0.0.1, Broadcast: Unspecified,
Generation: 10
Addresses, Flags: None
Destination: Unspecified, Local: 127.0.0.1, Broadcast: Unspecified,
Generation: 12
Protocol iso, MTU: Unlimited, Generation: 11, Route table: 0
Flags: None
Addresses, Flags: Is-Default Is-Primary
Destination: Unspecified, Local: 49.0004.1000.0000.0001,
Broadcast: Unspecified, Generation: 14

Logical interface lo0.16385 (Index 65) (SNMP ifIndex 76) (Generation 4)
Flags: SNMP-Traps Encapsulation: Unspecified
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
Protocol inet, MTU: Unlimited, Generation: 12, Route table: 1
Flags: None

```

```

show interfaces extensive (Loopback) user@host> show interfaces lo0 extensive
Physical interface: lo0, Enabled, Physical link is Up
Interface index: 6, SNMP ifIndex: 6, Generation: 4
Type: Loopback, Link-level type: Unspecified, MTU: Unlimited,
Clocking: Unspecified, Speed: Unspecified
Device flags : Present Running Loopback
Interface flags: SNMP-Traps
Link type : Unspecified
Link flags : None
Physical info : Unspecified
Hold-times : Up 0 ms, Down 0 ms
Current address: Unspecified, Hardware address: Unspecified
Alternate link address: Unspecified
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:

```

Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
Policed discards: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
Resource errors: 0

Logical interface lo0.0 (Index 64) (SNMP ifIndex 16) (Generation 3)
Flags: SNMP-Traps Encapsulation: Unspecified
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
Protocol inet, MTU: Unlimited, Generation: 10, Route table: 0
Flags: None
Addresses, Flags: Is-Default Is-Primary
Destination: Unspecified, Local: 10.0.0.1, Broadcast: Unspecified,
Generation: 10
Addresses, Flags: None
Destination: Unspecified, Local: 127.0.0.1, Broadcast: Unspecified,
Generation: 12
Protocol iso, MTU: Unlimited, Generation: 11, Route table: 0
Flags: None
Addresses, Flags: Is-Default Is-Primary
Destination: Unspecified, Local: 49.0004.1000.0000.0001,
Broadcast: Unspecified, Generation: 14

Logical interface lo0.16385 (Index 65) (SNMP ifIndex 76) (Generation 4)
Flags: SNMP-Traps Encapsulation: Unspecified
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
Protocol inet, MTU: Unlimited, Generation: 12, Route table: 1
Flags: None

show interfaces (Aggregated Ethernet)

Syntax	<pre>show interfaces <i>aenumber</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 only) Display status information about the specified aggregated Fast Ethernet or Gigabit Ethernet interface.
Options	<p><i>aenumber</i>—Display standard information about the specified aggregated Fast Ethernet or Gigabit Ethernet interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
List of Sample Output	<p>show interfaces (Aggregated Ethernet) on page 105</p> <p>show interfaces brief (Aggregated Ethernet) on page 106</p> <p>show interfaces detail (Aggregated Ethernet) on page 106</p> <p>show interfaces extensive (Aggregated Ethernet) on page 107</p> <p>show interfaces extensive (Aggregated Ethernet with VLAN Stacking) on page 108</p>
Output Fields	Table 15 on page 101 lists the output fields for the show interfaces (Aggregated Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 15: Aggregated Ethernet show interfaces Output Fields

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

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

Field Name	Field Description	Level of Output
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Minimum links needed	Number of child links that must be operational for the aggregate interface to be operational.	All levels
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under Common Output Fields Description.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interfaces Flags" section under Common Output Fields Description.	All levels
Current address	Configured MAC address.	detail extensive
Hardware address	Hardware MAC address.	detail extensive
Last flapped	Date, time, and how long ago the interface went from down to up or up to down. The format is Last flapped: year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive

Table 15: Aggregated Ethernet show interfaces 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. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface • Output packets—Number of packets transmitted on the interface. 	detail extensive
Input errors	<p>Input errors on the interface:</p> <ul style="list-style-type: none"> • Errors—Sum of 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. • Giants—Number of frames received that are larger than the giant threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • Resource errors—Sum of transmit drops. 	detail extensive
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 is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	detail extensive
IPv6 transit statistics	<p>Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive

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

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

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

Field Name	Field Description	Level of Output
Protocol	Protocol family configured on the logical interface. Possible values are described in the “Protocol Field” section under Common Output Fields Description.	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under Common Output Fields Description.	detail extensive none
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	detail extensive none
Addresses, Flags	Information about address flags. Possible values are described in the “Addresses Flags” section under Common Output Fields Description.	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

```

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

  Logical interface ae0.0 (Index 72) (SNMP ifIndex 60)
  Flags: SNMP-Traps 16384 Encapsulation: ENET2
  Statistics
  Bundle:
    Packets      pps      Bytes      bps
    Input :      0        0          0        0
    Output :      0        0          0        0

```

```

Protocol inet, MTU: 1500
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.100.1/24, Local: 10.100.1.2, Broadcast: 10.100.1.255

```

**show interfaces brief
(Aggregated Ethernet)**

```

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

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

```

**show interfaces detail
(Aggregated Ethernet)**

```

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

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

Logical interface ae0.0 (Index 72) (SNMP ifIndex 60) (Generation 18)
Flags: SNMP-Traps 16384 Encapsulation: ENET2
Statistics Packets pps Bytes bps
Bundle:
Input : 0 0 0 0
Output: 0 0 0 0
Link:
fe-0/1/0.0
Input : 0 0 0 0
Output: 0 0 0 0
fe-0/1/2.0
Input : 0 0 0 0
Output: 0 0 0 0
fe-0/1/3.0
Input : 0 0 0 0
Output: 0 0 0 0
Marker Statistics: Marker Rx Resp Tx Unknown Rx Illegal Rx
fe-0/1/0.0 0 0 0 0
fe-0/1/2.0 0 0 0 0

```

```

    fe-0/1/3.0          0          0          0          0
Protocol inet, MTU: 1500, Generation: 37, Route table: 0
Flags: Is-Primary, Mac-Validate-Strict
Mac-Validate Failures: Packets: 0, Bytes: 0
Destination: 10.100.1/24, Local: 10.100.1.2, Broadcast: 10.100.1.255,
Generation: 49

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

0 best-effort          7375          7375          0

1 expedited-fo          0          0          0

2 assured-forw          0          0          0

3 network-cont         2268          2268          0

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

```

Addresses, Flags: Is-Preferred Is-Primary
 Destination: 10.100.1/24, Local: 10.100.1.2, Broadcast: 10.100.1.255,
 Generation: 49

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

```
user@host> show interfaces ae0 detail
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 155, SNMP ifIndex: 48, Generation: 186
  Link-level type: 52, MTU: 1518, Speed: 2000mbps, Loopback: Disabled, Source
  filtering: Disabled,
  Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:12:1e:19:3f:f0, Hardware address: 00:12:1e:19:3f:f0
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :           2406875           40152 bps
    Output bytes  :           1124470           22056 bps
    Input packets :             5307             5 pps
    Output packets:            13295            21 pps
  IPv6 transit statistics:
    Input bytes   :             0
    Output bytes  :             0
    Input packets :             0
    Output packets:             0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Giants: 0, Policed discards:
  0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
  0
  Ingress queues: 4 supported, 4 in use
  Queue counters:      Queued packets  Transmitted packets      Dropped packets

