



Junos[®] OS

Broadband Subscriber Management Wholesale Feature Guide

Release
15.1



Modified: 2015-06-03

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

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

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

Junos[®] OS Broadband Subscriber Management Wholesale Feature Guide

15.1

Copyright © 2016, 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	xv
	Documentation and Release Notes	xv
	Supported Platforms	xv
	Using the Examples in This Manual	xv
	Merging a Full Example	xvi
	Merging a Snippet	xvi
	Documentation Conventions	xvii
	Documentation Feedback	xix
	Requesting Technical Support	xix
	Self-Help Online Tools and Resources	xix
	Opening a Case with JTAC	xx
Part 1	Configuring DHCP Layer 3 Wholesale Networks	
Chapter 1	Subscriber Management DHCP Layer 3 Wholesale Overview	3
	Layer 2 and Layer 3 Wholesale Overview	3
	Wholesale Network Configuration Options and Considerations	4
	DHCP Layer 3 Wholesale Configuration Interface Support	5
	Layer 3 Wholesale Configuration DHCP Support	5
	Subscriber to Logical System and Routing Instance Relationship	6
	RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview	6
Chapter 2	Configuring DHCPv4 Layer 3 Wholesale Networks	9
	Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements	9
	DHCPv4 Layer 3 Wholesale Network Topology Overview	10
	Configuring Loopback Interfaces for the DHCPv4 Layer 3 Wholesale Solution	11
	Configuring VLANs for the DHCPv4 Layer 3 Wholesale Network Solution	12
	Configuring Static Customer VLANs for the DHCPv4 Layer 3 Wholesale Network Solution	12
	Configuring Dynamic VLANs for the DHCPv4 Layer 3 Wholesale Network Solution	13
	Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution	15
	Configuring RADIUS Server Access	15
	Configuring a DHCP Wholesaler Access Profile	16
	Configuring DHCP Retailer Access Profiles	16

Chapter 3

Configuring Dynamic Profiles for the DHCPv4 Layer 3 Wholesale Network	
Solution	17
Configuring a Wholesale Dynamic Profile for use in the DHCPv4	
Solution	18
Configuring a Dynamic Profile for use by a Retailer in the DHCPv4	
Solution	19
Configuring Separate Routing Instances for DHCPv4 Service Retailers	20
Configure Default Forwarding Options for the DHCPv4 Wholesale Network	
Solution	22
Example: Wholesaler Dynamic Profile for a DHCPv4 Wholesale Network	24
Example: Retailer Dynamic Profile for a DHCPv4 Wholesale Network	24
Example: Default Forwarding Options Configuration for the DHCPv4 Wholesale	
Network	25
Example: Retailer Routing Instances for a DHCPv4 Wholesale Network	26
Configuring DHCPv6 Layer 3 Wholesale Networks	29
Broadband Subscriber Management DHCPv6 Layer 3 Wholesale Topology and	
Configuration Elements	29
DHCPv6 Layer 3 Wholesale Network Topology Overview	31
Configuring Loopback Interfaces for the DHCPv6 Layer 3 Wholesale Solution	32
Configuring VLANs for the DHCPv6 Layer 3 Wholesale Network Solution	32
Configuring Static Customer VLANs for the DHCPv6 Layer 3 Wholesale	
Network Solution	33
Configuring Dynamic Customer VLANs for the DHCPv6 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 DHCPv6 Layer 3 Wholesale Network	
Solution	37
Configuring a Wholesale Dynamic Profile for use in the DHCPv6	
Solution	37
Configuring a Dynamic Profile for use by Each Retailer in the DHCPv6	
Solution	38
Configuring Separate Routing Instances for DHCPv6 Service Retailers	40
Configuring Address Server Elements for the DHCPv6 Layer 3 Wholesale	
Solution	40
Configuring a DHCPv6 Address Assignment Pool	40
Configuring Extended DHCPv6 Local Server	42
Example: Retailer Dynamic Profile for a DHCPv6 Wholesale Network	43
Example: Retailer Routing Instances for a DHCPv6 Wholesale Network	44
Example: DHCPv6 Address Assignment Pool That Provides Full 128-bit IPv6	
Addresses for a DHCPv6 Wholesale Network	44
Example: DHCPv6 Address Assignment Pool That Provides 74-bit IPv6 Prefixes	
for a DHCPv6 Wholesale Network	45
Example: Extended DHCPv6 Local Server for a DHCPv6 Wholesale Network	45

Part 2	Configuring PPPoE Layer 3 Wholesale Networks	
Chapter 4	Subscriber Management PPPoE Wholesale Overview	49
	Layer 2 and Layer 3 Wholesale Overview	49
	PPPoE Layer 3 Wholesale Configuration Interface Support	50
	Subscriber to Logical System and Routing Instance Relationship	50
	RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview	51
Chapter 5	Configuring PPPoE Layer 3 Wholesale Networks	53
	Broadband Subscriber Management PPPoE Layer 3 Wholesale Topology and Configuration Elements	53
	PPPoE Layer 3 Wholesale Network Topology Overview	55
	Configuring Loopback Interfaces for the PPPoE Layer 3 Wholesale Solution	55
	Configuring Static Customer VLANs for the PPPoE Layer 3 Wholesale Network Solution	57
	Configuring Access Components for the PPPoE Wholesale Network Solution	57
	Configuring RADIUS Server Access	58
	Configuring a PPPoE Wholesaler Access Profile	58
	Configuring PPPoE Retailer Access Profiles	59
	Configuring Dynamic Profiles for the PPPoE Layer 3 Wholesale Network Solution	59
	Configuring a Wholesale Dynamic Profile for use in the PPPoE Solution	60
	Configuring Separate Routing Instances for PPPoE Service Retailers	61
	Example: Wholesaler Dynamic Profile for a PPPoE Wholesale Network	62
	Example: Retailer Routing Instances for a PPPoE Wholesale Network	63
Part 3	Configuring Layer 2 Wholesale Networks	
Chapter 6	Subscriber Management Layer 2 Wholesale Overview	67
	Layer 2 and Layer 3 Wholesale Overview	67
	Wholesale Network Configuration Options and Considerations	68
	RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview	69
Chapter 7	Configuring Layer 2 Wholesale Networks	71
	Broadband Subscriber Management Layer 2 Wholesale Topology and Configuration Elements	71
	Layer 2 Wholesale Network Topology Overview	73
	Configuring a Retail Dynamic Profile for Use in the Layer 2 Wholesale Solution	75
	Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution	76
	Configuring VLAN Interfaces for the Layer 2 Wholesale Solution	78
	Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces	79
	Configuring NNI ISP-Facing Interfaces for the Layer 2 Wholesale Solution	80
	Configuring Direct ISP-Facing Interfaces for the Layer 2 Wholesale Solution	81
	Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers	81

Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service	
Retailers	84
Configuring Access Components for the Layer 2 Wholesale Network	
Solution	86
Configuring RADIUS Server Access	86
Configuring a Layer 2 Wholesaler Access Profile	86
Example: Retailer Dynamic Profile for a Layer 2 Wholesale Network	87
Example: Access Interface for a Layer 2 Wholesale Network	87
Example: Retailer Access Routing Instances for a Layer 2 Wholesale	
Network	88
Example: Retailer NNI ISP-Facing Interfaces for a Layer 2 Wholesale	
Network	89
Example: Retailer Direct ISP-Facing Interface for a Layer 2 Wholesale	
Network	89

Part 4

Chapter 8

Configuration Statements and Operational Commands

Configuration Statements	93
accept	97
access-profile	98
accounting-server	99
active-server-group	100
address	101
address-assignment (Address-Assignment Pools)	105
authentication (DHCP Local Server)	106
authentication (DHCP Relay Agent)	107
authentication-order	108
authentication-server	109
auto-configure	110
connectivity-type	112
core-facing	113
demux0 (Dynamic Interface)	114
demux-options (Dynamic Interface)	115
demux-source (Dynamic IP Demux Interface)	116
demux-source (Dynamic Underlying Interface)	117
demux-source (Underlying Interface)	118
dhcp-attributes (Address-Assignment Pools)	119
dhcp-local-server	120
dhcp-relay	126
dhcpv6 (DHCP Local Server)	133
domain-name (Address-Assignment Pools)	136
dynamic-profile (DHCP Local Server)	137
dynamic-profile (DHCP Relay Agent)	138
dynamic-profile (Dynamic PPPoE)	139
dynamic-profile (Stacked VLAN)	140
dynamic-profile (VLAN)	141
dynamic-profiles	142
encapsulation (Dynamic Interfaces)	150
encapsulation (Logical Interface)	153

encapsulation (Physical Interface)	157
family	162
family (Address-Assignment Pools)	166
family (Dynamic Demux Interface)	167
family (Dynamic PPPoE)	168
family (Dynamic Standard Interface)	169
flexible-vlan-tagging	171
forwarding-options	172
grace-period	173
group (DHCP Local Server)	174
group (DHCP Relay Agent)	177
input-vlan-map (Dynamic Interfaces)	180
interface (DHCP Local Server)	181
interface (DHCP Relay Agent)	183
interface (Dynamic Routing Instances)	184
interface (Routing Instances)	185
interface-mac-limit (VPLS)	186
interfaces	187
interfaces (Static and Dynamic Subscribers)	188
instance-role	192
instance-type	193
ip-address-first	195
keepalives (Dynamic Profiles)	196
mac-validate (Dynamic IP Demux Interface)	197
multicast-replication	198
no-local-switching	198
no-tunnel-services	199
maximum-lease-time	200
output-vlan-map (Dynamic Interfaces)	201
pap (Dynamic PPP)	201
pool (Address-Assignment Pools)	202
pool-match-order	203
pop (Dynamic VLANs)	203
pppoe-options (Dynamic PPPoE)	204
pppoe-underlying-options (Static and Dynamic Subscribers)	205
ppp-options (Dynamic PPP)	206
prefix (Address-Assignment Pools)	207
profile (Access)	208
protocols	212
proxy-arp	215
proxy-arp (Dynamic Profiles)	216
push (Dynamic VLANs)	216
radius (Access Profile)	217
radius-server	219
range (Address-Assignment Pools)	220
ranges (Dynamic VLAN)	221
route-distinguisher	222
routing-instances (Dynamic Profiles)	224
routing-instances (Multiple Routing Entities)	226

secret	227
server (Dynamic PPPoE)	227
server-group	228
services (System Services)	229
site (VPLS Multihoming for FEC 128)	231
site-identifier (VPLS)	232
site-range	232
stacked-vlan-ranges	233
stacked-vlan-tagging	234
system	234
traceoptions (DHCP)	235
underlying-interface (demux0)	237
underlying-interface (Dynamic PPPoE)	238
unit	239
unit (Dynamic Demux Interface)	245
unit (Dynamic Profiles Standard Interface)	247
unnumbered-address (Dynamic PPPoE)	249
unnumbered-address (Dynamic Profiles)	250
unnumbered-address (Ethernet)	252
user-prefix (DHCP Local Server)	253
vlan-id (Dynamic VLANs)	254
vlan-id (VLAN ID to Be Bound to a Logical Interface)	255
vlan-model	256
vlan-ranges	257
vlan-tags	258
vlan-tags (Stacked VLAN Tags)	259
vpls (Routing Instance)	261
vrf-export	263
vrf-import	264
vrf-target	265
Chapter 9	
Operational Commands	267
AAA and Address Assignment Pool Commands	267
show network-access aaa statistics	268
show network-access aaa statistics authentication	273
show network-access aaa subscribers	276
show network-access address-assignment pool	279
DHCPv4 Local Server Commands	280
show dhcp server binding	281
show dhcp server statistics	287
clear dhcp server binding	290
clear dhcp server statistics	293
DHCPv4 Relay Commands	294
show dhcp relay binding	295
show dhcp relay statistics	300
clear dhcp relay binding	303
clear dhcp relay statistics	305

DHCPv6 Local Server Commands	307
show dhcpv6 server binding	308
show dhcpv6 server statistics	314
clear dhcpv6 server binding	317
clear dhcpv6 server statistics	319
Interface Commands	319
show interfaces (Loopback)	320
show interfaces (Aggregated Ethernet)	327
show interfaces (Fast Ethernet)	338
show interfaces (Gigabit Ethernet)	354
show interfaces (PPPoE)	378
show interfaces demux0 (Demux Interfaces)	388
show interfaces filters	397
show interfaces routing	399
show interfaces routing-instance	405
show ppp interface	407
Subscriber Commands	416
show subscribers	417
show subscribers summary	437
VPLS Commands	442
show vpls connections	443
show vpls flood event-queue	454
show vpls flood instance	456
show vpls flood route	458
show vpls mac-table	460
show vpls statistics	465

Part 5

Index

Index	471
-------------	-----

List of Figures

Part 1	Configuring DHCP Layer 3 Wholesale Networks	
Chapter 2	Configuring DHCPv4 Layer 3 Wholesale Networks	9
	Figure 1: Basic Subscriber Management Layer 3 Wholesale Solution Topology	10
	Figure 2: DHCPv4 Layer 3 Wholesale Network Reference Topology	11
Chapter 3	Configuring DHCPv6 Layer 3 Wholesale Networks	29
	Figure 3: Basic Subscriber Management DHCPv6 Layer 3 Wholesale Solution Topology	30
	Figure 4: DHCPv6 Layer 3 Wholesale Network Reference Topology	31
Part 2	Configuring PPPoE Layer 3 Wholesale Networks	
Chapter 5	Configuring PPPoE Layer 3 Wholesale Networks	53
	Figure 5: Basic Subscriber Management PPPoE Layer 3 Wholesale Solution Topology	54
	Figure 6: PPPoE Layer 3 Wholesale Network Reference Topology	55
Part 3	Configuring Layer 2 Wholesale Networks	
Chapter 7	Configuring Layer 2 Wholesale Networks	71
	Figure 7: Basic Subscriber Management Layer 2 Wholesale Solution Topology	72
	Figure 8: Layer 2 Wholesale Network Reference Topology	74

List of Tables

	About the Documentation	xv
	Table 1: Notice Icons	xvii
	Table 2: Text and Syntax Conventions	xvii
Part 1	Configuring DHCP Layer 3 Wholesale Networks	
Chapter 1	Subscriber Management DHCP Layer 3 Wholesale Overview	3
	Table 3: Required Juniper Networks VSAs for the Broadband Subscriber Management Wholesale Network Solution	6
Part 2	Configuring PPPoE Layer 3 Wholesale Networks	
Chapter 4	Subscriber Management PPPoE Wholesale Overview	49
	Table 4: Required Juniper Networks VSAs for the Broadband Subscriber Management Wholesale Network Solution	51
Part 3	Configuring Layer 2 Wholesale Networks	
Chapter 6	Subscriber Management Layer 2 Wholesale Overview	67
	Table 5: Required Juniper Networks VSAs for the Broadband Subscriber Management Wholesale Network Solution	69
Chapter 7	Configuring Layer 2 Wholesale Networks	71
	Table 6: Rewrite Operations on Single-Tagged and Dual-Tagged Frames	76
	Table 7: Applying Rewrite Operations to VLAN Maps	77
	Table 8: Encapsulation Combinations for Layer 2 Wholesale Interfaces	79
Part 4	Configuration Statements and Operational Commands	
Chapter 9	Operational Commands	267
	Table 9: show network-access aaa statistics Output Fields	268
	Table 10: show network-access aaa statistics authentication Output Fields	273
	Table 11: show network-access aaa subscribers Output Fields	276
	Table 12: show network-access address-assignment pool Output Fields	279
	Table 13: show dhcp server binding Output Fields	282
	Table 14: show dhcp server statistics Output Fields	288
	Table 15: show dhcp relay binding Output Fields	296
	Table 16: show dhcp relay statistics Output Fields	301
	Table 17: clear dhcp relay statistics Output Fields	306
	Table 18: show dhcpv6 server binding Output Fields	309
	Table 19: show dhcpv6 server statistics Output Fields	315
	Table 20: Loopback show interfaces Output Fields	320

Table 21: Aggregated Ethernet show interfaces Output Fields	327
Table 22: show interfaces Fast Ethernet Output Fields	338
Table 23: show interfaces (Gigabit Ethernet) Output Fields	355
Table 24: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type	369
Table 25: show interfaces (PPPoE) Output Fields	378
Table 26: show interfaces demux0 (Demux Interfaces) Output Fields	388
Table 27: show interfaces filters Output Fields	397
Table 28: show interfaces routing Output Fields	399
Table 29: show ppp interface Output Fields	407
Table 30: show subscribers Output Fields	420
Table 31: show subscribers summary Output Fields	438
Table 32: show vpls connections Output Fields	444
Table 33: show vpls flood event-queue Output Fields	454
Table 34: show vpls flood instance Output Fields	456
Table 35: show vpls flood route Output Fields	458
Table 36: show vpls mac-table Output fields	461
Table 37: show vpls statistics Output Fields	465

About the Documentation

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

Documentation and Release Notes

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

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

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

Supported Platforms

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

- MX Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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


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

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

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

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

Documentation Conventions

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

Table 1: Notice Icons

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

Table 2 on page xvii 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

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

Convention	Description	Examples
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none">Introduces or emphasizes important new terms.Identifies guide 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 CLI User Guide</i>RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none">To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.The console port is labeled CONSOLE.
< > (angle brackets)	Encloses 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)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none">In the Logical Interfaces box, select All Interfaces.To cancel the configuration, click Cancel.

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

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <http://www.juniper.net/techpubs/feedback/>.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software 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 Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>

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

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

Opening a Case with JTAC

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

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

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

PART 1

Configuring DHCP Layer 3 Wholesale Networks

- [Subscriber Management DHCP Layer 3 Wholesale Overview on page 3](#)
- [Configuring DHCPv4 Layer 3 Wholesale Networks on page 9](#)
- [Configuring DHCPv6 Layer 3 Wholesale Networks on page 29](#)

CHAPTER 1

Subscriber Management DHCP Layer 3 Wholesale Overview

- [Layer 2 and Layer 3 Wholesale Overview on page 3](#)
- [Wholesale Network Configuration Options and Considerations on page 4](#)
- [DHCP Layer 3 Wholesale Configuration Interface Support on page 5](#)
- [Layer 3 Wholesale Configuration DHCP Support on page 5](#)
- [Subscriber to Logical System and Routing Instance Relationship on page 6](#)
- [RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview on page 6](#)

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 9](#)
- [Broadband Subscriber Management PPPoE Layer 3 Wholesale Topology and Configuration Elements on page 53](#)
- [Broadband Subscriber Management Layer 2 Wholesale Topology and Configuration Elements on page 71](#)

Wholesale Network Configuration Options and Considerations

You can configure a wholesale network any number of ways using Juniper Networks hardware and Junos OS software. For information about subscriber management hardware support, see the *Junos OS Broadband Subscriber Management and Services Library*. The general configuration options, and considerations for each, are provided in the following table:

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 Library for Routing Devices*.

Related Documentation

- *Junos OS Network Interfaces Library for Routing Devices*
- *Subscriber Interfaces and Demultiplexing Overview* in the *Junos OS Broadband Subscriber Management and Services Library*.
- *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles* in the *Junos OS Broadband Subscriber Management and Services Library*.
- *Configuring a Subscriber Interface Using a Set of Static IP Demux Interfaces* in the *Junos OS Broadband Subscriber Management and Services Library*.

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 Broadband Subscriber Management and Services Library*.

Related Documentation

- *Extended DHCP Relay Agent Overview* in the *Junos OS Broadband Subscriber Management and Services Library*.
- *DHCP Relay Proxy Overview* in the *Junos OS Broadband Subscriber Management and Services Library*.
- *Extended DHCP Local Server Overview* in the *Junos OS Broadband Subscriber Management and Services Library*.

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 Library for Routing Devices*.

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 3 on page 6](#) is required for a wholesale network to function.

Table 3: 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 3: 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 Broadband Subscriber Management and Services Library.*

CHAPTER 2

Configuring DHCPv4 Layer 3 Wholesale Networks


- Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements on page 9
- DHCPv4 Layer 3 Wholesale Network Topology Overview on page 10
- Configuring Loopback Interfaces for the DHCPv4 Layer 3 Wholesale Solution on page 11
- Configuring VLANs for the DHCPv4 Layer 3 Wholesale Network Solution on page 12
- Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution on page 15
- Configuring Dynamic Profiles for the DHCPv4 Layer 3 Wholesale Network Solution on page 17
- Configuring Separate Routing Instances for DHCPv4 Service Retailers on page 20
- Configure Default Forwarding Options for the DHCPv4 Wholesale Network Solution on page 22
- Example: Wholesaler Dynamic Profile for a DHCPv4 Wholesale Network on page 24
- Example: Retailer Dynamic Profile for a DHCPv4 Wholesale Network on page 24
- Example: Default Forwarding Options Configuration for the DHCPv4 Wholesale Network on page 25
- Example: Retailer Routing Instances for a DHCPv4 Wholesale Network on page 26

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 1 on page 10](#) illustrates a basic Layer 3 wholesale topology model from which you can expand.

Figure 1: 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

Related Documentation

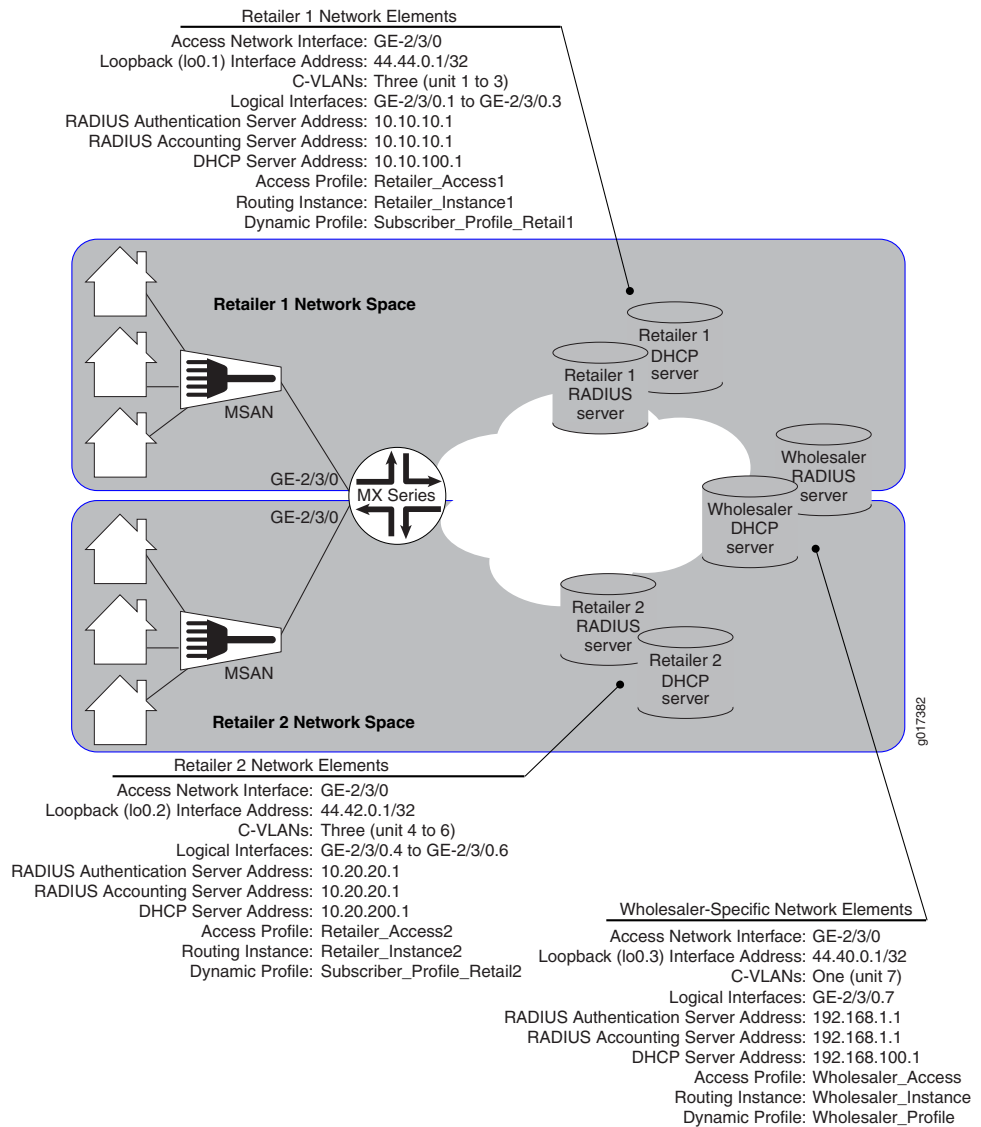
- [Layer 2 and Layer 3 Wholesale Overview on page 3](#)
- [DHCPv4 Layer 3 Wholesale Network Topology Overview on page 10](#)

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 2 on page 11](#) provides the reference topology for this configuration example.

Figure 2: DHCPv4 Layer 3 Wholesale Network Reference Topology



Related Documentation

- [Layer 2 and Layer 3 Wholesale Overview on page 3](#)
- [Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements on page 9](#)

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 • [Junos OS Network Interfaces Library for Routing Devices](#)

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 12](#)
- [Configuring Dynamic VLANs for the DHCPv4 Layer 3 Wholesale Network Solution on page 13](#)

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 15](#)
- [Configuring a DHCP Wholesaler Access Profile on page 16](#)
- [Configuring DHCP Retailer Access Profiles on page 16](#)

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 18](#)
- [Configuring a Dynamic Profile for use by a Retailer in the DHCPv4 Solution on page 19](#)

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

- [\[edit routing-instances\] Hierarchy Level](#)

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](#)

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;
          }
        }
      }
    }
  }
}
```

- Related Documentation**
- [Configuring Dynamic Profiles for the DHCPv4 Layer 3 Wholesale Network Solution on page 17](#)

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";
    }
}
}
}
}

```

Related Documentation • [Configuring Dynamic Profiles for the DHCPv4 Layer 3 Wholesale Network Solution on page 17](#)

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;
}
}
```

**Related
Documentation**

- [Configure Default Forwarding Options for the DHCPv4 Wholesale Network Solution on page 22](#)

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;  
    }  
    }  
    }  
    }  
}
```

Related Documentation

- [Configuring Separate Routing Instances for DHCPv4 Service Retailers on page 20](#)

CHAPTER 3

Configuring DHCPv6 Layer 3 Wholesale Networks

- Broadband Subscriber Management DHCPv6 Layer 3 Wholesale Topology and Configuration Elements on page 29
- DHCPv6 Layer 3 Wholesale Network Topology Overview on page 31
- Configuring Loopback Interfaces for the DHCPv6 Layer 3 Wholesale Solution on page 32
- Configuring VLANs for the DHCPv6 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 DHCPv6 Layer 3 Wholesale Network Solution on page 37
- Configuring Separate Routing Instances for DHCPv6 Service Retailers on page 40
- Configuring Address Server Elements for the DHCPv6 Layer 3 Wholesale Solution on page 40
- Example: Retailer Dynamic Profile for a DHCPv6 Wholesale Network on page 43
- Example: Retailer Routing Instances for a DHCPv6 Wholesale Network on page 44
- Example: DHCPv6 Address Assignment Pool That Provides Full 128-bit IPV6 Addresses for a DHCPv6 Wholesale Network on page 44
- Example: DHCPv6 Address Assignment Pool That Provides 74-bit IPV6 Prefixes for a DHCPv6 Wholesale Network on page 45
- Example: Extended DHCPv6 Local Server for a DHCPv6 Wholesale Network on page 45

Broadband Subscriber Management DHCPv6 Layer 3 Wholesale Topology and Configuration Elements


The network topology for the subscriber management DHCPv6 Layer 3 wholesale solution includes configuring separate routing instances for individual retailers that use a portion of the router. This solution uses a DHCPv6 local server configuration.



NOTE: Only DHCPv6 local server is currently supported for DHCPv6 Layer 3 wholesale configuration.

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

Figure 3: Basic Subscriber Management DHCPv6 Layer 3 Wholesale Solution Topology



A DHCPv6 Layer 3 wholesale network solution can use various combinations of the following configuration elements:

- Subscriber network VLAN configuration
- DHCPv6 configuration (local server only)
- 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

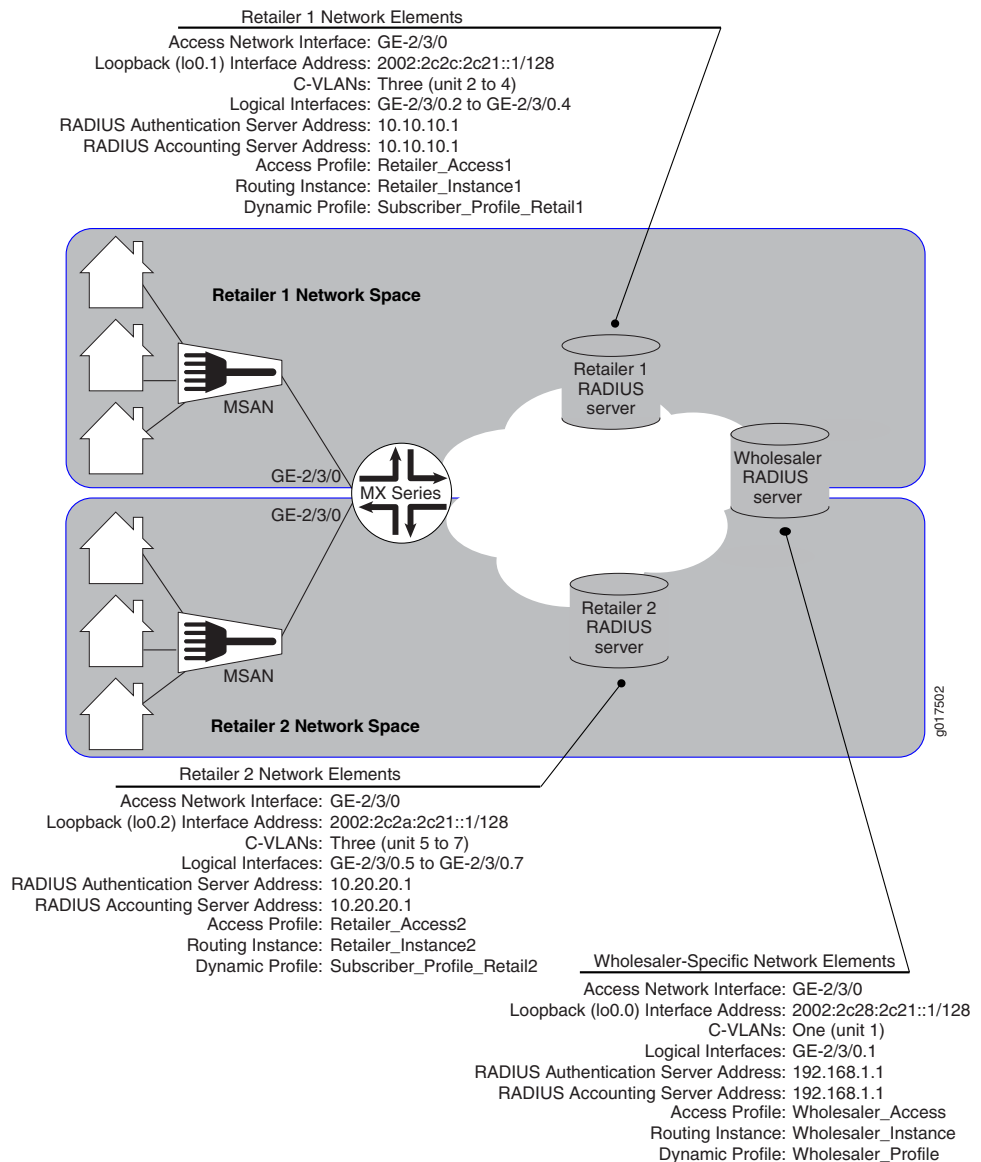
Related Documentation

- [Layer 2 and Layer 3 Wholesale Overview on page 3](#)
- [DHCPv6 Layer 3 Wholesale Network Topology Overview on page 31](#)

DHCPv6 Layer 3 Wholesale Network Topology Overview

This configuration explains how to configure a simple DHCPv6 Layer 3 wholesale subscriber access network. This solution incorporates two retailers sharing resources on a wholesaler router. [Figure 4 on page 31](#) provides the reference topology for this configuration example.

Figure 4: DHCPv6 Layer 3 Wholesale Network Reference Topology



- Related Documentation**
- [Layer 2 and Layer 3 Wholesale Overview on page 3](#)
 - [Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements on page 9](#)

Configuring Loopback Interfaces for the DHCPv6 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 loopback interface that you want to use for the wholesaler.

```
[edit interfaces lo0]  
user@host# edit unit 0
```

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

```
[edit interfaces lo0 unit 0]  
user@host# edit family inet6
```

4. Specify the wholesale loopback interface address.

```
[edit interfaces lo0 unit 0]  
user@host# set address 2002:2c28:2c21::1/128
```

5. Edit the unit for a retail loopback interface.

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

6. Edit the retail loopback interface family.

```
[edit interfaces lo0 unit 1]  
user@host# edit family inet6
```

7. Specify the retail loopback interface address.

```
[edit interfaces lo0 unit 1]  
user@host# set address 2002:2c2c:2c21::1/128
```

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

Configuring VLANs for the DHCPv6 Layer 3 Wholesale Network Solution

You can configure either static or dynamic customer VLANs for use in the DHCPv6 wholesale network solution.

- [Configuring Static Customer VLANs for the DHCPv6 Layer 3 Wholesale Network Solution on page 33](#)
- [Configuring Dynamic Customer VLANs for the DHCPv6 Layer 3 Wholesale Network Solution on page 33](#)

Configuring Static Customer VLANs for the DHCPv6 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 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 IPv6 demux interfaces.

```
[edit interfaces ge-2/3/0 unit 1]
user@host# set demux-source inet6
```

6. Edit the family for the first VLAN.

```
[edit interfaces ge-2/3/0 unit 1]
user@host# edit family inet6
```

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

```
[edit interfaces ge-2/3/0 unit 1 family inet6]
user@host# set unnumbered-address lo0.1 preferred-source-address
2002:2c28:2c21::1/128
```

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

Configuring Dynamic Customer VLANs for the DHCPv6 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. Specify that you want to create IPv6 demux interfaces.

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit"]
user@host# set demux-source inet6
```

- e. 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.

- f. 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.

- g. Access the family type.

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit"]
user@host# edit family inet6
```

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

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit" family inet6]
user@host# set unnumbered-address lo.0 preferred-source-address
2002:2c28:2c21::1/128
```

2. Associate the dynamic profile with the interface on which you want the VLANs 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 inet6
```

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 wholesale 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
```


- 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
```

- Specify that you want to configure RADIUS support.

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

- 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
```

- 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
```

- Configure any desired options for the RADIUS server.

See *Configuring RADIUS Server Options for Subscriber Access*.

- Configure subscriber accounting (RADIUS accounting).

See *Configuring Per-Subscriber Session Accounting*.

Configuring Dynamic Profiles for the DHCPv6 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 DHCPv6 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 DHCPv6 Solution on page 37](#)
- [Configuring a Dynamic Profile for use by Each Retailer in the DHCPv6 Solution on page 38](#)

Configuring a Wholesale Dynamic Profile for use in the DHCPv6 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 Wholesaler_Profile]
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 Wholesaler_Profile 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 Wholesaler_Profile 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 Wholesaler_Profile interfaces demux0 unit
"$junos-interface-unit"]
user@host# edit family inet6
```

- b. Configure the unnumbered address for the family.

```
[edit dynamic-profiles Wholesaler_Profile 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 Wholesaler_Profile interfaces demux0 unit
"$junos-interface-unit"]
user@host# set demux-source $junos-subscriber-ipv6-address
```

Configuring a Dynamic Profile for use by Each Retailer in the DHCPv6 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 inet6
```

- b. Configure the unnumbered address and preferred source address for the family.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 demux0 unit "$junos-interface-unit"
family inet6]
user@host# set unnumbered-address $junos-loopback-interface
preferred-source-address $junos-preferred-source-address
```

- c. Configure the variable that identifies the demux interface on the logical interface.

The variable is dynamically replaced with the IPv6 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-ipv6-address
```

Configuring Separate Routing Instances for DHCPv6 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 Retailer_Instance1
```

2. Specify the routing instance type for the retailer.

```
[edit routing-instances "Retailer_Instance1"]
user@host# set instance-type vrf
```

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

```
[edit routing-instances "Retailer_Instance1"]
user@host# set access-profile Retailer_Access1
```

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

```
[edit routing-instances "Retailer_Instance1"]
user@host# set interface ge-11/1/9.10
```

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

```
[edit routing-instances "Retailer_Instance1"]
user@host# set interface lo0.1
```



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

6. Repeat this procedure for other retailers.

Related Documentation

- [\[edit routing-instances\] Hierarchy Level](#)

Configuring Address Server Elements for the DHCPv6 Layer 3 Wholesale Solution

- [Configuring a DHCPv6 Address Assignment Pool on page 40](#)
- [Configuring Extended DHCPv6 Local Server on page 42](#)

Configuring a DHCPv6 Address Assignment Pool

Address assignment pools enable you to specify groups of IPv6 addresses that different client applications can share. In this configuration, the extended DHCPv6 local server configuration uses the address pool to provide addresses to subscribers that are accessing the network. You must create separate address assignment pools for each retailer routing instance.

You can create address assignment pools that provide full 128 bit IPv6 addresses or pools that provide prefixes of a specified length.

To configure an address assignment pool that provides full 128 -bit IPv6 addresses:

1. Create and name an address assignment pool.