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

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

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

  Logical interface ae0.451 (Index 69) (SNMP ifIndex 167) (Generation 601)
  Flags: SNMP-Traps VLAN-Tag [ 0x8100.451 ] Encapsulation: VLAN-VPLS
  Statistics      Packets      pps      Bytes      bps
  Bundle:
    Input :           289           0       25685       376
    Output:          1698           4      130375      3096
  Link:
```

```

ge-1/2/0.451
  Input :          289          0          25685          376
  Output:           0          0              0              0
ge-1/2/1.451
  Input :           0          0              0              0
  Output:        1698          4        130375        3096
Marker Statistics:  Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
ge-1/2/0.451              0              0              0              0
ge-1/2/1.451              0              0              0              0
Protocol vpls, MTU: 1518, Generation: 849, Route table: 3
Flags: Is-Primary

Logical interface ae0.452 (Index 70) (SNMP ifIndex 170) (Generation 602)
Flags: SNMP-Traps VLAN-Tag [ 0x8100.452 ] Encapsulation: VLAN-VPLS
Statistics      Packets      pps      Bytes      bps
Bundle:
  Input :         293          1          26003          1072
  Output:        1694          3        130057          2400
Link:
ge-1/2/0.452
  Input :         293          1          26003          1072
  Output:        1694          3        130057          2400
ge-1/2/1.452
  Input :           0          0              0              0
  Output:           0          0              0              0
Marker Statistics:  Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
ge-1/2/0.452              0              0              0              0
ge-1/2/1.452              0              0              0              0
Protocol vpls, MTU: 1518, Generation: 850, Route table: 3
Flags: None
...

```

show interfaces (Fast Ethernet)

Syntax	<pre>show interfaces <i>interface-type</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display status information about the specified Fast Ethernet interface.
Options	<p><i>interface-type</i>—On M Series and T Series routers, the interface type is <i>fe-fpc/pic/port</i>. On the J Series routers, the interface type is <i>fe-pim/O/port</i>.</p> <p><i>brief detail extensive terse</i>—(Optional) Display the specified level of output.</p> <p><i>descriptions</i>—(Optional) Display interface description strings.</p> <p><i>media</i>—(Optional) Display media-specific information about network interfaces.</p> <p><i>snmp-index snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><i>statistics</i>—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
List of Sample Output	<p>show interfaces (Fast Ethernet) on page 123</p> <p>show interfaces brief (Fast Ethernet) on page 124</p> <p>show interfaces detail (Fast Ethernet) on page 124</p> <p>show interfaces extensive (Fast Ethernet) on page 124</p>
Output Fields	<p>Table 16 on page 110 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 16: show interfaces Fast Ethernet Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under Common Output Fields Description.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none

Table 16: show interfaces Fast Ethernet Output Fields (*continued*)

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

Table 16: 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 16: 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 16: 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 16: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
OTN FEC statistics	<p>The forward error correction (FEC) counters provide the following statistics:</p> <ul style="list-style-type: none"> • Corrected Errors—The count of corrected errors in the last second. • Corrected Error Ratio—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits. 	
PCS statistics	<p>(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.</p> <ul style="list-style-type: none"> • Bit errors—High bit error rate. Indicates the number of bit errors when the PCS receiver is operating in normal mode. • Errored blocks—Loss of block lock. The number of errored blocks when PCS receiver is operating in normal mode. 	detail extensive

Table 16: 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 16: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the router from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local router (which the router is rejecting). • Input SA rejects—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • Output packet count—Number of packets that the filter has given to the MAC hardware. • Output packet pad count—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • Output packet error count—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • CAM destination filters, CAM source filters—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0. 	extensive
PMA PHY	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • PHY Lock—Phase-locked loop • PHY Light—Loss of optical signal 	extensive

Table 16: 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 16: 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 16: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner status—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the attached Ethernet device, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is None. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Local resolution—Information from the link partner: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive). 	extensive
Received path trace, Transmitted path trace	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.</p>	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive

Table 16: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under Common Output Fields Description.	All levels
VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags. <ul style="list-style-type: none"> • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. • push-pop—An outer VLAN tag is pushed in front of the existing VLAN tag, and the outer VLAN tag of the incoming frame is 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 16: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Demux:	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: <ul style="list-style-type: none"> Source Family Inet Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the "Protocol Field" section under Common Output Fields Description.	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. <ul style="list-style-type: none"> Input bytes, Output bytes—Number of bytes received and transmitted on the interface set Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the router.	extensive
Transit statistics	Number and rate of bytes and packets transiting the switch. <p>NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the Output bytes and Output packets interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.</p>	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. Possible values are described in the "Family Flags" section under Common Output Fields Description.	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none

Table 16: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	detail extensive none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	detail extensive
Output Filters	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	detail extensive
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under Common Output Fields Description.	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about address flag (possible values are described in the “Addresses Flags” section under Common Output Fields Description.	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

```

show interfaces user@host> show interfaces fe-0/0/0
(Fast Ethernet) 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, 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: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Active alarms : None
Active defects : None
MAC statistics:
  Total octets          Receive      Transmit
  Total packets         0             1
  Unicast packets       0             0
  Broadcast packets     0             1
  Multicast packets     0             0
  CRC/Align errors      0             0
  FIFO errors           0             0
  MAC control frames    0             0
  MAC pause frames      0             0
  Oversized frames      0
  Jabber frames         0
  Fragment frames       0
  VLAN tagged frames    0
  Code violations       0
Filter statistics:
  Input packet count    0
  Input packet rejects  0
  Input DA rejects      0
  Input SA rejects      0
  Output packet count   1
  Output packet pad count 0
  Output packet error count 0
  CAM destination filters: 1, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
    Link partner: Full-duplex, Flow control: None, Remote fault: Ok
  Local resolution:
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Bandwidth      Buffer Priority  Limit
               %      bps   %      usec
0 best-effort   95    950000000 95      0    low  none
3 network-control 5    500000000 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> otn-options { bytes { transmit-payload-type <i>number</i>; } } <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 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
List of Sample Output	<p>show interfaces (Gigabit Ethernet) on page 141</p> <p>show interfaces (Gigabit Ethernet on MX Series Router) on page 141</p> <p>show interfaces brief (Gigabit Ethernet) on page 141</p> <p>show interfaces detail (Gigabit Ethernet) on page 142</p> <p>show interfaces extensive (Gigabit Ethernet IQ2) on page 143</p> <p>show interfaces (Gigabit Ethernet Unnumbered Interface) on page 146</p>
Output Fields	See Table 17 on page 127 for the output fields for the show interfaces (Gigabit Ethernet) command. For Gigabit Ethernet IQ and IQE PICs, the traffic and MAC statistics vary by interface type. For more information, see Table 18 on page 140 .