```
[edit]
user@host# edit access address-assignment pool AddressPool_1
```

2. Edit the address pool family.

```
[edit access address-assignment pool AddressPool_1]
user@host# edit family inet6
```

3. Define the IPv6 network prefix.

```
[edit access address-pool AddressPool_1 family inet6]
user@host# set prefix 2121::0/64
```

4. Define a named address range for the pool of IPv6 addresses.

```
[edit access address-assignment pool AddressPool_1 family inet6]
user@host# set range Range1 low 2121::a/128
user@host# set range Range1 high 2121::7ffe/128
```

5. (Optional) Edit the family DHCP attributes.

```
[edit access address-assignment pool AddressPool_1 family inet6]
user@host# edit dhcp-attributes
```

6. (Optional) Set the maximum lease time.

```
[edit access address-assignment pool AddressPool_1 family inet dhcp-attributes]
user@host# set maximum-lease-time 3600
```

7. (Optional) Set the grace period.

```
[edit access address-assignment pool AddressPool_1 family inet dhcp-attributes]
user@host# set grace-period 60
```

To configure an address assignment pool that provides shorter, 74-bit IPv6 prefixes:

1. Create and name an address assignment pool.

```
[edit]
user@host# edit access address-assignment pool AddressPool_2
```

2. Edit the address pool family.

```
[edit access address-assignment pool AddressPool_2]
user@host# edit family inet6
```

3. Define the IPv6 network prefix.

```
[edit access address-pool AddressPool_2 family inet6]
user@host# set prefix 2222::0/64
```

4. Define a named address range limit for the pool of IPv6 addresses.

```
[edit access address-assignment pool AddressPool_2 family inet6]
user@host# set range BitLimit prefix-length 74
```

5. (Optional) Edit the family DHCP attributes.

```
[edit access address-assignment pool AddressPool_2 family inet6]
user@host# edit dhcp-attributes
```

6. (Optional) Set the maximum lease time.

```
[edit access address-assignment pool AddressPool_2 family inet dhcp-attributes]
user@host# set maximum-lease-time 3600
```

7. (Optional) Set the grace period.

```
[edit access address-assignment pool AddressPool_2 family inet dhcp-attributes]
user@host# set grace-period 60
```

Configuring Extended DHCPv6 Local Server

You can enable the MX Series router to function as an extended DHCPv6 local server. The extended DHCPv6 local server provides IPv6 addresses and other configuration information to a subscriber logging into the network. You must configure extended DHCPv6 local server for the wholesaler (default) routing instance and also for each retailer routing instance.

To configure the DHCPv6 local server:

1. Edit the routing system services.

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

2. Edit the DHCPv6 local server.

```
[edit system services]
user@host# edit dhcp-local-server
```

3. Define the DHCP pool match order.

```
[edit system services dhcp-local-server]
user@host# set pool-match-order ip-address-first
```

4. Set the authentication password.

```
[edit system services dhcp-local-server]
user@host# set authentication password auth-psswrđ
```

5. (Optional) Edit the values you want included with the username.

```
[edit system services dhcp-local-server]
user@host# edit authentication username-include
```

6. (Optional) Set the values you want included with the username.

```
[edit system services dhcp-local-server username-include]
user@host# set domain-name yourcompany.com
user@host# set user-prefix user-defined-prefix
```

7. Access the DHCPv6-specific service configuration.

```
[edit system services dhcp-local-server]
user@host# edit dhcpv6
```

8. Create and name a DHCPv6 local server group.

```
[edit system services dhcp-local-server dhcpv6]
```

```
user@host# edit group dhcp-ls-group
```

9. Specify a dynamic profile that you want the DHCPv6 local server group to use.

```
[edit system services dhcp-local-server dhcpv6 group dhcp-ls-group]
```

```
user@host# set dynamic-profile Wholesaler_Profile
```

10. Assign interfaces to the group.

```
[edit system services dhcp-local-server dhcpv6 group dhcp-ls-group]
```

```
user@host# set interface ge-1/3/0.1 upto ge-1/3/0.5
```

11. Edit the DHCPv6 local server trace options.

```
[edit system processes dhcp-service]
```

```
user@host# edit traceoptions
```

12. Specify a log file into which you want trace option information to be saved.

```
[edit system processes dhcp-service traceoptions]
```

```
user@host# set file dhcp-server-msgs.log
```

13. Specify the DHCPv6 local server message operations that you want saved in the log file.

```
[edit system processes dhcp-service traceoptions]
```

```
user@host# set flag all
```

- Related Documentation**
- *Address-Assignment Pools Overview*
 - *DHCPv6 Local Server Overview*

Example: Retailer Dynamic Profile for a DHCPv6 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 inet6 {
            demux-source {
              "$junos-subscriber-ip-address";
            }
            unnumbered-address "$junos-loopback-interface" preferred-source-address
              "$junos-preferred-source-address";
          }
        }
      }
    }
  }
}
```

- Related Documentation**
- [Configuring Dynamic Profiles for the DHCPv6 Layer 3 Wholesale Network Solution on page 37](#)

Example: Retailer Routing Instances for a DHCPv6 Wholesale Network

```
routing-instances {
  Retailer_Instance1 {
    instance-type vrf;
    access-profile Retailer_Access1;
    interface ge-11/1/9.10;
    interface lo0.1;
    route-distinguisher 1:1;
  }
  Retailer_Instance2 {
    instance-type vrf;
    access-profile Retailer_Access2;
    interface ge-7/1/9.10;
    interface lo0.2;
  }
}
```

- Related Documentation**
- [Configuring Separate Routing Instances for DHCPv6 Service Retailers on page 40](#)

Example: DHCPv6 Address Assignment Pool That Provides Full 128-bit IPV6 Addresses for a DHCPv6 Wholesale Network

```
access {
  address-assignment {
    pool AddressPool_1 {
      family inet6 {
        prefix 2121::0/64;
        range Range1 {
          low 2121::a/128;
          high 2121::7ffe/128;
        }
        dhcp-attributes {
          maximum-lease-time 3600;
          grace-period 60;
        }
      }
    }
  }
}
```

- Related Documentation**
- [DHCPv6 Address Assignment Pool That Provides Full 128-bit IPV6 Addresses for a DHCPv6 Wholesale Network](#)

Example: DHCPv6 Address Assignment Pool That Provides 74-bit IPV6 Prefixes for a DHCPv6 Wholesale Network

```
access {
  address-assignment {
    pool AddressPool_2 {
      family inet6 {
        prefix 2222::0/64;
        range BitLimit prefix-length 74;
        dhcp-attributes {
          maximum-lease-time 3600;
          grace-period 60;
        }
      }
    }
  }
}
```

Related Documentation • [Configuring Address Server Elements for the DHCPv6 Layer 3 Wholesale Solution on page 40](#)

Example: Extended DHCPv6 Local Server for a DHCPv6 Wholesale Network

```
system {
  services {
    dhcp-local-server {
      traceoptions {
        file dhcp-server-msgs.log;
        flag all;
      }
    }
    dhcpv6 {
      group dhcp-ls-group {
        dynamic-profile Wholesaler_Profile;
        interface ge-1/3/0.1 {
          upto ge-1/3/0.5;
        }
      }
    }
    pool-match-order {
      ip-address-first;
    }
    authentication {
      password auth-psswrld;
      username-include {
        domain-name yourcompany.com;
        user-prefix user-defined-prefix;
      }
    }
  }
}
```

- Related Documentation**
- [Configuring Address Server Elements for the DHCPv6 Layer 3 Wholesale Solution on page 40](#)

PART 2

Configuring PPPoE Layer 3 Wholesale Networks

- [Subscriber Management PPPoE Wholesale Overview on page 49](#)
- [Configuring PPPoE Layer 3 Wholesale Networks on page 53](#)

CHAPTER 4

Subscriber Management PPPoE Wholesale Overview

- [Layer 2 and Layer 3 Wholesale Overview on page 49](#)
- [PPPoE Layer 3 Wholesale Configuration Interface Support on page 50](#)
- [Subscriber to Logical System and Routing Instance Relationship on page 50](#)
- [RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview on page 51](#)

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 9](#)
- [Broadband Subscriber Management PPPoE Layer 3 Wholesale Topology and Configuration Elements on page 53](#)
- [Broadband Subscriber Management Layer 2 Wholesale Topology and Configuration Elements on page 71](#)

PPPoE Layer 3 Wholesale Configuration Interface Support

PPPoE Layer 3 wholesale requires the use of PPP interfaces. This means that you must specify the PP0 interface when configuring Layer 3 wholesaling in a PPPoE network.

For general additional information about configuring PPPoE interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

Related Documentation

- *Junos OS Network Interfaces Library for Routing Devices*
- *Configuring a PPPoE Dynamic Profile in the Junos OS Broadband Subscriber Management and Services Library.*
- *Configuring Dynamic PPPoE Subscriber Interfaces in the Junos OS Broadband Subscriber Management and Services Library.*

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 Library for Routing Devices*.

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 3 on page 6](#) 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
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 Broadband Subscriber Management and Services Library.*

CHAPTER 5

Configuring PPPoE Layer 3 Wholesale Networks


- [Broadband Subscriber Management PPPoE Layer 3 Wholesale Topology and Configuration Elements on page 53](#)
- [PPPoE Layer 3 Wholesale Network Topology Overview on page 55](#)
- [Configuring Loopback Interfaces for the PPPoE Layer 3 Wholesale Solution on page 55](#)
- [Configuring Static Customer VLANs for the PPPoE Layer 3 Wholesale Network Solution on page 57](#)
- [Configuring Access Components for the PPPoE Wholesale Network Solution on page 57](#)
- [Configuring Dynamic Profiles for the PPPoE Layer 3 Wholesale Network Solution on page 59](#)
- [Configuring Separate Routing Instances for PPPoE Service Retailers on page 61](#)
- [Example: Wholesaler Dynamic Profile for a PPPoE Wholesale Network on page 62](#)
- [Example: Retailer Routing Instances for a PPPoE Wholesale Network on page 63](#)

Broadband Subscriber Management PPPoE Layer 3 Wholesale Topology and Configuration Elements

The network topology for the subscriber management PPPoE Layer 3 wholesale solution includes configuring separate routing instances for individual retailers that use a portion of the router.

To explain the concept, but to limit complexity, this solution provides a configuration with one wholesaler and only two retailers. [Figure 5 on page 54](#) illustrates a basic PPPoE Layer 3 wholesale topology model from which you can expand.

Figure 5: Basic Subscriber Management PPPoE Layer 3 Wholesale Solution Topology



When you are configuring a PPPoE Layer 3 wholesale network solution, the following configuration elements are required:

- Subscriber network VLAN configuration
- Addressing server or addressing server access configuration
- RADIUS server access configuration
- Dynamic profile configuration for default (wholesaler) access
- Routing instance configuration for individual retailers
- Group configuration and forwarding options for the network
- Core network configuration

This implementation of PPPoE Layer 3 wholesale supports the following:

- Dynamic PPPoE interface creation.
- Static VLAN use only.
- AAA server assignment of subscribers to different routing instances within the same (default) logical system only.

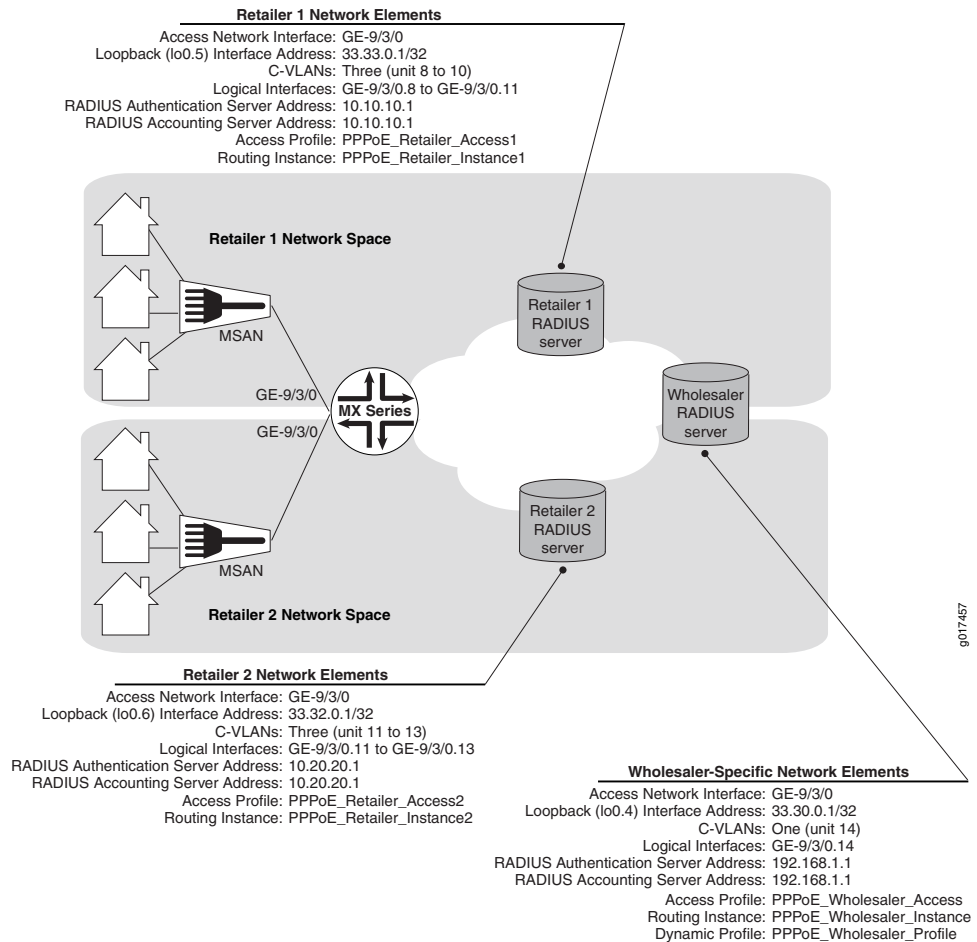
Related Documentation

- [Layer 2 and Layer 3 Wholesale Overview on page 3](#)
- [PPPoE Layer 3 Wholesale Network Topology Overview on page 55](#)

PPPoE Layer 3 Wholesale Network Topology Overview

This configuration explains how to configure a simple PPPoE Layer 3 wholesale subscriber access network. This solution incorporates two retailers sharing resources on a wholesaler router. [Figure 6 on page 55](#) provides the reference topology for this configuration example.

Figure 6: PPPoE Layer 3 Wholesale Network Reference Topology



Related Documentation

- [Layer 2 and Layer 3 Wholesale Overview on page 3](#)
- [Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements on page 9](#)

Configuring Loopback Interfaces for the PPPoE 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.



NOTE: If you do not configure the loopback interface, the routing platform chooses the first interface to come online as the default. If you configure more than one address on the loopback interface, we recommend that you configure one to be the primary address to ensure that it is selected for use with unnumbered interfaces. By default, the primary address is used as the source address when packets originate from the interface.

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 4
```

3. Edit the wholesale loopback interface family.

```
[edit interfaces lo0 unit 4]
user@host# edit family inet
```

4. Specify the wholesale loopback interface address.

```
[edit interfaces lo0 unit 4 family inet]
user@host# set address 33.30.0.1/32
```

5. (Optional) Specify the loopback interface address as the primary loopback interface.

```
[edit interfaces lo0 unit 4 family inet]
user@host# set address 33.30.0.2/32 primary
```

6. Edit the unit for a retail loopback interface.

```
[edit interfaces lo0]
user@host# edit unit 5
```

7. Edit the retail loopback interface family.

```
[edit interfaces lo0 unit 5]
user@host# edit family inet
```

8. Specify the retail loopback interface address.

```
[edit interfaces lo0 unit 5 family inet]
user@host# set address 33.33.0.1/32
```

9. (Optional) Specify the loopback interface address as the primary loopback interface.

```
[edit interfaces lo0 unit 5 family inet]
user@host# set address 33.33.0.2/32 primary
```

10. Repeat steps 7 through 10 for additional retailers, making sure to use unique unit and address values for each retailer loopback interface.

**Related
Documentation**

- *Junos OS Network Interfaces Library for Routing Devices*

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

In this example configuration, the access interface (**ge-9/3/0**) connects to a device (that is, a DSLAM) on the access side of the network. You can define static customer VLANs (C-VLANs) for use by the wholesaler and any access network subscribers.

To configure the customer VLANs:

1. Edit the access side interface.

```
[edit]
user@host# edit interfaces ge-9/3/0
```

2. Specify the use of flexible VLAN tagging.

```
[edit interfaces ge-9/3/0]
user@host# set flexible-vlan-tagging
```

3. Edit the interface unit for the wholesaler VLAN.

```
[edit interfaces ge-9/3/0]
user@host# edit unit 14
```

4. Specify the type of encapsulation that you want the wholesaler VLAN to use.

```
[edit interfaces ge-9/3/0 unit 14]
user@host# set encapsulation ppp-over-ether
```

5. (Optional) Specify that you want the wholesaler VLAN to use Proxy ARP.

```
[edit interfaces ge-9/3/0 unit 14]
user@host# set proxy-arp
```

6. Define a unique VLAN ID for the wholesaler VLAN.

```
[edit interfaces ge-9/3/0 unit 14]
user@host# set vlan-id 14
```

7. Specify the dynamic profile that you want the wholesaler VLAN to use.

```
[edit interfaces ge-9/3/0 unit 14]
user@host# set pppoe-underlying-options dynamic-profile PPPoE_Wholesaler_Profile
```

Configuring Access Components for the PPPoE Wholesale Network Solution

When configuring a wholesale network, you must configure several components globally. This configuration provides access to RADIUS servers (if used) 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 58](#)
- [Configuring a PPPoE Wholesaler Access Profile on page 58](#)
- [Configuring PPPoE Retailer Access Profiles on page 59](#)

Configuring RADIUS Server Access

You can globally define any RADIUS servers in your network that either the wholesale 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 PPPoE 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 PPPoE_Wholesaler_Access
```

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

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

3. Specify that you want to configure RADIUS support.

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

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

```
[edit access profile PPPoE_Wholesaler_Access 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 PPPoE_Wholesaler_Access 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 PPPoE 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 PPPoE_Retailer_Access1
```

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

```
[edit access profile PPPoE_Retailer_Access1]
user@host# set authentication-order radius
```

3. Specify that you want to configure RADIUS support.

```
[edit access profile PPPoE_Retailer_Access1]
user@host# edit radius
```

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

```
[edit access profile PPPoE_Retailer_Access1 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 PPPoE_Retailer_Access1 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 PPPoE 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 PPPoE 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. Subscribers are assigned by the wholesaler RADIUS server to a particular retailer routing instance and can then be redirected to that retailer network space.

- [Configuring a Wholesale Dynamic Profile for use in the PPPoE Solution on page 60](#)

Configuring a Wholesale Dynamic Profile for use in the PPPoE Solution

You can configure a basic access profile to initially manage PPPoE 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 PPPoE_Wholesaler_Profile
```

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

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

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

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

4. Specify that you want to configure the **pp0** interface in the dynamic profile.

```
[edit dynamic-profiles PPPoE_Wholesaler_Profile]
user@host# edit interfaces pp0
```

5. Configure the unit for the **pp0** interface.

- a. Configure the variable for the unit number of the **pp0** interface.

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

```
[edit dynamic-profiles PPPoE_Wholesaler_Profile interfaces pp0]
user@host# edit unit $junos-interface-unit
```

- b. Configure PAP or CHAP (or both) to function on the interface.

```
[edit dynamic-profiles PPPoE_Wholesaler_Profile interfaces pp0 unit
"$junos-interface-unit"]
user@host# set ppp-options chap pap
```

- c. Configure the variable for the underlying interface of the pp0 interfaces.

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

```
[edit dynamic-profiles PPPoE_Wholesaler_Profile interfaces pp0 unit
"$junos-interface-unit"]
user@host# set pppoe-options underlying-interface $junos-underlying-interface
```


- d. Configure the router to act as a PPPoE server.

```
[edit dynamic-profiles PPPoE_Wholesaler_Profile interfaces pp0 unit
"$junos-interface-unit"]
user@host# set pppoe-options server
```

6. (Optional) Modify the PPPoE keepalive interval.

```
[edit dynamic-profiles PPPoE_Wholesaler_Profile interfaces pp0 unit
"$junos-interface-unit"]
user@host# set keepalives interval 15
```

7. Configure the family for the **pp0** interface.

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



NOTE: You can specify **inet** for IPv4 and **inet6** for IPv6. However, this solution provides the IPv4 configuration only.

```
[edit dynamic-profiles PPPoE_Wholesaler_Profile interfaces pp0 unit
"$junos-interface-unit"]
user@host# edit family inet
```

- b. Configure the unnumbered address for the family.

```
[edit dynamic-profiles PPPoE_Wholesaler_Profile interfaces pp0 unit
"$junos-interface-unit" family inet]
user@host# set unnumbered-address $junos-loopback-interface
```

Configuring Separate Routing Instances for PPPoE Service Retailers

As the owner of the system, the wholesaler 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 PPPoE_Retailer_Instance1
```

2. Specify the routing instance type for the retailer.

```
[edit routing-instances "PPPoE_Retailer_Instance1"]
user@host# set instance-type vrf
```

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

```
[edit routing-instances "PPPoE_Retailer_Instance1"]
user@host# set access-profile PPPoE_Retailer_Access1
```

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

```
[edit routing-instances "PPPoE_Retailer_Instance1"]
user@host# set interface ge-11/1/9.10
```

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

```
[edit routing-instances "PPPoE_Retailer_Instance1"]
user@host# set interface lo0.5
```



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

6. Specify an identifier to distinguish the VPN to which the route belongs.

```
[edit routing-instances "PPPoE_Retailer_Instance1"]
user@host# set route-distinguisher 1:1
```

7. Specify how routes are imported into the local PE router's VPN routing table from the remote PE router.

```
[edit routing-instances "PPPoE_Retailer_Instance1"]
user@host# set vrf-import policyImport
```

8. Specify which routes are exported from the local instance table to the remote PE router.

```
[edit routing-instances "PPPoE_Retailer_Instance1"]
user@host# set vrf-export policyExport
```

9. Repeat this procedure for other retailers.

Related Documentation

- [\[edit routing-instances\] Hierarchy Level](#)

Example: Wholesaler Dynamic Profile for a PPPoE Wholesale Network

This example specifies a dynamic profile name of *PPPoE_Wholesaler_Profile*, uses pp0 interfaces, and references the predefined input firewall filter.

```
PPPoE_Wholesaler_Profile {
  routing-instances {
    "$junos-routing-instance" {
      interface "$junos-interface-name";
    }
  }
  interfaces {
    pp0 {
      unit "$junos-interface-unit" {
        ppp-options {
          chap;
          pap;
        }
        pppoe-options {
          underlying-interface "$junos-underlying-interface";
          server;
        }
      }
      keepalives interval 15;
      family inet {
        filter {
          input "$junos-input-filter";
        }
      }
    }
  }
}
```

```

        output "$junos-output-filter";
    }
    unnumbered-address "$junos-loopback-interface";
}
}
}
}
}
}

```

- Related Documentation**
- [Configuring Dynamic Profiles for the PPPoE Layer 3 Wholesale Network Solution on page 59](#)

Example: Retailer Routing Instances for a PPPoE Wholesale Network

```

routing-instances {
  PPPoE_Retailer_Instance1 {
    instance-type vrf;
    access-profile PPPoE_Retailer_Access1;
    interface ge-11/1/9.10;
    interface lo0.5;
    route-distinguisher 1:1;
    vrf-import policyImport;
    vrf-export policyExport;
  }
  Retailer_Instance2 {
    instance-type vrf;
    access-profile PPPoE_Retailer_Access2;
    interface ge-11/1/9.10;
    interface lo0.6;
    route-distinguisher 2:2;
    vrf-import policyImport;
    vrf-export policyExport;
  }
}

```

- Related Documentation**
- [Configuring Separate Routing Instances for PPPoE Service Retailers on page 61](#)

PART 3

Configuring Layer 2 Wholesale Networks

- [Subscriber Management Layer 2 Wholesale Overview on page 67](#)
- [Configuring Layer 2 Wholesale Networks on page 71](#)

CHAPTER 6

Subscriber Management Layer 2 Wholesale Overview

- [Layer 2 and Layer 3 Wholesale Overview on page 67](#)
- [Wholesale Network Configuration Options and Considerations on page 68](#)
- [RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview on page 69](#)

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 9](#)
- [Broadband Subscriber Management PPPoE Layer 3 Wholesale Topology and Configuration Elements on page 53](#)
- [Broadband Subscriber Management Layer 2 Wholesale Topology and Configuration Elements on page 71](#)

Wholesale Network Configuration Options and Considerations

You can configure a wholesale network any number of ways using Juniper Networks hardware and Junos OS software. For information about subscriber management hardware support, see the *Junos OS Broadband Subscriber Management and Services Library*. The general configuration options, and considerations for each, are provided in the following table:

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).
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.

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 3 on page 6](#) is required for a wholesale network to function.

Table 5: 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
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 Broadband Subscriber Management and Services Library.*

CHAPTER 7

Configuring Layer 2 Wholesale Networks

- [Broadband Subscriber Management Layer 2 Wholesale Topology and Configuration Elements on page 71](#)
- [Layer 2 Wholesale Network Topology Overview on page 73](#)
- [Configuring a Retail Dynamic Profile for Use in the Layer 2 Wholesale Solution on page 75](#)
- [Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution on page 76](#)
- [Configuring VLAN Interfaces for the Layer 2 Wholesale Solution on page 78](#)
- [Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces on page 79](#)
- [Configuring NNI ISP-Facing Interfaces for the Layer 2 Wholesale Solution on page 80](#)
- [Configuring Direct ISP-Facing Interfaces for the Layer 2 Wholesale Solution on page 81](#)
- [Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers on page 81](#)
- [Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers on page 84](#)
- [Configuring Access Components for the Layer 2 Wholesale Network Solution on page 86](#)
- [Example: Retailer Dynamic Profile for a Layer 2 Wholesale Network on page 87](#)
- [Example: Access Interface for a Layer 2 Wholesale Network on page 87](#)
- [Example: Retailer Access Routing Instances for a Layer 2 Wholesale Network on page 88](#)
- [Example: Retailer NNI ISP-Facing Interfaces for a Layer 2 Wholesale Network on page 89](#)
- [Example: Retailer Direct ISP-Facing Interface for a Layer 2 Wholesale Network on page 89](#)


Broadband Subscriber Management Layer 2 Wholesale Topology and Configuration Elements

The network topology for the subscriber management Layer 2 wholesale solution includes configuring separate routing instances for individual retailers that use a portion of the router. This solution uses a Virtual Private LAN Service (VPLS) configuration.

Layer 2 wholesale networks are supported on MPC/MIC interfaces.

To explain the concept but limit complexity, this solution provides a configuration with one wholesaler and only two retailers. Figure 7 on page 72 illustrates a basic Layer 2 wholesale topology model from which you can expand.

Figure 7: Basic Subscriber Management Layer 2 Wholesale Solution Topology



- A Wholesaler Access PE Router Network Elements**
 Access Network Interface: GE-2/3/0
 RADIUS Authentication Server Address: 10.10.10.1
 RADIUS Accounting Server Address: 10.10.10.1
 Access Profile: AccessProfile
 Routing Instances: Retailer_Instance1
 Retailer_Instance2
 Dynamic Profile: 1.2_Access_Profile

- B Wholesaler Direct ISP-Facing Interface**
 Interface facing ISP Retailer 1: GE-1/1/0.1
 VPLS Routing Instances: Retailer_Instance1
- C Wholesaler NNI-1-ISP-Facing Interface**
 Interface facing ISP Retailer 1: GE-1/1/0.0
 VPLS Routing Instances: Retailer_Instance1
- D Wholesaler NNI-2-ISP-Facing Interface**
 Interface facing ISP Retailer 2: GE-2/2/0.0
 VPLS Routing Instances: Retailer_Instance2

80-4212

When you are configuring a Layer 2 wholesale network solution, the following configuration elements are required:

- Subscriber access dynamic VLAN configuration including dynamic profile configuration for retailer routing instances
- Routing instance configuration for individual retailers on provider edge (PE) routers and network-to-network interface (NNI) routers.
- VLAN interface configuration
- RADIUS server access configuration
- Core network configuration

**Related
Documentation**

- [Layer 2 and Layer 3 Wholesale Overview on page 3](#)
- [Layer 2 Wholesale Network Topology Overview on page 73](#)

Layer 2 Wholesale Network Topology Overview

This configuration explains how to configure a simple Layer 2 wholesale subscriber access network. This solution illustrates two Internet Service Provider (ISP) retailers sharing access to a wholesaler network. The wholesaler network contains a Layer 2 Network access router and two Virtual Private LAN Service (VPLS) network-to-network interface (NNI) routers.



NOTE: You can have more than one ISP router connecting to a single VPLS NNI router with VPLS interfaces configured with routing instances specific to each different ISP-facing interfaces.


The example also shows two different connection options from one subscriber access router to one of the individual ISP access routers. One connection option uses an interface on the subscriber access router to connect directly to the ISP access router. Another connection option uses two routers: a subscriber access router and another NNI router that connects to the ISP access router.



NOTE: When using the NNI router connection option, use a standard BGP or MPLS configuration between the subscriber access routers and the edge router that connects to the ISP access routers. See the *Junos OS Routing Protocols Library for Routing Devices* for information about BGP configuration. See the *Junos OS MPLS Applications Library for Routing Devices* for information about MPLS configuration.

[Figure 8 on page 74](#) provides the reference topology for this configuration example.

Figure 8: Layer 2 Wholesale Network Reference Topology



Related Documentation

- [Layer 2 and Layer 3 Wholesale Overview on page 3](#)
- [Broadband Subscriber Management Layer 2 Wholesale Topology and Configuration Elements on page 71](#)

Configuring a Retail Dynamic Profile for Use in the Layer 2 Wholesale 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. Define the dynamic interfaces variable for the dynamic profile.

```
[edit dynamic-profiles Subscriber_Profile_Retail1]
user@host# set interfaces $junos-interface-ifd-name
```

5. Define the dynamic interface unit variable for the dynamic profile.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name"]
user@host# set unit $junos-interface-unit
```

6. (Optional) Define the VLAN encapsulation for the dynamic interfaces.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name"
unit "$junos-interface-unit"]
user@host# set encapsulation vlan-vpls
```



NOTE: If you choose not to specify an encapsulation for the logical interface, you must specify encapsulation for the physical interface.

7. Define the VLAN tag variables for the dynamic profile:



NOTE: This solution example uses stacked VLAN tagging. However, you can also specify single-tag VLANs. For additional information about configuring dynamic VLANs, see the *Junos OS Broadband Subscriber Management and Services Library*.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name"
unit "$junos-interface-unit"]
user@host# set vlan-tags outer $junos-stacked-vlan-id inner $junos-vlan-id
```

8. Define the input and output VLAN maps. See [“Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution” on page 76](#) for details.

9. Specify the unit family as **vpls** at the `[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" family]` hierarchy level.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name"
  unit "$junos-interface-unit"]
user@host# set family vpls
```

Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution

Stacking and rewriting VLAN tags allows you to use an additional (outer) VLAN tag to differentiate between routers in the Layer 2 wholesale network. A frame can be received on an interface, or it can be internal to the system (as a result of the **input-vlan-map** statement).

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

You can configure the following single-action VLAN rewrite operations:

- **pop**—Remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.
- **push**—Add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.
- **swap**—Replace the inner VLAN tag of the incoming frame with a user-specified VLAN tag value.

You configure VLAN rewrite operations for logical interfaces in the input VLAN map for incoming frames and in the output VLAN map for outgoing frames.

You can include both the **input-vlan-map** and **output-vlan-map** statements at the **[edit dynamic-profiles profile-name interface "\$junos-interface-ifd-name" unit "\$junos-interface-unit]** hierarchy level.

The type of VLAN rewrite operation permitted depends upon whether the frame is single-tagged or dual-tagged. [Table 6 on page 76](#) shows supported rewrite operations and whether they can be applied to single-tagged frames or dual-tagged frames. The table also indicates the number of tags being added or removed during the operation.

Table 6: Rewrite Operations on Single-Tagged and Dual-Tagged Frames

Rewrite Operation	Single-Tagged	Dual-Tagged	Number of Tags
pop	Yes	Yes	– 1
push	Yes	Yes	+1
swap	Yes	Yes	0

Depending on the VLAN rewrite operation, you configure the rewrite operation for the interface in the input VLAN map, the output VLAN map, or both. [Table 7 on page 77](#) shows what rewrite operation combinations you can configure. “None” means that no rewrite operation is specified for the VLAN map.

Table 7: Applying Rewrite Operations to VLAN Maps

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

To configure the input VLAN map:



NOTE: You configure the `input-vlan-map` statement only when there is a need either to push an outer tag on a single-tagged subscriber packet or to modify the outer tag in a subscriber dual-tagged packet.

1. Include the `input-vlan-map` statement.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name"
 unit "$junos-interface-unit"]
user@host# edit input-vlan-map
```

2. Specify the action that you want the input VLAN map to take.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name"
 unit "$junos-interface-unit" input-vlan-map]
user@host# set push
```

3. Include the `vlan-id` statement along with the `$junos-vlan-map-id` dynamic variable.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name"
 unit "$junos-interface-unit" input-vlan-map]
user@host# set vlan-id $junos-vlan-map-id
```

To configure the output VLAN map:



NOTE: You configure the `output-vlan-map` statement only when there is a need to either pop or modify the outer tag found in a dual-tagged packet meant for the subscriber.

1. Include the `output-vlan-map` statement.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name"
 unit "$junos-interface-unit"]
user@host# edit output-vlan-map
```

2. Specify the action that you want the output VLAN map to take.

```
[edit dynamic-profiles Subscriber_Profile_Retail] interfaces "$junos-interface-ifd-name"
  unit "$junos-interface-unit" output-vlan-map]
user@host# set pop
```

You must know whether the VLAN rewrite operation is valid and is applied to the input VLAN map or the output VLAN map. You must also know whether the rewrite operation requires you to include statements to configure the inner and outer tag protocol identifiers (TPIDs) and inner and outer VLAN IDs in the input VLAN map or output VLAN map. For information about configuring inner and outer TPIDs and inner and outer VLAN IDs, see *Configuring Inner and Outer TPIDs and VLAN IDs*.

Configuring VLAN Interfaces for the Layer 2 Wholesale Solution

Clients access the Layer 2 Wholesale network through a specific interface. After they access this interface, and when they are authenticated, VLANs are dynamically created to carry the client traffic.

To configure a VLAN interface for dynamic access of clients:

1. Access the physical interface that you want to use for dynamically creating VLAN interfaces.

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

2. Specify the desired VLAN tagging.



NOTE: This example uses flexible VLAN tagging to simultaneously support transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port.

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

3. Specify that you want to automatically configure VLAN interfaces.

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

4. Specify that you want to configure single VLANs.

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

5. Define the VLAN ranges for the configuration.

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

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

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

7. Specify that any type of VLAN Ethernet packet is accepted by the interface.

```
[edit interfaces ge-2/3/0 auto-configure vlan-ranges dynamic-profile
  "Subscriber_Profile_Retail1"]
user@host# set accept any
```

8. Repeat steps for any other interfaces that you want to use for creating VLANs.
9. Specify the encapsulation type for the VLAN interfaces.

```
[edit interfaces ge-2/3/0]
user@host# edit encapsulation flexible-ethernet-services
```

Related Documentation • [Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces on page 79](#)

Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces

Each dynamic VLAN interface in a Layer 2 wholesale network must use encapsulation. You can configure encapsulation dynamically for each VLAN interface by using the **encapsulation** statement at the **[edit dynamic-profiles *profile-name* interface "\$junos-interface-ifd-name" unit "\$junos-interface-unit"]** hierarchy level or configure encapsulation for the physical interfaces at the **[edit interfaces *interface-name*]** hierarchy level for each dynamically created VLAN interface to use. However, how you choose to configure (or not configure) encapsulation at the **[edit dynamic-profiles *profile-name* interface "\$junos-interface-ifd-name" unit "\$junos-interface-unit"]** hierarchy level affects how you configure encapsulation at the **[edit interfaces *interface-name*]** hierarchy level.

[Table 8 on page 79](#) provides the valid encapsulation combinations for both dynamic profiles and physical interfaces in the Layer 2 wholesale network.

Table 8: Encapsulation Combinations for Layer 2 Wholesale Interfaces

Dynamic Profile Encapsulation	Physical Interface Encapsulation	Usage Notes
vlan-vpls	vlan-vpls	Using the vlan-vpls encapsulation type in both the dynamic profile and when configuring the physical interface limits the VLAN ID value to a number greater than or equal to 512.
vlan-vpls	flexible-ethernet-services	Using the flexible-ethernet-services encapsulation type removes any VLAN ID value limitation.
vlan-vpls	extended-vlan-vpls	The extended-vlan-vpls encapsulation type can support multiple TPIDs. Using this encapsulation type removes any VLAN ID value limitation.
No encapsulation type	extended-vlan-vpls	The extended-vlan-vpls encapsulation type can support multiple TPIDs. Using this encapsulation type removes any VLAN ID value limitation.

To configure encapsulation for Layer 2 wholesale VLAN interfaces:

1. (Optional) Define the VLAN encapsulation for the dynamic interfaces.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name"
  unit "$junos-interface-unit"]
user@host# set encapsulation encapsulation-type
```

- Specify the encapsulation type for the physical VLAN interface.