Table 17: show interfaces Gigabit Ethernet Output Fields

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

Table 17: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Link flags	Information about the link. Possible values are described in the “Links Flags” section under Common Output Fields Description.	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Schedulers	(Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the show interfaces (10-Gigabit Ethernet) command.</p>	detail extensive

Table 17: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 17: 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 17: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
OTN 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 17: 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. • 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 17: 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 17: 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 17: 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 17: 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	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.</p>	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive

Table 17: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 17: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. • pop-swap—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none
Demux:	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> • Source Family Inet • Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under Common Output Fields Description.	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Dynamic Profile	(MX Series routers with Trio MPCs only) Name of the dynamic profile that was used to create this interface configured with 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 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
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

Table 17: 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 specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface set • Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the router.	extensive
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 Common Output Fields Description.	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	detail extensive none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	detail extensive
Output Filters	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	detail extensive
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the "Addresses Flags" section under Common Output Fields Description.	detail extensive none

Table 17: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about address flag (possible values are described in the “Addresses Flags” section under Common Output Fields Description.	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Table 18: 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      user@host> show interfaces ge-3/0/2
(Gigabit Ethernet)  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
                      Input packets : 0
                      Output packets: 0
                      Protocol ccc, MTU: 1522
                      Flags: Is-Primary

show interfaces      user@host> show interfaces ge-2/2/2
(Gigabit Ethernet on  Physical interface: ge-2/2/2, Enabled, Physical link is Up
MX Series Router)   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
                      Input packets : 0
                      Output packets: 0
                      Protocol aenet, AE bundle: ae0.0    Link Index: 4

show interfaces brief user@host> show interfaces ge-3/0/2 brief
(Gigabit Ethernet)  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
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 extensive ge-7/1/3
Physical interface: ge-7/1/3, Enabled, Physical link is Up
Interface index: 170, SNMP ifIndex: 70, Generation: 171
Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4004000
Link flags : None
CoS queues : 8 supported, 4 maximum usable queues
Schedulers : 256
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:14:f6:30:5e:74, Hardware address: 00:14:f6:30:5e:74
Last flapped : 2007-11-07 21:31:41 PST (02:03:33 ago)
Statistics last cleared: Never
Traffic statistics:
```

```

Input bytes :          38910844056          7952 bps
Output bytes :          7174605          8464 bps
Input packets:          418398473          11 pps
Output packets:          78903          12 pps
IPv6 transit statistics:
  Input bytes :          0
  Output bytes :          0
  Input packets:          0
  Output packets:          0
Ingress traffic statistics at Packet Forwarding Engine:
Input bytes :          38910799145          7952 bps
Input packets:          418397956          11 pps
Drop bytes :          0          0 bps
Drop packets:          0          0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort          418390823          418390823          0

  1 expedited-fo          0          0          0

  2 assured-forw          0          0          0

  3 network-cont          7133          7133          0

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

  0 best-effort          1031          1031          0

  1 expedited-fo          0          0          0

  2 assured-forw          0          0          0

  3 network-cont          77872          77872          0

Active alarms : None
Active defects : None
MAC statistics:
  Total octets          Receive          Transmit
  Total packets          38910844056          7174605
  Unicast packets          418398473          78903
  Broadcast packets          408021893366          1026
  Multicast packets          10          12
  CRC/Align errors          418398217          77865
  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 17 on page 127](#).

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 demux0 (Demux Interfaces)

Syntax	show interfaces demux0 <i>logical-interface-number</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i> > <statistics>
Release Information	Command introduced in Junos OS Release 9.0.
Description	(MX Series and M Series routers only) Display status information about the specified demux interface.
Options	<p>none—Display standard information about the specified demux interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
List of Sample Output	show interfaces (Demux) on page 153 show interfaces (PPPoE over Aggregated Ethernet) on page 154 show interfaces extensive (Targeted Distribution for Aggregated Ethernet Links) on page 154
Output Fields	Table 19 on page 147 lists the output fields for the show interfaces (demux interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 19: Demux show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	brief detail extensive none
Interface index	Index number of the physical interface, which reflects its initialization sequence.	brief detail extensive none
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under Common Output Fields Description.	brief detail extensive none

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

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

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

Field Name	Field Description	Level of Output
Hardware address	Hardware MAC address.	detail extensive
Alternate link address	Backup address of the link.	detail extensive
Last flapped	Date, time, and how long ago the interface went from down to up. The format is 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
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. • IPv6 transit statistics—Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Input errors	<p>Input errors on the interface whose definitions are as follows:</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. • Giants—Number of frames received that are larger than the giant packet 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. • Resource errors—Sum of transmit drops. 	extensive
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	none

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

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: <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. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Output Rate	Output rate in bps and pps.	none
Logical Interface		
Logical interface	Name of the logical interface.	brief detail extensive none
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
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under Common Output Fields Description.	brief detail extensive none
Encapsulation	Encapsulation on the logical interface.	brief extensive none
Demux	Specific IP demultiplexing (demux) values: <ul style="list-style-type: none"> • Underlying interface—The underlying interface that the demux interface uses. • Index—Index number of the logical interface. • Family—Protocol family configured on the logical interface. • Source prefixes, total—Total number of source prefixes for the underlying interface. • Destination prefixes, total—Total number of destination prefixes for the underlying interface. • Prefix—inet family prefix. 	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface.	brief

Table 19: Demux show interfaces 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 specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. • Input packets, Output packets—Number of packets received and transmitted on the interface set. • IPv6 transit statistics—Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Local statistics	<p>Number of transit bytes and packets received and transmitted on the local 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. 	detail extensive
Transit statistics	<p>Number and rate of bytes and packets transiting the switch.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
IPv6 Transit statistics	<p>Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Input packets	Number of packets received on the interface.	none
Output packets	Number of packets transmitted on the interface.	none
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under Common Output Fields Description.	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none

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

Field Name	Field Description	Level of Output
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
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under Common Output Fields Description.	detail extensive none
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under Common Output Fields Description.	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive statistics none
Local	IP address of the logical interface.	detail extensive terse none
Remote	IP address of the remote interface.	terse
Broadcast	Broadcast address of the logical interlace.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link	Name of the physical interfaces for member links in an aggregated Ethernet bundle for a PPPoE over aggregated Ethernet configuration. PPPoE traffic goes out on these interfaces.	detail extensive none
Dynamic-profile	Name of the PPPoE dynamic profile assigned to the underlying interface.	detail extensive none
Service Name Table	Name of the PPPoE service name table assigned to the PPPoE underlying interface.	detail extensive none
Max Sessions	Maximum number of dynamic PPPoE logical interfaces that the router can activate on the underlying interface.	detail extensive none
Duplicate Protection	State of duplicate protection: On or Off . Duplicate protection prevents the activation of another dynamic PPPoE logical interface on the same underlying interface when a dynamic PPPoE logical interface for a client with the same MAC address is already active on that interface.	detail extensive none
AC Name	Name of the access concentrator.	detail extensive none

Sample Output

```

show interfaces user@host> show interfaces demux0
(Demux) Physical interface: demux0, Enabled, Physical link is Up
          Interface index: 128, SNMP ifIndex: 79, Generation: 129
          Type: Software-Pseudo, Link-level type: Unspecified, MTU: 9192, Clocking: 1,
          Speed: Unspecified
          Device flags   : Present Running
          Interface flags: Point-To-Point SNMP-Traps
          Link type      : Full-Duplex
          Link flags     : None
          Physical info  : Unspecified
          Hold-times     : Up 0 ms, Down 0 ms
          Current address: Unspecified, Hardware address: Unspecified
          Alternate link address: Unspecified
          Last flapped   : Never
          Statistics last cleared: Never
          Traffic statistics:
            Input bytes   :                0                0 bps
            Output bytes  :                0                0 bps
            Input packets :                0                0 pps
            Output packets:                0                0 pps
          IPv6 transit statistics:
            Input bytes   :                0
            Output bytes  :                0
            Input packets :                0
            Output packets:                0
          Input errors:
            Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
            Policed discards: 0, Resource errors: 0
          Output errors:
            Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
            Resource errors: 0

Logical interface demux0.0 (Index 87) (SNMP ifIndex 84) (Generation 312)
  Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
  Demux:
    Underlying interface: ge-2/0/1.0 (Index 74)
  Family Inet Source prefixes, total 1
  Prefix: 1.1.1/24
  Traffic statistics:
    Input bytes   :                0
    Output bytes  :             1554
    Input packets :                0
    Output packets:             37
  IPv6 transit statistics:
    Input bytes   :                0
    Output bytes  :                0
    Input packets :                0
    Output packets:                0
  Local statistics:
    Input bytes   :                0
    Output bytes  :             1554
    Input packets :                0
    Output packets:             37
  Transit statistics:
    Input bytes   :                0                0 bps
    Output bytes  :                0                0 bps
    Input packets :                0                0 pps
    Output packets:                0                0 pps