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



NOTE: If you choose not to specify an encapsulation for the logical interface, you must specify `extended-vlan-vpls` encapsulation for the physical interface.

Related Documentation

- [Configuring a Retail Dynamic Profile for Use in the Layer 2 Wholesale Solution on page 75](#)
- [Configuring VLAN Interfaces for the Layer 2 Wholesale Solution on page 78](#)

Configuring NNI ISP-Facing Interfaces for the Layer 2 Wholesale Solution

You must configure separate, ISP-facing interfaces on each NNI ISP-facing router that connect to individual retailer ISP access routers in the Layer 2 Wholesale solution.



NOTE: On the network-to-network (NNI) or egress interfaces of provider edge (PE) routers, you cannot configure the inner-range `vid1—vid2` option with the `vlan-tags` statement for ISP-facing interfaces.

To configure an NNI ISP-facing interface:

- Access the physical interface that you want to use to access the retailer ISP network.

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

- Specify the encapsulation type for the VLAN interfaces.

```
[edit interfaces ge-1/1/0]
user@host# edit encapsulation ethernet-vpls
```

- Specify the interface unit that you want ISP clients to use.

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

- Repeat these steps for any other NNI ISP-facing interfaces that you want to use. In this example, you must also configure interface `ge-2/2/0.0`.

Related Documentation

- [Configuring Direct ISP-Facing Interfaces for the Layer 2 Wholesale Solution on page 81](#)
- [Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers on page 81](#)

Configuring Direct ISP-Facing Interfaces for the Layer 2 Wholesale Solution

When connecting a subscriber access router directly to an ISP access router, you must define any ISP-facing interfaces that connect to the retailer ISP access routers as core-facing interfaces.

To configure a direct ISP-facing interface:

1. Access the physical interface that you want to use to access the retailer ISP network.

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

2. Specify the encapsulation type for the VLAN interfaces.

```
[edit interfaces ge-1/1/0]
user@host# edit encapsulation ethernet-vpls
```

3. Specify the interface unit that you want ISP clients to use.

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

4. Specify the unit family.

```
[edit interfaces ge-1/1/0 unit 1]
user@host# set family vpls
```

5. Define the interface as core-facing to ensure that the network does not improperly treat the interface as a client interface..

```
[edit interfaces ge-1/1/0 unit 1 family vpls]
user@host# set core-facing
```

6. Repeat steps for any other direct ISP-facing interfaces that you want to use..

Related Documentation

- [Configuring NNI ISP-Facing Interfaces for the Layer 2 Wholesale Solution on page 80](#)
- [Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers on page 81](#)

Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers

As the owner of the system, the wholesaler 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.

When creating separate routing instances, it is important to understand the role that the router plays in the Layer 2 Wholesale network and specify that role (either access or NNI) in the routing instance configuration. If the router connects directly to an ISP network (or ISP-controlled device), you must configure the routing instances as an NNI routing instance. See [“Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers” on page 84](#).

To define an access retailer routing instance:

1. Create the retailer routing instance.

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

2. Specify the VLAN model that you want the retailer to follow.

```
[edit routing-instances RetailerInstance1]
user@host# set vlan-model one-to-one
```

3. Specify the role that you want the routing instance to take.

```
[edit routing-instances RetailerInstance1]
user@host# set instance-role access
```

4. Specify the routing instance type for the retailer.

```
[edit routing-instances RetailerInstance1]
user@host# set instance-type l2backhaul-vpn
```

5. Specify the access interface for the retailer.

```
[edit routing-instances RetailerInstance1]
user@host# set interface ge-2/3/0.0
```

6. Specify that access ports in this VLAN domain do not forward packets to each other.

```
[edit routing-instances RetailerInstance1]
user@host# set no-local-switching
```

7. Specify a unique identifier attached to a route that enables you to distinguish to which VPN the route belongs.

```
[edit routing-instances RetailerInstance1]
user@host# set route-distinguisher 10.10.1.1
```

8. (Optional) Specify a VRF target community.

```
[edit routing-instances RetailerInstance1]
user@host# set vrf-target target:100:1
```



NOTE: The purpose of the `vrf-target` statement is to simplify the configuration by allowing you to configure most statements at the `[edit routing-instances]` hierarchy level.

9. Define the VPLS protocol for the routing instance.

- a. Access the routing instance
- protocols**
- hierarchy.

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

- b. Enable VPLS on the routing instance.

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

- c. Specify the maximum number of sites allowed for the VPLS domain.

```
[edit routing-instances RetailerInstance1 protocols vpls]
user@host# set site-range 10
```

- d. Specify the size of the VPLS MAC address table for the routing instance.

```
[edit routing-instances RetailerInstance1 protocols vpls]
user@host# set mac-table-size 6000
```

- e. Specify the maximum number of MAC addresses that can be learned by the VPLS routing instance.

```
[edit routing-instances RetailerInstance1 protocols vpls]
user@host# set interface-mac-limit 2000
```

- f. (Optional) Specify the
- no-tunnel-services**
- statement if the router does not have a Tunnel Services PIC.

```
[edit routing-instances RetailerInstance1 protocols vpls]
user@host# set no-tunnel-services
```

- g. Specify a site name.

```
[edit routing-instances RetailerInstance1 protocols vpls]
user@host# set site A-PE
```

- h. Specify a site identifier.

```
[edit routing-instances RetailerInstance1 protocols vpls site A-PE]
user@host# set site-identifier 1
```

10. Repeat this procedure for other retailers. In this example, you must configure a routing instance for Retailer 2.

Related Documentation

- [Configuring VPLS Routing Instances](#)
- [Configuring NNI ISP-Facing Interfaces for the Layer 2 Wholesale Solution on page 80](#)
- [Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers on page 84](#)

Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers

As the owner of the system, the wholesaler 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.

When creating separate routing instances, it is important to understand the role that the router plays in the Layer 2 Wholesale network and specify that role (either access or NNI) in the routing instance configuration. If the router connects to the access portion of the network (for example, to an MSAN device), you must configure the routing instances as an access routing instance. See [“Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers” on page 81](#).

To define a retailer routing instance:

1. Create the retailer routing instance.

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

2. Specify the VLAN model that you want the retailer to follow.

```
[edit routing-instances RetailerInstance1]
user@host# set vlan-model one-to-one
```

3. Specify the role that you want the routing instance to take.

```
[edit routing-instances RetailerInstance1]
user@host# set instance-role nni
```

4. Specify the routing instance type for the retailer.

```
[edit routing-instances RetailerInstance1]
user@host# set instance-type l2backhaul-vpn
```

5. Define the NNI ISP-facing interface for this retailer.

```
[edit routing-instances RetailerInstance1]
user@host# set interface ge-1/1/0.0
```

6. Specify that access ports in this VLAN domain do not forward packets to each other.

```
[edit routing-instances RetailerInstance1]
user@host# set no-local-switching
```

7. Specify a unique identifier attached to a route that enables you to distinguish to which VPN the route belongs.

```
[edit routing-instances RetailerInstance1]
user@host# set route-distinguisher 10.10.1.1
```

8. (Optional) Specify a VRF target community.

```
[edit routing-instances RetailerInstance1]
user@host# set vrf-target target:100:1
```




NOTE: The purpose of the `vrf-target` statement is to simplify the configuration by allowing you to configure most statements at the `[edit routing-instances]` hierarchy level.

9. Define the VPLS protocol for the routing instance.

- a. Access the routing instance **protocols** hierarchy.

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

- b. Enable VPLS on the routing instance.

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

- c. Specify the maximum number of sites allowed for the VPLS domain.

```
[edit routing-instances RetailerInstance1 protocols vpls]
user@host# set site-range 1000
```

- d. (Optional) Specify the **no-tunnel-services** statement if the router does not have a Tunnel Services PIC.

```
[edit routing-instances RetailerInstance1 protocols vpls]
user@host# set no-tunnel-services
```

- e. Specify a site name.

```
[edit routing-instances RetailerInstance1 protocols vpls]
user@host# set site A-PE
```

- f. Specify a site identifier.

```
[edit routing-instances RetailerInstance1 protocols vpls site A-PE]
user@host# set site-identifier 1
```

- g. Define the connectivity of the VPLS routing instance as **permanent** to keep the VPLS connection up until specifically taken down.

```
[edit routing-instances RetailerInstance1 protocols vpls]
user@host# set connectivity-type permanent
```

10. Repeat this procedure for other retailers.

**Related
Documentation**

- [Configuring VPLS Routing Instances](#)
- [Configuring VLAN Interfaces for the Layer 2 Wholesale Solution on page 78](#)
- [Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers on page 81](#)

Configuring Access Components for the Layer 2 Wholesale Network Solution

When configuring a wholesale network, you must configure several components globally. This configuration provides access to RADIUS servers (if used) 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 86](#)
- [Configuring a Layer 2 Wholesaler Access Profile on page 86](#)

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 Layer 2 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 AccessProfile
```

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

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

3. Specify that you want to configure RADIUS support.

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

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

```
[edit access profile AccessProfile 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 AccessProfile 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*.

Example: Retailer Dynamic Profile for a Layer 2 Wholesale Network

```
dynamic-profiles {
  Subscriber_Profile_Retail1 {
    routing-instances {
      "$junos-routing-instance" {
        interface "$junos-interface-name";
      }
    }
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-interface-unit" {
          encapsulation vlan-vpls;
          vlan-tags outer "$junos-stacked-vlan-id" inner "$junos-vlan-id";
          input-vlan-map {
            swap;
            vlan-id "$junos-vlan-map-id";
          }
          output-vlan-map swap;
          family vpls;
        }
      }
    }
  }
}
```

- Related Documentation**
- [Layer 2 and Layer 3 Wholesale Overview on page 3](#)
 - [Layer 2 Wholesale Network Topology Overview on page 73](#)
 - [Configuring a Retail Dynamic Profile for Use in the Layer 2 Wholesale Solution on page 75](#)

Example: Access Interface for a Layer 2 Wholesale Network

```
interfaces {
  ge-2/3/0 {
    flexible-vlan-tagging;
    auto-configure {
      stacked-vlan-ranges {
        dynamic-profile Subscriber_Profile_Retail1 {
```

```
        accept any;
        ranges {
            any,any;
        }
    }
    access-profile AccessProfile;
}
}
encapsulation flexible-ethernet-services;
}
```

**Related
Documentation**

- [Layer 2 and Layer 3 Wholesale Overview on page 3](#)
- [Layer 2 Wholesale Network Topology Overview on page 73](#)
- [Configuring VLAN Interfaces for the Layer 2 Wholesale Solution on page 78](#)

Example: Retailer Access Routing Instances for a Layer 2 Wholesale Network

You need to create a routing instance for each retailer to keep routing information for different retailers separate and to define servers and forwarding options specific to each retailer.

There are two types of routing instances that you can create: access or NNI. The following code snippets show how to configure separate access routing instances for two retailers: Retailer_Instance1 and Retailer_Instance2.

```
routing-instances {
  Retailer_Instance1 {
    vlan-model one-to-one;
    instance-role access;
    instance-type l2backhaul-vpn;
    interface ge-1/1/0.0
    no-local-switching;
    route-distinguisher 10.10.1.1:1;
    vrf-target target:100:1;
    protocols {
      vpls {
        site-range 10;
        mac-table-size {
          6000;
        }
        interface-mac-limit {
          2000;
        }
        no-tunnel-services;
        site A-PE {
          site-identifier 1;
        }
      }
    }
  }
  Retailer_Instance2 {
    vlan-model one-to-one;
```

```

instance-role access;
instance-type l2backhaul-vpn;
interface ge-2/2/0.0
no-local-switching;
route-distinguisher 10.10.1.1:2;
vrf-target target:300:1;
protocols {
  vpls {
    site-range 1000;
    no-tunnel-services;
    site A-PE {
      site-identifier 1;
    }
  }
}
}
}

```

- Related Documentation**
- [Layer 2 and Layer 3 Wholesale Overview on page 3](#)
 - [Layer 2 Wholesale Network Topology Overview on page 73](#)
 - [Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers on page 81](#)

Example: Retailer NNI ISP-Facing Interfaces for a Layer 2 Wholesale Network

```

interfaces {
  ge-1/1/0 {
    description Retailer 1 NNI ISP-facing interface;
    encapsulation ethernet-vpls;
    unit 0 {
    }
  }
  interfaces {
    ge-2/2/0 {
      description Retailer 2 NNI ISP-facing interface;
      encapsulation ethernet-vpls;
      unit 0;
    }
  }
}

```

- Related Documentation**
- [Layer 2 and Layer 3 Wholesale Overview on page 3](#)
 - [Layer 2 Wholesale Network Topology Overview on page 73](#)
 - [Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers on page 84](#)

Example: Retailer Direct ISP-Facing Interface for a Layer 2 Wholesale Network

```

interfaces {
  ge-1/1/0 {
    description Retailer 1 Direct ISP-facing interface;
    encapsulation ethernet-vpls;
    unit 1
  }
}

```

```
        family vpls {  
            core-facing;  
        }  
    }  
}
```

**Related
Documentation**

- [Layer 2 and Layer 3 Wholesale Overview on page 3](#)
- [Layer 2 Wholesale Network Topology Overview on page 73](#)
- [Configuring Direct ISP-Facing Interfaces for the Layer 2 Wholesale Solution on page 81](#)

PART 4

Configuration Statements and Operational Commands

- [Configuration Statements on page 93](#)
- [Operational Commands on page 267](#)

CHAPTER 8

Configuration Statements



- [accept](#) on page 97
- [access-profile](#) on page 98
- [accounting-server](#) on page 99
- [active-server-group](#) on page 100
- [address](#) on page 101
- [address-assignment](#) (Address-Assignment Pools) on page 105
- [authentication](#) (DHCP Local Server) on page 106
- [authentication](#) (DHCP Relay Agent) on page 107
- [authentication-order](#) on page 108
- [authentication-server](#) on page 109
- [auto-configure](#) on page 110
- [connectivity-type](#) on page 112
- [core-facing](#) on page 113
- [demux0](#) (Dynamic Interface) on page 114
- [demux-options](#) (Dynamic Interface) on page 115
- [demux-source](#) (Dynamic IP Demux Interface) on page 116
- [demux-source](#) (Dynamic Underlying Interface) on page 117
- [demux-source](#) (Underlying Interface) on page 118
- [dhcp-attributes](#) (Address-Assignment Pools) on page 119
- [dhcp-local-server](#) on page 120
- [dhcp-relay](#) on page 126
- [dhcpv6](#) (DHCP Local Server) on page 133
- [domain-name](#) (Address-Assignment Pools) on page 136
- [dynamic-profile](#) (DHCP Local Server) on page 137
- [dynamic-profile](#) (DHCP Relay Agent) on page 138
- [dynamic-profile](#) (Dynamic PPPoE) on page 139
- [dynamic-profile](#) (Stacked VLAN) on page 140
- [dynamic-profile](#) (VLAN) on page 141

- [dynamic-profiles](#) on page 142
- [encapsulation \(Dynamic Interfaces\)](#) on page 150
- [encapsulation \(Logical Interface\)](#) on page 153
- [encapsulation \(Physical Interface\)](#) on page 157
- [family](#) on page 162
- [family \(Address-Assignment Pools\)](#) on page 166
- [family \(Dynamic Demux Interface\)](#) on page 167
- [family \(Dynamic PPPoE\)](#) on page 168
- [family \(Dynamic Standard Interface\)](#) on page 169
- [flexible-vlan-tagging](#) on page 171
- [forwarding-options](#) on page 172
- [grace-period](#) on page 173
- [group \(DHCP Local Server\)](#) on page 174
- [group \(DHCP Relay Agent\)](#) on page 177
- [input-vlan-map \(Dynamic Interfaces\)](#) on page 180
- [interface \(DHCP Local Server\)](#) on page 181
- [interface \(DHCP Relay Agent\)](#) on page 183
- [interface \(Dynamic Routing Instances\)](#) on page 184
- [interface \(Routing Instances\)](#) on page 185
- [interface-mac-limit \(VPLS\)](#) on page 186
- [interfaces](#) on page 187
- [interfaces \(Static and Dynamic Subscribers\)](#) on page 188
- [instance-role](#) on page 192
- [instance-type](#) on page 193
- [ip-address-first](#) on page 195
- [keepalives \(Dynamic Profiles\)](#) on page 196
- [mac-validate \(Dynamic IP Demux Interface\)](#) on page 197
- [multicast-replication](#) on page 198
- [no-local-switching](#) on page 198
- [no-tunnel-services](#) on page 199
- [maximum-lease-time](#) on page 200
- [output-vlan-map \(Dynamic Interfaces\)](#) on page 201
- [pap \(Dynamic PPP\)](#) on page 201
- [pool \(Address-Assignment Pools\)](#) on page 202
- [pool-match-order](#) on page 203
- [pop \(Dynamic VLANs\)](#) on page 203
- [pppoe-options \(Dynamic PPPoE\)](#) on page 204

- [pppoe-underlying-options \(Static and Dynamic Subscribers\)](#) on page 205
- [ppp-options \(Dynamic PPP\)](#) on page 206
- [prefix \(Address-Assignment Pools\)](#) on page 207
- [profile \(Access\)](#) on page 208
- [protocols](#) on page 212
- [proxy-arp](#) on page 215
- [proxy-arp \(Dynamic Profiles\)](#) on page 216
- [push \(Dynamic VLANs\)](#) on page 216
- [radius \(Access Profile\)](#) on page 217
- [radius-server](#) on page 219
- [range \(Address-Assignment Pools\)](#) on page 220
- [ranges \(Dynamic VLAN\)](#) on page 221
- [route-distinguisher](#) on page 222
- [routing-instances \(Dynamic Profiles\)](#) on page 224
- [routing-instances \(Multiple Routing Entities\)](#) on page 226
- [secret](#) on page 227
- [server \(Dynamic PPPoE\)](#) on page 227
- [server-group](#) on page 228
- [services \(System Services\)](#) on page 229
- [site \(VPLS Multihoming for FEC 128\)](#) on page 231
- [site-identifier \(VPLS\)](#) on page 232
- [site-range](#) on page 232
- [stacked-vlan-ranges](#) on page 233
- [stacked-vlan-tagging](#) on page 234
- [system](#) on page 234
- [traceoptions \(DHCP\)](#) on page 235
- [underlying-interface \(demux0\)](#) on page 237
- [underlying-interface \(Dynamic PPPoE\)](#) on page 238
- [unit](#) on page 239
- [unit \(Dynamic Demux Interface\)](#) on page 245
- [unit \(Dynamic Profiles Standard Interface\)](#) on page 247
- [unnumbered-address \(Dynamic PPPoE\)](#) on page 249
- [unnumbered-address \(Dynamic Profiles\)](#) on page 250
- [unnumbered-address \(Ethernet\)](#) on page 252
- [user-prefix \(DHCP Local Server\)](#) on page 253
- [vlan-id \(Dynamic VLANs\)](#) on page 254
- [vlan-id \(VLAN ID to Be Bound to a Logical Interface\)](#) on page 255

- [vlan-model on page 256](#)
- [vlan-ranges on page 257](#)
- [vlan-tags on page 258](#)
- [vlan-tags \(Stacked VLAN Tags\) on page 259](#)
- [vpls \(Routing Instance\) on page 261](#)
- [vrf-export on page 263](#)
- [vrf-import on page 264](#)
- [vrf-target on page 265](#)

accept

Syntax	<code>accept (any dhcp-v4 dhcp-v6 inet inet6 pppoe);</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> auto-configure stacked-vlan-ranges dynamic-profile <i>profile-name</i>], [edit interfaces <i>interface-name</i> auto-configure vlan-ranges dynamic-profile <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5. dhcp-v4 option added in Junos OS Release 10.0. dhcp-v6 , inet6 and pppoe options added in Junos OS Release 10.2. any option added in Junos OS Release 10.4.
Description	Specify the type of VLAN Ethernet packet accepted by an interface that is associated with a VLAN dynamic profile or stacked VLAN dynamic profile.
Options	<p>any—Any packet type. Specifies that any incoming packets trigger the dynamic creation of a VLAN with properties determined by the auto-configure interface configuration stanza and associated profile attributes. This option is used when configuring wholesaling in a Layer 2 network.</p> <p>dhcp-v4—IPv4 DHCP packet type. Specifies that incoming IPv4 DHCP discover packets trigger the dynamic creation of a VLAN with properties determined by the auto-configure interface configuration stanza and associated profile attributes</p> <p>.....</p> <p> NOTE: The DHCP-specific mac-address and option-82 options are rejected if the accept statement is not set to dhcp-v4.</p> <p>.....</p> <p>dhcp-v6—IPv6 DHCP packet type. Specifies that incoming IPv6 DHCP discover packets trigger the dynamic creation of a VLAN with properties determined by the auto-configure interface configuration stanza and associated profile attributes.</p> <p>inet—IPv4 Ethernet and ARP packet type.</p> <p>inet6—IPv6 Ethernet packet type.</p> <p>pppoe—Point-to-Point Protocol over Ethernet packet type.</p> <p>.....</p> <p> NOTE: The pppoe VLAN Ethernet packet type option is supported only for MPC/MIC interfaces.</p> <p>.....</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs](#)
 - [Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs](#)
 - [Configuring VLAN Interfaces for the Layer 2 Wholesale Solution on page 78](#)
 - [Configuring Subscriber Packet Types to Trigger VLAN Authentication](#)

access-profile

Syntax	<code>access-profile <i>profile-name</i>;</code>
Hierarchy Level	<code>[edit],</code> <code>[edit forwarding-options dhcp-relay]</code> <code>[edit forwarding-options dhcp-relay group <i>group-name</i>]</code> <code>[edit forwarding-options dhcp-relay dhcpv6]</code> <code>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i>]</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>]</code> <code>[edit interfaces <i>interface-name</i> auto-configure vlan-ranges],</code> <code>[edit interfaces <i>interface-name</i> auto-configure stacked-vlan-ranges],</code> <code>[edit routing-instances <i>routing-instances-name</i>]</code> <code>[edit system services dhcp-local-server]</code> <code>[edit system services dhcp-local-server group <i>group-name</i>]</code> <code>[edit system services dhcp-local-server dhcpv6]</code> <code>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	After you have created the access profile that specifies authentication and accounting parameters, you must specify where the profile is used. Authentication and accounting will not run unless you specify the profile. You can attach access profiles globally at the [edit] hierarchy level, or you can apply them to DHCP clients or subscribers, VLANs, or to a routing instance.
Options	<i>profile-name</i> —Name of the access profile that you configured at the [edit access profile name] hierarchy level.
Required Privilege Level	access —To view this statement in the configuration. access-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Attaching Access Profiles• Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces• Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution on page 15• Configuring Access Components for the PPPoE Wholesale Network Solution on page 57

accounting-server

Syntax	<code>accounting-server [<i>ip-address</i>];</code>
Hierarchy Level	[edit access profile <i>profile-name</i> radius]
Release Information	Statement introduced in Junos OS Release 9.1.
Description	Specify a list of the RADIUS accounting servers used for accounting for DHCP, L2TP, and PPP clients.
Options	<i>ip-address</i> —IP version 4 (IPv4) address.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Authentication and Accounting Parameters for Subscriber Access</i>

active-server-group

Syntax	<code>active-server-group <i>server-group-name</i>;</code>
Hierarchy Level	<pre> [edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit forwarding-options dhcp-relay group <i>group-name</i>], [edit forwarding-options dhcp-relay group <i>group-name</i> dhcpv6], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6], [edit logical-systems <i>logical-system-name</i> forwarding-options group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options group <i>group-name</i> dhcpv6], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> dhcpv6], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay] [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i>]</pre>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Apply a DHCP relay agent configuration to the named group of DHCP server addresses. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.</p> <p>A group-specific configuration overrides a global option.</p>
Options	<i>server-group-name</i> —Name of the group of DHCP or DHCPv6 server addresses to which the DHCP or DHCPv6 relay agent configuration applies.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Extended DHCP Relay Agent Overview</i> • <i>Configuring Active Server Groups</i> • <i>Configuring Group-Specific DHCP Relay Options</i> • dhcp-relay on page 126

address

```

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

```

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure the interface address.



NOTE: The vrrp High Availability functionality is not available on the QFX Series.

Options *address*—Address of the interface.

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



NOTE: If you configure the same address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration. The remaining address configurations are ignored and can leave interfaces without an address. Interfaces that do not have an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

For example, in the following configuration the address configuration of interface xe-0/0/1.0 is ignored:

```
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.1/24;
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.1.1/24;
      }
    }
  }
}
```

For more information on configuring the same address on multiple interfaces, see *Configuring the Interface Address*.

The remaining statements are explained separately.



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

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

**Related
Documentation**

- *Configuring the Protocol Family*
- *Junos OS Administration Library for Routing Devices*
- *family*
- *negotiate-address*
- [unnumbered-address \(Ethernet\) on page 252](#)
- *Junos OS Administration Library for Routing Devices*

address-assignment (Address-Assignment Pools)

Syntax

```

address-assignment {
    abated-utilization percentage;
    abated-utilization-v6 percentage;
    high-utilization percentage;
    high-utilization-v6 percentage;
    neighbor-discovery-router-advertisement ndra-pool-name;
    pool pool-name {
        family family {
            dhcp-attributes {
                protocol-specific attributes;
            }
            host hostname {
                hardware-address mac-address;
                ip-address ip-address;
            }
            network ip-prefix / <prefix-length>;
            prefix ipv6-prefix;
            range range-name {
                high upper-limit;
                low lower-limit;
                prefix-length prefix-length;
            }
        }
        link pool-name;
    }
}

```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Configure address-assignment pools that can be used by different client applications.



NOTE: Support for subordinate statements is platform-specific. See individual statement topics for support information.

Options *pool-name*—Name assigned to an address-assignment pool.

The remaining statements are explained separately.

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

Related Documentation

- *Address-Assignment Pools Overview*
- *Configuring Address-Assignment Pools*

- *Configuring an Address-Assignment Pool for L2TP LNS with Inline Services*

authentication (DHCP Local Server)

Syntax

```
authentication {  
  password password-string;  
  username-include {  
    circuit-type;  
    client-id;  
    delimiter delimiter-character;  
    domain-name domain-name-string;  
    interface-name ;  
    logical-system-name;  
    mac-address;  
    option-60;  
    option-82 <circuit-id> <remote-id>;  
    relay-agent-interface-id;  
    relay-agent-remote-id;  
    relay-agent-subscriber-id;  
    routing-instance-name;  
    user-prefix user-prefix-string;  
  }  
}
```

Hierarchy Level

[edit system services [dhcp-local-server](#)],
[edit system services dhcp-local-server [dhcipv6](#)],
[edit system services dhcp-local-server dhcipv6 [group group-name](#)],
[edit system services dhcp-local-server [group group-name](#)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services [dhcp-local-server ...](#)],
[edit logical-systems *logical-system-name* system services [dhcp-local-server ...](#)],
[edit routing-instances *routing-instance-name* system services [dhcp-local-server ...](#)]

Release Information

Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure the parameters the router sends to the external AAA server. A group configuration takes precedence over a global DHCP relay or DHCP local server configuration.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.


Related Documentation

- *Using External AAA Authentication Services with DHCP*

authentication (DHCP Relay Agent)

Syntax	<pre> authentication { password <i>password-string</i>; username-include { circuit-type; client-id; delimiter <i>delimiter-character</i>; domain-name <i>domain-name-string</i>; interface-name; logical-system-name; mac-address; option-60; option-82 <circuit-id> <remote-id>; relay-agent-interface-id; relay-agent-remote-id; relay-agent-subscriber-id; routing-instance-name; user-prefix <i>user-prefix-string</i>; } }</pre>
Hierarchy Level	<pre> [edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i>], [edit forwarding-options dhcp-relay group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</pre>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p>
Description	<p>Configure the parameters the router sends to the external AAA server. A group configuration takes precedence over a global DHCP relay configuration. Use the statement at the [edit...dhcpv6] hierarchy levels to configure DHCPv6 support.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • dhcp-relay on page 126 • <i>Using External AAA Authentication Services with DHCP</i>

authentication-order

Syntax	<code>authentication-order [<i>authentication-methods</i>];</code>
Hierarchy Level	<code>[edit access <i>profile</i> <i>profile-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. none option added in Junos OS Release 11.2.
Description	Set the order in which the Junos OS tries different authentication methods when verifying that a client can access the router or switch. For each login attempt, the software tries the authentication methods in order, from first to last.
Default	<code>password</code>
Options	<i>authentication-methods</i> <ul style="list-style-type: none">• none—Grants authentication without examining the client credentials. Can be used, for example, when the Diameter function Gx-Plus is employed for notification during subscriber provisioning.• password—Verify the client using the information configured at the <code>[edit access profile <i>profile-name</i> client <i>client-name</i>]</code> hierarchy level.• radius—Verify the client using RADIUS authentication services.
<div> NOTE: For subscriber access management, you must always specify the radius method. Subscriber access management does not support the password option (the default), and authentication fails when no method is specified.</div>	
Required Privilege Level	<code>admin</code> —To view this statement in the configuration. <code>admin-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring CHAP Authentication with RADIUS</i>• <i>Specifying the Authentication and Accounting Methods for Subscriber Access</i>• <i>Configuring Access Profiles for L2TP or PPP Parameters</i>

authentication-server



Syntax	authentication-server [<i>ip-address</i>];
Hierarchy Level	[edit access profile <i>profile-name</i> radius]
Release Information	Statement introduced in Junos OS Release 9.1.
Description	Specify a list of the RADIUS authentication servers used to authenticate DHCP, L2TP, and PPP clients. The servers in the list are also used as RADIUS dynamic-request servers, from which the router accepts and processes RADIUS disconnect requests, CoA requests, and dynamic service activations and deactivations.
Options	<i>ip-address</i> —IPv4 address.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring RADIUS Server Parameters for Subscriber Access</i>

auto-configure

```
Syntax auto-configure {
    vlan-ranges {
        access-profile profile-name;
        authentication {
            packet-types [packet-types];
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                mac-address;
                option-18;
                option-37;
                option-82 <circuit-id> <remote-id>;
                radius-realm radius-realm-string;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name {
            accept (any | dhcp-v4 | dhcp-v6 | inet | inet6 | pppoe);
            ranges (any | low-tag)–(any | high-tag);
        }
        override;
    }
    stacked-vlan-ranges {
        access-profile profile-name;
        authentication {
            packet-types [packet-types];
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                mac-address;
                option-18;
                option-37;
                option-82 <circuit-id> <remote-id>;
                radius-realm radius-realm-string;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name {
            accept (any | dhcp-v4 | dhcp-v6 | inet | inet6 | pppoe);
            ranges (any | low-tag-high-tag), (any | low-tag-high-tag);
        }
        override;
    }
    remove-when-no-subscribers;
}
```

Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Enable the configuration of dynamic, auto-sensed VLANs. The remaining statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs</i>• <i>Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs</i>

connectivity-type

Syntax	connectivity-type (ce irb permanent);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls], [edit routing-instances <i>routing-instance-name</i> protocols vpls]
Release Information	Statement introduced in Junos OS Release 9.1. irb option introduced in Junos OS Release 9.3. permanent option introduced in Junos OS Release 10.4.
Description	Specify when a VPLS connection is taken down depending on whether or not the interface for the VPLS routing instance is customer-facing or integrated routing and bridging (IRB).
<div>  <p>NOTE: The connectivity-type statement is not supported for FEC 129 VPLS (also known as LDP VPLS with BGP-based autodiscovery).</p> </div>	
Default	ce
Options	<p>ce—Require that for the VPLS connection to be up, the customer-facing interface for the VPLS routing instance must also be up. If the customer-facing interface fails, the VPLS connection is taken down.</p> <p>irb—Allow a VPLS connection to remain up so long as an IRB interface is configured for the VPLS routing instance.</p> <p>permanent—Allow a VPLS connection to remain up until specifically taken down. This option is reserved for use in configuring Layer 2 Wholesale subscriber networks. See the <i>Broadband Subscriber Management Solutions Guide</i> for details about configuring a Layer 2 Wholesale network.</p>
<div>  <p>NOTE: To specifically take down a VPLS routing instance that is using the permanent option, all associated static logical interfaces must also be down.</p> </div>	
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring VPLS Routing Instances Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers on page 84

core-facing

Syntax	core-facing;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced in Junos OS Release 11.2.
Description	Specifies that the VLAN is physically connected to a core-facing ISP router and ensures that the network does not improperly treat the interface as a client interface. When specified, the interface is inserted into the core-facing default mesh group where traffic from pseudowires that belong to the default mesh group is not forwarded on the core-facing link.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Broadband Subscriber Management and Services Library</i>

demux0 (Dynamic Interface)

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

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

Release Information Statement introduced in Junos OS Release 9.3.

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

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

The remaining statements are explained separately.

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

- Related Documentation**
- *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*
 - *Demultiplexing Interface Overview*

demux-options (Dynamic Interface)

Syntax	demux-options { underlying-interface interface-name }
Hierarchy Level	[edit dynamic-profiles profile-name interfaces demux0 interface-name unit logical-unit-number]
Release Information	Statement introduced in Junos OS Release 9.3.
Description	Configure logical demultiplexing (demux) interface options in a dynamic profile. The remaining statement is explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles</i> • <i>Demultiplexing Interface Overview</i>

demux-source (Dynamic IP Demux Interface)

Syntax	<code>demux-source { source-address; }</code>
Hierarchy Level	[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family]
Release Information	Statement introduced in Junos OS Release 9.3.
Description	Configure a logical demultiplexing (demux) source address for a subscriber in a dynamic profile.
Options	source-address —Either the specific source address you want to assign to the subscriber interface or the source address variable. For IPv4, specify \$junos-subscriber-ip-address ; for IPv6, specify \$junos-subscriber-ipv6-address . The source address for the interface is dynamically supplied by DHCP when the subscriber accesses the router.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles</i>• <i>Demultiplexing Interface Overview</i>

demux-source (Dynamic Underlying Interface)

Syntax	<code>demux-source <i>family</i>;</code>
Hierarchy Level	[edit <code>dynamic-profiles interfaces interface-name unit logical-unit-number</code>]
Release Information	Statement introduced in Junos OS Release 9.6.
Description	Configure the logical demultiplexing (demux) source family type on the IP demux underlying interface within a dynamic profile.



NOTE: The IP demux interface feature currently supports only Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet underlying interfaces.

Options	<p><i>family</i>—Protocol family:</p> <ul style="list-style-type: none"> • inet—Internet Protocol version 4 suite • inet6—Internet Protocol version 6 suite
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

demux-source (Underlying Interface)

Syntax	<code>demux-source <i>family</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 9.0. Support for aggregated Ethernet added in Junos OS Release 9.4.
Description	Configure the logical demultiplexing (demux) source family type on the IP demux underlying interface.



NOTE: The IP demux interface feature currently supports only Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet underlying interfaces.

Options	<i>family</i> —Protocol family: <ul style="list-style-type: none">• inet—Internet Protocol version 4 suite• inet6—Internet Protocol version 6 suite
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring an IP Demultiplexing Interface</i>• <i>Configuring a VLAN Demultiplexing Interface</i>

dhcp-attributes (Address-Assignment Pools)

Syntax	<pre> dhcp-attributes { boot-file <i>filename</i>; boot-server (<i>address</i> <i>hostname</i>); dns-server [<i>ipv6-address</i>]; domain-name <i>domain-name</i>; grace-period <i>seconds</i>; maximum-lease-time <i>seconds</i>; name-server [<i>server-list</i>]; netbios-node-type <i>node-type</i>; option { [(<i>id-number</i> <i>option-type</i> <i>option-value</i>) (<i>id-number</i> <i>array</i> <i>option-type</i> <i>option-value</i>)]; } option-match { option-82 { circuit-id <i>value</i> <i>range</i> <i>named-range</i>; remote-id <i>value</i> <i>range</i> <i>named-range</i>; } } preferred-lifetime <i>seconds</i>; router [<i>router-address</i>]; server-identifier <i>ip4-address</i>; sip-server-address [<i>ipv6-address</i>]; sip-server-domain-name <i>domain-name</i>; t1-percentage <i>percentage</i>; t2-percentage <i>percentage</i>; tftp-server <i>address</i>; valid-lifetime <i>seconds</i>; wins-server [<i>servers</i>]; } </pre>
Hierarchy Level	[edit access address-assignment pool <i>pool-name</i> family <i>family</i>]
Release Information	<p>Statement introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 12.3 for EX Series switches.</p>
Description	<p>Configure address pools that can be used by different client applications.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Address-Assignment Pools Overview • DHCP Attributes for Address-Assignment Pools • Configuring Address-Assignment Pools • Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address

dhcp-local-server

```
Syntax  dhcp-local-server {
        access-profile profile-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                logical-system-name;
                mac-address;
                option-60;
                option-82 <circuit-id> <remote-id>;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dhcpv6 {
            access-profile profile-name;
            authentication {
                ...
            }
            group group-name {
                access-profile profile-name;
                authentication {
                    ...
                }
            }
            interface interface-name {
                access-profile profile-name;
                exclude;
                liveness-detection {
                    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                    method {
                        bfd {
                            version (0 | 1 | automatic);
                            minimum-interval milliseconds;
                            minimum-receive-interval milliseconds;
                            multiplier number;
                            no-adaptation;
                            transmit-interval {
                                minimum-interval milliseconds;
                                threshold milliseconds;
                            }
                        }
                        detection-time {
                            threshold milliseconds;
                        }
                        session-mode (automatic | multihop | singlehop);
                        holddown-interval milliseconds;
                    }
                }
            }
        }
        overrides {
```

```

        interface-client-limit number;
        multi-address-embedded-option-response;
        process-inform {
            pool pool-name;
        }
        rapid-commit;
    }
    service-profile dynamic-profile-name;
    trace;
    upto upto-interface-name;
}
overrides {
    delegated-pool;
    interface-client-limit number;
    multi-address-embedded-option-response;
    process-inform {
        pool pool-name;
    }
    rapid-commit;
}
route-suppression;
service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
overrides {
    delegated-pool;
    include-option-82 {
        forcerenew;
        nak;
    }
    interface-client-limit number;
    multi-address-embedded-option-response;
    process-inform {
        pool pool-name;
    }
    rapid-commit;
}

```

```
}
reconfigure {
  attempts attempt-count;
  clear-on-abort;
  strict;
  timeout timeout-value;
  token token-value;
  trigger {
    radius-disconnect;
  }
}
route-suppression;
service-profile dynamic-profile-name;
}
duplicate-clients-in-subnet (incoming-interface | option-82);
dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
primary-profile-name>;
forward-snooped-clients (all-interfaces | configured-interfaces |
  non-configured-interfaces);
group group-name {
  authentication {
    ...
  }
  dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
primary-profile-name>;
  interface interface-name {
    exclude;
    liveness-detection {
      failure-action (clear-binding | clear-binding-if-interface-up | log-only);
      method {
        bfd {
          version (0 | 1 | automatic);
          minimum-interval milliseconds;
          minimum-receive-interval milliseconds;
          multiplier number;
          no-adaptation;
          transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
          }
          detection-time {
            threshold milliseconds;
          }
          session-mode (automatic | multihop | singlehop);
          holddown-interval milliseconds;
        }
      }
    }
  }
}
overrides {
  client-discover-match (option60-and-option82 | incoming-interface);
  include-option-82 {
    forcerenew;
    nak;
  }
  interface-client-limit number;
  process-inform {
```

```

        pool pool-name;
    }
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
overrides {
    client-discover-match (option60-and-option82 | incoming-interface);
    include-option-82 {
        forcerenew;
        nak;
    }
    interface-client-limit number;
    process-inform {
        pool pool-name;
    }
}
requested-ip-network-match subnet-mask
route-suppression;
service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
overrides {
    client-discover-match <option60-and-option82 | incoming-interface>;
    interface-client-limit number;
    process-inform {
        pool pool-name;
    }
}
pool-match-order {
    external-authority;
    ip-address-first;
    option-82;
}
reconfigure {

```

```

    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
        radius-disconnect;
    }
}
requested-ip-network-match subnet-mask;
route-suppression;
service-profile dynamic-profile-name;
}

```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services],
 [edit logical-systems *logical-system-name* system services],
 [edit routing-instances *routing-instance-name* system services],
 [edit system services]

Release Information Statement introduced in Junos OS Release 9.0.
 Statement introduced in Junos OS Release 12.1 for EX Series switches.
 Statement introduced in Junos OS Release 13.2X51 for the QFX Series.
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure Dynamic Host Configuration Protocol (DHCP) local server options on the router or switch and enable the router or switch to function as an extended DHCP local server. The DHCP local server receives DHCP request and reply packets from DHCP clients and then responds with an IP address and other optional configuration information to the client.

The extended DHCP local server is incompatible with the DHCP server on J Series routers and so is not supported on J Series routers. Also, the DHCP local server and the DHCP/BOOTP relay server, which are configured under the **[edit forwarding-options helpers]** hierarchy level, cannot both be enabled on the router or switch at the same time. The extended DHCP local server is fully compatible with the extended DHCP relay feature.

The **dhcipv6** stanza configures the router or switch to support Dynamic Host Configuration Protocol for IPv6 (DHCPv6). The DHCPv6 local server is fully compatible with the extended DHCP local server and the extended DHCP relay feature.



NOTE: When you configure the **dhcp-local-server** statement at the routing instance hierarchy level, you must use a routing instance type of **virtual-router**.

The remaining statements are explained separately.

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

- Related Documentation**
- *Extended DHCP Local Server Overview*
 - *DHCPv6 Local Server Overview*

dhcp-relay

```
Syntax  dhcp-relay {
        access-profile profile-name;
        active-server-group server-group-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                logical-system-name;
                mac-address;
                option-60;
                option-82 <circuit-id> <remote-id>;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
    }
    dhcpv6 {
        access-profile profile-name;
        active-server-group server-group-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                logical-system-name;
                relay-agent-interface-id;
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        forward-only {
            logical-system <current | default | logical-system-name>;
            routing-instance <current | default | routing-instance-name>;
        }
        forward-only-replies;
        group group-name {
            access-profile profile-name;
            active-server-group server-group-name;
            authentication {
                ...
            }
        }
    }
}
```

```

dynamic-profile profile-name {
  ...
}
forward-only {
  ...
}
forward-only {
  ...
}
interface interface-name {
  access-profile profile-name;
  exclude;
  liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
      bfd {
        version (0 | 1 | automatic);
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        detection-time {
          threshold milliseconds;
        }
        session-mode (automatic | multihop | singlehop);
        holddown-interval milliseconds;
      }
    }
  }
  overrides {
    ...
  }
  relay-option {
    ...
  }
  service-profile dynamic-profile-name;
  trace;
  upto upto-interface-name;
}
route-suppression:
service-profile dynamic-profile-name;
overrides {
  ...
}
relay-agent-interface-id {
  ...
}
relay-agent-remote-id {
  ...
}
relay-option {
  ...
}

```

```
    }
    route-suppression;
    server-response-time seconds;
    service-profile dynamic-profile-name;
  }
  liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
      bfd {
        version (0 | 1 | automatic);
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        detection-time {
          threshold milliseconds;
        }
        session-mode (automatic | multihop | singlehop);
        holddown-interval milliseconds;
      }
    }
  }
  overrides {
    allow-snooped-clients;
    delay-authentication;
    delete-binding-on-renegotiation;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    send-release-on-delete;
  }
  relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
  }
  relay-agent-remote-id {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
  }
  server-group {
    server-group-name {
      server-ip-address;
    }
  }
  duplicate-clients-in-subnet (incoming-interface | option-82):
  dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
  }
  forward-only {
```

```

    logical-system <current | default | logical-system-name>;
    routing-instance <current | default | routing-instance-name>;
}
forward-only-replies;
forward-snooped-clients (all-interfaces | configured-interfaces |
    non-configured-interfaces);
group group-name {
    access-profile profile-name;
    active-server-group server-group-name;
    authentication {
        ...
    }
    dynamic-profile profile-name {
        ...
    }
}
forward-only {
    logical-system <current | default | logical-system-name>;
    routing-instance <current | default | routing-instance-name>;
}
forward-only {
    logical-system <current | default | logical-system-name>;
    routing-instance <current | default | routing-instance-name>;
}
interface interface-name {
    access-profile profile-name;
    exclude;
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                detection-time {
                    threshold milliseconds;
                }
                session-mode (automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
        }
    }
    overrides {
        ...
    }
    service-profile dynamic-profile-name;
    trace;
    upto upto-interface-name;
}
overrides {
    ...
}

```

```
    }
    relay-option {
        ...
    }
    relay-option-82 {
        ...
    }
    route-suppression:
    service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
overrides {
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    client-discover-match (option60-and-option82 | incoming-interface);
    delay-authentication;
    delete-binding-on-renegotiation;
    disable-relay;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-bind-on-request;
    proxy-mode;
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        relay-server-group group-name;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
```

```

        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
    starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
}
}
relay-option-82 {
    circuit-id {
        prefix prefix;
        use-interface-description (logical | device);
    }
    remote-id {
        prefix prefix;
        use-interface-description (logical | device);
    }
}
}
}
server-group {
    server-group-name {
        server-ip-address;
    }
}
route-suppression:
server-response-time seconds;
service-profile dynamic-profile-name;
}

```

Hierarchy Level [edit forwarding-options],
 [edit logical-systems *logical-system-name* forwarding-options],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options],
 [edit routing-instances *routing-instance-name* forwarding-options]

Release Information Statement introduced in Junos OS Release 8.3.
 Statement introduced in Junos OS Release 12.1 for EX Series switches.
 Statement introduced in Junos OS Release 13.2X51 for the QFX Series.
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the router or switch and enable the router (or switch) 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.

DHCP relay supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication or client authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.

The extended DHCP and DHCPv6 relay agent options configured with the **dhcp-relay** and **dhcpv6** statements are incompatible with the DHCP/BOOTP relay agent options configured with the **bootp** statement. As a result, the extended DHCP or DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot both be enabled on the router (or switch) at the same time.

The remaining statements are explained separately.

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

Related Documentation

- *Extended DHCP Relay Agent Overview*
- *DHCPv6 Relay Agent Overview*
- *DHCP Relay Proxy Overview*
- *Using External AAA Authentication Services with DHCP*

dhcpx6 (DHCP Local Server)

```
Syntax  dchpx6 {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name-string;
            logical-system-name;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix-string;
        }
    }
    group group-name {
        access-profile profile-name;
        authentication {
            ...
        }
        interface interface-name {
            access-profile profile-name;
            exclude;
            liveness-detection {
                failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                method {
                    bfd {
                        version (0 | 1 | automatic);
                        minimum-interval milliseconds;
                        minimum-receive-interval milliseconds;
                        multiplier number;
                        no-adaptation;
                        transmit-interval {
                            minimum-interval milliseconds;
                            threshold milliseconds;
                        }
                        detection-time {
                            threshold milliseconds;
                        }
                    }
                    session-mode (automatic | multihop | singlehop);
                    holddown-interval milliseconds;
                }
            }
        }
    }
    overrides {
        delete-binding-on-renegotiation;
        interface-client-limit number;
        multi-address-embedded-option-response;
        process-inform {
            pool pool-name;
        }
    }
}
```

```
        rapid-commit;
    }
    service-profile dynamic-profile-name;
    trace;
    upto upto-interface-name;
}
overrides {
    delegated-pool;
    delete-binding-on-renegotiation;
    interface-client-limit number;
    multi-address-embedded-option-response;
    process-inform {
        pool pool-name;
    }
    rapid-commit;
}
route-suppression;
service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
        }
        session-mode (automatic | multihop | singlehop);
        holddown-interval milliseconds;
    }
}
}
overrides {
    delegated-pool;
    delete-binding-on-renegotiation;
    delete-binding-on-renegotiation;
    interface-client-limit number;
    multi-address-embedded-option-response;
    process-inform {
        pool pool-name;
    }
    rapid-commit;
    reconfigure {
        attempts attempt-count;
        clear-on-abort;
        strict;
        timeout timeout-value;
        token token-value;
```

```

        trigger {
            radius-disconnect;
        }
    }
}
reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
        radius-disconnect;
    }
}
requested-ip-network-match subnet-mask;
route-suppression;
service-profile dynamic-profile-name;
}

```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services [dhcp-local-server](#)],
 [edit logical-systems *logical-system-name* system services [dhcp-local-server](#)],
 [edit routing-instances *routing-instance-name* system services [dhcp-local-server](#)],
 [edit system services [dhcp-local-server](#)]

Release Information Statement introduced in Junos OS Release 9.6.
 Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description Configure DHCPv6 local server options on the router or switch and enable the router or switch to function as a server for the DHCP protocol for IP version 6 (IPv6). The DHCPv6 local server sends and receives packets using the IPv6 protocol and informs IPv6 of the routing requirements of router clients. The local server works together with the AAA service framework to control subscriber access (or DHCP client access) and accounting.

The DHCPv6 local server is fully compatible with the extended DHCP local server and DHCP relay agent.

The remaining statements are explained separately.

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

Related Documentation • [DHCPv6 Local Server Overview](#)

domain-name (Address-Assignment Pools)

Syntax	<code>domain-name <i>domain-name</i>;</code>
Hierarchy Level	[edit access address-assignment pool <i>pool-name</i> family inet dhcp-attributes]
Release Information	Statement introduced in Junos OS Release 9.0.
Description	Configure the name of the domain in which clients search for a DHCP server host. This is the default domain name that is appended to hostnames that are not fully qualified. This is equivalent to DHCP option 15.
Options	<i>domain-name</i> —Name of the domain.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Address-Assignment Pools</i>


dynamic-profile (DHCP Local Server)

Syntax	<pre>dynamic-profile <i>profile-name</i> { aggregate-clients (merge replace); use-primary <i>primary-profile-name</i>; }</pre>
Hierarchy Level	<pre>[edit system services dhcp-local-server], [edit system services dhcp-local-server dhcpv6], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i>], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> interface <i>interface-name</i>], [edit system services dhcp-local-server group <i>group-name</i>], [edit system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...]</pre>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Options aggregate-clients and use-primary introduced in Junos OS Release 9.3.</p> <p>Support at the [edit ... interface] hierarchy levels introduced in Junos OS Release 11.2.</p>
Description	Specify the dynamic profile that is attached to all interfaces, a named group of interfaces, or a specific interface.
Options	<p><i>profile-name</i>—Name of the dynamic profile.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces</i> • <i>Configuring a Default Subscriber Service</i>

dynamic-profile (DHCP Relay Agent)

Syntax	<pre>dynamic-profile <i>profile-name</i> { aggregate-clients (merge replace); use-primary <i>primary-profile-name</i>; }</pre>
Hierarchy Level	<pre>[edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> interface <i>interface-name</i>], [edit forwarding-options dhcp-relay group <i>group-name</i>], [edit forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</pre>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Specify the dynamic profile that is attached to all interfaces, to a named group of interfaces, or to a specific interface.</p> <p>M120 and M320 routers do not support DHCPv6.</p>
Options	<p><i>profile-name</i>—Name of the dynamic profile.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• dhcp-relay on page 126• <i>Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces</i>• <i>Grouping Interfaces with Common DHCP Configurations</i>• <i>Configuring a Default Subscriber Service</i>

dynamic-profile (Dynamic PPPoE)

Syntax	<code>dynamic-profile <i>profile-name</i>;</code>
Hierarchy Level	<p>[edit dynamic-profiles <i>profile-name</i> interfaces demux0 unit <i>logical-unit-number</i> family pppoe],</p> <p>[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family pppoe],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family pppoe],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> pppoe-underlying-options],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family pppoe],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> pppoe-underlying-options]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.1.</p> <p>Support for the [edit ... family pppoe] hierarchies introduced in Junos OS Release 11.2.</p>
Description	<p>Attach a PPPoE dynamic profile to an underlying Ethernet interface. This underlying interface is configured with either the encapsulation ppp-over-ether statement or the family pppoe statement; the two statements are mutually exclusive. When the router creates a dynamic PPPoE logical interface on the underlying interface, it uses the information in the dynamic profile to determine the properties of the dynamic PPPoE logical interface.</p>
<div>  NOTE: The [edit ... family pppoe] hierarchies are supported only on MX Series routers with MPCs. </div>	
Options	<p><i>profile-name</i>—Name of a previously configured PPPoE dynamic profile, up to 64 characters in length, defined at the [edit dynamic-profiles <i>profile-name</i> interfaces pp0] hierarchy level.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces Configuring the PPPoE Family for an Underlying Interface Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview

dynamic-profile (Stacked VLAN)

Syntax	<code>dynamic-profile <i>profile-name</i> { accept (any dhcp-v4 dhcp-v6 inet inet6 pppoe); ranges (any <i>low-tag-high-tag</i>), (any <i>low-tag-high-tag</i>); }</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> auto-configure stacked-vlan-ranges]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Configure a dynamic profile for use when configuring dynamic stacked VLANs.
Options	<p><i>profile-name</i>—Name of the dynamic profile that you want to use when configuring dynamic stacked VLANs.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Dynamic Profiles Overview</i>• <i>Configuring a Basic Dynamic Profile</i>• <i>Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs</i>

dynamic-profile (VLAN)

Syntax	dynamic-profile <i>profile-name</i> { accept (any dhcp-v4 dhcp-v6 inet inet6 pppoe); ranges (any low-tag)–(any high-tag); }
Hierarchy Level	[edit interfaces <i>interface-name</i> auto-configure vlan-ranges]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Configure a dynamic profile for use when configuring dynamic VLANs.
Options	<p><i>profile-name</i>—Name of the dynamic profile that you want to use when configuring dynamic VLANs.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Dynamic Profiles Overview</i> • <i>Configuring a Basic Dynamic Profile</i> • <i>Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs</i>

dynamic-profiles

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

```

    term term-name {
        from {
            match-conditions;
        }
        then {
            action;
            action-modifiers;
        }
        only-at-create;
    }
}
filter filter-name {
    enhanced-mode-override;
    fast-lookup-filter;
    instance-shared;
    interface-shared;
    interface-specific;
    term term-name {
        from {
            match-conditions;
        }
        then {
            action;
            action-modifiers;
        }
    }
}
}
hierarchical-policer uid {
    aggregate {
        if-exceeding {
            bandwidth-limit-limit bps;
            burst-size-limit bytes;
        }
        then {
            policer-action;
        }
    }
}
premium {
    if-exceeding {
        bandwidth-limit bps;
        burst-size-limit bytes;
    }
    then {
        policer-action;
    }
}
}
policer uid {
    filter-specific;
    if-exceeding {
        (bandwidth-limit bps | bandwidth-percent percentage);
        burst-size-limit bytes;
    }
    logical-bandwidth-policer;
    logical-interface-policer;
}

```

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

```

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

```

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

```

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

```

```

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

```

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

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

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

Options *profile-name*—Name of the dynamic profile; string of up to 80 alphanumeric characters.

The remaining statements are explained separately.

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

Related Documentation

- *Configuring a Basic Dynamic Profile*
- *Configuring Dynamic VLANs Based on Agent Circuit Identifier Information*
- *Dynamic Profiles Overview*

encapsulation (Dynamic Interfaces)

Syntax	<code>encapsulation (atm-ccc-cell-relay atm-ccc-vc-mux atm-cisco-nlpid atm-tcc-vc-mux atm-mlppp-llc atm-nlpid atm-ppp-llc atm-ppp-vc-mux atm-snap atm-tcc-snap atm-vc-mux ether-over-atm-llc ether-vpls-over-atm-llc ether-vpls-over-fr ether-vpls-over-ppp ethernet frame-relay-ccc frame-relay-ppp frame-relay-tcc frame-relay-ether-type frame-relay-ether-type-tcc multilink-frame-relay-end-to-end multilink-ppp ppp-over-ether ppp-over-ether-over-atm-llc vlan-bridge vlan-ccc vlan-vci-ccc vlan-tcc vlan-vpls);</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Dynamic interface configuration of the logical link-layer encapsulation type.
Options	<p>atm-ccc-cell-relay—Use ATM cell-relay encapsulation.</p> <p>atm-ccc-vc-mux—Use ATM virtual circuit (VC) multiplex encapsulation on circuit cross-connect (CCC) circuits. When you use this encapsulation type, you can configure the ccc family only.</p> <p>atm-cisco-nlpid—Use Cisco ATM network layer protocol ID (NLPID) encapsulation. When you use this encapsulation type, you can configure the inet family only.</p> <p>atm-mlppp-llc—For ATM2 IQ interfaces only, use Multilink Point-to-Point Protocol (MLPPP) over AAL5 LLC. For this encapsulation type, your router must be equipped with a link services or voice services PIC. MLPPP over ATM encapsulation is not supported on ATM2 IQ OC48 interfaces.</p> <p>atm-nlpid—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the inet family only.</p> <p>atm-ppp-llc—For ATM2 IQ interfaces only, use PPP over AAL5 LLC encapsulation.</p> <p>atm-ppp-vc-mux—For ATM2 IQ interfaces only, use PPP over ATM AAL5 multiplex encapsulation.</p> <p>atm-snap—Use ATM subnetwork attachment point (SNAP) encapsulation.</p> <p>atm-tcc-snap—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.</p> <p>atm-tcc-vc-mux—Use ATM VC multiplex encapsulation on TCC circuits. When you use this encapsulation type, you can configure the tcc family only.</p> <p>atm-vc-mux—Use ATM VC multiplex encapsulation. When you use this encapsulation type, you can configure the inet family only.</p> <p>ether-over-atm-llc—For interfaces that carry IPv4 traffic, use Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces.</p>

ether-vpls-over-atm-llc—For ATM2 IQ interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

ether-vpls-over-fr—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Frame Relay encapsulation to support Bridged Ethernet over Frame Relay encapsulated TDM interfaces for VPLS applications, as per *Multiprotocol Interconnect over Frame Relay* (RFC 2427 [1490]).

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

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

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

extended-vlan-vpls—Use extended virtual LAN (VLAN) VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901.



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

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

frame-relay-ppp—Use PPP over Frame Relay circuits. When you use this encapsulation type, you can configure the **ppp** family only.

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

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

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

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

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

ppp-over-ether—You use PPP over Ethernet encapsulation to configure an underlying Ethernet interface for a dynamic PPPoE logical interface.

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

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

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

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

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

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

Related Documentation	<ul style="list-style-type: none">• Configuring a Retail Dynamic Profile for Use in the Layer 2 Wholesale Solution on page 75• <i>Configuring PPP over ATM2 Encapsulation Overview</i>
------------------------------	---

encapsulation (Logical Interface)

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

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

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

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

ether-vpls-over-atm-llc—For ATM2 IQ interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

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



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

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

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

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

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



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

ethernet-vpls-fr—Use in a VPLS setup when a CE device is connected to a PE device over a time-division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

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

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

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

frame-relay-ppp—Use PPP over Frame Relay circuits. When you use this encapsulation type, you can configure the **ppp** family only.

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

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

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

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

ppp-over-ether—You use PPP over Ethernet encapsulation to configure an underlying Ethernet interface for a dynamic PPPoE logical interface.

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

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

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

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

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

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

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

Related Documentation

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

encapsulation (Physical Interface)

Syntax	encapsulation (atm-ccc-cell-relay atm-pvc cisco-hdlc cisco-hdlc-ccc cisco-hdlc-tcc ethernet-bridge ethernet-ccc ethernet-over-atm ethernet-tcc ethernet-vpls ethernet-vpls-fr ether-vpls-over-atm-llc ethernet-vpls-ppp extended-frame-relay-ccc extended-frame-relay-ether-type-tcc extended-frame-relay-tcc extended-vlan-bridge extended-vlan-ccc extended-vlan-tcc extended-vlan-vpls flexible-ethernet-services flexible-frame-relay frame-relay frame-relay-ccc frame-relay-ether-type frame-relay-ether-type-tcc frame-relay-port-ccc frame-relay-tcc generic-services multilink-frame-relay-uni-nni ppp ppp-ccc ppp-tcc vlan-ccc vlan-vci-ccc vlan-vpls);
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces rlsq <i>number:number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers (flexible-ethernet-services , ethernet-ccc , and ethernet-tcc options only).
Description	Specify the physical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.
Default	For SONET interfaces— ppp For ATM interfaces— atm-pvc For DS1, E1, and T1 interfaces— ppp For DS3, E3, and T3 interfaces— ppp For all channelized interfaces— ppp
Options	atm-ccc-cell-relay —Use ATM cell-relay encapsulation. atm-pvc —Use ATM PVC encapsulation. cisco-hdlc —Use Cisco-compatible High-Level Data Link Control (HDLC) framing. cisco-hdlc-ccc —Use Cisco-compatible HDLC framing on CCC circuits. cisco-hdlc-tcc —Use Cisco-compatible HDLC framing on TCC circuits for connecting different media. ethernet-bridge —Use Ethernet bridge encapsulation on Ethernet interfaces that have bridging enabled and that must accept all packets. ethernet-ccc —Use Ethernet CCC encapsulation on Ethernet interfaces that must accept packets carrying standard Tag Protocol ID (TPID) values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, CCC is not supported. ethernet-over-atm —For interfaces that carry IPv4 traffic, use Ethernet over ATM encapsulation. When you use this encapsulation type, you cannot configure multipoint

interfaces. As defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, this encapsulation type allows ATM interfaces to connect to devices that support only bridge protocol data units (BPDUs). Junos OS does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet LLC/SNAP frames with IP or ARP in the payload, and drops the rest. For packets destined to the Ethernet LAN, a route lookup is done using the destination IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and MAC header, and the packet is forwarded to the ATM interface.

ethernet-tcc—For interfaces that carry IPv4 traffic, use Ethernet TCC encapsulation on interfaces that must accept packets carrying standard TPID values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard TPID values. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

ethernet-vpls-fr—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

ethernet-vpls-ppp—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 PPP connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use it to forward the packet into a given VPLS instance.

ether-vpls-over-atm-llc—For ATM intelligent queuing (IQ) interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

extended-frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC.

extended-frame-relay-ether-type-tcc—Use extended Frame Relay ether type TCC for Cisco-compatible Frame Relay for DLCIs 1 through 1022. This encapsulation type is used for circuits with different media on either side of the connection.

extended-frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect different media. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC.

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

extended-vlan-ccc—Use extended VLAN encapsulation on CCC circuits with Gigabit Ethernet and 4-port Fast Ethernet interfaces that must accept packets carrying 802.1Q values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC is not supported.

extended-vlan-tcc—For interfaces that carry IPv4 traffic, use extended VLAN encapsulation on TCC circuits with Gigabit Ethernet interfaces on which you want to use 802.1Q tagging. For 4-port Gigabit Ethernet PICs, extended VLAN TCC is not supported.

extended-vlan-vpls—Use extended VLAN VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



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

flexible-ethernet-services—For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. Aggregated Ethernet bundles can use this encapsulation type. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

flexible-frame-relay—For IQ interfaces only, use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.

frame-relay—Use Frame Relay encapsulation.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with the Cisco Frame Relay.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media.

frame-relay-port-ccc—Use Frame Relay port CCC encapsulation to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect different media.

generic-services—Use generic services encapsulation for services with a hierarchical scheduler.

multilink-frame-relay-uni-nni—Use MLFR UNI NNI encapsulation. This encapsulation is used on link services, voice services interfaces functioning as FRF.16 bundles, and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

ppp—Use serial PPP encapsulation.

ppp-ccc—Use serial PPP encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

ppp-tcc—Use serial PPP encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-ccc—Use Ethernet VLAN encapsulation on CCC circuits.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only. All logical interfaces configured on the Ethernet interface must also have the encapsulation type set to **vlan-vci-ccc**.

vlan-vpls—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



NOTE:

- Label-switched interfaces (LSIs) do not support VLAN VPLS encapsulation. Therefore, you can only use VLAN VPLS encapsulation on a PE-router-to-CE-router interface and not a core-facing interface.
 - Starting with Junos OS release 13.3, a commit error occurs when you configure **vlan-vpls** encapsulation on a physical interface and configure family **inet** on one of the logical units. Previously, it was possible to commit this invalid configuration.
-

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

**Related
Documentation**

- *Configuring Interface Encapsulation on Physical Interfaces*
- *Configuring CCC Encapsulation for Layer 2 VPNs*
- *Configuring Layer 2 Switching Cross-Connects Using CCC*
- *Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits*
- *Configuring ATM Interface Encapsulation*
- *Configuring ATM-to-Ethernet Interworking*
- *Configuring VLAN and Extended VLAN Encapsulation*
- *OBSOLETE: Configuring Extended VLAN Encapsulation*
- [Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces on page 79](#)
- *Configuring Interfaces for Layer 2 Circuits*
- *Configuring Interface Encapsulation on PTX Series Packet Transport Routers*
- *Configuring MPLS LSP Tunnel Cross-Connects Using CCC*
- *Configuring TCC*
- *Configuring VPLS Interface Encapsulation*
- *Configuring Interfaces for VPLS Routing*
- *Defining the Encapsulation for Switching Cross-Connects*

family

Syntax family *family* {
 accounting {
 destination-class-usage;
 source-class-usage {
 (input | output | input output);
 }
 }
 access-concentrator *name*;
 address *address* {
 ... *the address subhierarchy appears after the main* [edit interfaces *interface-name* unit
 logical-unit-number family *family-name*] *hierarchy* ...
 }
 bundle *interface-name*;
 core-facing;
 demux-destination {
 destination-prefix;
 }
 demux-source {
 source-prefix;
 }
 direct-connect;
 duplicate-protection;
 dynamic-profile *profile-name*;
 filter {
 group *filter-group-number*;
 input *filter-name*;
 input-list [*filter-names*];
 output *filter-name*;
 output-list [*filter-names*];
 }
 interface-mode (access | trunk);
 ipsec-sa *sa-name*;
 keep-address-and-control;
 mac-validate (loose | strict);
 max-sessions *number*;
 max-sessions-vsa-ignore;
 mtu *bytes*;
 multicast-only;
 negotiate-address;
 no-redirects;
 policer {
 arp *policer-template-name*;
 input *policer-template-name*;
 output *policer-template-name*;
 }
 primary;
 protocols [inet iso mpls];
 proxy inet-address *address*;
 receive-options-packets;
 receive-ttl-exceeded;
 remote (inet-address *address* | mac-address *address*);
 rpf-check {

```

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

```

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

Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Option max-sessions-vs-a-ignore introduced in Junos OS Release 11.4.
Description	Configure protocol family information for the logical interface.



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

Options *family*—Protocol family:

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



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

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

The remaining statements are explained separately.

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

Related Documentation • *Configuring the Protocol Family*

family (Address-Assignment Pools)

Syntax `family family {
 dhcp-attributes {
 [protocol-specific attributes]
 }
 host hostname {
 hardware-address mac-address;
 ip-address ip-address;
 }
 network ip-prefix / <prefix-length>;
 prefix ipv6-prefix;
 range range-name {
 high upper-limit;
 low lower-limit;
 prefix-length prefix-length;
 }
 }`

Hierarchy Level [edit access address-assignment **pool** *pool-name*]

Release Information Statement introduced in Junos OS Release 9.0.
 Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description Configure the protocol family for the address-assignment pool.



NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options *family*—Protocol family:

- **inet**—Internet Protocol version 4 suite
- **inet6**—Internet Protocol version 6 suite

The remaining statements are explained separately.

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

Related Documentation • *Address-Assignment Pools Overview*
 • *Configuring Address-Assignment Pools*

family (Dynamic Demux Interface)

Syntax `family family {`
 `access-concentrator name;`
 `address address;`
 `demux-source {`
 `source-address;`
 `}`
 `direct-connect;`
 `duplicate-protection;`
 `dynamic-profile profile-name;`
 `filter {`
 `input filter-name;`
 `output filter-name;`
 `}`
 `mac-validate (loose | strict);`
 `max-sessions number;`
 `max-sessions-vsa-ignore;`
 `service-name-table table-name;`
 `short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max`
 `maximum-seconds>;`
 `unnumbered-address interface-name <preferred-source-address address>;`
 `}`

Hierarchy Level [edit `dynamic-profiles profile-name interfaces demux0 unit logical-unit-number`]

Release Information Statement introduced in Junos OS Release 9.3.
`pppoe` option added in Junos OS Release 11.2.

Description Configure protocol family information for the logical interface.



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

Options `family`—Protocol family:

- `inet`—Internet Protocol version 4 suite
- `inet6`—Internet Protocol version 6 suite
- `pppoe`—(MX Series routers with MPCs only) Point-to-Point Protocol over Ethernet

The remaining statements are explained separately.

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

Related Documentation

- *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*
- *Subscriber Interfaces and Demultiplexing Overview*

family (Dynamic PPPoE)

Syntax `family family {
 unnumbered-address interface-name;
 address address;
 service {
 input {
 service-set service-set-name {
 service-filter filter-name;
 }
 post-service-filter filter-name;
 }
 output {
 service-set service-set-name {
 service-filter filter-name;
 }
 }
 }
 filter {
 input filter-name {
 precedence precedence;
 }
 output filter-name {
 precedence precedence;
 }
 }
 }
 }`

Hierarchy Level [edit `dynamic-profiles` *profile-name* `interfaces` pp0 unit "\$junos-interface-unit"]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure protocol family information for the logical interface.

Options *family*—Protocol family:

- **inet**—Internet Protocol version 4 suite
- **inet6**—Internet Protocol version 6 suite

The remaining statements are explained separately.

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

Related Documentation

- *Configuring a PPPoE Dynamic Profile*
- *Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview*

family (Dynamic Standard Interface)

```

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

```

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

Release Information Statement introduced in Junos OS Release 9.2.

pppoe option added in Junos OS Release 11.2.

Description Configure protocol family information for the logical interface.



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

Options *family*—Protocol family:

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

The remaining statements are explained separately.

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

Related Documentation

- *Example: Configuring Static Routing on Logical Systems*
- *Configuring the Protocol Family*

flexible-vlan-tagging

Syntax	flexible-vlan-tagging;
Hierarchy Level	[edit interfaces aex], [edit interfaces ge- <i>fpc/pic/port</i>], [edit interfaces et- <i>fpc/pic/port</i>], [edit interfaces ps0], [edit interfaces xe- <i>fpc/pic/port</i>]
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Support for aggregated Ethernet added in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.</p>
Description	<p>Support simultaneous transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port, and on pseudowire logical interfaces.</p> <p>This statement is supported on M Series and T Series routers, for Fast Ethernet and Gigabit Ethernet interfaces only on Gigabit Ethernet IQ2 and IQ2-E, IQ, and IQE PICs, and for aggregated Ethernet interfaces with member links in IQ2, IQ2-E, and IQ PICs or in MX Series DPCs, or on Ethernet interfaces for PTX Series Packet Transport Routers or 100-Gigabit Ethernet Type 5 PIC with CFP. This statement is supported on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series switches.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Mixed Tagging</i> • <i>Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers</i>

forwarding-options

Syntax	forwarding-options { ... }
Hierarchy Level	[edit] [edit routing-instance <i>routing-instance-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure traffic forwarding. The statements that apply to services interfaces are explained separately. For other statements, see the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i> .