```

```
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Protocol inet, MTU: 1500, Generation: 395, Route table: 0
  Flags: Is-Primary, Mac-Validate-Strict
  Mac-Validate Failures: Packets: 0, Bytes: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 11.1.1/24, Local: 11.1.1.1, Broadcast: 11.1.1.255,
    Generation: 434
```

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

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

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

```
user@host> show interfaces demux0.1073741824 extensive
Logical interface demux0.1073741824 (Index 75) (SNMP ifIndex 558) (Generation
346)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2
  Demux:
    Underlying interface: ae0 (Index 201)
  Link:
    ge-1/0/0
    ge-1/1/0
    ge-2/0/7
    ge-2/0/8
  Targeting summary:
    ge-1/1/0, primary, Physical link is Up
    ge-2/0/8, backup, Physical link is Up
  Bandwidth: 1000mbps
```

show interfaces filters

Syntax	show interfaces filters <interface-name>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced on PTX Series Packet Transport Switches for Junos OS Release 12.1.
Description	Display all firewall filters that are installed on each interface in a system.
Options	none —Display filter information about all interfaces. interface-name —(Optional) Display filter information about a particular interface.
Additional Information	For information about how to configure firewall filters, see the Junos OS Policy Framework Configuration Guide . For related operational mode commands, see the Junos OS Routing Protocols and Policies Command Reference .
Required Privilege Level	view
List of Sample Output	show interfaces filters on page 156 show interfaces filters interface-name on page 156 show interfaces filters (PTX Series Packet Transport Switches) on page 156
Output Fields	Table 20 on page 155 lists the output fields for the show interfaces filters command. Output fields are listed in the approximate order in which they appear.

Table 20: show interfaces filters Output Fields

Field Name	Field Description
Interface	Name of the interface.
Admin	Interface state: up or down .
Link	Link state: up or down .
Proto	Protocol configured on the interface.
Input Filter	Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.
Output Filter	Names of any firewall filters to be evaluated when packets are transmitted on the interface, including any filters attached through activation of dynamic service.

Sample Output

```

show interfaces filters user@host> show interfaces filters
Interface      Admin Link Proto Input Filter      Output Filter
ge-0/0/0        up    up
ge-0/0/0.0      up    up    inet
                                   iso

ge-5/0/0        up    up
ge-5/0/0.0      up    up    any
                                   inet
                                   multiservice

gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
vt-0/3/0        up    up
at-1/0/0        up    up
at-1/0/0.0      up    up    inet
                                   iso

at-1/1/0        up    down
at-1/1/0.0      up    down inet
                                   iso

....

show interfaces filters user@host> show interfaces filters so-2/1/0
interface-name Interface      Admin Link Proto Input Filter      Output Filter
so-2/1/0        up    down
so-2/1/0.0      up    down inet    goop    outfilter
                                   iso
                                   inet6 v6in    v6out

user@host > show interfaces filters ge-3/0/1
Interface      Admin Link Proto Input Filter      Output Filter
ge-3/0/1        up    up
ge-3/0/1.0      up    up    inet    F1-ge-3/0/1.0-in    F2-ge-3/0/1.0-out
                                   inet    F3-ge-3/0/1.0-in

show interfaces filters user@host > show interfaces filters em0
(PTX Series Packet Interface      Admin Link Proto Input Filter      Output Filter
Transport Switches)   em0        up    up
em0.0              up    up    inet

```

show interfaces routing

Syntax	show interfaces routing <brief detail> <interface-name> <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display the state of the router's interfaces. Use this command for performing router diagnostics only, when you are determining whether the routing protocols and the Junos OS differ about the state of an interface.
Options	<p>none—Display standard information about the state of all router interfaces on all logical systems.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>interface-name—(Optional) Name of a specific 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>
Additional Information	For information about how to configure routing protocols, see the Junos OS Routing Protocols Configuration Guide . For information about related operational mode commands for routing instances and protocols, see the Junos OS Routing Protocols and Policies Command Reference .
Required Privilege Level	view
List of Sample Output	show interfaces routing brief on page 158 show interfaces routing brief (TX Matrix Plus Router) on page 159 show interfaces routing detail on page 159 show interfaces routing detail (TX Matrix Plus Router) on page 160
Output Fields	Table 21 on page 157 lists the output fields for the show interfaces routing command. Output fields are listed in the approximate order in which they appear.

Table 21: show interfaces routing Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the physical interface.	none brief
State	State of the physical interface: Up or Down .	none brief
Addresses	Protocols and addresses configured on the interface.	none brief
Index	Interface index number, which reflects its initialization sequence.	detail

Table 21: show interfaces routing Output Fields (*continued*)

Field Name	Field Description	Level of Output
Refcount	Number of references to the interface in the routing software.	detail
State	State (Up or Down) and type of interface.	detail
Change	Reflects one or more of the following recent changes to the interface: <ul style="list-style-type: none"> • Add—The interface was just added. • Address—The interface's link-layer address has changed. • Delete—The interface is being deleted. • Encapsulation—The type of encapsulation on the interface has changed. • Metric—The interface's metric value has changed. • MTU—The interface's maximim transmission unit size has changed. • UpDown—The interface has made an up or down transition. 	detail
Up/down transitions	Number of times the interface has gone from Down to Up .	detail
Link layer	Describes the link layer of the interface.	detail
Encapsulation	Encapsulation on the interface.	detail
Bandwidth	Speed at which the interface is running.	detail
Protocol address	Information about the configuration of protocols on the interface: <ul style="list-style-type: none"> • Address—Address configured on the interface for the protocol type. • State—State (Up or down) and type of interface. • Change—Reflects one or more of the following recent changes to the interface: <ul style="list-style-type: none"> • Add—The interface was just added. • Address—The interface's address has changed. • Broadcast—The interface's broadcast address has changed. • Delete—The interface is being deleted. • Netmask—The interface's netmask has changed. • UpDown—The interface has made an up or down transition. • Preference—Preference value for the route for this address. • Metric—Metric value on the interface for the protocol type. • MTU—Maximim transmission unit value of the interface. • Local address—On a point-to-point link, the address of the local side of the link. Not used for multicast links. • Destination—For a point-to-point link, the address of the remote side of the link. For multicast links, the network address. 	detail

Sample Output

```
show interfaces user@host> show interfaces routing brief
routing brief
```