NOTE: The `next-hop-group` statement is present in the `forwarding-options` stanza for a routing instance, but the `next-hop-group` statement is not allowed in a routing instance. In other words, in a routing instance, `[edit routing-instances routing-instance-name forwarding-options next-hop-group]` is not supported. You will get an error message if you try to commit this type of configuration.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>[edit forwarding-options] Hierarchy Level</i>• <i>Configuring Flow Monitoring</i>• <i>Configuring Traffic Sampling</i>

grace-period

Syntax	<code>grace-period <i>seconds</i>;</code>
Hierarchy Level	[edit access address-assignment pool <i>pool-name</i> family (inet inet6) dhcp-attributes]
Release Information	Statement introduced in Junos OS Release 9.0.
Description	Configure the amount of time that the client retains the address lease after the lease expires. The address cannot be reassigned to another client during the grace period.
Options	<i>seconds</i> —Number of seconds the lease is retained. Range: 0 through 4,294,967,295 seconds Default: 0 (no grace period)
Required Privilege Level	admin —To view this statement in the configuration. admin-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Address-Assignment Pools</i>

group (DHCP Local Server)

```
Syntax  group group-name {
        access-profile profile-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                logical-system-name;
                mac-address;
                option-60;
                option-82 <circuit-id> <remote-id>;
                relay-agent-interface-id
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
            primary-profile-name>;
        interface interface-name {
            access-profile profile-name;
            exclude;
            overrides {
                client-discover-match <option60-and-option82>;
                interface-client-limit number;
                process-inform {
                    pool pool-name;
                }
                rapid-commit;
            }
            service-profile dynamic-profile-name;
            trace;
            upto upto-interface-name;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                }
                detection-time {
                    threshold milliseconds;
                }
            }
        }
    }
```

```

    }
    session-mode(automatic | multihop | singlehop);
    holddown-interval milliseconds;
  }
}
}
overrides {
  client-discover-match <option60-and-option82>;
  delegated-pool;
  delete-binding-on-renegotiation;
  interface-client-limit number;
  process-inform {
    pool pool-name;
  }
  rapid-commit;
}
reconfigure {
  attempts attempt-count;
  clear-on-abort;
  strict;
  timeout timeout-value;
  token token-value;
  trigger {
    radius-disconnect;
  }
}
route-suppression;
service-profile dynamic-profile-name;
}

```

Hierarchy Level	[edit system services dhcp-local-server], [edit system services dhcp-local-server dhcpv6], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server ...], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...]
Release Information	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.1 for EX Series switches.
Description	Configure a group of interfaces that have a common configuration, such as authentication parameters. A group must contain at least one interface.
Options	<p><i>group-name</i>—Name of the group.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

- Related Documentation**
- *Extended DHCP Local Server Overview*
 - *Grouping Interfaces with Common DHCP Configurations*
 - *Using External AAA Authentication Services with DHCP*
 - *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*

group (DHCP Relay Agent)

```
Syntax  group group-name {
        access-profile profile-name;
        active-server-group server-group-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                logical-system-name;
                mac-address;
                option-60;
                option-82 [circuit-id] [remote-id];
                relay-agent-interface-id;
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        forward-only {
            logical-system <current | default | logical-system-name>;
            routing-instance <current | default | routing-instance-name>;
        }
        interface interface-name {
            access-profile profile-name;
            exclude;
            liveness-detection {
                failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                method {
                    bfd {
                        version (0 | 1 | automatic);
                        minimum-interval milliseconds;
                        minimum-receive-interval milliseconds;
                        multiplier number;
                        no-adaptation;
                        transmit-interval {
                            minimum-interval milliseconds;
                            threshold milliseconds;
                        }
                        detection-time {
                            threshold milliseconds;
                        }
                    }
                    session-mode (automatic | multihop | singlehop);
                    holddown-interval milliseconds;
                }
            }
        }
    }
```

```
    }
    overrides {
        ...
    }
    service-profile dynamic-profile-name;
    trace;
    upto upto-interface-name;
}
overrides {
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    client-discover-match <option60-and-option82>;
    disable-relay;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-arp;
    no-bind-on-request;
    proxy-mode;
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
}
relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
}
relay-option-60 {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
}
relay-option {
    option-number option-number;
    vendor-option {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
    (equals | starts-with) (ascii match-string | hexadecimal match-hex) {
        (default-relay-server-group server-group-name |
        default-local-server-group local-server-group-name |
        drop);
    }
    (default-relay-server-group server-group-name |
    default-local-server-group local-server-group-name |
    drop);
}
}
relay-option-82 {
    circuit-id {
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82;
```

```

    }
    remote-id {
        prefix prefix;
        use-interface-description (logical | device);
    }
}
route-suppression;
service-profile dynamic-profile-name;


```

Hierarchy Level	<p>[edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3. Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Specify the name of a group of interfaces that have a common DHCP or DHCPv6 relay agent configuration. A group must contain at least one interface. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.</p>
Options	<p><i>group-name</i>—Name of a group of interfaces that have a common DHCP or DHCPv6 relay agent configuration.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • dhcp-relay on page 126 • <i>dhcp-relay (EX Series Switches only)</i> • <i>Extended DHCP Relay Agent Overview</i> • <i>Understanding the Extended DHCP Relay Agent for EX Series Switches</i> • <i>Configuring an Extended DHCP Relay Server on EX Series Switches (CLI Procedure)</i> • <i>Configuring Group-Specific DHCP Relay Options</i> • <i>Grouping Interfaces with Common DHCP Configurations</i> • <i>Using External AAA Authentication Services with DHCP</i> • <i>Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces</i>

input-vlan-map (Dynamic Interfaces)

Syntax	<pre>input-vlan-map { inner-tag-protocol-id <i>tpid</i>; inner-vlan-id <i>number</i>; (push swap); tag-protocol-id <i>tpid</i>; vlan-id <i>number</i>; }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	<p>For dynamic interfaces, define the rewrite profile to be applied to incoming frames on this logical interface.</p> <p>The statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution on page 76

interface (DHCP Local Server)

Syntax	<pre> interface <i>interface-name</i> { access-profile <i>profile-name</i>; exclude; overrides { client-discover-match <option60-and-option82 incoming-interface>; interface-client-limit <i>number</i>; rapid-commit; } service-profile <i>dynamic-profile-name</i>; trace; upto <i>upto-interface-name</i>; } </pre>
Hierarchy Level	<p>[edit system services dhcp-local-server group <i>group-name</i>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server ...],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Options upto and exclude introduced in Junos OS Release 9.1.</p>
Description	<p>Specify one or more interfaces, or a range of interfaces, that are within a specified group on which the DHCP local server is enabled. You can repeat the interface <i>interface-name</i> statement to specify multiple interfaces within a group, but you cannot specify the same interface in more than one group. Also, you cannot use an interface that is being used by the DHCP relay agent.</p>
<div>  <p>NOTE: DHCP values are supported in Integrated Routing and Bridging (IRB) configurations. When you configure an IRB interface in a network that is using DHCP, the DHCP information (for example, authentication, address assignment, and so on) is propagated in the associated bridge domain. This enables the DHCP server to configure client IP addresses residing within the bridge domain. IRB currently only supports static DHCP configurations. For additional information about how to configure IRB, see <i>Configuring Integrated Routing and Bridging for Bridge Domains</i>.</p> </div>	
Options	<p>exclude—Exclude an interface or a range of interfaces from the group. This option and the overrides option are mutually exclusive.</p> <p>interface-name—Name of the interface. You can repeat this option multiple times.</p>

upto-interface-name—Upper end of the range of interfaces; the lower end of the range is the ***interface-name*** entry. The interface device name of the ***upto-interface-name*** must be the same as the device name of the ***interface-name***.

The remaining statements are explained separately.

Required Privilege Level	system—To view this statement in the configuration.
	system-control—To add this statement to the configuration.
Related Documentation	• <i>Extended DHCP Local Server Overview</i>
	• <i>Grouping Interfaces with Common DHCP Configurations</i>
	• <i>Using External AAA Authentication Services with DHCP</i>

interface (DHCP Relay Agent)

Syntax interface *interface-name* {
 `access-profile` *profile-name*;
 exclude;
 overrides {
 allow-snooped-clients;
 always-write-giaddr;
 always-write-option-82;
 client-discover-match <option60-and-option82 | incoming-interface>;
 disable-relay;
 interface-client-limit *number*;
 layer2-unicast-replies;
 no-allow-snooped-clients;
 proxy-mode;
 replace-ip-source-with;
 send-release-on-delete;
 trust-option-82;
 }
 service-profile *dynamic-profile-name*;
 trace;
 upto *upto-interface-name*;
 }

Hierarchy Level [edit forwarding-options dhcp-relay dhcpv6 `group` *group-name*],
 [edit forwarding-options dhcp-relay `group` *group-name*],
 [edit logical-systems *logical-system-name* forwarding-options `dhcp-relay` ...],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
 forwarding-options `dhcp-relay` ...],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay ...]

Release Information Statement introduced in Junos OS Release 8.3.
 Options **upto** and **exclude** introduced in Junos OS Release 9.1.
 Support at the [edit ... **dhcpv6**] hierarchy levels introduced in Junos OS Release 11.4.
 Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Specify one or more interfaces, or a range of interfaces, that are within a specified group on which the DHCP or DHCPv6 relay agent is enabled. You can repeat the **interface** *interface-name* statement to specify multiple interfaces within a group, but you cannot specify the same interface in more than one group. Also, you cannot use an interface that is being used by the DHCP local server. Use the statement at the [edit ... **dhcpv6**] hierarchy levels to configure DHCPv6 support.

EX Series switches do not support DHCPv6.



NOTE: DHCP values are supported in Integrated Routing and Bridging (IRB) configurations. When you configure an IRB interface in a network that is using DHCP, the DHCP information (for example, authentication, address assignment, and so on) is propagated in the associated bridge domain. This enables the DHCP server to configure client IP addresses residing within the bridge domain. IRB currently only supports static DHCP configurations. For

additional information about how to configure IRB, see *Configuring Integrated Routing and Bridging for Bridge Domains*.

.....

Options	<p>exclude—Exclude an interface or a range of interfaces from the group. This option and the overrides option are mutually exclusive.</p> <p>interface-name—Name of the interface. You can repeat this option multiple times.</p> <p>overrides—Override the specified default configuration settings for the interface. The overrides statement is described separately.</p> <p>upto-interface-name—Upper end of the range of interfaces; the lower end of the range is the interface-name entry. The interface device name of the upto-interface-name must be the same as the device name of the interface-name.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Extended DHCP Relay Agent Overview</i>• dhcp-relay on page 126• <i>Grouping Interfaces with Common DHCP Configurations</i>• <i>Using External AAA Authentication Services with DHCP</i>

interface (Dynamic Routing Instances)

Syntax	interface <i>interface-name</i> ;
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced in Junos OS Release 9.6.
Description	Assign the specified interface to the dynamically created routing instance.
Options	<p>interface-name—The interface name variable (<i>\$junos-interface-name</i>). The interface name variable is dynamically replaced with the interface the accessing client uses when connecting to the router.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>[edit routing-instances] Hierarchy Level</i>

interface (Routing Instances)

Syntax	<pre>interface <i>interface-name</i> { description <i>text</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 13.2 for MX 3D Series routers.</p>
Description	Specify the interface over which the VPN traffic travels between the PE device and CE device. You configure the interface on the PE device. If the value vrf is specified for the instance-type statement included in the routing instance configuration, this statement is required.
Options	<p><i>interface-name</i>—Name of the interface.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Routing Instances on PE Routers in VPNs</i> • <i>Configuring EVPN Routing Instances</i> • <i>Configuring EVPN Routing Instances on EX9200 Switches</i> • <i>interface (VPLS Routing Instances)</i>

interface-mac-limit (VPLS)

Syntax	<code>interface-mac-limit <i>limit</i> { packet-action drop; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i> interfaces <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols evpn], [edit routing-instances <i>routing-instance-name</i> protocols evpn interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i> interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Support for EVPNs introduced in Junos OS Release 13.2 on MX 3D Series routers. Support for EVPNs introduced in Junos OS Release 14.2 on EX Series switches.
Description	<p>Specify the maximum number of media access control (MAC) addresses that can be learned by the EVPN or VPLS routing instance. You can configure the same limit for all interfaces configured for a routing instance. You can also configure a limit for a specific interface.</p> <p>Starting with Junos OS Release 12.3R4, if you do not configure the parameter to limit the number of MAC addresses to be learned by a VPLS instance, the default value is not effective. Instead, if you do not include the interface-mac-limit option at the [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i> interfaces <i>interface-name</i>], hierarchy level, this setting is not present in the configuration with the default value of 1024 addresses. If you upgrade a router running a Junos OS release earlier than Release 12.3R4 to Release 12.3R4 or later, you must configure the interface-mac-limit option with a valid value for it to be saved in the configuration.</p>
Options	<p>limit—Number of MAC addresses that can be learned from each interface.</p> <p>Range: 16 through 65,536 MAC addresses</p> <p>Default: 1024 addresses</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring EVPN Routing Instances</i>• <i>Configuring EVPN Routing Instances on EX9200 Switches</i>• <i>Configuring VPLS Routing Instances</i>• <i>interface</i>• <i>mac-table-size</i>

interfaces

Syntax	interfaces { ... }
Hierarchy Level	[edit]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure interfaces on the router.
Default	The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Physical Interface Configuration Statements Overview</i>• <i>Configuring Aggregated Ethernet Link Protection</i>

interfaces (Static and Dynamic Subscribers)

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



```

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

```

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

```

    }
    output filter-name {
        precedence precedence;
        shared-name filter-shared-name;
    }
}
}
}
}
}
}

```

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

Release Information Statement introduced in Junos OS Release 9.2.

Description Define interfaces for dynamic profiles.

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



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


The remaining statements are explained separately.

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

Related Documentation

- *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*
- *Configuring Dynamic PPPoE Subscriber Interfaces*
- *Configuring Dynamic VLANs Based on Agent Circuit Identifier Information*
- *DHCP Subscriber Interface Overview*
- *Configuring Subscribers over Static Interfaces*
- *Demultiplexing Interface Overview*

instance-role

Syntax	<code>instance-role (access nni);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced in Junos OS Release 11.2.
Description	Define the role of the routing instance in a Layer 2 Wholesale network.
Options	<p>access—Defines the connectivity role of the routing instance in a Layer 2 Wholesale network as an access routing instance. When defined for this role, the same process occurs as in a Layer 3 Wholesale network—when the first packet is received from a given client, authentication for the client initiates with an external entity (for example, RADIUS). If authentication is successful, a logical interface is created with the appropriate outer and inner VLAN tags for that client.</p> <p>nni—Defines the connectivity role of the routing instance in a Layer 2 Wholesale network as a network to network interface (NNI) routing instance. When defined for this role, only outer VLAN tags are learned. In addition, when the NNI routing instance receives a response from the ISP, the packets are forwarded to the appropriate client, provided the packet has the same two tags that were verified during authentication.</p>
<div> NOTE: If you connect an access node or MSAN device to a router participating in the Layer 2 Wholesale network in an NNI role, you must create a new routing instance of type <code>l2backhaul-vpn</code> with an instance role of type <code>access</code> for that connection.</div>	
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers on page 81• Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers on page 84• Subscriber Management Overview

instance-type

Syntax	<code>instance-type type;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. virtual-switch and layer2-control options introduced in Junos OS Release 8.4. Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. evpn option introduced in Junos OS Release 13.2 for MX 3D Series routers. forwarding option introduced in Junos OS Release 14.2 for PTX Series Packet Transport Routers.
Description	Define the type of routing instance.

Options



NOTE: On ACX Series routers, you can configure only the forwarding, virtual router, and VRF routing instances.

type—Can be one of the following:

- **evpn**—(MX Series routers only) Enable an Ethernet VPN (EVPN) on the routing instance. You cannot configure the **evpn** option under the `[edit logical-systems logical-system-name routing-instances routing-instance-name instance-type]` hierarchy level.
- **forwarding**—Provide support for filter-based forwarding, where interfaces are not associated with instances. All interfaces belong to the default instance. Other instances are used for populating RPD learned routes. For this instance type, there is no one-to-one mapping between an interface and a routing instance. All interfaces belong to the default instance inet.0.
- **l2backhaul-vpn**—Provide support for Layer 2 wholesale VLAN packets with no existing corresponding logical interface. When using this instance, the router learns both the outer tag and inner tag of the incoming packets, when the **instance-role** statement is defined as **access**, or the outer VLAN tag only, when the **instance-role** statement is defined as **nni**.
- **l2vpn**—Enable a Layer 2 VPN on the routing instance. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.
- **layer2-control**—(MX Series routers only) Provide support for RSTP or MSTP in customer edge interfaces of a VPLS routing instance. This instance type cannot be used if the customer edge interface is multihomed to two provider edge interfaces. If the customer

edge interface is multihomed to two provider edge interfaces, use the default BPDU tunneling.

- **no-forwarding**—This is the default routing instance. Do not create a corresponding forwarding instance. Use this routing instance type when a separation of routing table information is required. There is no corresponding forwarding table. All routes are installed into the default forwarding table. IS-IS instances are strictly nonforwarding instance types.
- **virtual-router**—Enable a virtual router routing instance. This instance type is similar to a VPN routing and forwarding instance type, but used for non-VPN-related applications. You must configure the **interface** statement for this type of routing instance. You do not need to configure the **route-distinguisher**, **vrf-import**, and **vrf-export** statements.
- **virtual-switch**—(MX Series routers and EX9200 switches only) Provide support for Layer 2 bridging. Use this routing instance type to isolate a LAN segment with its Spanning Tree Protocol (STP) instance and to separate its VLAN identifier space.
- **vpls**—Enable VPLS on the routing instance. Use this routing instance type for point-to-multipoint LAN implementations between a set of sites in a VPN. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.
- **vrf**—VPN routing and forwarding (VRF) instance. Provides support for Layer 3 VPNs, where interface routes for each instance go into the corresponding forwarding table only. Required to create a Layer 3 VPN. Create a VRF table (*instance-name.inet.0*) that contains the routes originating from and destined for a particular Layer 3 VPN. For this instance type, there is a one-to-one mapping between an interface and a routing instance. Each VRF instance corresponds with a forwarding table. Routes on an interface go into the corresponding forwarding table. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.

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

- | | |
|------------------------------|--|
| Related Documentation | <ul style="list-style-type: none">• <i>Configuring Routing Instances on PE Routers in VPNs</i>• <i>Configuring EVPN Routing Instances</i>• <i>Configuring EVPN Routing Instances on EX9200 Switches</i>• <i>Configuring Virtual Router Routing Instances</i>• <i>Example: Configuring Filter-Based Forwarding on the Source Address</i>• <i>Example: Configuring Filter-Based Forwarding on Logical Systems</i>• <i>Layer 2 Routing Instance Types</i> |
|------------------------------|--|

ip-address-first

Syntax	ip-address-first;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server pool-match-order], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server pool-match-order], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server pool-match-order], [edit system services dhcp-local-server pool-match-order]
Release Information	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.1 for EX Series switches.
Description	Configure the extended DHCP local server to use the IP address method to determine which address-assignment pool to use. The local server uses the IP address in the gateway IP address if one is present in the DHCP client PDU. If no gateway IP address is present, the local server uses the IP address of the receiving interface to find the address-assignment pool. The DHCP local server uses this method by default when no method is explicitly specified.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use</i> • <i>Extended DHCP Local Server Overview</i> • <i>Address-Assignment Pools Overview</i>


keepalives (Dynamic Profiles)

Syntax	<pre>keepalives { interval <i>seconds</i>; }</pre>
Hierarchy Level	<pre>[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit <i>logical-unit-number</i>] [edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit"] [edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit"]</pre>
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Support at the <code>[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit"]</code> hierarchy level introduced in Junos OS Release 10.1.</p> <p>Support at the <code>[edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit"]</code> hierarchy level introduced in Junos OS Release 12.2.</p>
Description	Specify the keepalive interval in a PPP dynamic profile.
Default	Sending of keepalives is enabled by default.
Options	<p>interval <i>seconds</i>—The time in seconds between successive keepalive requests.</p> <p>Range: 1 through 32767 seconds</p> <p>Default: 30 seconds for LNS-based PPP sessions. 10 seconds for all other PPP sessions.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Dynamic Profiles Overview</i>• <i>Configuring Dynamic Authentication for PPP Subscribers</i>• <i>Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface</i>

mac-validate (Dynamic IP Demux Interface)

Syntax	mac-validate (loose strict);
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces demux0 unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced in Junos OS Release 9.3.
Description	Enable IP and MAC address validation for dynamic IP demux interfaces in a dynamic profile.
Options	<p>loose—Forwards incoming packets when both the IP source address and the MAC source address match one of the trusted address tuples. Drops packets when the IP source address matches one of the trusted tuples, but the MAC address does not match the MAC address of the tuple. Continues to forward incoming packets when the source address of the incoming packet does not match any of the trusted IP addresses.</p> <p>strict—Forwards incoming packets when both the IP source address and the MAC source address match one of the trusted address tuples. Drops packets when the MAC address does not match the tuple's MAC source address, or when IP source address of the incoming packet does not match any of the trusted IP addresses.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring MAC Address Validation for Subscriber Interfaces</i>

multicast-replication

Syntax	multicast-replication { ingress; local-latency-fairness; }
Hierarchy Level	[edit forwarding-options]
Release Information	Statement introduced in Junos OS Release 15.1 for MX Series routers.
Description	Configure the mode of multicast replication that helps to optimize multicast latency.
<div> NOTE: The multicast-replication statement is supported only on platforms with the enhanced-ip mode enabled.</div>	
Default	This statement is disabled by default.
Options	ingress —Complete ingress replication of the multicast data packets where all the egress Packet Forwarding Engines receive packets from the ingress Packet Forwarding Engines directly. local-latency-fairness —Complete parallel replication of the multicast data packets.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• forwarding-options on page 172

no-local-switching

Syntax	no-local-switching
Hierarchy Level	[edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches.
Description	Specify that access ports in this VLAN domain do not forward packets to each other. You use this statement with primary VLANs and isolated secondary VLANs.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

no-tunnel-services

Syntax	no-tunnel-services;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols vpls static-vpls], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls], [edit protocols vpls static-vpls], [edit routing-instances <i>routing-instance-name</i> protocols vpls]
Release Information	Statement introduced in Junos OS Release 7.6. Support for static VPLS added in Junos OS Release 10.2.
Description	Configure VPLS on a router without a Tunnel Services PIC. Configuring the no-tunnel-services statement creates a label-switched interface (LSI) to provide VPLS functionality. An LSI MPLS label is used as the inner label for VPLS. This label maps to a VPLS routing instance. On the PE router, the LSI label is stripped and then mapped to a logical LSI interface. The Layer 2 Ethernet frame is then forwarded using the LSI interface to the correct VPLS routing instance.



NOTE: In VPLS documentation, the word *Router* in terms such as *PR Router* is used to refer to any device that provides routing functions.

Label-switched interfaces configured with the **no-tunnel-services** statement are not supported with GRE tunnels.



NOTE: Although visible in the CLI, the **no-tunnel-services** statement is not supported on DPC cards at the [edit logical-systems *logical-system-name* protocols vpls static-vpls] and the [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls] hierarchy levels.

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

Related Documentation	<ul style="list-style-type: none"> Configuring VPLS Without a Tunnel Services PIC Configuring Static Pseudowires for VPLS Configuring EXP-Based Traffic Classification for VPLS
------------------------------	--

maximum-lease-time

Syntax	<code>maximum-lease-time seconds;</code>
Hierarchy Level	<code>[edit access address-assignment pool <i>pool-name</i> family (inet inet6) dhcp-attributes]</code>
Release Information	Statement introduced in Junos OS Release 9.0.
Description	Specify the maximum length of time, in seconds, that the lease is held for a client if the client does not renew the lease. This is equivalent to DHCP option 51. The maximum-lease-time is mutually exclusive with both the preferred-lifetime and the valid-lifetime , and cannot be configured with either timer.
Options	seconds —Maximum number of seconds the lease can be held. Range: 30 through 4,294,967,295 seconds Default: 86,400 (24 hours)
Required Privilege Level	admin —To view this statement in the configuration. admin-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Address-Assignment Pools</i>• <i>DHCP Attributes for Address-Assignment Pools</i>• <i>preferred-lifetime (Address-Assignment Pools)</i>• <i>valid-lifetime (Address-Assignment Pools)</i>

output-vlan-map (Dynamic Interfaces)

Syntax	<pre>output-vlan-map { inner-tag-protocol-id <i>tpid</i>; inner-vlan-id <i>number</i>; (pop swap); tag-protocol-id <i>tpid</i>; vlan-id <i>number</i>; }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	<p>For dynamic interfaces, define the rewrite profile to be applied to outgoing frames on this logical interface.</p> <p>The statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution on page 76

pap (Dynamic PPP)

Syntax	pap;
Hierarchy Level	<p>[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit" ppp-options],</p> <p>[edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit" ppp-options]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Support at the [edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit" ppp-options] hierarchy level introduced in Junos OS Release 12.2.</p>
Description	Specify PAP authentication in a PPP dynamic profile.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Dynamic Profiles Overview • Configuring Dynamic Authentication for PPP Subscribers • Attaching Dynamic Profiles to Static PPP Subscriber Interfaces • Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface

pool (Address-Assignment Pools)

Syntax `pool pool-name {
 family family {
 dhcp-attributes {
 [protocol-specific attributes]
 }
 host hostname {
 hardware-address mac-address;
 ip-address ip-address;
 }
 network ip-prefix / <prefix-length>;
 prefix ipv6-prefix;
 range range-name {
 high upper-limit;
 low lower-limit;
 prefix-length prefix-length;
 }
 }
 link pool-name;
 }`

Hierarchy Level [edit access [address-assignment](#)]

Release Information Statement introduced in Junos OS Release 9.0.
 Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Configure the name of an address-assignment pool.



NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options *pool-name*—Name assigned to the address-assignment pool.

The remaining statements are explained separately.

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

Related Documentation • *Address-Assignment Pools Overview*
 • *Configuring Address-Assignment Pools*

pool-match-order

Syntax	pool-match-order { external-authority; ip-address-first; option-82; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server], [edit system services dhcp-local-server]
Release Information	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.1.
Description	Configure the order in which the DHCP local server uses information in the DHCP client PDU to determine how to obtain an address for the client. The remaining statements are explained separately.
Default	DHCP local server uses the ip-address-first method to determine which address pool to use.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use</i> • <i>Extended DHCP Local Server Overview</i>

pop (Dynamic VLANs)

Syntax	pop;
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	For dynamic VLAN interfaces, specify the VLAN rewrite operation to remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Removing a VLAN Tag</i> • <i>Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution on page 76</i>

pppoe-options (Dynamic PPPoE)

Syntax	<pre>pppoe-options { underlying-interface <i>interface-name</i>; server; }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit"]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	<p>Configure the underlying interface and PPPoE server mode for a dynamic PPPoE logical interface in a dynamic profile.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring a PPPoE Dynamic Profile</i>• <i>Configuring Dynamic PPPoE Subscriber Interfaces</i>

pppoe-underlying-options (Static and Dynamic Subscribers)

Syntax	<pre>pppoe-underlying-options { access-concentrator <i>name</i>; dynamic-profile <i>profile-name</i>; direct-connect duplicate-protection; max-sessions <i>number</i>; max-sessions-vsa-ignore; service-name-table <i>table-name</i>; short-cycle-protection <lockout-time-min <i>minimum-seconds</i>> <lockout-time-max <i>maximum-seconds</i>> <filter [<i>aci</i>]>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 10.0.
Description	<p>Configure PPPoE-specific interface properties for the underlying interface on which the router creates a static or dynamic PPPoE logical interface. The underlying interface must be configured with PPPoE (ppp-over-ether) encapsulation.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring PPPoE</i> (for static interfaces) • <i>Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces</i> • <i>Assigning a Service Name Table to a PPPoE Underlying Interface</i>

ppp-options (Dynamic PPP)

Syntax	<pre>ppp-options { authentication [<i>authentication-protocols</i>]; mtu (<i>size</i> use-lower-layer); mru <i>size</i>; chap { challenge-length minimum <i>minimum-length</i> maximum <i>maximum-length</i>; } initiate-ncp (ip ipv6 dual-stack-passive) on-demand-ip-address; pap; }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit"], [edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit"]
Release Information	Statement introduced in Junos OS Release 9.5. Support at the [edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit"] hierarchy level introduced in Junos OS Release 12.2.
Description	Configure PPP-specific interface properties in a dynamic profile. The remaining statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Dynamic Profiles Overview</i>• <i>Configuring Dynamic Authentication for PPP Subscribers</i>• <i>Attaching Dynamic Profiles to Static PPP Subscriber Interfaces</i>• <i>Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface</i>

prefix (Address-Assignment Pools)

Syntax	<code>prefix <i>ipv6-prefix</i>;</code>
Hierarchy Level	[edit access address-assignment <code>pool <i>pool-name</i></code> family inet6]
Release Information	Statement introduced in Junos OS Release 10.0. Statement introduced in Junos OS Release 12.3 for EX Series switches.
Description	Specify the IPv6 prefix for the IPv6 address-assignment pool. This statement is mandatory for IPv6 address-assignment pools.
Options	<i>ipv6-prefix</i> —The IPv6 prefix.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Address-Assignment Pools Overview</i>• <i>Configuring Address-Assignment Pools</i>

profile (Access)

Syntax `profile profile-name {`
 `accounting {`
 `address-change-immediate-update`
 `accounting-stop-on-access-deny;`
 `accounting-stop-on-failure;`
 `ancp-speed-change-immediate-update;`
 `coa-immediate-update;`
 `coa-no-override service-class-attribute;`
 `duplication;`
 `duplication-filter;`
 `duplication-vrf {`
 `access-profile-name profile-name;`
 `vrf-name vrf-name;`
 `}`
 `immediate-update;`
 `order [accounting-method];`
 `send-acct-status-on-config-change;`
 `statistics (time | volume-time);`
 `update-interval minutes;`
 `wait-for-acct-on-ack;`
 `}`
 `authentication-order [authentication-methods];`
 `client client-name {`
 `chap-secret chap-secret;`
 `group-profile profile-name;`
 `ike {`
 `allowed-proxy-pair {`
 `remote remote-proxy-address local local-proxy-address;`
 `}`
 `pre-shared-key (ascii-text character-string | hexadecimal hexadecimal-digits);`
 `ike-policy policy-name;`
 `interface-id string-value;`
 `}`
 `l2tp {`
 `aaa-access-profile profile-name;`
 `interface-id interface-id;`
 `lcp-renegotiation;`
 `local-chap;`
 `maximum-sessions-per-tunnel number;`
 `multilink {`
 `drop-timeout milliseconds;`
 `fragment-threshold bytes;`
 `}`
 `override-result-code session-out-of-resource;`
 `ppp-authentication (chap | pap);`
 `ppp-profile profile-name;`
 `shared-secret shared-secret;`
 `}`
 `pap-password pap-password;`
 `ppp {`
 `cell-overhead;`
 `encapsulation-overhead bytes;`

```

    framed-ip-address ip-address;
    framed-pool framed-pool;
    idle-timeout seconds;
    interface-id interface-id;
    keepalive seconds;
    primary-dns primary-dns;
    primary-wins primary-wins;
    secondary-dns secondary-dns;
    secondary-wins secondary-wins;
  }
  user-group-profile profile-name;
}
domain-name-server;
domain-name-server-inet;
domain-name-server-inet6;
preauthentication-order preauthentication-method;
provisioning-order (gx-plus | jsr);
radius {
  accounting-server [ ip-address ];
  attributes {
    exclude {
      ...
    }
    ignore {
      framed-ip-netmask;
      input-filter;
      logical-system:routing-instance;
      output-filter;
    }
  }
}
authentication-server [ ip-address ];
options {
  accounting-session-id-format (decimal | description);
  calling-station-id-delimiter delimiter-character;
  calling-station-id-format {
    agent-circuit-id;
    agent-remote-id;
    interface-description;
    nas-identifier;
  }
  chap-challenge-in-request-authenticator;
  client-accounting-algorithm (direct | round-robin);
  client-authentication-algorithm (direct | round-robin);
  coa-dynamic-variable-validation;
  ethernet-port-type-virtual;
  interface-description-format {
    exclude-adapter;
    exclude-sub-interface;
  }
  juniper-dsl-attributes;
  nas-identifier identifier-value;
  nas-port-extended-format {
    adapter-width width;
    ae-width width;
    port-width width;
    slot-width width;
  }
}

```

```
    stacked-vlan-width width;  
    vlan-width width;  
    atm {  
        adapter-width width;  
        port-width width;  
        slot-width width;  
        vci-width width;  
        vpi-width width;  
    }  
}  
nas-port-id-delimiter delimiter-character;  
nas-port-id-format {  
    agent-circuit-id;  
    agent-remote-id;  
    interface-description;  
    interface-text-description;  
    nas-identifier;  
    order {  
        agent-circuit-id;  
        agent-remote-id;  
        interface-description;  
        interface-text-description;  
        nas-identifier;  
        postpend-vlan-tags;  
    }  
    postpend-vlan-tags;  
}  
nas-port-type {  
    ethernet {  
        port-type;  
    }  
}  
revert-interval interval;  
vlan-nas-port-stacked-format;  
}  
preauthentication-server ip-address;  
}  
radius-server server-address {  
    accounting-port port-number;  
    accounting-retry number;  
    accounting-timeout seconds;  
    dynamic-request-port;  
    port port-number;  
    preauthentication-port port-number;  
    preauthentication-secret password;  
    retry attempts;  
    routing-instance routing-instance-name;  
    secret password;  
    max-outstanding-requests value;  
    source-address source-address;  
    timeout seconds;  
}  
service {  
    accounting-order (activation-protocol | radius);  
}  
session-options {
```

```

        client-idle-timeout minutes;
        client-session-timeout minutes;
    }
}

```

Hierarchy Level	[edit access]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure PPP CHAP, or a profile and its subscriber access, L2TP, or PPP properties.
Options	<p><i>profile-name</i>—Name of the profile.</p> <p>For CHAP, the name serves as the mapping between peer identifiers and CHAP secret keys. This entity is queried for the secret key whenever a CHAP challenge or response is received.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the PPP Authentication Protocol</i> • <i>Configuring Access Profiles for L2TP or PPP Parameters</i> • <i>Configuring L2TP Properties for a Client-Specific Profile</i> • <i>Configuring an L2TP LNS with Inline Service Interfaces</i> • <i>Configuring PPP Properties for a Client-Specific Profile</i> • <i>Configuring Service Accounting with JSRC</i> • <i>AAA Service Framework Overview</i> • show network-access aaa statistics on page 268 • <i>clear network-access aaa statistics</i>

protocols

```
Syntax protocols {
    bgp {
        ... bgp-configuration ...
    }
    isis {
        ... isis-configuration ...
    }
    ldp {
        ... ldp-configuration ...
    }
    mpls {
        ... mpls-configuration ...
    }
    msdp {
        ... msdp-configuration ...
    }
    mstp {
        ... mstp-configuration ...
    }
    ospf {
        domain-id domain-id;
        domain-vpn-tag number;
        route-type-community (iana | vendor);
        traffic-engineering {
            <advertise-unnumbered-interfaces>;
            <credibility-protocol-preference>;
            ignore-lsp-metrics;
            multicast-rpf-routes;
            no-topology;
            shortcuts {
                lsp-metric-into-summary;
            }
        }
        ... ospf-configuration ...
    }
    ospf3 {
        domain-id domain-id;
        domain-vpn-tag number;
        route-type-community (iana | vendor);
        traffic-engineering {
            <advertise-unnumbered-interfaces>;
            <credibility-protocol-preference>;
            ignore-lsp-metrics;
            multicast-rpf-routes;
            no-topology;
            shortcuts {
                lsp-metric-into-summary;
            }
        }
        ... ospf3-configuration ...
    }
    pim {
```



```

    ... pim-configuration ...
  }
  rip {
    ... rip-configuration ...
  }
  ripng {
    ... ripng-configuration ...
  }
  rstp {
    rstp-configuration;
  }
  rsvp {
    ... rsvp-configuration ...
  }
  vstp {
    vstp configuration;
  }
  vpls {
    vpls configuration;
  }
}

```

Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Support for RIPng introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. mpls and rsvp options added in Junos OS Release 15.1. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the protocol for a routing instance. You can configure multiple instances of many protocol types. Not all protocols are supported on the switches. See the switch CLI.

- Options**
- bgp**—Specify BGP as the protocol for a routing instance.
 - isis**—Specify IS-IS as the protocol for a routing instance.
 - ldp**—Specify LDP as the protocol for a routing instance or for a virtual router instance.
 - l2vpn**—Specify Layer 2 VPN as the protocol for a routing instance.
 - mpls**—Specify MPLS as the protocol for a routing instance.
 - msdp**—Specify the Multicast Source Discovery Protocol (MSDP) for a routing instance.
 - mstp**—Specify the Multiple Spanning Tree Protocol (MSTP) for a virtual switch routing instance.
 - ospf**—Specify OSPF as the protocol for a routing instance.
 - ospf3**—Specify OSPF version 3 (OSPFv3) as the protocol for a routing instance.




NOTE: OSPFv3 supports the **no-forwarding**, **virtual-router**, and **vrf** routing instance types only.

- pim**—Specify the Protocol Independent Multicast (PIM) protocol for a routing instance.
- rip**—Specify RIP as the protocol for a routing instance.
- ripng**—Specify RIP next generation (RIPng) as the protocol for a routing instance.
- rstp**—Specify the Rapid Spanning Tree Protocol (RSTP) for a virtual switch routing instance.
- rsvp**—Specify the RSVP for a routing instance.
- vstp**—Specify the VLAN Spanning Tree Protocol (VSTP) for a virtual switch routing instance.
- vpls**—Specify VPLS as the protocol for a routing instance.

- | | |
|---------------------------------|---|
| Required Privilege Level | routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Example: Configuring Multiple Routing Instances of OSPF</i> |

proxy-arp

Syntax	<code>proxy-arp (restricted unrestricted);</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.6 for EX Series switches. restricted added in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 12.2 for the QFX Series.
Description	For Ethernet interfaces only, configure the router or switch to respond to any ARP request, as long as the router or switch has an active route to the ARP request's target address.
<div>  <p>NOTE: You must configure the IP address and the inet family for the interface when you enable proxy ARP.</p> </div>	
Default	Proxy ARP is not enabled. The router or switch responds to an ARP request only if the destination IP address is its own.
Options	<ul style="list-style-type: none"> • none—The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address. • restricted—(Optional) The router or switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are in the same subnet. The router or switch must also have a route to the target IP address. • unrestricted—(Optional) The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.
	Default: unrestricted
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Restricted and Unrestricted Proxy ARP</i> • <i>Configuring Proxy ARP (CLI Procedure)</i> • <i>Example: Configuring Proxy ARP on an EX Series Switch</i> • <i>Configuring Gratuitous ARP</i>

proxy-arp (Dynamic Profiles)

Syntax	proxy-arp;
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	For Ethernet interfaces only, configure the router to respond to any ARP request, as long as the router has an active route to the target address of the ARP request.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Restricted and Unrestricted Proxy ARP</i>• <i>Configuring Gratuitous ARP</i>

push (Dynamic VLANs)

Syntax	push;
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	For dynamic VLAN interfaces, specify the VLAN rewrite operation to add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag. If you include the push statement in the configuration, you must also include the <i>pop</i> statement at the [edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map] hierarchy level.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution on page 76

radius (Access Profile)

```
Syntax  radius {
    accounting-server [ ip-address ];
    attributes {
        exclude
        ...
    }
    ignore {
        framed-ip-netmask;
        input-filter;
        logical-system-routing-instance;
        output-filter;
    }
}
authentication-server [ ip-address ];
options {
    accounting-session-id-format (decimal | description);
    calling-station-id-delimiter delimiter-character;
    calling-station-id-format {
        agent-circuit-id;
        agent-remote-id;
        interface-description;
        nas-identifier;
    }
    chap-challenge-in-request-authenticator;
    client-accounting-algorithm (direct | round-robin);
    client-authentication-algorithm (direct | round-robin);
    coa-dynamic-variable-validation;
    ethernet-port-type-virtual;
    interface-description-format {
        exclude-adapter;
        exclude-sub-interface;
    }
    ip-address-change-notify message;
    juniper-dsl-attributes;
    nas-identifier identifier-value;
    nas-port-extended-format {
        adapter-width width;
        ae-width width;
        port-width width;
        slot-width width;
        stacked-vlan-width width;
        vlan-width width;
        atm {
            adapter-width width;
            port-width width;
            slot-width width;
            vci-width width;
            vpi-width width;
        }
    }
    nas-port-id-delimiter delimiter-character;
    nas-port-id-format {
```

```
agent-circuit-id;
agent-remote-id;
interface-description;
interface-text-description;
nas-identifier;
order {
    agent-circuit-id;
    agent-remote-id;
    interface-description;
    interface-text-description;
    nas-identifier;
    postpend-vlan-tags;
}
postpend-vlan-tags;
}
nas-port-type {
    ethernet {
        port-type;
    }
}
revert-interval interval;
vlan-nas-port-stacked-format;
}
preauthentication-server ip-address;
}
```

Hierarchy Level [edit access [profile](#) *profile-name*]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description Configure the RADIUS parameters that the router uses for AAA authentication and accounting for subscribers.

The remaining statements are explained separately.

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

Related Documentation

- *Configuring RADIUS Server Parameters for Subscriber Access*
- *RADIUS Server Options for Subscriber Access*

radius-server

Syntax	<pre>radius-server server-address { accounting-port <i>port-number</i>; accounting-retry <i>number</i>; accounting-timeout <i>seconds</i>; dynamic-request-port; port <i>port-number</i>; preauthentication-port <i>port-number</i>; preauthentication-secret <i>password</i>; retry <i>attempts</i>; routing-instance <i>routing-instance-name</i>; secret <i>password</i>; max-outstanding-requests <i>value</i>; source-address <i>source-address</i>; timeout <i>seconds</i>; }</pre>
Hierarchy Level	<p>[edit access],</p> <p>[edit access profile <i>profile-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>dynamic-request-port option added in Junos OS Release 14.2R1 for MX Series routers.</p> <p>preauthentication-port and preauthentication-secret options added in Junos OS Release 14.1X51 for MX Series routers.</p>
Description	<p>Configure RADIUS for subscriber access management, L2TP, or PPP.</p> <p>To configure multiple RADIUS servers, include multiple radius-server statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.</p>
Options	<p>server-address—Address of the RADIUS authentication server.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring RADIUS Authentication for L2TP</i> • <i>Configuring the PPP Authentication Protocol</i> • <i>Configuring RADIUS Server Authentication</i> • <i>Configuring Authentication and Accounting Parameters for Subscriber Access</i> • show network-access aaa statistics on page 268 • <i>clear network-access aaa statistics</i>

range (Address-Assignment Pools)

Syntax	<pre>range range-name { high upper-limit; low lower-limit; prefix-length prefix-length; }</pre>
Hierarchy Level	[edit access address-assignment pool <i>pool-name</i> family (inet inet6)]
Release Information	Statement introduced in Junos OS Release 9.0. IPv6 support introduced in Junos OS Release 10.0. Statement introduced in Junos OS Release 12.3 for EX Series switches.
Description	Configure a named range of IPv4 addresses or IPv6 prefixes, used within an address-assignment pool.
Options	<p>high upper-limit—Upper limit of an address range or IPv6 prefix range.</p> <p>low lower-limit—Lower limit of an address range or IPv6 prefix range.</p> <p>prefix-length prefix-length—Assigned length of the IPv6 prefix.</p> <p>range-name—Name assigned to the range of IPv4 addresses or IPv6 prefixes.</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Address-Assignment Pools Overview</i>• <i>Configuring Address-Assignment Pools</i>

ranges (Dynamic VLAN)

Syntax	<code>ranges (any <i>low-tag</i>)-(any <i>high-tag</i>);</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> auto-configure vlan-ranges dynamic-profile <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Configure VLAN ranges for dynamic, auto-sensed VLANs.
Options	<p>any—The entire VLAN range.</p> <p><i>low-tag</i>—The lower limit of the VLAN range.</p> <p><i>high-tag</i>—The upper limit of the VLAN range.</p> <p>Range: 1 through 4094</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs</i>

route-distinguisher

Syntax	<code>route-distinguisher (as-number:id ip-address:id);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn mesh-group <i>mesh-group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn mesh-group <i>mesh-group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Support at [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>] hierarchy level introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Support at [edit routing-instances <i>routing-instance-name</i> protocols l2vpn mesh-group <i>mesh-group-name</i>] hierarchy level introduced in Junos OS Release 13.2.</p> <p>Statement introduced in Junos OS Release 14.1X53-D30 for QFX Series switches.</p>
Description	<p>Specify an identifier attached to a route, enabling you to distinguish to which VPN or virtual private LAN service (VPLS) the route belongs. Each routing instance must have a unique route distinguisher (RD) associated with it. The RD is used to place bounds around a VPN so that the same IP address prefixes can be used in different VPNs without having them overlap. If the instance type is vrf, the route-distinguisher statement is required.</p> <p>For Layer 2 VPNs and VPLS, if you configure the l2vpn-use-bgp-rules statement, you must configure a unique RD for each PE router participating in the routing instance.</p> <p>For other types of VPNs, we recommend that you use a unique RD for each provider edge (PE) router participating in specific routing instance. Although you can use the same RD on all PE routers for the same VPN routing instance, if you use a unique RD, you can determine the customer edge (CE) router from which a route originated within the VPN.</p> <p>For Layer 2 VPNs and VPLSs, if you configure mesh groups, the RD in each mesh group must be unique.</p>



CAUTION: We strongly recommend that if you change an RD that has already been configured, make the change during a maintenance window, as follows:

1. Deactivate the routing instance.
2. Change the RD.
3. Activate the routing instance.

This is not required if you are configuring the RD for the first time.

Options *as-number:number*—*as-number* is an assigned AS number, and *number* is any 2-byte or 4-byte value. The AS number can be from 1 through 4,294,967,295. If the AS number is a 2-byte value, the administrative number is a 4-byte value. If the AS number is a 4-byte value, the administrative number is a 2-byte value. An RD consisting of a 4-byte AS number and a 2-byte administrative number is defined as a type 2 RD in RFC 4364 *BGP/MPLS IP VPNs*.



NOTE: In Junos OS Release 9.1 and later, the numeric range for AS numbers is extended to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. All releases of Junos OS support 2-byte AS numbers. To configure an RD that includes a 4-byte AS number, append the letter “L” to the end of the AS number. For example, an RD with the 4-byte AS number 7,765,000 and an administrative number of 1,000 is represented as 7765000L:1000.

In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For example, the 4-byte AS number of 65,546 in the plain-number format is represented as 1.10 in AS dot notation format.

ip-address:id—IP address (*ip-address* is a 4-byte value) within your assigned prefix range and a 2-byte value for the *id*. The IP address can be any globally unique unicast address.

Range: 0 through 4,294,967,295 ($2^{32} - 1$). If the router you are configuring is a BGP peer of a router that does not support 4-byte AS numbers, you need to configure a local AS number. For more information, see *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview*.

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

Related Documentation

- *Example: Configuring BGP Route Target Filtering for VPNs*
- *Example: Configuring FEC 129 BGP Autodiscovery for VPWS*
- *Configuring EVPN Routing Instances*
- *Configuring Routing Instances on PE Routers in VPNs*
- *Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)*
- *Configuring an MPLS-Based Layer 3 VPN (CLI Procedure)*
- *l2vpn-use-bgp-rules*

routing-instances (Dynamic Profiles)

```

Syntax  routing-instances routing-instance-name {
        interface interface-name;
        multicast-snooping-options {
        }
        routing-options {
            access {
                route prefix {
                    metric route-cost;
                    next-hop next-hop;
                    preference route-distance;
                    tag route-tag;
                }
            }
            access-internal {
                route subscriber-ip-address {
                    qualified-next-hop underlying-interface {
                        mac-address address;
                    }
                }
            }
        }
        multicast {
            interface interface-name {
                no-qos-adjust;
            }
        }
        rib routing-table-name {
            access {
                route prefix {
                    metric route-cost;
                    next-hop next-hop;
                    preference route-distance;
                    tag route-tag;
                }
            }
            access-internal {
                route subscriber-ip-address {
                    qualified-next-hop underlying-interface {
                        mac-address address;
                    }
                }
            }
        }
    }

```

Hierarchy Level [edit [dynamic-profiles](#)]
 [edit logical-systems *logical-system-name*]

Release Information Statement introduced in Junos OS Release 9.6.
 Support at the **logical-systems** hierarchy level was introduced in Junos OS Release 14.2.

Description Dynamically configure an additional routing entity for a router.

Options *routing-instance-name*—The routing instance variable (*\$junos-routing-instance*). The routing instance variable is dynamically replaced with the routing instance the accessing client uses when connecting to the router.

The remaining statements are explained separately.

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

Related • [Configuring a Dynamic Profile for use by a Retailer in the DHCPv4 Solution on page 19](#)
Documentation

routing-instances (Multiple Routing Entities)

Syntax	<code>routing-instances <i>routing-instance-name</i> { ... }</code>
Hierarchy Level	[edit], [edit logical-systems <i>logical-system-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Configure an additional routing entity for a router. You can create multiple instances of BGP, IS-IS, OSPF, OSPFv3, and RIP for a router. You can also create multiple routing instances for separating routing tables, routing policies, and interfaces for individual wholesale subscribers (retailers) in a Layer 3 wholesale network.</p> <p>Each routing instance consist of the following:</p> <ul style="list-style-type: none">• A set of routing tables• A set of interfaces that belong to these routing tables• A set of routing option configurations <p>Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name my-instance, its corresponding IP unicast table is my-instance.inet.0. All routes for my-instance are installed into my-instance.inet.0.</p> <p>Routes are installed into the default routing instance inet.0 by default, unless a routing instance is specified.</p> <p>In Junos OS Release 9.0 and later, you can no longer specify a routing-instance name of <i>master</i>, <i>default</i>, or <i>bgp</i> or include special characters within the name of a routing instance.</p> <p>In Junos OS Release 9.6 and later, you can include a slash (/) in a routing-instance name only if a logical system is not configured. That is, you cannot include the slash character in a routing-instance name if a logical system other than the default is explicitly configured. Routing-instance names, further, are restricted from having the form <code>__.*__</code> (beginning and ending with underscores). The colon : character cannot be used when multitopology routing (MTR) is enabled.</p>
Default	Routing instances are disabled for the router.
Options	<i>routing-instance-name</i> —Name of the routing instance. This must be a non-reserved string of not more than 128 characters.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Interprovider Layer 3 VPN Option A</i>• <i>Example: Configuring Interprovider Layer 3 VPN Option B</i>

- *Example: Configuring Interprovider Layer 3 VPN Option C*

secret

Syntax	<code>secret password;</code>
Hierarchy Level	[edit access profile <i>profile-name</i> radius-server <i>server-address</i>], [edit access radius-disconnect <i>client-address</i>], [edit access radius-server <i>server-address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure the password to use with the RADIUS server. The secret password used by the local router or switch must match that used by the server.
Options	password —Password to use; it can include spaces if the character string is enclosed in quotation marks.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Authentication and Accounting Parameters for Subscriber Access</i> • <i>Configuring Router or Switch Interaction with RADIUS Servers</i> • <i>Example: Configuring CHAP Authentication with RADIUS</i> • <i>Configuring RADIUS Authentication for L2TP</i> • <i>Configuring the RADIUS Disconnect Server for L2TP</i>

server (Dynamic PPPoE)

Syntax	<code>server;</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit" pppoe-options]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	In a dynamic profile, configure the router to act as a PPPoE server, also known as a remote access concentrator, when a PPPoE logical interface is dynamically created.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring a PPPoE Dynamic Profile</i> • <i>Subscriber Interfaces and PPPoE Overview</i>

server-group

Syntax	<pre>server-group { server-group-name { server-ip-address; } }</pre>
Hierarchy Level	[edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6]
Release Information	Statement introduced in Junos OS Release 8.3. Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.1 for EX Series switches.
Description	Specify the name of a group of DHCP server addresses for use by the extended DHCP relay agent. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.
Options	server-group-name —Name of the group of DHCP or DHCPv6 server addresses. server-ip-address —IP address of the DHCP server belonging to this named server group. Use IPv6 addresses when configuring DHCPv6 support. You can configure a maximum of five IP addresses in each named server group.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• dhcp-relay on page 126• <i>Extended DHCP Relay Agent Overview</i>• <i>Configuring Server Groups</i>

services (System Services)

```
Syntax  services {
    dhcp { \* DHCP not supported on a DCF
        dhcp_services;
    }
    finger {
        connection-limit limit;
        rate-limit limit;
    }
    ftp {
        authentication-order [authentication-methods];
        connection-limit limit;
        rate-limit limit;
    }
    service-deployment {
        servers address {
            port-number port-number;
        }
        source-address address;
    }
    ssh {
        authentication-order [authentication-methods];
        connection-limit limit;
        protocol-version [v1 v2];
        rate-limit limit;
        root-login (allow | deny | deny-password);
    }
    telnet {
        authentication-order [authentication-methods];
        connection-limit limit;
        rate-limit limit;
    }
    web-management {
        http {
            interfaces [ names ];
            port port;
        }
        https {
            interfaces [ names ];
            local-certificate name;
            port port;
        }
        session {
            idle-timeout [ minutes ];
            session-limit [ limit ];
        }
    }
    xnm-clear-text {
        connection-limit limit;
        rate-limit limit;
    }
    xnm-ssl {
        connection-limit limit;
    }
}
```

```
    local-certificate name;  
    rate-limit limit;  
    ssl-renegotiation;  
  }  
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure the router or switch so that users on remote systems can access the local router or switch through the DHCP server, finger, rlogin, SSH, telnet, Web management, Junos XML protocol clear-text, Junos XML protocol SSL, and network utilities or enable Junos OS to work with the Session and Resource Control (SRC) software.

The remaining statements are explained separately.

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

Related Documentation

- *Configuring clear-text or SSL Service for Junos XML Protocol Client Applications*
- *Configuring the Junos OS to Work with SRC Software*

site (VPLS Multihoming for FEC 128)

Syntax	<pre> site <i>site-name</i> { active-interface (any primary <i>interface-name</i>); best-site; interface <i>interface-name</i> { interface-mac-limit <i>limit</i>; } mesh-group <i>mesh-group-name</i>; multi-homing; site-identifier <i>identifier</i>; site-preference <i>preference-value</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the site name and site identifier for a site. Allows you to configure a remote site ID for remote sites.
Options	<p><i>site-name</i>—Name of the site.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring VPLS Routing Instances</i>

site-identifier (VPLS)

Syntax	<code>site-identifier <i>identifier</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the numerical identifier for the local VPLS site.
Options	<i>identifier</i> —Specify the numerical identifier for the local VPLS site. The identifier must be an unsigned 16-bit number greater than zero.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring VPLS Routing Instances</i>

site-range

Syntax	<code>site-range <i>number</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls], [edit routing-instances <i>routing-instance-name</i> protocols vpls]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify an upper limit on the maximum site identifier that can be accepted to allow a pseudowire to be brought up. Pseudowires cannot be established to sites with site identifiers greater than the configured site range. If you issue the show vpls connections command, such sites are displayed as OR (out of range).
Options	<i>number</i> —Maximum number of site identifiers. We recommend using the default value. Range: 1 through 65,534 Default: 65,534
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring VPLS Routing Instances</i>

stacked-vlan-ranges

Syntax

```
stacked-vlan-ranges {
  access-profile profile-name;
  authentication {
    packet-types [packet-types];
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-name;
      mac-address;
      option-18
      option-37
      option-82;
      radius-realm radius-realm-string;
      user-prefix user-prefix-string;
    }
  }
  dynamic-profile profile-name {
    accept (any | dhcp-v4 | inet);
    ranges (any | low-tag-high-tag), (any | low-tag-high-tag);
  }
  override;
}
```

Hierarchy Level [edit interfaces *interface-name* [auto-configure](#)]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure multiple VLANs. Each VLAN is assigned a VLAN ID number from the range.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs*
- *Configuring Interfaces to Support Both Single and Stacked VLANs*

stacked-vlan-tagging

Syntax	stacked-vlan-tagging;
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Description	<p>For Gigabit Ethernet IQ interfaces, Gigabit Ethernet, 10-Gigabit Ethernet LAN/WAN PIC, and 100-Gigabit Ethernet Type 5 PIC with CFP, enable stacked VLAN tagging for all logical interfaces on the physical interface.</p> <p>For pseudowire subscriber interfaces, enable stacked VLAN tagging for logical interfaces on the pseudowire service.</p>



NOTE: This statement is not supported on PTX1000 routers.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview</i>• vlan-tags (Stacked VLAN Tags) on page 259

system

Syntax	system { ... }
Hierarchy Level	[edit]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure system management properties.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>System Management Configuration Statements</i>

traceoptions (DHCP)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i> > <size <i>maximum-file-size</i>> <world-readable no-world-readable>; flag <i>flag</i>; level (all error info notice verbose warning); no-remote-trace; } </pre>
Hierarchy Level	[edit system processes dhcp-service]
Release Information	<p>Statement introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Define global tracing operations for extended DHCP local server and extended DHCP relay agent processes.</p> <p>This statement replaces the deprecated traceoptions statements at the [edit forwarding-options dhcp-relay] and [edit system services dhcp-local-server] hierarchy levels.</p>
Options	<p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000</p> <p>Default: 3 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements:</p> <ul style="list-style-type: none"> • all—Trace all events. • auth—Trace authentication events. • database—Trace database events. • fwd—Trace firewall process events. • general—Trace miscellaneous events. • ha—Trace high availability-related events. • interface—Trace interface operations. • io—Trace I/O operations. • liveness-detection—Trace liveness detection operations. • packet—Trace packet and option decoding operations. • performance—Trace performance measurement operations.

- **profile**—Trace profile operations.
- **rpd**—Trace routing protocol process events.
- **rtsock**—Trace routing socket operations.
- **security-persistence**—Trace security persistence events.
- **session-db**—Trace session database events.
- **state**—Trace changes in state.
- **statistics**—Trace baseline statistics.
- **ui**—Trace user interface operations.

level—Level of tracing to perform; also known as severity level. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:

- **all**—Match messages of all levels.
- **error**—Match error messages.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure **verbose**, messages at all higher levels are traced. Therefore, the result is the same as when you configure **all**.
- **warning**—Match warning messages.

Default: error

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access, allowing only the user **root** and users who have the Junos OS **maintenance** permission to access the trace files.

size *maximum-file-size*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (***maximum-file-sizek***), megabytes (***maximum-file-sizem***), or gigabytes (***maximum-file-sizeg***). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Range: 10,240 through 1,073,741,824


Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

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

- Related Documentation**
- *Tracing Extended DHCP Operations*

underlying-interface (demux0)

Syntax	<code>underlying-interface <i>underlying-interface-name</i>;</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces demux0 <i>interface-name</i> unit <i>unit logical-unit-number</i> demux-options]
Release Information	Statement introduced in Junos OS Release 9.3. Support for aggregated Ethernet introduced in Junos OS Release 9.4.
Description	Configure the underlying interface on which the demultiplexing (demux) interface is running.
Options	<p><i>underlying-interface-name</i>—Either the specific name of the interface on which the DHCP discover packet arrives or one of the following interface variables:</p> <ul style="list-style-type: none"> • \$junos-underlying-interface when configuring dynamic IP demux interfaces. • \$junos-interface-ifd-name when configuring dynamic VLAN demux interfaces. <p>The variable is used to specify the underlying interface when a new demux interface is dynamically created. The variable is dynamically replaced with the underlying interface that DHCP supplies when the subscriber logs in.</p>
<div>  <p>NOTE: Logical demux interfaces are currently supported on Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet interfaces.</p> </div>	
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles</i> • <i>Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles</i> • <i>Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview</i>

underlying-interface (Dynamic PPPoE)

Syntax	<code>underlying-interface <i>interface-name</i>;</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit" ppoe-options]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	In a dynamic profile, configure the underlying interface on which the router creates the dynamic PPPoE logical interface.
Options	<i>interface-name</i> —Variable used to specify the name of the underlying interface on which the PPPoE logical interface is dynamically created. In the underlying-interface <i>interface-name</i> statement for dynamic PPPoE logical interfaces, you must use the predefined variable \$junos-underlying-interface in place of <i>interface-name</i> . When the router creates the dynamic PPPoE interface, the \$junos-underlying-interface predefined variable is dynamically replaced with the name of the underlying interface supplied by the network when the subscriber logs in.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring a PPPoE Dynamic Profile</i>• <i>Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview</i>

unit

```

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

```

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

```

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

```

```
}
short-sequence;
targeted-distribution;
transmit-weight number;
(traps | no-traps);
trunk-bandwidth rate;
trunk-id number;
tunnel {
    backup-destination address;
    destination address;
    key number;
    routing-instance {
        destination routing-instance-name;
    }
    source source-address;
    ttl number;
}
vci vpi-identifier.vci-identifier;
vci-range start start-vci end end-vci;
vpi vpi-identifier;
vlan-id number;
vlan-id-range number-number;
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
family family {
    accounting {
        destination-class-usage;
        source-class-usage {
            (input | output | input output);
        }
    }
}
access-concentrator name;
address address {
    ... the address subhierarchy appears after the main [edit interfaces interface-name unit
       logical-unit-number family family-name] hierarchy ...
}
bundle interface-name;
core-facing;
demux-destination {
    destination-prefix;
}
demux-source {
    source-prefix;
}
direct-connect;
duplicate-protection;
dynamic-profile profile-name;
filter {
    group filter-group-number;
    input filter-name;
    input-list [filter-names];
    output filter-name;
    output-list [filter-names];
}
interface-mode (access | trunk);
ipsec-sa sa-name;
keep-address-and-control;
```

```

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

```

```

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

```

Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i>], [edit interfaces interface-set <i>interface-set-name</i> interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.
Options	<p>logical-unit-number—Number of the logical unit.</p> <p>Range: 0 through 1,073,741,823 for demux and PPPoE static interfaces. 0 through 16,385 for all other static interface types.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

unit (Dynamic Demux Interface)

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

Hierarchy Level [edit `dynamic-profiles profile-name interfaces demux0`]

Release Information Statement introduced in Junos OS Release 9.3.

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

Options *logical-unit-number*—Either the specific unit number of the interface or the unit number variable (`$junos-interface-unit`). The variable is used to specify the unit of the interface when a new demux interface is dynamically created. The static unit number variable is dynamically replaced with the unit number that DHCP supplies when the subscriber logs in.

The remaining statements are explained separately.

Required Privilege	interface—To view this statement in the configuration.
Level	interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles</i>

unit (Dynamic Profiles Standard Interface)

```

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

```

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

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

Release Information Statement introduced in Junos OS Release 9.2.

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

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

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

The remaining statements are explained separately.

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

Related Documentation

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

unnumbered-address (Dynamic PPPoE)

Syntax unnumbered-address *interface-name*;

Hierarchy Level [edit **dynamic-profiles** *profile-name* **interfaces** pp0 unit "\$junos-interface-unit" **family** inet]

Release Information Statement introduced in Junos OS Release 10.1.

Description For dynamic PPPoE interfaces, enable the local address to be derived from the specified interface. Configuring unnumbered Ethernet interfaces enables IP processing on the interface without assigning an explicit IP address to the interface.

Options *interface-name*—Interface from which the local address is derived. The interface name must include a logical unit number and must have a configured address.

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

Related Documentation

- *Configuring a PPPoE Dynamic Profile*
- *Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview*

unnumbered-address (Dynamic Profiles)

Syntax	<code>unnumbered-address interface-name <preferred-source-address address>;</code>
Hierarchy Level	[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family family], [edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family]
Release Information	Statement introduced in Junos OS Release 9.2. \$junos-preferred-source-address variable support added in Junos OS Release 9.6. Support for the \$junos-loopback-interface predefined variable introduced in Junos OS Release 9.6.
Description	<p>For Ethernet interfaces, enable the local address to be derived from the specified interface. Configuring unnumbered Ethernet interfaces enables IP processing on the interface without assigning an explicit IP address to the interface. To configure unnumbered address dynamically, include the \$junos-loopback-interface-address predefined variable.</p> <p>You can configure unnumbered address support on Ethernet interfaces for IPv4 and IPv6 address families.</p>
Options	interface-name —Name of the interface from which the local address is derived. The specified interface must have a logical unit number, a configured IP address, and must not be an unnumbered interface. This value can be a specific interface name or the \$junos-loopback-interface dynamic variable.

When defining the **unnumbered-address** statement using a static interface, keep the following in mind:

- If you choose to include the **routing-instance** statement at the [edit [dynamic-profiles](#)] hierarchy level, that statement must be configured with a valid, static routing instance value. In addition, whatever static unnumbered interface you specify must belong to that routing instance.
- If you choose to not include the **routing-instance** statement at the [edit [dynamic-profiles](#)] hierarchy level, the unnumbered-address statement uses the default routing instance. The use of the default routing instance requires that the unnumbered interface be configured statically and that it reside in the default routing instance.

When defining the **unnumbered-address** statement using the **\$junos-loopback-interface** dynamic variable, keep the following in mind:

- To use the **\$junos-loopback-interface** dynamic variable, the dynamic profile must also contain the **routing-instance** statement configured with the **\$junos-routing-instance** dynamic variable at the [edit [dynamic-profiles](#)] hierarchy level.
- The applied loopback interface is based on the dynamically obtained routing instance of the subscriber.

address—(Optional) Secondary IP address of the donor interface. Configuring the preferred source address enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet interfaces in your network. This value can be a static IP address, the **\$junos-preferred-source-address** dynamic variable for the inet family, or **\$junos-preferred-source-ipv6-address** dynamic variable for the inet6 family.

When defining the **preferred-source-address** value using a static IP address, keep the following in mind:

- The unnumbered interface must be statically configured.
- The IP address specified as the **preferred-source-address** must be configured in the specified unnumbered interface.

When defining the **preferred-source-address** value using the **\$junos-preferred-source-address** or **\$junos-preferred-source-ipv6-address** dynamic variables, keep the following in mind:

- You must configure the **unnumbered-address** statement using the **\$junos-loopback-interface** dynamic variable.
- You must configure the **routing-instance** statement using the **\$junos-routing-instance** dynamic variable at the **[edit dynamic-profiles]** hierarchy level.
- The preferred source address chosen is based on the dynamically applied loopback address which is in turn derived from the dynamically obtained routing instance of the subscriber. The configured loopback address with the closest network match to the user IP address is selected as the preferred source address.

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

Related Documentation	<ul style="list-style-type: none"> • <i>Dynamic Profiles Overview</i>
------------------------------	--

unnumbered-address (Ethernet)

Syntax	<code>unnumbered-address interface-name <preferred-source-address address>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced in Junos OS Release 8.2. preferred-source-address option introduced in Junos OS Release 9.0.
Description	For Ethernet interfaces, enable the local address to be derived from the specified interface. Configuring an unnumbered Ethernet interface enables IP processing on the interface without assigning an explicit IP address to the interface.
Options	interface-name —Name of the interface from which the local address is derived. The specified interface must have a logical unit number and a configured IP address, and must not be an unnumbered interface. The preferred-source-address statement is explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring an Unnumbered Interface</i>• address on page 101• <i>Junos System Basics Configuration Guide</i>

user-prefix (DHCP Local Server)

Syntax	<code>user-prefix <i>user-prefix-string</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit system services dhcp-local-server authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit system services dhcp-local-server group group-name authentication username-include]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Specify the user prefix that is concatenated with the username during the subscriber authentication or DHCP client authentication process.
Options	<i>user-prefix-string</i> —User prefix string.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

- Related Documentation**
- *Using External AAA Authentication Services with DHCP*

vlan-id (Dynamic VLANs)

Syntax	<code>vlan-id number;</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	<p>For dynamic VLAN interfaces, specify the line VLAN identifiers to be rewritten at the input or output interface.</p> <p>You cannot include the vlan-id statement with the swap statement, swap-push statement, push-push statement, or push-swap statement at the [edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map] hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the vlan-id statement that you include at the [edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.</p>
Options	number —A valid VLAN identifier. When used for input VLAN maps, you can specify the <code>\$junos-vlan-map-id</code> predefined variable to dynamically obtain the VLAN identifier.
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Rewriting the VLAN Tag on Tagged Frames</i>• <i>Binding VLAN IDs to Logical Interfaces</i>

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

Syntax	<code>vlan-id <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For Fast Ethernet, Gigabit Ethernet, and Aggregated Ethernet interfaces only, bind a 802.1Q VLAN tag ID to a logical interface.
Options	<p><i>number</i>—A valid VLAN identifier.</p> <p>Range: For aggregated Ethernet, 4-port, 8-port, and 12-port Fast Ethernet PICs, and for management and internal Ethernet interfaces, 1 through 1023.</p> <p>For 48-port Fast Ethernet and Gigabit Ethernet PICs, 1 through 4094.</p> <p>VLAN ID 0 is reserved for tagging the priority of frames.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Mixed Tagging</i>

vlan-model

Syntax	vlan-model one-to-one;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced in Junos OS Release 11.2.
Description	Define the network VLAN model.
Options	one-to-one —Specify that any received, dual-tagged VLAN packet triggers the provisioning process in a Layer 2 Wholesale network. Using this option, the router learns VLAN tags for each individual client. The router learns both the outer tag and inner tag of the incoming packets, when the instance-role statement is defined as access , or the outer VLAN tag only, when the instance-role statement is defined as nni .
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers on page 81• Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers on page 84

vlan-ranges

```
Syntax  vlan-ranges {
        access-profile profile-name;
        authentication {
            packet-types [packet-types];
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                mac-address;
                option-82 <circuit-id> <remote-id>;
                radius-realm radius-realm-string;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name {
            accept (any | dhcp-v4 | inet);
            ranges (any | low-tag)–(any | high-tag);
        }
        override;
    }
```

Hierarchy Level [edit interfaces *interface-name* [auto-configure](#)]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure multiple VLANs. Each VLAN is assigned a VLAN ID number from the range.

The remaining statements are explained separately.

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

Related Documentation

- *Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs*
- *Configuring Interfaces to Support Both Single and Stacked VLANs*

vlan-tags

Syntax	<code>vlan-tags outer [<i>tpid</i>].<i>vlan-id</i> [inner [<i>tpid</i>].<i>vlan-id</i>];</code>
Hierarchy Level	<code>[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.5. VLAN demux interface support introduced in Junos OS Release 10.2.
Description	For Gigabit Ethernet IQ and IQE interfaces only, binds TPIDs and 802.1Q VLAN tag IDs to a logical interface. You must include the stacked-vlan-tagging statement at the <code>[edit interfaces <i>interface-name</i>]</code> hierarchy level.



NOTE: The inner-range *vid1–vid2* option is supported on IQE PICs only.

Options	<p>inner [<i>tpid</i>].<i>vlan-id</i>—A TPID (optional) and a valid VLAN identifier in the format <i>tpid.vlan-id</i>. When used in the dynamic-profiles hierarchy, specify the <code>\$junos-vlan-id</code> predefined variable to dynamically obtain the VLAN ID.</p> <p>Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.</p> <p>outer [<i>tpid</i>].<i>vlan-id</i>—A TPID (optional) and a valid VLAN identifier in the format <i>tpid.vlan-id</i>. When used in the dynamic-profiles hierarchy, specify the <code>\$junos-stacked-vlan-id</code> predefined variable.</p> <p>Range: For VLAN ID, 1 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces. VLAN ID 0 is reserved for tagging the priority of frames.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Dual VLAN Tags• stacked-vlan-tagging on page 234

vlan-tags (Stacked VLAN Tags)

Syntax	<code>vlan-tags inner <i>tpid.vlan-id</i> inner-range <i>vid1—vid2</i> outer <i>tpid.vlan-id</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.
Description	For Gigabit Ethernet IQ and IQE interfaces only, bind TPIDs and 802.1Q VLAN tag IDs to a logical interface.
Options	inner <i>tpid.vlan-id</i> —A TPID and a valid VLAN identifier. Range: (most routers) For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames. For PTX Series, VLAN ID 0 is not supported. inner-range <i>tpid. vid1—vid2</i> —Specify a TPID and a range of VLAN IDs where <i>vid1</i> is the start of the range and <i>vid2</i> is the end of the range.



NOTE: On the network-to-network (NNI) or egress interfaces of provider edge (PE) routers, you cannot configure the inner-range *tpid. vid1—vid2* option with the `vlan-tags` statement for ISP-facing interfaces.

Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

outer *tpid.vlan-id*—A TPID and a valid VLAN identifier.

Range: (most routers) For VLAN ID, 1 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces. VLAN ID 0 is reserved for tagging the priority of frames. For PTX Series, VLAN ID 0 is not supported.



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

The following examples illustrate this further:

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

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

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

Related Documentation

- *Configuring Dual VLAN Tags*
- *Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers*
- [stacked-vlan-tagging on page 234](#)

vpls (Routing Instance)

```

Syntax  vpls {
        active-interface {
            any;
            primary interface-name;
        }
        community COMM;
        connectivity-type (ce | irb);
        control-word;
        encapsulation-type ethernet;
        ignore-encapsulation-mismatch;
        ignore-mtu-mismatch;
        interface interface-name;
        interface-mac-limit limit;
        label-block-size size;
        mac-flush [ explicit-mac-flush-message-options ];
        mac-table-aging-time time;
        mac-table-size size;
        mesh-group mesh-group-name {
            l2vpn-id (as-number:id | ip-address:id);
            local-switching;
            mac-flush [ explicit-mac-flush-message-options ];
            neighbor address {...}
            peer-as all;
            pseudowire-status-tlv;
            route-distinguisher (as-number:id | ip-address:id);
            vpls-id number;
            vrf-export [ policy-names ];
            vrf-import [ policy-names ];
            vrf-target {
                community;
                import community-name;
                export community-name;
            }
        }
        mtu mtu;
        no-control-word;
        no-tunnel-services;
        site site-name {
            active-interface interface-name {
                any;
                primary preference-value;
            }
            best-site;
            interface interface-name {
                interface-mac-limit limit;
            }
            mesh-group mesh-group-name;
            multi-homing;
            site-identifier identifier;
            site-preference preference-value {
                backup;
                primary;
            }
        }
    }

```