Interface	State	Addresses
so-5/0/3.0	Down	ISO enabled
so-5/0/2.0	Up	MPLS enabled ISO enabled INET 192.168.2.120 INET enabled
so-5/0/1.0	Up	MPLS enabled ISO enabled INET 192.168.2.130 INET enabled
at-1/0/0.3	Up	CCC enabled
at-1/0/0.2	Up	CCC enabled
at-1/0/0.0	Up	ISO enabled INET 192.168.90.10 INET enabled
lo0.0	Up	ISO 47.0005.80ff.f800.0000.0108.0001.1921.6800.5061.00 ISO enabled INET 127.0.0.1
fxp1.0	Up	
fxp0.0	Up	INET 192.168.6.90

**show interfaces
routing brief (TX Matrix
Plus Router)**

user@host> show interfaces routing brief

Interface	State	Addresses
...		
ge-23/0/4.0	Up	INET 2.9.1.1 ISO enabled MPLS enabled
ge-23/0/3.0	Up	INET 2.8.1.1 ISO enabled MPLS enabled
ge-23/0/2.0	Up	INET 2.7.1.1 ISO enabled MPLS enabled
ge-23/0/1.0	Up	INET 2.6.1.1 ISO enabled MPLS enabled
ge-23/0/0.0	Up	INET 2.5.1.1 ISO enabled MPLS enabled
ge-31/0/7.599	Up	INET 2.14.10.93
ge-31/0/7.598	Up	INET 2.14.10.89
ge-31/0/7.597	Up	INET 2.14.10.85
ge-31/0/7.596	Up	INET 2.14.10.81
ge-31/0/7.595	Up	INET 2.14.10.77
ge-31/0/7.594	Up	INET 2.14.10.73
...		
ixgbe1.0	Up	INET 10.34.0.4 INET 162.0.0.4 INET6 fe80::200:1ff:fe22:4 INET6 fec0::a:22:0:4
ixgbe0.0	Up	INET 10.34.0.4 INET 162.0.0.4 INET6 fe80::200:ff:fe22:4 INET6 fec0::a:22:0:4
em0.0	Up	INET 192.168.178.11

**show interfaces
routing detail**

user@host> show interfaces routing detail

so-5/0/3.0

Index: 15, Refcount: 2, State: Up <Broadcast PointToPoint Multicast> Change:<>

Metric: 0, Up/down transitions: 0, Full-duplex

```

Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
ISO address (null)
  State: <Broadcast PointToPoint Multicast> Change: <>
  Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
so-5/0/2.0
  Index: 14, Refcount: 7, State: <Up Broadcast PointToPoint Multicast> Change:<>

Metric: 0, Up/down transitions: 0, Full-duplex
Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
MPLS address (null)
  State: <Up Broadcast PointToPoint Multicast> Change: <>
  Preference: 0 (120 down), Metric: 0, MTU: 4458 bytes
ISO address (null)
  State: <Up Broadcast PointToPoint Multicast> Change: <>
  Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
INET address 192.168.2.120
  State: <Up Broadcast PointToPoint Multicast Localup> Change: <>
  Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
  Local address: 192.168.2.120
  Destination: 192.168.2.110/32
INET address (null)
  State: <Up Broadcast PointToPoint Multicast> Change: <>
  Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
...

```

**show interfaces
routing detail (TX
Matrix Plus Router)**

```

user@host> show interfaces routing detail
ge-23/0/4.0
  Index: 77, Refcount: 5, State: <Up Broadcast Multicast> Change: <>
  0 metric, 0 up/down transitions, reth state 0, full-duplex
  Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
  Link address #0 0.1d.b5.14.da.2d
  INET address 2.9.1.1
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 2.9.1.3
    Destination: 2.9.1.0/30
    System flags: <Is-Preferred Is-Primary>
  ISO address (null)
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1497 bytes
    System flags: <>
  MPLS address (null)
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1488 bytes
    System flags: <>
ge-23/0/3.0
  Index: 76, Refcount: 5, State: <Up Broadcast Multicast> Change: <>
  0 metric, 0 up/down transitions, reth state 0, full-duplex
  Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
  Link address #0 0.1d.b5.14.da.2c
  INET address 2.8.1.1
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 2.8.1.3
    Destination: 2.8.1.0/30
    System flags: <Is-Preferred Is-Primary>
  ISO address (null)
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1497 bytes
    System flags: <>
  MPLS address (null)

```



```

    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1488 bytes
    System flags: <>
ge-23/0/2.0
    Index: 75, Refcount: 5, State: <Up Broadcast Multicast> Change: <>
    0 metric, 0 up/down transitions, reth state 0, full-duplex
    Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
    Link address #0 0.1d.b5.14.da.2b
    INET address 2.7.1.1
        State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
        Preference 0, metric 0, MTU 1500 bytes
        Broadcast address 2.7.1.3
        Destination: 2.7.1.0/30
        System flags: <Is-Preferred Is-Primary>
    ISO address (null)
        State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
        Preference 0, metric 0, MTU 1497 bytes
        System flags: <>
    MPLS address (null)
        State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
        Preference 0, metric 0, MTU 1488 bytes
        System flags: <>
ge-23/0/1.0
    Index: 74, Refcount: 5, State: <Up Broadcast Multicast> Change: <>
    0 metric, 0 up/down transitions, reth state 0, full-duplex
    Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
    Link address #0 0.1d.b5.14.da.2a
    INET address 2.6.1.1
        State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
        Preference 0, metric 0, MTU 1500 bytes
        Broadcast address 2.6.1.3
    ...
ixgbe1.0
    Index: 5, Refcount: 5, State: <Up Broadcast Multicast> Change: <>
    0 metric, 0 up/down transitions, reth state 0, full-duplex
    Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
    Link address #0 2.0.1.22.0.4
    INET address 10.34.0.4
        State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
        Preference 0, metric 0, MTU 1500 bytes
        Broadcast address 10.255.255.255
        Destination: 10.0.0.0/8
        System flags: <Is-Preferred>
    INET address 162.0.0.4
        State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
        Preference 0, metric 0, MTU 1500 bytes
        Broadcast address 191.255.255.255
        Destination: 128.0.0.0/2
        System flags: <Primary Is-Preferred Is-Primary>
    INET6 address fe80::200:1ff:fe22:4
        State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
        Preference 0, metric 0, MTU 1500 bytes
        Destination: fe80::/64
        System flags: <Is-Preferred>
    INET6 address fec0::a:22:0:4
        State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
        Preference 0, metric 0, MTU 1500 bytes
        Destination: fec0::/64
        System flags: <Is-Preferred Is-Primary>
ixgbe0.0
    Index: 4, Refcount: 5, State: <Up Broadcast Multicast> Change: <>

```

```
0 metric, 0 up/down transitions, reth state 0, full-duplex
Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
Link address #0 2.0.0.22.0.4
INET address 10.34.0.4
  State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
  Preference 0, metric 0, MTU 1500 bytes
  Broadcast address 10.255.255.255
  Destination: 10.0.0.0/8
  System flags: <Is-Preferred>
INET address 162.0.0.4
  State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
  Preference 0, metric 0, MTU 1500 bytes
  Broadcast address 191.255.255.255
  Destination: 128.0.0.0/2
  System flags: <Primary Is-Default Is-Preferred Is-Primary>
INET6 address fe80::200:ff:fe22:4
  State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
  Preference 0, metric 0, MTU 1500 bytes
  Destination: fe80::/64
  System flags: <Is-Preferred>
INET6 address fec0::a:22:0:4
  State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
  Preference 0, metric 0, MTU 1500 bytes
  Destination: fec0::/64
  System flags: <Is-Default Is-Preferred Is-Primary>
em0.0
Index: 3, Refcount: 2, State: <Up Broadcast Multicast> Change: <>
0 metric, 0 up/down transitions, reth state 0, full-duplex
Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 100Mbps
Link address #0 0.80.f9.26.0.c0
INET address 192.168.178.11
  State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
  Preference 0, metric 0, MTU 1500 bytes
  Broadcast address 192.168.178.127
  Destination: 192.168.178.0/25
  System flags: <Is-Preferred Is-Primary>
```

CHAPTER 11

Subscriber Management Subscriber CLI Commands

show subscribers

Syntax `show subscribers`
 `<address address>`
 `<client-type client-type>`
 `<interface interface>`
 `<logical-system logical-system>`
 `<mac-address mac-address>`
 `<profile-name profile-name>`
 `<routing-instance routing-instance>`
 `<stacked-vlan-id stacked-vlan-id>`
 `<subscriber-state subscriber-state>`
 `<vlan-id vlan-id>`
 `<count | detail | extensive | summary (all | logical-system logical-system | routing-instance routing-instance) | terse>`

Release Information Command introduced in Junos OS Release 9.3.
 Command introduced in Junos OS Release 9.3 for EX Series switches.
 client-type, **mac-address**, **subscriber-state**, **extensive**, and **summary** options introduced in Junos OS Release 10.2.
 count option usage with other options introduced in Junos OS Release 10.2.
 Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display information for active subscribers.

Options **address**—(Optional) Display subscribers whose IP address matches the specified address.

client-type—(Optional) Display subscribers whose client type matches the specified client type (DHCP, L2TP, PPP, PPPOE, VLAN, or static).

count—(Optional) Display the count of total subscribers and active subscribers for any specified option. You can use the **count** option alone or with the **address**, **client-type**, **interface**, **logical-system**, **mac-address**, **profile-name**, **routing-instance**, **stacked-vlan-id**, **subscriber-state**, or **vlan-id** options.

id—(Optional) Display a specific subscriber session whose session id matches the specified subscriber ID. You can display subscriber IDs by using the **show subscribers extensive** or the **show subscribers interface extensive** commands.

interface—(Optional) Display subscribers whose interface matches the specified interface.

logical-system—(Optional) Display subscribers whose logical system matches the specified logical system.

mac-address—(Optional) Display subscribers whose MAC address matches the specified MAC address.

profile-name—(Optional) Display subscribers whose dynamic profile matches the specified profile name.

routing-instance—(Optional) Display subscribers whose routing instance matches the specified routing instance.

subscriber-state—(Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).

vlan-id—(Optional) Display subscribers whose VLAN ID matches the specified VLAN ID.

stacked-vlan-id—(Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.

detail | extensive | summary | terse—(Optional) Display the specified level of output.



NOTE: Due to display limitations, logical system and routing instance output values are truncated when necessary.

Required Privilege Level	view
List of Sample Output	show subscribers (IPv4) on page 168 show subscribers (IPv6) on page 168 show subscribers (IPv4 and IPv6 Dual Stack) on page 168 show subscribers (LNS on MX Series Routers) on page 168 show subscribers detail (IPv4) on page 169 show subscribers detail (IPv6) on page 169 show subscribers detail (IPv6 Static Demux Interface) on page 169 show subscribers detail (L2TP LNS Subscribers on MX Series Routers) on page 169 show subscribers detail (Tunneled Subscriber) on page 170 show subscribers interface on page 170 show subscribers logical-system on page 170 show subscribers count on page 171 show subscribers routing-instance inst1 count on page 171 show subscribers vlan-id on page 171 show subscribers vlan-id detail on page 171 show subscribers stacked-vlan-id detail on page 171 show subscribers stacked-vlan-id vlan-id detail (Combined Output) on page 171 show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface) on page 171 show subscribers client-type dhcp detail on page 172 show subscribers extensive on page 172 show subscribers extensive (L2TP LNS Subscribers on MX Series Routers) on page 172 show subscribers summary on page 173 show subscribers summary all on page 173 show subscribers terse on page 173
Output Fields	Table 22 on page 166 lists the output fields for the show subscribers command. Output fields are listed in the approximate order in which they appear.

Table 22: show subscribers Output Fields

Field Name	Field Description
User Name	Name of subscriber.
Type	Subscriber client type (DHCP, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN).
IP Address	Subscriber IPv4 address.
IP Netmask	Subscriber IP netmask.
IPv6 Address	Subscriber IPv6 address, or multiple addresses.
IPv6 Prefix	Subscriber IPv6 prefix.
IPv6 Address Pool	Subscriber IPv6 address pool. The IPv6 address pool is used to allocate IPv6 prefixes to the DHCPv6 clients.
IPv6 Network Prefix Length	Length of the network portion of the IPv6 address.
IPv6 Prefix Length	Length of the subscriber IPv6 prefix.
Logical System	Logical system associated with the subscriber.
Routing Instance	Routing instance associated with the subscriber.
Interface	Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface. The * character indicates a continuation of addresses for the same session.
Interface Type	Whether the subscriber interface is Static or Dynamic .
Dynamic Profile Name	Dynamic profile used for the subscriber.
MAC Address	MAC address associated with the subscriber.
State	Current state of the subscriber session (Init , Configured , Active , Terminating , Tunneled).
VLAN Id	VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> .
Stacked VLAN Id	Stacked VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> .
RADIUS Accounting ID	RADIUS accounting ID associated with the subscriber.
Agent Circuit ID	Option 82 agent circuit ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.
Agent Remote ID	Option 82 agent remote ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.

Table 22: show subscribers Output Fields (*continued*)

Field Name	Field Description
DHCP Relay IP Address	IP address used by the DHCP relay agent.
Login Time	Date and time at which the subscriber logged in.
DHCP Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options, as defined in RFC 2132.
Session ID	ID number for a subscriber service session.
Service Sessions	Number of service sessions (that is, a service activated using RADIUS CoA) associated with the subscribers.
Service Session Name	Service session profile name.
Session Timeout (seconds)	Number of seconds of access provided to the subscriber before the session is automatically terminated.
Idle Timeout (seconds)	Number of seconds subscriber can be idle before the session is automatically terminated.
ADF IPv4 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv4 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
IPv4 Input Filter Name	Name assigned to the IPv4 input filter (client or service session).
IPv4 Output Filter Name	Name assigned to the IPv4 output filter (client or service session).
IPv6 Input Filter Name	Name assigned to the IPv6 input filter (client or service session).
IPv6 Output Filter Name	Name assigned to the IPv6 output filter (client or service session).
IFL Input Filter Name	Name assigned to the logical interface input filter (client or service session).
IFL Output Filter Name	Name assigned to the logical interface output filter (client or service session).