```
    }  
  }  
  site-range number;  
  traceoptions {  
    file filename <files number> <size size> <world-readable | no-world-readable>;  
    flag flag <flag-modifier> <disable>;  
  }  
  tunnel-services {  
    devices device-names;  
    primary primary-device-name;  
  }  
}
```

Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. The mac-flush option was added in Junos OS Release 10.0.
Description	Configure a virtual private LAN service (VPLS) routing instance. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring VPLS Routing Instances</i>

vrf-export

Syntax	<code>vrf-export [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>] [edit routing-instances <i>routing-instance-name</i>] [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>] [edit switch-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D30 for QFX Series switches.</p>
Description	<p>Specify how routes are exported from the local PE router's VRF table (<i>routing-instance-name</i>.inet.0) to the remote PE router. If the value vrf is specified for the instance-type statement included in the routing instance configuration, this statement is required.</p> <p>You can configure multiple export policies on the PE router or PE switch.</p>
Default	If the instance-type is vrf , vrf-export is a required statement. The default action is to reject.
Options	<i>policy-names</i> —Names for the export policies.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Implementing EVPN-VXLAN for Data Centers</i> • instance-type on page 193 • <i>Configuring Policies for the VRF Table on PE Routers in VPNs</i>

vrf-import

Syntax	<code>vrf-import [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> vpls mesh-group <i>mesh-group-name</i>]</code> <code>[edit routing-instances <i>routing-instance-name</i>]</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>]</code> <code>[edit switch-options]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D30 for QFX Series switches.
Description	<p>Specify how routes are imported into the virtual routing and forwarding (VRF) table (<i>routing-instance-name</i>.inet.0) of the local provider edge (PE) router or switch from the remote PE router. If the value vrf is specified for the instance-type statement included in the routing instance configuration, this statement is required.</p> <p>You can configure multiple import policies on the PE router or switch.</p>
Default	If the instance type is vrf , vrf-import is a required statement. The default action is to accept.
Options	<i>policy-names</i> —Names for the import policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Implementing EVPN-VXLAN for Data Centers</i>• instance-type on page 193• <i>Configuring Policies for the VRF Table on PE Routers in VPNs</i>

vrf-target

Syntax	<pre>vrf-target { community; auto import community-name; export community-name; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn mesh-group <i>mesh-group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>], [edit routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols l2vpn mesh-group <i>mesh-group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>], [edit switch-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D30 for QFX Series switches. auto option was also added at this time.</p>
Description	<p>Specify a virtual routing and forwarding (VRF) target community. If you configure the community option only, default VRF import and export policies are generated that accept and tag routes with the specified target community. The purpose of the vrf-target statement is to simplify the configuration by allowing you to configure most statements at the [edit routing-instances] hierarchy level. In effect, this statement configures a single policy for import and a single policy for export to replace the per-VRF policies for every community.</p> <p>You can still create more complex policies by explicitly configuring VRF import and export policies using the import and export options.</p>
Options	<p>community—Community name.</p> <p>auto—Automatically derives the route target (RT) for supported QFX Series switches.</p> <p>import community-name—Allowed communities accepted from neighbors.</p> <p>export community-name—Allowed communities sent to neighbors.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Implementing EVPN-VXLAN for Data Centers Configuring Policies for the VRF Table on PE Routers in VPNs

- *Example: Configuring FEC 129 BGP Autodiscovery for VPWS*

CHAPTER 9

Operational Commands

- [AAA and Address Assignment Pool Commands on page 267](#)
- [DHCPv4 Local Server Commands on page 280](#)
- [DHCPv4 Relay Commands on page 294](#)
- [DHCPv6 Local Server Commands on page 307](#)
- [Interface Commands on page 319](#)
- [Subscriber Commands on page 416](#)
- [VPLS Commands on page 442](#)

AAA and Address Assignment Pool Commands

- [show network-access aaa statistics](#)
- [show network-access aaa statistics authentication](#)
- [show network-access aaa subscribers](#)
- [show network-access address-assignment pool](#)

show network-access aaa statistics

Syntax	<pre>show network-access aaa statistics <accounting (detail)> <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> <p>Option detail introduced in Junos OS Release 13.3.</p>
Description	Display AAA accounting, address-assignment, dynamic request statistics, and RADIUS settings and statistics.
Options	<p>accounting (detail)—(Optional) Display AAA accounting statistics. The detail keyword displays additional accounting information</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
Related Documentation	<ul style="list-style-type: none"> <i>Verifying and Managing Subscriber AAA Information</i>
List of Sample Output	<p>show network-access aaa statistics accounting on page 271</p> <p>show network-access aaa statistics accounting detail on page 271</p> <p>show network-access aaa statistics address-assignment client on page 271</p> <p>show network-access aaa statistics address-assignment pool on page 271</p> <p>show network-access aaa statistics dynamic-requests on page 272</p> <p>show network-access aaa statistics radius on page 272</p>
Output Fields	<p>Table 9 on page 268 lists the output fields for the show network-access aaa statistics command. Output fields are listed in the approximate order in which they appear.</p>

Table 9: 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 on requests	Number of accounting on requests sent from a client to a RADIUS accounting server.

Table 9: show network-access aaa statistics Output Fields (*continued*)

Field Name	Field Description
Accounting start requests	Number of accounting start requests sent from a client to a RADIUS accounting server.
Accounting interim requests	Number of accounting interim requests sent from a client to a RADIUS accounting server.
Accounting stop requests	Number of accounting stop requests sent from a client to a RADIUS accounting 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.
Timed out requests	Number of accounting requests to the accounting server that timed out.
Accounting on responses	Number of accounting on requests acknowledged by the RADIUS accounting server.
Accounting start responses	Number of accounting start requests acknowledged by the RADIUS accounting server.
Accounting interim responses	Number of accounting interim requests acknowledged by the RADIUS accounting server.
Accounting stop responses	Number of accounting stop requests acknowledged by the RADIUS accounting server.
Accounting rollover requests	Number of accounting requests coming to a RADIUS accounting server after a previous server timing out.
Accounting unknown requests	Number of unknown accounting requests sent from a client to a RADIUS accounting server (for example, the header has invalid or unsupported information).
Accounting pending account requests	Number of accounting requests sent from a client to a RADIUS accounting server that are waiting for a response from the server.
Accounting malformed responses	Number of accounting responses from a RADIUS accounting server that have invalid or unexpected attributes.
Accounting retransmissions	Number of accounting requests made by a client to the RADIUS sever that were retransmitted.
Accounting bad authenticators	Number of accounting responses from a RADIUS accounting server that have an incorrect authenticator (for example, the client and server RADIUS secret do not match).
Accounting packets dropped	Number of accounting responses from a RADIUS accounting server that are dropped by a client.

Table 9: show network-access aaa statistics Output Fields (*continued*)

Field Name	Field Description
Client	Client type; for example, DHCP, Mobile IP, PPP.
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.

Table 9: show network-access aaa statistics Output Fields (*continued*)

Field Name	Field Description
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 aaa statistics accounting

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

show network-access aaa statistics accounting detail

```
user@host> show network-access aaa statistics accounting detail
Accounting module statistics
  Requests received: 261
    Accounting on requests: 261
    Accounting start requests: 0
    Accounting interim requests: 0
    Accounting stop requests: 0
  Accounting response failures: 0
  Accounting response success: 0
    Accounting on responses: 0
    Accounting start responses: 0
    Accounting interim responses: 0
    Accounting stop responses: 0
  Timed out requests: 260
  Accounting rollover requests: 0
  Accounting unknown responses: 0
  Accounting pending account requests: 1
  Accounting malformed responses: 0
  Accounting retransmissions: 783
  Accounting bad authenticators: 0
  Accounting packets dropped: 0
```

show network-access aaa statistics address-assignment client

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

show network-access aaa statistics address-assignment pool

```
user@host> show network-access aaa statistics address-assignment pool isp_1
Address-assignment statistics
  Pool Name: isp_1
  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 aaa statistics dynamic-requests

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

show network-access aaa statistics radius

```
user@host> show network-access aaa statistics radius
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
Related Documentation	<ul style="list-style-type: none"> <i>Verifying and Managing Subscriber AAA Information</i>
List of Sample Output	show network-access aaa statistics authentication on page 275 show network-access aaa statistics authentication detail on page 275
Output Fields	Table 10 on page 273 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 10: 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
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
Timed out requests	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 10: 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 aaa statistics authentication

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

show network-access aaa statistics authentication detail

```
user@host> show network-access aaa statistics authentication detail
Authentication module statistics
  Requests received: 2118
  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
  Timed out requests: 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
Related Documentation	<ul style="list-style-type: none"> • <i>Verifying and Managing Subscriber AAA Information</i>
List of Sample Output	<p>show network-access aaa subscribers logical-system on page 277</p> <p>show network-access aaa subscribers logical-system routing-instance on page 277</p> <p>show network-access aaa subscribers statistics username on page 278</p> <p>show network-access aaa subscribers username on page 278</p>
Output Fields	Table 11 on page 276 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 11: 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.

Table 11: show network-access aaa subscribers Output Fields (*continued*)

Field Name	Field Description
INTERIM sent successfully	Number of accounting interim requests generated by the AAA framework for this subscriber.
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-cmbrg-12
conley@isp3.com    dhcp            default:isp-gtown-r3-00
0020df980102.2334 dhcp            isp-bos-metro-16:isp-cmbrg-12

```

show network-access aaa subscribers logical-system routing-instance

```

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

```

show network-access aaa subscribers statistics username

```
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
  Successful 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

```
user@host> show network-access aaa subscribers username fred@isp5.net
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               time         6000 secs  on/volume
GamingBurst       time         6000 secs  on/volume
```

show network-access address-assignment pool

Syntax	<code>show network-access address-assignment pool <i>pool-name</i></code> <code><logical-system <i>logical-system-name</i>></code> <code><routing-instance <i>routing-instance-name</i>></code>
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 279
Output Fields	Table 12 on page 279 lists the output fields for the show network-access address-assignment pool command. Output fields are listed in the approximate order in which they appear.

Table 12: 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

show network-access address-assignment pool

```

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

```

DHCPv4 Local Server Commands

- `show dhcp server binding`
- `show dhcp server statistics`
- `clear dhcp server binding`
- `clear dhcp server statistics`

show dhcp server binding

Syntax `show dhcp server binding`
`<address>`
`<interfaces-vlan><brief | detail | summary>`
`<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 Display the address bindings in the client table on the extended Dynamic Host Configuration Protocol (DHCP) local server.



NOTE: If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the `clear dhcp server binding` command before you delete the DHCP server configuration.

Options *address*—(Optional) Display DHCP binding information for a specific client identified by one of the following entries:

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

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`.

interface interface-name—(Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.

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

interfaces-wildcard—(Optional) The set of interfaces on which to show the binding state information. This option supports the use of the wildcard character (*).

logical-system logical-system-name—(Optional) Display information about active client bindings for DHCP clients on the specified logical system.

routing-instance routing-instance-name—(Optional) Display information about active client bindings for DHCP clients on the specified routing instance.

Required Privilege Level view

Related Documentation

- *Clearing DHCP Bindings for Subscriber Access*
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*
- [clear dhcp server binding on page 290](#)

List of Sample Output

[show dhcp server binding on page 284](#)
[show dhcp server binding detail on page 285](#)
[show dhcp server binding detail \(ACI Interface Set Configured\) on page 285](#)
[show dhcp server binding interface <vlan-id> on page 286](#)
[show dhcp server binding interface <svlan-id> on page 286](#)
[show dhcp server binding <ip-address> on page 286](#)
[show dhcp server binding <session-id> on page 286](#)
[show dhcp server binding summary on page 286](#)
[show dhcp server binding <interfaces-vlan> on page 286](#)
[show dhcp server binding <interfaces-wildcard> on page 286](#)

Output Fields [Table 13 on page 282](#) lists the output fields for the **show dhcp server binding** command. Output fields are listed in the approximate order in which they appear.

Table 13: 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

Table 13: show dhcp server binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
Lease Start	Date and time at which the client's IP address lease started.	detail
Lease time violated	Lease time violation has occurred.	detail
Last Packet Received	Date and time at which the router received the last packet.	detail
Incoming Client Interface	Client's incoming interface.	detail
Client Interface Svlan Id	S-VLAN ID of the client's incoming interface.	detail
Client Interface Vlan Id	VLAN ID of the client's incoming interface.	detail
Demux Interface	Name of the IP demultiplexing (demux) interface.	detail
Server IP Address or Server Identifier	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

Table 13: show dhcp server binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Liveness Detection State	<p>State of the liveness detection status for a subscriber's Bidirectional Forwarding Detection (BFD) protocol session:</p> <p>NOTE: This output field displays status only when liveness detection has been explicitly configured for a subscriber and the liveness detection protocol is actively functioning for that subscriber.</p> <ul style="list-style-type: none"> DOWN—Liveness detection has been enabled for a subscriber but the broadband network gateway (BNG) detects that the liveness detection session for the BFD protocol is in the DOWN state. A liveness detection session that was previously in an UP state has transitioned to a DOWN state, beginning with a liveness detection failure, and ending with the deletion of the client binding. The DOWN state is reported only during this transition period of time. UNKNOWN—Liveness detection has been enabled for a subscriber but the actual liveness detection state has not yet been determined. The UNKNOWN state is reported after a DHCP subscriber initially logs in while the underlying liveness detection protocol handshake, such as BFD, is still processing and the BFD session has not yet reached the UP state. UP—Liveness detection has been enabled for a subscriber, and the BNG and the subscriber or client have <i>both</i> determined that the liveness detection session for the BFD protocol is in the UP state. WENT_DOWN—State is functionally equivalent to the DOWN state. A liveness detection session that was previously in an UP state has transitioned to a DOWN state implying a liveness detection failure. The WENT_DOWN state applies to the internal distribution of the liveness detection mechanism between the Junos DHCP Daemon for Subscriber Services (JDHCPd), the BFD plug-in within the Broadband Edge Subscriber Management Daemon (BBE-SMGD), and the Packet Forwarding Engine. 	detail
ACI Interface Set Name	Internally generated name of the dynamic agent circuit identifier (ACI) interface set.	detail
ACI Interface Set Index	Index number of the dynamic ACI interface set.	detail
ACI Interface Set Session ID	Identifier of the dynamic ACI interface set entry in the session database.	detail

Sample Output

show dhcp server binding

```
user@host> show dhcp server 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
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
  Liveness Detection State: UP
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
  Lease time violated:    yes
  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
  Liveness Detection State: UP

```

show dhcp server binding detail (ACI Interface Set Configured)

```

user@host> show dhcp server binding detail
Client IP Address: 100.20.22.14
  Hardware Address:      00:00:64:34:01:02
  State:                 BOUND(LOCAL_SERVER_STATE_BOUND)
  Lease Expires:         2012-03-13 09:53:32 PDT
  Lease Expires in:      82660 seconds
  Lease Start:           2012-03-12 10:23:32 PDT
  Last Packet Received:  2012-03-12 10:23:32 PDT
  Incoming Client Interface: demux0.1073741827
  Client Interface Svlan Id: 1802
  Client Interface Vlan Id: 302
  Demux Interface:       demux0.1073741832
  Server Identifier:     100.20.200.202
  Session Id:            11
  Client Pool Name:      poolA
  Client Profile Name:    DEMUXprofile
  Liveness Detection State: UP

```

ACI Interface Set Name: aci-1002-demux0.1073741827
ACI Interface Set Index: 2
ACI Interface Set Session ID: 6

show dhcp server binding interface <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 <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>

```
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>

```
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 <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 <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	<p>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.</p> <p>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.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear dhcp server statistics on page 293
List of Sample Output	show dhcp server statistics on page 288
Output Fields	Table 14 on page 288 lists the output fields for the show dhcp server statistics command. Output fields are listed in the approximate order in which they appear.

Table 14: 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 • Lease Time Violation—Number of packets discarded because of a lease time violation • 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                  1

```

Lease Time Violation	1
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.



NOTE: If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the `clear dhcp server binding` command before you delete the DHCP server configuration.

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 281](#)

List of Sample Output

- [clear dhcp server binding <ip-address> on page 291](#)
- [clear dhcp server binding all on page 291](#)
- [clear dhcp server binding interface on page 292](#)
- [clear dhcp server binding <interfaces-vlan> on page 292](#)
- [clear dhcp server binding <interfaces-wildcard> on page 292](#)

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 <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP local server bindings over a specific interface:

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


clear dhcp server statistics

Syntax	<code>clear dhcp server statistics</code> <code><logical-system <i>logical-system-name</i>></code> <code><routing-instance <i>routing-instance-name</i>></code>
Release Information	Command introduced in Junos OS Release 9.0.
Description	Clear all extended Dynamic Host Configuration Protocol (DHCP) local server statistics.
Options	<p>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.</p> <p>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.</p>
Required Privilege Level	view
List of Sample Output	clear dhcp server statistics on page 293
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                1
    Lease Time Violation 1

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

DHCPv4 Relay Commands

- [show dhcp relay binding](#)
- [show dhcp relay statistics](#)
- [clear dhcp relay binding](#)
- [clear dhcp relay statistics](#)

show dhcp relay binding

Syntax	<pre> show dhcp relay binding <address> <brief> <detail> <interface interface-name> <interfaces-vlan> <interfaces-wildcard> <ip-address mac-address> <logical-system logical-system-name> <routing-instance routing-instance-name> <summary> </pre>
Release Information	<p>Command introduced in Junos OS Release 8.3.</p> <p>Options interface and mac-address added in Junos OS Release 8.4.</p> <p>Options interfaces-vlan and interfaces-wildcard added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.</p>
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 interface-name—(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 logical-system-name—(Optional) Perform this operation on the specified logical system.</p> <p>routing-instance routing-instance-name—(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 303](#)

List of Sample Output

- [show dhcp relay binding on page 297](#)
- [show dhcp relay binding detail on page 298](#)
- [show dhcp relay binding interface on page 298](#)
- [show dhcp relay binding interface vlan-id on page 298](#)
- [show dhcp relay binding interface svlan-id on page 298](#)
- [show dhcp relay binding ip-address on page 299](#)
- [show dhcp relay binding mac-address on page 299](#)
- [show dhcp relay binding session-id on page 299](#)
- [show dhcp relay binding <interfaces-vlan> on page 299](#)
- [show dhcp relay binding <interfaces-wildcard> on page 299](#)
- [show dhcp relay binding summary on page 299](#)

Output Fields Table 15 on page 296 lists the output fields for the **show dhcp relay binding** command. Output fields are listed in the approximate order in which they appear.

Table 15: 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
Generated Remote ID	Remote ID generated by the Option 82 Agent Remote ID (suboption 1)	detail
Hardware address	Hardware address of the DHCP client.	briefdetail
Expires	Number of seconds in which the lease expires.	briefdetail

Table 15: show dhcp relay binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
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
Lease time violated	Lease time violation has occurred.	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
  Lease time violated: yes
  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
  Generated Remote ID: host:ge-1/0/0:100
```

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 vlan-id

```
user@host> show dhcp relay 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 relay binding interface svlan-id

```
user@host> show dhcp relay 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 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 <interfaces-vlan>

```

user@host> show dhcp relay 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 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	<code>show 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. Command introduced in Junos OS Release 12.1 for EX Series switches. Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.
Description	Display 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 displayed 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 displayed for the default routing instance.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear dhcp relay statistics on page 305
List of Sample Output	show dhcp relay statistics on page 302
Output Fields	Table 16 on page 301 lists the output fields for the <code>show dhcp relay statistics</code> command. Output fields are listed in the approximate order in which they appear.

Table 16: 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. • Lease Time Violation—Number of packets discarded because of a lease time violation • 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
External Server Response	State of the external DHCP server responsiveness.
Packets forwarded	<p>Number of packets forwarded.</p> <ul style="list-style-type: none"> • BOOTREQUEST—Number of BOOTREQUEST protocol data units (PDUs) forwarded • BOOTREPLY—Number of BOOTREPLY protocol data units (PDUs) forwarded

Table 16: show dhcp relay statistics Output Fields (*continued*)

Field Name	Field Description
External Server Response	State of the external DHCP server responsiveness.

Sample Output

show dhcp relay statistics

```

user@host> show dhcp relay statistics
Packets dropped:
    Total                               34
    Bad hardware address                 1
    Bad opcode                           1
    Bad options                           3
    Invalid server address                5
    Lease Time Violation                 1
    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

Packets forwarded:
    Total                                 4
    BOOTREQUEST                           2
    BOOTREPLY                             2

External Server Response:
    State                                Responding

```

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	<p>Command introduced in Junos OS Release 8.3.</p> <p>Options all and interface added in Junos OS Release 8.4.</p> <p>Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.</p>
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"> <i>Clearing DHCP Bindings for Subscriber Access</i> show dhcp relay binding on page 295

List of Sample Output [clear dhcp relay binding on page 304](#)
[clear dhcp relay binding all on page 304](#)
[clear dhcp relay binding interface on page 304](#)
[clear dhcp relay binding <interfaces-vlan> on page 304](#)
[clear dhcp relay binding <interfaces-wildcard> on page 304](#)

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	<pre>clear 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.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.</p>
Description	Clear all 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 cleared 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 cleared for the default routing instance.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show dhcp relay statistics on page 300
List of Sample Output	clear dhcp relay statistics on page 306
Output Fields	Table 17 on page 306 lists the output fields for the clear dhcp relay statistics command.

Table 17: 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. • Lease Time Violation—Number of packets discarded because of a lease time violation • 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 • DHC PNACK—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          1
  Lease Time Violated 1

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
```

DHCPv6 Local Server Commands

- `show dhcpv6 server binding`
- `show dhcpv6 server statistics`
- `clear dhcpv6 server binding`
- `clear dhcpv6 server statistics`

show dhcpv6 server binding

Syntax	<pre>show dhcpv6 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 9.6. 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 for IPv6 (DHCPv6) local server.
Options	<p>address—(Optional) One of the following identifiers for the DHCPv6 client whose binding state you want to show:</p> <ul style="list-style-type: none">• <i>CID</i>—The specified Client ID (CID).• <i>ipv6-prefix</i>—The specified IPv6 prefix.• <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 dhcpv6 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) Interface VLAN ID or S-VLAN ID interface on which to show binding state information.</p> <p>interfaces-wildcard—(Optional) Set of interfaces on which to show 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 DHCPv6 clients on the specified logical system.</p> <p>routing-instance routing-instance-name—(Optional) Display information about active client bindings for DHCPv6 clients on the specified routing instance.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>Clearing DHCP Bindings for Subscriber Access</i>• clear dhcpv6 server binding on page 317

List of Sample Output

- [show dhcpv6 server binding on page 310](#)
- [show dhcpv6 server binding detail on page 310](#)
- [show dhcpv6 server binding interface on page 311](#)
- [show dhcpv6 server binding interface detail on page 311](#)
- [show dhcpv6 server binding \(IPv6 Prefix\) on page 312](#)
- [show dhcpv6 server binding \(Session ID\) on page 312](#)
- [show dhcpv6 server binding \(Interfaces VLAN\) on page 312](#)
- [show dhcpv6 server binding \(Interfaces Wildcard\) on page 312](#)
- [show dhcpv6 server binding \(Interfaces Wildcard\) on page 312](#)
- [show dhcpv6 server binding summary on page 313](#)

Output Fields [Table 18 on page 309](#) lists the output fields for the **show dhcpv6 server binding** command. Output fields are listed in the approximate order in which they appear.

Table 18: show dhcpv6 server binding Output Fields

Field Name	Field Description	Level of Output
<i>number clients</i> , (<i>number init</i> , <i>number bound</i> , <i>number selecting</i> , <i>number requesting</i> , <i>number renewing</i> , <i>number releasing</i>)	Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state.	summary
Prefix	Client's DHCPv6 prefix, or prefix used to support multiple address assignment.	brief detail
Session Id	Session ID of the subscriber session.	brief detail
Expires	Number of seconds in which lease expires.	brief detail
State	State of the address binding table on the extended DHCPv6 local server: <ul style="list-style-type: none"> • BOUND—Client has active IP address lease. • INIT—Initial state. • RECONFIGURE—Server has sent reconfigure message to client. • RELEASE—Client is releasing IP address lease. • RENEWING—Client sending request to renew IP address lease. • REQUESTING—Client requesting a DHCPv6 server. • SELECTING—Client receiving offers from DHCPv6 servers. 	brief detail
Interface	Interface on which the DHCPv6 request was received.	brief
Client IPv6 Address	Client's IPv6 address.	detail
Client IPv6 Prefix	Client's IPv6 prefix.	detail
Client DUID	Client's DHCP Unique Identifier (DUID).	brief detail
Lease expires	Date and time at which the client's IP address lease expires.	detail

Table 18: show dhcpv6 server binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Lease expires in	Number of seconds in which lease expires.	detail
Preferred Lease Expires	Date and UTC time at which the client's IPv6 prefix expires.	detail
Preferred Lease Expires in	Number of seconds at which client's IPv6 prefix expires.	detail
Lease Start	Date and time at which the client's address lease was obtained.	detail
Lease time violated	Lease time violation has occurred.	detail
Incoming Client Interface	Client's incoming interface.	detail
Server IP Address	IP address of DHCPv6 server.	detail
Server Interface	Interface of DHCPv6 server.	detail
Client Pool Name	Address pool used to assign IPv6 address.	detail
Client Prefix Pool Name	Address pool used to assign IPv6 prefix.	detail
Client Id length	Length of the DHCPv6 client ID, in bytes.	detail
Client Id	ID of the DHCPv6 client.	detail

Sample Output

show dhcpv6 server binding

```

user@host> show dhcpv6 server binding
Prefix          Session Id Expires State Interface Client DUID
2001:bd8:1111:2222::/64 6 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:01
2001:bd8:1111:2222::/64 7 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
2001:bd8:1111:2222::/64 8 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03
2001:bd8:1111:2222::/64 9 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:04
2001:bd8:1111:2222::/64 10 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:05
2002::1/74 11 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:06

```

show dhcpv6 server binding detail

```

user@host> show dhcpv6 server binding detail

```

```

Session Id: 6
  Client IPv6 Prefix:      2001:bd8:1111:2222::/64
  Client DUID:             LL_TIME0x1-0x2e159c0-00:10:94:00:00:01

  State:
  BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
    Lease Expires:         2009-07-21 10:41:15 PDT
    Lease Expires in:      86308 seconds
    Preferred Lease Expires: 2012-07-24 00:18:14 UTC
    Preferred Lease Expires in: 600 seconds
    Lease Start:           2009-07-20 10:41:15 PDT
    Lease time violated:   yes
    Incoming Client Interface: ge-1/0/0.0
    Server Ip Address:     0.0.0.0
    Server Interface:      none
    Client Id Length:      14
    Client Id:
    /0x00010001/0x02e159c0/0x00109400/0x0001

```

```

Session Id: 7
  Client IPv6 Address:     2002::1/128
  Client IPv6 Prefix:      2001:bd8:1111:2222::/64
  Client DUID:             LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

  State:
  BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
    Lease Expires:         2009-07-21 10:41:15 PDT
    Lease Expires in:      86308 seconds
    Preferred Lease Expires: 2012-07-24 00:18:14 UTC
    Preferred Lease Expires in: 600 seconds
    Lease Start:           2009-07-20 10:41:15 PDT
    Incoming Client Interface: ge-1/0/0.0
    Server Ip Address:     0.0.0.0
    Client Pool Name:      bos-v6-pool
    Client Prefix Pool Name: bos-v6-prefix-pool
    Client Id Length:      14
    Client Id:
    /0x00010001/0x02e159c0/0x00109400/0x0002

```

show dhcpv6 server binding interface

```

user@host> show dhcpv6 server binding interface ge-1/0/0:10-101
Prefix      Session Id Expires State Interface Client DUID
2001:bd8:1111:2222::/64 1      86055   BOUND   ge-1/0/0.100
LL_TIME0x1-0x4b0a53b9-00:10:94:00:00:01

```

show dhcpv6 server binding interface detail

```

user@host> show dhcpv6 server binding interface ge-1/0/0:10-101 detail
Session Id: 7
  Client IPv6 Prefix:      2001:bd8:1111:2222::/64
  Client DUID:             LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

  State:                   BOUND(bound)
  Lease Expires:           2009-07-21 10:41:15 PDT
  Lease Expires in:        86136 seconds
  Preferred Lease Expires: 2012-07-24 00:18:14 UTC
  Preferred Lease Expires in: 600 seconds
  Lease Start:             2009-07-20 10:41:15 PDT
  Incoming Client Interface: ge-1/0/0.0
  Server Ip Address:       0.0.0.0

```

```

Server Interface:          none
Client Id Length:         14
Client Id:
/0x00010001/0x02e159c0/0x00109400/0x0002

```

show dhcpv6 server binding (IPv6 Prefix)

```

user@host> show dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005
detail
Session Id: 7
Client IPv6 Prefix:      2001:bd8:1111:2222::/64
Client DUID:             LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

State:                   BOUND(bound)
Lease Expires:           2009-07-21 10:41:15 PDT
Lease Expires in:        86136 seconds
Preferred Lease Expires: 2012-07-24 00:18:14 UTC
Preferred Lease Expires in: 600 seconds
Lease Start:             2009-07-20 10:41:15 PDT
Incoming Client Interface: ge-1/0/0.0
Server Ip Address:       0.0.0.0
Server Interface:        none
Client Id Length:        14
Client Id:
/0x00010001/0x02e159c0/0x00109400/0x0002

```

show dhcpv6 server binding (Session ID)

```

user@host> show dhcpv6 server binding 8
Prefix      Session Id Expires State Interface Client DUID
2001:DB8::/32 8      86235 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03

```

show dhcpv6 server binding (Interfaces VLAN)

```

user@host> show dhcpv6 server binding ge-1/0/0:100-200
Prefix      Session Id Expires State Interface Client DUID
2001:DB8::/32 11      87583 BOUND ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32 12      87583 BOUND ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

show dhcpv6 server binding (Interfaces Wildcard)

```

user@host> show dhcpv6 server binding demux0
Prefix      Session Id Expires State Interface Client DUID
2001:DB8::/32 30      79681 BOUND demux0.1073741824
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32 31      79681 BOUND demux0.1073741825
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:CB9::/32 32      79681 BOUND demux0.1073741826
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

show dhcpv6 server binding (Interfaces Wildcard)

```

user@host> show dhcpv6 server binding ge-1/3/*
Prefix      Session Id Expires State Interface Client DUID
2001:DB8::/32 22      79681 BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32 33      79681 BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

```
2001:CB9::/32      24      79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 server binding summary

```
user@host> show dhcpv6 server binding summary
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

show dhcpv6 server statistics

Syntax	show dhcpv6 server statistics <logical-system <i>logical-system-name</i>> <routing-instance <i>routing-instance-name</i>>
Release Information	Command introduced in Junos OS Release 9.6.
Description	Display extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.
Options	logical-system <i>logical-system-name</i> —(Optional) Display information about extended DHCPv6 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 DHCPv6 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 dhcpv6 server statistics on page 319
List of Sample Output	show dhcpv6 server statistics on page 315
Output Fields	Table 19 on page 315 lists the output fields for the show dhcpv6 server statistics command. Output fields are listed in the approximate order in which they appear.

Table 19: show dhcpv6 server statistics Output Fields

Field Name	Field Description
Packets dropped	<p>Number of packets discarded by the extended DHCPv6 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 DHCPv6 local server • Strict Reconfigure—Number of solicit messages discarded because the client does not support reconfiguration • 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 • Lease Time Violation—Number of packets discarded because of a lease time violation • 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 DHCPv6 local server could not send
Messages received	<p>Number of DHCPv6 messages received.</p> <ul style="list-style-type: none"> • DHCPV6_CONFIRM—Number of DHCPv6 CONFIRM PDUs received. • DHCPV6_DECLINE—Number of DHCPv6 DECLINE PDUs received. • DHCPV6_INFORMATION_REQUEST—Number of DHCPv6 INFORMATION-REQUEST PDUs received. • DHCPV6_REBIND—Number of DHCPv6 REBIND PDUs received. • DHCPV6_RELAY_FORW—Number of DHCPv6 RELAY-FORW PDUs received. • DHCPV6_RELAY_REPL—Number of DHCPv6 RELAY-REPL PDUs received. • DHCPV6_RELEASE—Number of DHCPv6 RELEASE PDUs received. • DHCPV6_RENEW—Number of DHCPv6 RENEW PDUs received. • DHCPV6_REQUEST—Number of DHCPv6 REQUEST PDUs received. • DHCPV6_SOLICIT—Number of DHCPv6 SOLICIT PDUs received.
Messages sent	<p>Number of DHCPv6 messages sent.</p> <ul style="list-style-type: none"> • DHCPV6_ADVERTISE—Number of DHCPv6 ADVERTISE PDUs transmitted. • DHCPV6_REPLY—Number of DHCPv6 ADVERTISE PDUs transmitted. • DHC6_RECONFIGURE—Number of DHCPv6 RECONFIGURE PDUs transmitted.

Sample Output

show dhcpv6 server statistics

```
user@host> show dhcpv6 server statistics
```

Dhcpv6 Packets dropped:

Total	1
Lease Time Violation	1

Messages received:

DHCPV6_DECLINE	0
DHCPV6_SOLICIT	9
DHCPV6_INFORMATION_REQUEST	0
DHCPV6_RELEASE	0
DHCPV6_REQUEST	5
DHCPV6_CONFIRM	0
DHCPV6_RENEW	0
DHCPV6_REBIND	0
DHCPV6_RELAY_FORW	0
DHCPV6_RELAY_REPL	0

Messages sent:

DHCPV6_ADVERTISE	9
DHCPV6_REPLY	5
DHCPV6_RECONFIGURE	0

clear dhcpv6 server binding

Syntax	<pre>clear dhcpv6 server binding <address> <all> <interface interface-name> <interfaces-vlan> <interfaces-wildcard> <logical-system logical-system-name> <routing-instance routing-instance-name></pre>
Release Information	<p>Command introduced in Junos OS Release 9.6.</p> <p>Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.</p>
Description	Clear the binding state of a Dynamic Host Configuration Protocol for IPv6 (DHCPv6) client from the client table on the extended DHCPv6 local server.
Options	<p>address—(Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:</p> <ul style="list-style-type: none"> • <i>CID</i>—The specified Client ID (CID). • <i>ipv6-prefix</i>—The specified IPv6 prefix. • <i>session-id</i>—The specified session ID. <p>all—(Optional) Clear the binding state for all DHCPv6 clients.</p> <p>interface interface-name—(Optional) Clear the binding state for DHCPv6 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) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (*).</p> <p>logical-system logical-system-name—(Optional) Clear the binding state for DHCPv6 clients on the specified logical system.</p> <p>routing-instance routing-instance-name—(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • Clearing DHCP Bindings for Subscriber Access • show dhcpv6 server binding on page 308
List of Sample Output	<p>clear dhcpv6 server binding all on page 318</p> <p>clear dhcpv6 server binding <ipv6-prefix> on page 318</p>

[clear dhcpv6 server binding interface on page 318](#)
[clear dhcpv6 server binding <interfaces-vlan> on page 318](#)
[clear dhcpv6 server binding <interfaces-wildcard> on page 318](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear dhcpv6 server binding all

The following command clears all DHCPv6 local server bindings:

```
user@host> clear dhcpv6 server binding all
```

clear dhcpv6 server binding <ipv6-prefix>

The following command clears DHCPv6 local server bindings for a specific IPv6 prefix:

```
user@host> clear dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005
```

clear dhcpv6 server binding interface

The following command clears DHCPv6 local server bindings on a specific interface:

```
user@host> clear dhcpv6 server binding interface fe-0/0/2
```

clear dhcpv6 server binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCPv6 local server bindings on top of the underlying interface **ae0**, which clears DHCPv6 bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcpv6 server binding interface ae0
```

clear dhcpv6 server binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCPv6 local server bindings over a specific interface:

```
user@host> clear dhcpv6 server binding ge-1/0/0.*
```

clear dhcpv6 server statistics

Syntax	<code>clear dhcpv6 server statistics</code> <code><interface <i>interface-name</i>></code> <code><logical-system <i>logical-system-name</i>></code> <code><routing-instance <i>routing-instance-name</i>></code>
Release Information	Command introduced in Junos OS Release 9.6.
Description	Clear all extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.
Options	<p>logical-system <i>logical-system-name</i>—(Optional) Clear the statistics for DHCPv6 clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Clear the statistics for DHCPv6 clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show dhcpv6 server statistics on page 314
List of Sample Output	clear dhcpv6 server statistics on page 319
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear dhcpv6 server statistics

```
user@host> clear dhcpv6 server statistics
```

Interface Commands

- [show interfaces \(Loopback\)](#)
- [show interfaces \(Aggregated Ethernet\)](#)
- [show interfaces \(Fast Ethernet\)](#)
- [show interfaces \(Gigabit Ethernet\)](#)
- [show interfaces \(PPPoE\)](#)
- [show interfaces demux0 \(Demux Interfaces\)](#)
- [show interfaces filters](#)
- [show interfaces routing](#)
- [show interfaces routing-instance](#)
- [show ppp interface](#)

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 323](#)
[show interfaces brief \(Loopback\) on page 324](#)
[show interfaces detail \(Loopback\) on page 324](#)
[show interfaces extensive \(Loopback\) on page 325](#)

Output Fields [Table 20 on page 320](#) lists the output fields for the **show interfaces** (loopback) command. Output fields are listed in the approximate order in which they appear.