Table 22: show subscribers Output Fields (*continued*)

Field Name	Field Description
Subscribers by State	<p>Number of subscribers summarized by state. The summary information includes the following:</p> <ul style="list-style-type: none"> • Init—Number of subscriber currently in the initialization state. • Configured—Number of configured subscribers. • Active—Number of active subscribers. • Terminating—Number of subscribers currently terminating. • Terminated—Number of terminated subscribers. <p>Summary information includes subscriber counts per state and the total number of subscribers.</p>
Subscribers by Client Type	<p>Number of subscribers summarized by client type. Client types can include DHCP, VLAN, PPP, PPPOE, L2TP, and static. Summary information includes subscriber counts per client type and the total number of subscribers.</p>
Subscribers by LS:RI	<p>Number of subscribers summarized by logical system:routing instance (LS:RI) combination. Summary information includes subscriber counts per LS:RI and the total number of subscribers.</p>

Sample Output

show subscribers (IPv4)	<pre> user@host> show subscribers Interface IP Address/VLAN ID User Name LS:RI ge-1/3/0.1073741824 100 WHOLESALER-CLIENT default:default demux0.1073741824 100.0.0.10 RETAILER1-CLIENT test1:retailer1 demux0.1073741825 101.0.0.3 RETAILER1-CLIENT test1:retailer1 demux0.1073741826 102.0.0.3 RETAILER2-CLIENT test1:retailer2 </pre>
show subscribers (IPv6)	<pre> user@host> show subscribers Interface IP Address/VLAN ID User Name LS:RI ge-1/0/0.0 2001::c0:0:0:0/74 WHOLESALER-CLIENT default:default * 2002::1/128 subscriber-25 default:default </pre>
show subscribers (IPv4 and IPv6 Dual Stack)	<pre> user@host> show subscribers Interface IP Address/VLAN ID User Name LS:RI demux0.1073741834 0x8100.1002 0x8100.1 default:default demux0.1073741835 0x8100.1001 0x8100.1 default:default pp0.1073741836 61.1.1.1 dualstackuser1@ISP1.com default:ASP-1 * 2041:1:1:1::/48 * 2061:1:1:1:1::/64 pp0.1073741837 23.1.1.3 dualstackuser2@ISP1.com default:ASP-1 * 2001:1:2:5:5::/64 </pre>
show subscribers (LNS on MX Series Routers)	<pre> user@host> show subscribers Interface IP Address/VLAN ID User Name LS:RI si-4/0/0.1 192.168.4.1 xyz@example.com default:default </pre>


```

show subscribers user@host> show subscribers detail
detail (IPv4) Type: DHCP
IP Address: 100.20.9.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:95:00:00:98
State: Active
Radius Accounting ID: jnpr :2304
Session Timeout (seconds): 3600
Idle Timeout (seconds): 600
Login Time: 2009-08-25 14:43:52 PDT
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 08 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 36 2f
33 2d 37 2d 30 37 05 01 06 0f 21 2c
Service Sessions: 2

```

```

show subscribers user@host> show subscribers detail
detail (IPv6) Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2002:db2:ffff:1::/64
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:51:ff:ff:00:03
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00
00 00

```

```

show subscribers user@host> show subscribers detail
detail (IPv6 Static Type: STATIC-INTERFACE
Demux Interface) User Name: demux0.1@jnpr.net
IPv6 Prefix: 1:2:3:4:5:6:7:aa/128
Logical System: default
Routing Instance: default
Interface: demux0.1
Interface type: Static
Dynamic Profile Name: junos-default-profile
State: Active
Radius Accounting ID: 185
Login Time: 2010-05-18 14:33:56 EDT

```

```

show subscribers user@host> show subscribers detail
detail (L2TP LNS Type: L2TP
Subscribers on MX User Name: user1@jnpr.net
Series Routers) IP Address: 10.1.32.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824

```

Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST

**show subscribers
detail (Tunneled
Subscriber)** user@host> show subscribers detail
Type: PPPoE
User Name: user1@example.com
Logical System: default
Routing Instance: default
Interface: pp0.1
State: Active, Tunneled
Radius Accounting ID: 512

**show subscribers
interface** user@host> show subscribers interface demux0.1073741826 extensive
Type: VLAN
User Name: test1@test.com
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Dynamic
Dynamic Profile Name: profile-vdemux-relay-23qos
MAC Address: 00:00:6e:56:01:04
State: Active
Radius Accounting ID: 12
Session ID: 12
Stacked VLAN Id: 0x8100.1500
VLAN Id: 0x8100.2902
Login Time: 2011-10-20 16:21:59 EST

Type: DHCP
User Name: test1@test.com
IP Address: 172.16.200.6
IP Netmask: 255.255.255.0
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Static
MAC Address: 00:00:6e:56:01:04
State: Active
Radius Accounting ID: 21
Session ID: 21
Login Time: 2011-10-20 16:24:33 EST
Service Sessions: 2

Service Session ID: 25
Service Session Name: SUB-QOS
State: Active

Service Session ID: 26
Service Session Name: service-cb-content
State: Active
IPv4 Input Filter Name: content-cb-in-demux0.1073741826-in
IPv4 Output Filter Name: content-cb-out-demux0.1073741826-out

**show subscribers
logical-system** user@host> show subscribers logical-system test1 terse

	Interface	IP Address/VLAN ID	User Name	LS:RI
	demux0.1073741825	101.0.0.3	RETAILER1-CLIENT	test1:retailer1
	demux0.1073741826	102.0.0.3	RETAILER2-CLIENT	test1:retailer2


```

show subscribers count      user@host> show subscribers count
                               Total Subscribers: 188, Active Subscribers: 188

show subscribers routing-instance inst1 count
user@host> show subscribers routing-instance inst1 count
Total Subscribers: 188, Active Subscribers: 183

show subscribers vlan-id    user@host> show subscribers vlan-id 100
                               Interface      IP Address      User Name
                               ge-1/0/0.1073741824
                               ge-1/2/0.1073741825

show subscribers vlan-id detail
user@host> show subscribers vlan-id 100 detail
Type: VLAN
Interface: ge-1/0/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT

Type: VLAN
Interface: ge-1/2/0.1073741825
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT

show subscribers stacked-vlan-id detail
user@host> show subscribers stacked-vlan-id 101 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT

show subscribers stacked-vlan-id vlan-id detail (Combined Output)
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT

show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface)
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 interface ge-1/2/0.* detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active

```

Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT

show subscribers user@host> **show subscribers client-type dhcp detail**
client-type dhcp detail

Type: DHCP
IP Address: 100.20.9.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:95:00:00:98
State: Active
Radius Accounting ID: jnpr :2304
Login Time: 2009-08-25 14:43:52 PDT

Type: DHCP
IP Address: 100.20.10.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744383
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:94:00:01:f3
State: Active
Radius Accounting ID: jnpr :2560
Login Time: 2009-08-25 14:43:56 PDT

show subscribers user@host> **show subscribers extensive**
extensive

Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2002:db2:ffff:1::/64
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:51:ff:ff:00:03
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00
00 00
IPv6 Address Pool: pd_pool
IPv6 Network Prefix Length: 48

show subscribers user@host> **show subscribers extensive**
extensive (L2TP LNS Type: L2TP
Subscribers on MX User Name: user1@jnpr.net
Series Routers) IP Address: 10.1.32.58
 IP Netmask: 255.255.0.0

 Logical System: default
 Routing Instance: default
 Interface: si-5/2/0.1073749824
 Interface type: Dynamic

```

Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
IPv4 Input Filter Name: classify-si-5/2/0.1073749824-in
IPv4 Output Filter Name: classify-si-5/2/0.1073749824-out

```

show subscribers summary user@host> show subscribers summary

```

Subscribers by State
Init          3
Configured    2
Active       183
Terminating    2
Terminated     1

TOTAL        191

Subscribers by Client Type
DHCP         107
PPP           76
VLAN          8

TOTAL        191

```

show subscribers summary all user@host> show subscribers summary all

```

Subscribers by State
Init          3
Configured    2
Active       183
Terminating    2
Terminated     1

TOTAL        191

Subscribers by Client Type
DHCP         107
PPP           76
VLAN          8

TOTAL        191

Subscribers by LS:RI
default:default  1
default:ri1      28
default:ri2      16
ls1:default      22
ls1:riA          38
ls1:riB          44
logsysX:routinstY 42

TOTAL        191

```

show subscribers terse user@host> show subscribers summary terse

Interface	IP Address/VLAN ID	User Name	LS:RI
ge-1/3/0.1073741824	100		default:default
demux0.1073741824	100.0.0.10	WHOLESALE-CLIENT	default:default

demux0.1073741825	101.0.0.3	RETAILER1-CLIENT	test1:retailer1
demux0.1073741826	102.0.0.3	RETAILER2-CLIENT	test1:retailer2

PART 4

Index

- [Index on page 177](#)

Index

Symbols

#, comments in configuration statements.....	xiv
(), in syntax descriptions.....	xiv
< >, in syntax descriptions.....	xiv
[], in configuration statements.....	xiv
{ }, in configuration statements.....	xiv
(pipe), in syntax descriptions.....	xiv

A

AAA	
subscriber statistics	
displaying.....	54, 57
subscribers	
displaying.....	60
access	
Layer 3 wholesale	
DHCP.....	35
access profile	
retailer	
DHCP.....	36
wholesaler	
DHCP.....	36
address assignment pools	
displaying.....	63
aggregated Ethernet interfaces	
status information, displaying.....	101

B

binding state of DHCP client	
clearing.....	73, 88
displaying.....	66, 80
braces, in configuration statements.....	xiv
brackets	
angle, in syntax descriptions.....	xiv
square, in configuration statements.....	xiv
broadband access networks	
DHCP.....	17
broadband services router (BSR).....	9
high-speed Internet access support.....	9
IPTV support.....	10

network placement.....	10
overview.....	9
broadband subscriber management	
AAA service framework.....	18
DHCP.....	17
edge routers.....	9
Layer 3 wholesale topology	
DHCPv4.....	29
platform support.....	4
solution overview.....	3
supporting documentation.....	6
terms.....	5
VLAN architecture.....	15
BSR See broadband services router	

C

clear dhcp relay binding command.....	88
clear dhcp relay statistics command.....	90
clear dhcp server binding command.....	73
clear dhcp server statistics command.....	76
comments, in configuration statements.....	xiv
conventions	
text and syntax.....	xiii
curly braces, in configuration statements.....	xiv
customer support.....	xv
contacting JTAC.....	xv
customer VLAN	
overview.....	15

D

DHCP	
extended server binding	
clearing.....	73
displaying.....	66
extended server statistics	
clearing.....	76
displaying.....	70
relay binding	
clearing.....	88
relay binding state	
displaying.....	80
relay statistics	
clearing.....	90
displaying.....	85

DHCP client		
binding state		
clearing.....	73, 88	
displaying.....	66, 80	
statistics		
clearing.....	76, 90	
displaying.....	70	
documentation		
comments on.....	xiv	
dynamic profiles		
DHCPv4 retailer example.....	46	
DHCPv4 wholesaler example.....	45	
retailer		
DHCPv4.....	39	
wholesale network		
DHCPv4.....	37	
wholesaler		
DHCPv4.....	38	
E		
edge router placement		
multiedge network.....	11	
single-edge network.....	10	
Ethernet interfaces		
status information, displaying		
aggregated.....	101	
Fast Ethernet.....	110	
Gigabit Ethernet.....	126	
F		
Fast Ethernet interfaces		
status information, displaying.....	110	
firewall filters		
displaying.....	155	
font conventions.....	xiii	
forwarding options		
configuring.....	42	
DHCP relay agent.....	80, 85, 88, 90	
DHCPv4 example.....	46	
G		
Gigabit Ethernet interfaces		
demultiplexing interface information,		
displaying.....	147	
status information, displaying.....	126	
H		
hybrid customer VLAN.....	16	
I		
interfaces		
loopback		
DHCPv4 Layer 3 wholesale.....	31	
L		
Layer 2 wholesale		
configuration options		
DHCPv4.....	22	
DHCPv6.....	22	
overview.....	21	
Layer 3 wholesale		
access		
DHCP.....	35	
RADIUS server.....	35	
access profile		
DHCP.....	36	
basic topology		
DHCPv4.....	29	
configuration elements		
DHCPv4.....	29	
configuration options		
DHCPv4.....	22	
DHCPv6.....	22	
PPPoE.....	22	
DHCP support.....	23	
DHCPv4		
forwarding options example.....	46	
dynamic profiles		
DHCPv4.....	37, 38, 39	
DHCPv4 retailer example.....	46	
DHCPv4 wholesaler example.....	45	
dynamic VLAN		
DHCPv4.....	33	
forwarding options		
configuring.....	42	
interface support.....	23	
overview.....	21	
RADIUS VSAs.....	24	
reference topology		
DHCPv4.....	30	
routing instances		
DHCP.....	40	
DHCPv4 example.....	47	
static VLAN		
DHCPv4.....	32	
vlans		
DHCPv4.....	32	

- logical systems
 - subscriber relationship with.....24
- loopback interface
 - Layer 3 wholesale
 - DHCPv4.....31
 - status information, displaying.....94
- M**
- manuals
 - comments on.....xiv
- MSAN *See* multiservice access node
- multiservice access node (MSAN)
 - choosing.....12
 - delivery options.....13
 - overview.....11
 - VLAN interaction.....16
- P**
- parentheses, in syntax descriptions.....xiv
- R**
- RADIUS
 - wholesale VSA support.....24
- RADIUS server
 - access configuration
 - configuring.....35
- routing information
 - interfaces
 - state, displaying.....157
- routing instances
 - retailer
 - DHCPv4.....40
 - DHCPv4 example.....47
 - subscriber relationship with.....24
- S**
- service VLAN.....15
- show dhcp relay binding command.....80
- show dhcp relay statistics command.....85
- show dhcp server binding command.....66
- show dhcp server statistics command.....70
- show interfaces (Aggregated Ethernet)
 - command.....101
- show interfaces (Fast Ethernet) command.....110
- show interfaces (Gigabit Ethernet) command.....126
- show interfaces (Loopback) command.....94
- show interfaces demux0 (Demux Interfaces)
 - command.....147
- show interfaces filters command.....155
- show interfaces routing command.....157
- show network-access aaa statistics authentication
 - command.....57
- show network-access aaa statistics command.....54
- show network-access aaa subscribers
 - command.....60
- show network-access address-assignment pool
 - command.....63
- show subscribers command.....164
- subscriber access
 - subscriber information, displaying.....164
- subscribers
 - displaying.....164
- support, technical *See* technical support
- syntax conventions.....xiii
- T**
- technical support
 - contacting JTAC.....xv
- topology
 - Layer 3 wholesale network
 - DHCPv4.....29
- V**
- video services router (VSR).....9
 - network placement.....10
 - overview.....10
- VLAN
 - customer VLAN.....15
 - DHCPv4 wholesale
 - dynamic configuration.....33
 - static configuration.....32
 - Ethernet aggregation and.....16
 - hybrid.....16
 - multiservice access node interaction.....16
 - residential gateway interaction.....17
 - service VLAN.....15
- VSR *See* video services router
- W**
- wholesale *See* Layer 3 wholesale