Table 20: Loopback show interfaces Output Fields

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

Table 20: 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 <i>Common Output Fields Description</i> .	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 <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
Link type	Data transmission type.	detail extensive
Link flags	Information about the link. Possible values are described in the “Link Flags” section under <i>Common Output Fields Description</i> .	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: 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
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive

Table 20: 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 <i>Common Output Fields Description</i> .	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 20: 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 <i>Common Output Fields Description</i> .	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive
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 (Loopback)

```

user@host> show interfaces lo0
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
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding MAC Limiting</i> • <i>Ethernet Interfaces Feature Guide for Routing Devices</i>
List of Sample Output	<p>show interfaces (Aggregated Ethernet) on page 332</p> <p>show interfaces brief (Aggregated Ethernet) on page 333</p> <p>show interfaces detail (Aggregated Ethernet) on page 333</p> <p>show interfaces extensive (Aggregated Ethernet) on page 334</p> <p>show interfaces extensive (Aggregated Ethernet with VLAN Stacking) on page 335</p>
Output Fields	Table 21 on page 327 lists the output fields for the show interfaces (Aggregated Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 21: 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

Table 21: Aggregated Ethernet show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Enabled	State of the physical interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	All levels
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 <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interfaces Flags” section under <i>Common Output Fields Description</i> .	All levels
Current address	Configured MAC address.	detail extensive
Hardware address	Hardware MAC address.	detail extensive
Last flapped	Date, time, and how long ago the interface went from down to up or from up to down. The format is Last flapped: year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive

Table 21: 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 and rate, in bps, at which bytes are received on the interface. • Output bytes—Number of bytes and rate, in bps, at which bytes are transmitted on the interface. • Input packets—Number of packets and rate, in pps, at which packets are received on the interface. • Output packets—Number of packets and rate, in pps, at which packets are transmitted on the interface. 	detail extensive
Input errors	<p>Input errors on the interface:</p> <ul style="list-style-type: none"> • Errors—Sum of incoming frame aborts and frame check sequence (FCS) errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's random early detection (RED) mechanism. • Framing errors—Number of packets received with an invalid 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 were not of interest. Usually, this field reports protocols that Junos OS does not handle. • Resource errors—Sum of transmit drops. 	detail extensive
Output errors	<p>Output errors on the interface. 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), then the cable, the far-end system, or the PIC is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	detail extensive

Table 21: Aggregated Ethernet show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
Queue counters	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. <p>NOTE: In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the show interfaces command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenario, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs</p>	detail extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface (which reflects its initialization sequence).	detail extensive none
SNMP ifIndex	SNMP interface index number of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags Field" section under <i>Common Output Fields Description</i> .	All levels
VLAN-Tag	Tag Protocol Identifier (TPID) and VLAN identifier.	All levels
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

Table 21: Aggregated Ethernet show interfaces Output Fields (*continued*)

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

Table 21: Aggregated Ethernet show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
LACP Statistics	<p>LACP statistics for each aggregated interface.</p> <ul style="list-style-type: none"> • LACP Rx—LACP received counter that increments for each normal hello. • LACP Tx—Number of LACP transmit packet errors logged. • Unknown Rx—Number of unrecognized packet errors logged. • Illegal Rx—Number of invalid packets received. <p>NOTE: For LACP Rx and LACP Tx, Packet count is updated only on snmp timer expiry (30 secs).</p>	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. Possible values are described in the "Protocol Field" section under <i>Common Output Fields Description</i> .	brief
Protocol	Protocol family configured on the logical interface. Possible values are described in the "Protocol Field" section under <i>Common Output Fields Description</i> .	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive
Flags	Information about protocol family flags. Possible values are described in the "Family Flags Field" section under <i>Common Output Fields Description</i> .	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 <i>Common Output Fields Description</i> .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces (Aggregated Ethernet)

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



```

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

```

```

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

```

show interfaces brief (Aggregated Ethernet)

```

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

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

```

show interfaces detail (Aggregated Ethernet)

```

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

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

```

```

Logical interface ae0.0 (Index 72) (SNMP ifIndex 60) (Generation 18)
Flags: SNMP-Traps 16384 Encapsulation: ENET2
Statistics
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
Adaptive Statistics:
  Adaptive Adjusts:           0
  Adaptive Scans  :           0
  Adaptive Updates:           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
LACP info:          Role      System              System      Port      Port      Port
                  priority          identifier  priority  number    key

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

LACP Statistics:      LACP Rx      LACP Tx      Unknown Rx      Illegal Rx
  fe-1/0/3.0          3188          3186           0           0
Marker Statistics:    Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
  fe-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, Runts: 0, Giants: 0, Policed discards:
0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0

```

Ingress queues: 4 supported, 4 in use

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

Egress queues: 4 supported, 4 in use

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

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

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

Statistics	Packets	pps	Bytes	bps
Bundle:				
Input :	289	0	25685	376
Output:	1698	4	130375	3096

Link:

ge-1/2/0.451				
Input :	289	0	25685	376
Output:	0	0	0	0
ge-1/2/1.451				
Input :	0	0	0	0
Output:	1698	4	130375	3096

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

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

Flags: Is-Primary

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

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

Statistics	Packets	pps	Bytes	bps
Bundle:				
Input :	293	1	26003	1072
Output:	1694	3	130057	2400

Link:

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

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

show interfaces (Fast Ethernet)

Syntax	<pre>show interfaces <i>interface-type</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display status information about the specified Fast Ethernet interface.
Options	<p><i>interface-type</i>—On M Series and T Series routers, the interface type is <i>fe-fpc/pic/port</i>.</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 351</p> <p>show interfaces brief (Fast Ethernet) on page 352</p> <p>show interfaces detail (Fast Ethernet) on page 352</p> <p>show interfaces extensive (Fast Ethernet) on page 352</p>
Output Fields	Table 22 on page 338 lists the output fields for the show interfaces (Fast Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 22: show interfaces Fast Ethernet Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Table 22: show interfaces Fast Ethernet Output Fields (*continued*)

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

Table 22: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 <i>show interfaces (10-Gigabit Ethernet)</i> command.</p>	detail extensive

Table 22: 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 22: 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	<p>Total number of egress queues supported on the specified interface.</p> <p>NOTE: In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the show interfaces command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenario, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs</p>	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

Table 22: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Ingress queues	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	extensive
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	extensive
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link . <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	The forward error correction (FEC) counters provide the following statistics: <ul style="list-style-type: none"> • Corrected Errors—The count of corrected errors in the last second. • Corrected Error Ratio—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits. 	
PCS statistics	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. <ul style="list-style-type: none"> • Bit errors—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • Errored blocks—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode. 	detail extensive

Table 22: 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 <i>show interfaces (10-Gigabit Ethernet)</i> 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 22: show interfaces Fast Ethernet Output Fields (*continued*)

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

Table 22: 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 22: 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 22: show interfaces Fast Ethernet Output Fields (*continued*)

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

Table 22: show interfaces Fast Ethernet Output Fields (*continued*)

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

Table 22: show interfaces Fast Ethernet Output Fields (*continued*)

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

Table 22: show interfaces Fast Ethernet Output Fields (*continued*)

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

Sample Output

show interfaces (Fast Ethernet)

```

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

```

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

show interfaces brief (Fast Ethernet)

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

show interfaces detail (Fast Ethernet)

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

show interfaces extensive (Fast Ethernet)

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

```

Traffic statistics:
Input bytes :          0          0 bps
Output bytes :         42          0 bps
Input packets:         0          0 pps
Output packets:        1          0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Active alarms : None
Active defects : None
MAC statistics:
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	<code>show interfaces <i>ge-fpc/pic/port</i></code> <code><brief detail extensive terse></code> <code><descriptions></code> <code><media></code> <code><snmp-index <i>snmp-index</i>></code> <code><statistics></code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display status information about the specified Gigabit Ethernet interface.
Options	<p><code><i>ge-fpc/pic/port</i></code>—Display standard information about the specified Gigabit Ethernet interface.</p> <p><code>brief detail extensive terse</code>—(Optional) Display the specified level of output.</p> <p><code>descriptions</code>—(Optional) Display interface description strings.</p> <p><code>media</code>—(Optional) Display media-specific information about network interfaces.</p> <p><code>snmp-index <i>snmp-index</i></code>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><code>statistics</code>—(Optional) Display static interface statistics.</p>
Additional Information	In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration</i>
List of Sample Output	<p>show interfaces (Gigabit Ethernet) on page 370</p> <p>show interfaces (Gigabit Ethernet on MX Series Routers) on page 371</p> <p>show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration) on page 371</p> <p>show interfaces brief (Gigabit Ethernet) on page 371</p> <p>show interfaces detail (Gigabit Ethernet) on page 372</p> <p>show interfaces extensive (Gigabit Ethernet IQ2) on page 373</p> <p>show interfaces (Gigabit Ethernet Unnumbered Interface) on page 376</p> <p>show interfaces (ACI Interface Set Configured) on page 377</p>
Output Fields	<p>Table 23 on page 355 describes the output fields for the show interfaces (Gigabit Ethernet) command. Output fields are listed in the approximate order in which they appear. For Gigabit Ethernet IQ and IQE PICs, the traffic and MAC statistics vary by interface type. For more information, see Table 24 on page 369.</p>

Table 23: show interfaces (Gigabit Ethernet) Output Fields

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

Table 23: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

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

Table 23: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> Input bytes—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. Output bytes—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. Input packets—Number of packets received on the interface. Output packets—Number of packets transmitted on the interface. <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the <i>show interfaces (10-Gigabit Ethernet)</i> command.</p>	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> Errors—Sum of the incoming frame aborts and FCS errors. Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. Framing errors—Number of packets received with an invalid frame checksum (FCS). Runts—Number of frames received that are smaller than the runt threshold. Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the ignore-l3-incompletes statement. L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. Resource errors—Sum of transmit drops. 	extensive

Table 23: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the Drops field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> <ul style="list-style-type: none"> • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Egress queues	<p>Total number of egress queues supported on the specified interface.</p> <p>NOTE: In DPCs that are not of the enhanced type, such as DPC 40x 1GER, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the show interfaces command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenario, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs</p>	detail extensive

Table 23: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the Dropped packets field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>	detail extensive
Ingress queues	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	extensive
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
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
Interface transmit statistics	<p>(On MX Series devices) Status of the interface-transmit-statistics configuration: Enabled or Disabled.</p> <ul style="list-style-type: none"> • Enabled—When the interface-transmit-statistics statement is included in the configuration. If this is configured, the interface statistics show the actual transmitted load on the interface. • Disabled—When the interface-transmit-statistics statement is not included in the configuration. If this is not configured, the interface statistics show the offered load on the interface. 	detail extensive
OTN FEC statistics	<p>The forward error correction (FEC) counters provide the following statistics:</p> <ul style="list-style-type: none"> • Corrected Errors—Count of corrected errors in the last second. • Corrected Error Ratio—Corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits. 	detail extensive

Table 23: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
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—Number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode.• Errored blocks—Number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.	detail extensive

Table 23: 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 <i>show interfaces (10-Gigabit Ethernet)</i> 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 runs (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. <p>NOTE: The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the VLAN tagged frames field displays 0 when the show interfaces command is executed on a 20-port Gigabit Ethernet MIC. In other words, the number of VLAN tagged frames cannot be determined for the 20-port Gigabit Ethernet MIC.</p> • 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 23: 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 23: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

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

Table 23: 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 23: 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 23: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

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

Table 23: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value. • pop-swap—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none
Demux	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> • Source Family Inet • Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
ACI VLAN: Dynamic Profile	Name of the dynamic profile that defines the agent circuit identifier (ACI) interface set. If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information.	brief detail extensive none
Protocol	Protocol family. Possible values are described in the "Protocol Field" section under <i>Common Output Fields Description</i> .	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Neighbor Discovery Protocol (NDP) Queue Statistics	<p>NDP statistics for protocol inet6 under logical interface statistics.</p> <ul style="list-style-type: none"> • Max nh cache—Maximum interface neighbor discovery nexthop cache size. • New hold nh limit—Maximum number of new unresolved nexthops. • Curr nh cnt—Current number of resolved nexthops in the NDP queue. • Curr new hold cnt—Current number of unresolved nexthops in the NDP queue. • NH drop cnt—Number of NDP requests not serviced. 	All levels
Dynamic Profile	Name of the dynamic profile that was used to create this interface configured with a Point-to-Point Protocol over Ethernet (PPPoE) family.	detail extensive none
Service Name Table	Name of the service name table for the interface configured with a PPPoE family.	detail extensive none

Table 23: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Max Sessions	Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	detail extensive none
Duplicate Protection	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
Direct Connect	State of the configuration to ignore DSL Forum VSAs: On or Off . When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	detail extensive none
AC Name	Name of the access concentrator.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the router.	extensive
Transit statistics	<p>Number and rate of bytes and packets transiting the switch.</p> <p>NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the Output bytes and Output packets interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.</p>	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	detail extensive

Table 23: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

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

Table 24: 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.

Table 24: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type (*continued*)

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
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 include 6 bytes for the destination MAC address plus 4 bytes for VLAN plus 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 (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues    : 4 supported, 4 maximum usable queues
  Current address: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
  Last flapped  : 2006-08-10 17:25:10 PDT (00:01:08 ago)
  Input rate    : 0 bps (0 pps)
  Output rate   : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms : None
  Active defects : None

Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69)
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
  0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  Egress account overhead: 100
  Ingress account overhead: 90
  Input packets : 0
  Output packets: 0
  Protocol ccc, MTU: 1522
  Flags: Is-Primary

```

show interfaces (Gigabit Ethernet on MX Series Routers)

```

user@host> show interfaces ge-2/2/2
Physical interface: ge-2/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 188
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, MAC-REWRITE Error: None,
  Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers    : 0
  Current address: 00:1f:12:b7:d7:c0, Hardware address: 00:1f:12:b7:d6:76
  Last flapped   : 2008-09-05 16:44:30 PDT (3d 01:04 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  Logical interface ge-2/2/2.0 (Index 82) (SNMP ifIndex 219)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Input packets : 10232
    Output packets: 10294
    Protocol inet, MTU: 1500
      Flags: Sendbroadcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
    Protocol inet6, MTU: 1500
      Max nh cache: 4, New hold nh limit: 100000, Curr nh cnt: 4, Curr new hold
      cnt: 4, NH drop cnt: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 2001:db8::/32, Local: 2001:db8::5
      Addresses, Flags: Is-Preferred
        Destination: 2001:db8::/32, Local: 2001:db8:223:9cff:fe9f:3e78
    Protocol multiservice, MTU: Unlimited
      Flags: Is-Primary

```

show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration)

```

user@host> show interfaces ge-2/1/2 extensive | match "output|interface"
Physical interface: ge-2/1/2, Enabled, Physical link is Up
  Interface index: 151, SNMP ifIndex: 530, Generation: 154
  Interface flags: SNMP-Traps Internal: 0x4000
  Output bytes   :          240614363944          772721536 bps
  Output packets:          3538446506          1420444 pps
  Direction     : Output
  Interface transmit statistics: Enabled

  Logical interface ge-2/1/2.0 (Index 331) (SNMP ifIndex 955) (Generation 146)
    Output bytes   :          195560312716          522726272 bps
    Output packets:          4251311146          1420451 pps

```

show interfaces brief (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2 brief
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,

```

Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
 Remote fault: Online
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 Link flags : None

Logical interface ge-3/0/2.0
 Flags: SNMP-Traps 0x4000
 VLAN-Tag [0x8100.512 0x8100.513] In(pop-swap 0x8100.530) Out(swap-push
 0x8100.512 0x8100.513)
 Encapsulation: VLAN-CCC
 ccc

Logical interface ge-3/0/2.32767
 Flags: SNMP-Traps 0x4000 VLAN-Tag [0x0000.0] Encapsulation: ENET2

show interfaces detail (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2 detail
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35, Generation: 177
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags : None
  CoS queues : 4 supported, 4 maximum usable queues
  Hold-times : Up 0 ms, Down 0 ms
  Current address: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
  Last flapped : 2006-08-09 17:17:00 PDT (01:31:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes : 0 0 bps
    Output bytes : 0 0 bps
    Input packets: 0 0 pps
    Output packets: 0 0 pps
  Ingress traffic statistics at Packet Forwarding Engine:
    Input bytes : 0 0 bps
    Input packets: 0 0 pps
    Drop bytes : 0 0 bps
    Drop packets: 0 0 pps
  Ingress queues: 4 supported, 4 in use
  Queue counters:
    Queued packets Transmitted packets Dropped packets

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

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

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



```

3 network-cont                                0                                0                                0

Active alarms : None
Active defects : None

Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69) (Generation 140)
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530)
Out(swap-push 0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  Egress account overhead: 100
  Ingress account overhead: 90
  Traffic statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Local statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Transit statistics:
    Input bytes : 0 0 bps
    Output bytes : 0 0 bps
    Input packets: 0 0 pps
    Output packets: 0 0 pps
  Protocol ccc, MTU: 1522, Generation: 149, Route table: 0
  Flags: Is-Primary

Logical interface ge-3/0/2.32767 (Index 71) (SNMP ifIndex 70)
(Generation 139)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2
  Traffic statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Local statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Transit statistics:
    Input bytes : 0 0 bps
    Output bytes : 0 0 bps
    Input packets: 0 0 pps
    Output packets: 0 0 pps

```

show interfaces extensive (Gigabit Ethernet IQ2)

```

user@host> show interfaces ge-7/1/3 extensive
Physical interface: ge-7/1/3, Enabled, Physical link is Up
Interface index: 170, SNMP ifIndex: 70, Generation: 171
Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4004000
Link flags : None

```

```

CoS queues      : 8 supported, 4 maximum usable queues
Schedulers     : 256
Hold-times      : Up 0 ms, Down 0 ms
Current address: 00:14:f6:30:5e:74, Hardware address: 00:14:f6:30:5e:74
Last flapped    : 2007-11-07 21:31:41 PST (02:03:33 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes :          38910844056          7952 bps
Output bytes :           7174605          8464 bps
Input packets:         418398473           11 pps
Output packets:         78903           12 pps
IPv6 transit statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:              0
Ingress traffic statistics at Packet Forwarding Engine:
Input bytes :         38910799145          7952 bps
Input packets:         418397956           11 pps
Drop bytes :              0           0 bps
Drop packets:              0           0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

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

0 best-effort          418390823          418390823              0

1 expedited-fo              0              0              0

2 assured-forw              0              0              0

3 network-cont           7133           7133              0

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

0 best-effort          1031           1031              0

1 expedited-fo              0              0              0

2 assured-forw              0              0              0

3 network-cont          77872          77872              0

Active alarms : None
Active defects : None
MAC statistics:
Total octets          38910844056          7174605
Total packets         418398473          78903
Unicast packets       408021893366          1026
Broadcast packets           10           12
Multicast packets     418398217          77865
CRC/Align errors              0              0
FIFO errors              0              0

```

```

MAC control frames          0          0
MAC pause frames           0          0
Oversized frames           0
Jabber frames              0
Fragment frames            0
VLAN tagged frames         0
Code violations             0 OTN Received Overhead Bytes:
APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58
Payload Type: 0x08
OTN Transmitted Overhead Bytes:
APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
Payload Type: 0x08
Filter statistics:
  Input packet count        418398473
  Input packet rejects      479
  Input DA rejects         479
  Input SA rejects         0
  Output packet count              78903
  Output packet pad count         0
  Output packet error count      0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
    Link mode: Full-duplex, Flow control: Symmetric/Asymmetric,
    Remote fault: OK
  Local resolution:
    Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 7
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer      Priority      Limit
                           %      bps      %      usec
  0 best-effort           95      950000000  95      0
low  none
  3 network-control       5      500000000  5      0
low  none
  Direction : Input
  CoS transmit queue      Bandwidth      Buffer      Priority      Limit
                           %      bps      %      usec
  0 best-effort           95      950000000  95      0
low  none
  3 network-control       5      500000000  5      0
low  none

Logical interface ge-7/1/3.0 (Index 70) (SNMP ifIndex 85) (Generation 150)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
  Input bytes :      812400
  Output bytes :    1349206
  Input packets:      9429
  Output packets:     9449
IPv6 transit statistics:
  Input bytes :      0
  Output bytes :      0
  Input packets:      0
  Output packets:      0
Local statistics:
  Input bytes :      812400
  Output bytes :    1349206

```

```

Input packets:          9429
Output packets:         9449
Transit statistics:
Input bytes :           0          7440 bps
Output bytes :          0          7888 bps
Input packets:          0          10 pps
Output packets:         0          11 pps
IPv6 transit statistics:
Input bytes :           0
Output bytes :          0
Input packets:          0
Output packets:         0
Protocol inet, MTU: 1500, Generation: 169, Route table: 0
Flags: Is-Primary, Mac-Validate-Strict
Mac-Validate Failures: Packets: 0, Bytes: 0
Addresses, Flags: Is-Preferred Is-Primary
Input Filters: F1-ge-3/0/1.0-in, F3-ge-3/0/1.0-in
Output Filters: F2-ge-3/0/1.0-out (53)
Destination: 10.74.2/24, Local: 10.74.2.2, Broadcast: 10.74.2.255,
Generation: 196
Protocol multiservice, MTU: Unlimited, Generation: 170, Route table: 0
Flags: Is-Primary
Policer: Input: __default_arp_policer__

```

NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics displayed in the **show interfaces** command output might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the interface counters. For detailed information, see the description of the logical interface **Transit statistics** fields in [Table 23 on page 355](#).

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: 192.02.5

```

show interfaces (ACI Interface Set Configured)

```
user@host> show interfaces ge-1/0/0.4001
Logical interface ge-1/0/0.4001 (Index 340) (SNMP ifIndex 548)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4001 ] Encapsulation: PPP-over-
    Ethernet
    ACI VLAN:
      Dynamic Profile: aci-vlan-set-profile
    PPPoE:
      Dynamic Profile: aci-vlan-pppoe-profile,
      Service Name Table: None,
      Max Sessions: 32000, Max Sessions VSA Ignore: Off,
      Duplicate Protection: On, Short Cycle Protection: Off,
      Direct Connect: Off,
      AC Name: nbc
    Input packets : 9
    Output packets: 8
    Protocol multiservice, MTU: Unlimited
```

show interfaces (PPPoE)

Syntax	<pre>show interfaces pp0.logical <brief detail extensive terse> <descriptions> <media> <snmp-index snmp-index> <statistics></pre>
Release Information	Command introduced before Junos OS Release 7.4.
Description	(M120 routers, M320 routers, and MX Series routers only). Display status information about the PPPoE interface.
Options	<p>pp0.logical—Display standard status information about the PPPoE 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 PPPoE interfaces.</p> <p>snmp-index snmp-index—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display PPPoE interface statistics.</p>
Required Privilege Level	view
List of Sample Output	<p>show interfaces (PPPoE) on page 384</p> <p>show interfaces (PPPoE over Aggregated Ethernet) on page 384</p> <p>show interfaces brief (PPPoE) on page 384</p> <p>show interfaces detail (PPPoE) on page 385</p> <p>show interfaces extensive (PPPoE on M120 and M320 Routers) on page 386</p>
Output Fields	Table 25 on page 378 lists the output fields for the show interfaces (PPPoE) command. Output fields are listed in the approximate order in which they appear.

Table 25: show interfaces (PPPoE) Output Fields

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

Table 25: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Type	Physical interface type (PPPoE).	All levels
Link-level type	Encapsulation on the physical interface (PPPoE).	All levels
MTU	MTU size on the physical interface.	All levels
Clocking	Reference clock source. It can be Internal or External .	All levels
Speed	Speed at which the interface is running.	All levels
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> .	All levels
Link type	Physical interface link type: full duplex or half duplex .	All levels
Link flags	Information about the interface. Possible values are described in the "Link Flags" section under <i>Common Output Fields Description</i> .	All levels
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
Physical Info	Physical interface information.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive
Hardware address	MAC address of the hardware.	detail extensive
Alternate link address	Backup address of the link.	detail extensive
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. <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 25: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
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> <p>NOTE: These fields include dropped traffic and exception traffic, as those fields are not separately defined.</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 B chip Tx drops and IXP Tx net transmit drops. 	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), then the cable, the far-end system, or the 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 B chip Tx drops and IXP Tx net 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 none
SNMP ifIndex	Logical interface SNMP interface index number.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Table 25: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
Encapsulation	Type of encapsulation configured on the logical interface.	All levels
PPP parameters	PPP status: <ul style="list-style-type: none"> • LCP restart timer—Length of time (in milliseconds) between successive Link Control Protocol (LCP) configuration requests. • NCP restart timer—Length of time (in milliseconds) between successive Network Control Protocol (NCP) configuration requests. 	detail
PPPoE	PPPoE status: <ul style="list-style-type: none"> • State—State of the logical interface (up or down). • Session ID—PPPoE session ID. • Service name—Type of service required. Can be used to indicate an Internet service provider (ISP) name or a class or quality of service. • Configured AC name—Configured access concentrator name. • Auto-reconnect timeout—Time after which to try to reconnect after a PPPoE session is terminated, in seconds. • Idle Timeout—Length of time (in seconds) that a connection can be idle before disconnecting. • Underlying interface—Interface on which PPPoE is running. 	All levels
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.	All levels
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. This counter usually takes less than 1 second to stabilize.	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. <p>NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p> <ul style="list-style-type: none"> • 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	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. This counter usually takes less than 1 second to stabilize.	detail extensive

Table 25: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transit statistics	<p>Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize.</p> <p>NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p>	detail extensive
Keepalive settings	<p>(PPP and HDLC) Configured settings for keepalives.</p> <ul style="list-style-type: none"> interval seconds—The time in seconds between successive keepalive requests. The range is 10 seconds through 32,767 seconds, with a default of 10 seconds. down-count number—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is 1 through 255, with a default of 3. up-count number—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is 1 through 255, with a default of 1. 	detail extensive
Keepalive statistics	<p>(PPP and HDLC) Information about keepalive packets.</p> <ul style="list-style-type: none"> Input—Number of keepalive packets received by PPP. <ul style="list-style-type: none"> (last seen 00:00:00 ago)—Time the last keepalive packet was received, in the format <i>hh:mm:ss</i>. Output—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received. <ul style="list-style-type: none"> (last seen 00:00:00 ago)—Time the last keepalive packet was sent, in the format <i>hh:mm:ss</i>. <p>(MX Series routers with MPCs/MICs) When an MX Series router with MPCs/MICs is using PPP fast keepalive for a PPP link, the display does not include the number of keepalive packets received or sent, or the amount of time since the router received or sent the last keepalive packet.</p>	detail extensive
Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified
LCP state	<p>(PPP) Link Control Protocol state.</p> <ul style="list-style-type: none"> Conf-ack-received—Acknowledgement was received. Conf-ack-sent—Acknowledgement was sent. Conf-req-sent—Request was sent. Down—LCP negotiation is incomplete (not yet completed or has failed). Not-configured—LCP is not configured on the interface. Opened—LCP negotiation is successful. 	none detail extensive

Table 25: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
NCP state	(PPP) Network Control Protocol state. <ul style="list-style-type: none"> • Conf-ack-received—Acknowledgement was received. • Conf-ack-sent—Acknowledgement was sent. • Conf-req-sent—Request was sent. • Down—NCP negotiation is incomplete (not yet completed or has failed). • Not-configured—NCP is not configured on the interface. • Opened—NCP negotiation is successful. 	detail extensive none
CHAP state	(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction. <ul style="list-style-type: none"> • Chap-Chal-received—Challenge was received but response not yet sent. • Chap-Chal-sent—Challenge was sent. • Chap-Resp-received—Response was received for the challenge sent, but CHAP has not yet moved into the Success state. (Most likely with RADIUS authentication.) • Chap-Resp-sent—Response was sent for the challenge received. • Closed—CHAP authentication is incomplete. • Failure—CHAP authentication failed. • Not-configured—CHAP is not configured on the interface. • Success—CHAP authentication was successful. 	none detail extensive
Protocol	Protocol family configured on the logical interface.	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
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	Routing 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 the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
Addresses, Flags	Information about the addresses configured for the protocol family. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address.	detail extensive none

Sample Output

show interfaces (PPPoE)

```
user@host> show interfaces pp0
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 24
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

Logical interface pp0.0 (Index 72) (SNMP ifIndex 72)
  Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
  PPPoE:
    State: SessionDown, Session ID: None,
    Service name: None, Configured AC name: sapphire,
    Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
    Underlying interface: at-5/0/0.0 (Index 70)
  Input packets : 0
  Output packets: 0
  LCP state: Not-configured
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Closed
    Protocol inet, MTU: 100
    Flags: User-MTU, Negotiate-Address
```

show interfaces (PPPoE over Aggregated Ethernet)

```
user@host> show interfaces pp0.1073773821
Logical interface pp0.1073773821 (Index 80) (SNMP ifIndex 32584)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
  PPPoE:
    State: SessionUp, Session ID: 1,
    Session AC name: alcor, Remote MAC address: 00:10:94:00:00:01,
    Underlying interface: demux0.100 (Index 88)
  Link:
    ge-1/0/0.32767
    ge-1/0/1.32767
  Input packets : 6
  Output packets: 6
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Closed
  PAP state: Success
    Protocol inet, MTU: 1500
    Flags: Sendbroadcast-pkt-to-re
    Addresses, Flags: Is-Primary
    Local: 192.162.22.22
```

show interfaces brief (PPPoE)

```
user@host> show interfaces pp0 brief
Physical interface: pp0, Enabled, Physical link is Up
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532, Speed: Unspecified
  Device flags   : Present Running
```

Interface flags: Point-To-Point SNMP-Traps

Logical interface pp0.0

Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE

PPPoE:

State: SessionDown, Session ID: None,
Service name: None, Configured AC name: sapphire,
Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
Underlying interface: at-5/0/0.0 (Index 70)

inet

show interfaces detail (PPPoE)

user@host> show interfaces pp0 detail

Physical interface: pp0, Enabled, Physical link is Up

Interface index: 128, SNMP ifIndex: 24, Generation: 9

Type: PPPoE, Link-level type: PPPoE, MTU: 1532, 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

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

Logical interface pp0.0 (Index 72) (SNMP ifIndex 72) (Generation 14)

Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE

PPPoE:

State: SessionDown, Session ID: None,
Service name: None, Configured AC name: sapphire,
Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
Underlying interface: at-5/0/0.0 (Index 70)

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

LCP state: Not-configured

NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
mpls: Not-configured

CHAP state: Closed

Protocol inet, MTU: 100, Generation: 14, Route table: 0

Flags: User-MTU, Negotiate-Address

show interfaces extensive (PPPoE on M120 and M320 Routers)

```

user@host> show interfaces pp0 extensive
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 93, Generation: 129
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532, 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
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :          972192          0 bps
    Output bytes  :          975010          0 bps
    Input packets :          1338          0 pps
    Output packets:          1473          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, Runt: 0, Giants: 0, Policed discards:
0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0

Logical interface pp0.0 (Index 69) (SNMP ifIndex 96) (Generation 194)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
  PPPoE:
    State: SessionUp, Session ID: 26,
    Session AC name: None, AC MAC address: 00:17:cb:48:c8:12,
    Service name: None, Configured AC name: None,
    Auto-reconnect timeout: Never, Idle timeout: Never,
    Underlying interface: ge-3/0/1.0 (Index 67)
  Traffic statistics:
    Input bytes   :          252
    Output bytes  :          296
    Input packets :          7
    Output packets:          8
  IPv6 transit statistics:
    Input bytes   :          0
    Output bytes  :          0
    Input packets :          0
    Output packets:          0
  Local statistics:
    Input bytes   :          252
    Output bytes  :          296
    Input packets :          7
    Output packets:          8
  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
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive statistics:
  Input : 1 (last seen 00:00:00 ago)
  Output: 1 (last sent 00:00:03 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Closed
PAP state: Closed
  Protocol inet, MTU: 1492, Generation: 171, Route table: 0
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 192.0.12.2, Local: 192.0.12.1, Broadcast: Unspecified,
Generation: 206
```

show interfaces demux0 (Demux Interfaces)

Syntax	<pre>show interfaces demux0.logical-interface-number <brief detail extensive terse> <descriptions> <media> <snmp-index snmp-index> <statistics></pre>
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 snmp-index—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration
List of Sample Output	<p>show interfaces demux0 (Demux) on page 394</p> <p>show interfaces demux0 (PPPoE over Aggregated Ethernet) on page 395</p> <p>show interfaces demux0 extensive (Targeted Distribution for Aggregated Ethernet Links) on page 396</p> <p>show interfaces demux0 (ACI Interface Set Configured) on page 396</p>
Output Fields	Table 26 on page 388 lists the output fields for the show interfaces demux0 (Demux Interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 26: show interfaces demux0 (Demux Interfaces) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	brief detail extensive none

Table 26: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 <i>Common Output Fields Description</i> .	brief detail extensive none
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 <i>Common Output Fields Description</i> .	brief detail extensive none
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	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 <i>Common Output Fields Description</i> .	detail extensive none

Table 26: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
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. <p>NOTE: These fields include dropped traffic and exception traffic, as those fields are not separately defined.</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 26: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface whose definitions are as follows: <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
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 <i>Common Output Fields Description</i> .	brief detail extensive none
Encapsulation	Encapsulation on the logical interface.	brief extensive none

Table 26: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
ACI VLAN: Dynamic Profile	Name of the dynamic profile that defines the agent circuit identifier (ACI) interface set. If configured, the ACI interface set enables the underlying demux interface to create dynamic VLAN subscriber interfaces based on ACI information.	brief detail 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
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. • IPv6 transit statistics—Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. <p>NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p> <ul style="list-style-type: none"> • 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	Number of transit bytes and packets received and transmitted on the local interface. <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 26: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transit statistics	<p>Number and rate of bytes and packets transiting the switch.</p> <p>NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p> <ul style="list-style-type: none"> • 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> <p>NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p> <ul style="list-style-type: none"> • 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 <i>Common Output Fields Description</i> .	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
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 <i>Common Output Fields Description</i> .	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 <i>Common Output Fields Description</i> .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive statistics none

Table 26: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
Direct Connect	State of the configuration to ignore DSL Forum VSAs: On or Off . When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	detail extensive none
AC Name	Name of the access concentrator.	detail extensive none

Sample Output

show interfaces demux0 (Demux)

```

user@host> show interfaces demux0
Physical interface: demux0, Enabled, Physical link is Up
Interface index: 128, SNMP ifIndex: 79, Generation: 129
Type: Software-Pseudo, Link-level type: Unspecified, MTU: 9192, Clocking: 1,
Speed: Unspecified
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps
Link type      : Full-Duplex
Link flags     : None
Physical info  : Unspecified
Hold-times    : Up 0 ms, Down 0 ms
Current address: Unspecified, Hardware address: Unspecified
Alternate link address: Unspecified
Last flapped   : Never
Statistics last cleared: Never

```

```

Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
Policed discards: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
Resource errors: 0

Logical interface demux0.0 (Index 87) (SNMP ifIndex 84) (Generation 312)
Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
Demux:
Underlying interface: ge-2/0/1.0 (Index 74)
Family Inet Source prefixes, total 1
Prefix: 192.1.1/24
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: 192.1.1/24, Local: 192.1.1.1, Broadcast: 192.1.1.255,
Generation: 434

```

show interfaces demux0 (PPPoE over Aggregated Ethernet)

```

user@host> show interfaces demux0.100
Logical interface demux0.100 (Index 76) (SNMP ifIndex 61160)
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]
Encapsulation: ENET2

```

```
Demux:
  Underlying interface: ae0 (Index 199)
Link:
  ge-1/0/0
  ge-1/1/0
Input packets : 0
Output packets: 0
Protocol pppoe
  Dynamic Profile: pppoe-profile,
  Service Name Table: service-table1,
  Max Sessions: 100, Duplicate Protection: On,
  Direct Connect: Off,
  AC Name: pppoe-server-1
```

show interfaces demux0 extensive (Targeted Distribution for Aggregated Ethernet Links)

```
user@host> show interfaces demux0.1073741824 extensive
```

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

show interfaces demux0 (ACI Interface Set Configured)

```
user@host> show interfaces demux0.1073741827
Logical interface demux0.1073741827 (Index 346) (SNMP ifIndex 527)
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1802 0x8100.302 ] Encapsulation: ENET2
Demux: Source Family Inet
ACI VLAN:
  Dynamic Profile: aci-vlan-set-profile
Demux:
  Underlying interface: ge-1/0/0 (Index 138)
Input packets : 18
Output packets: 16
Protocol inet, MTU: 1500
  Flags: Sendbcst-pkt-to-re, Unnumbered
  Donor interface: lo0.0 (Index 322)
  Preferred source address: 192.0.200.202
  Addresses, Flags: Primary Is-Default Is-Primary
    Local: 10.4.12.119
Protocol pppoe
  Dynamic Profile: aci-vlan-pppoe-profile,
  Service Name Table: None,
  Max Sessions: 32000, Max Sessions VSA Ignore: Off,
  Duplicate Protection: On, Short Cycle Protection: Off,
  Direct Connect: Off,
  AC Name: nbc
```


show interfaces filters

Syntax	show interfaces filters <interface-name>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced on PTX Series Packet Transport Routers 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 <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i> . For related operational mode commands, see the CLI Explorer .
Required Privilege Level	view
List of Sample Output	show interfaces filters on page 398 show interfaces filters (Interface-Name) on page 398 show interfaces filters (PTX Series Packet Transport Routers) on page 398
Output Fields	Table 27 on page 397 lists the output fields for the show interfaces filters command. Output fields are listed in the approximate order in which they appear.

Table 27: 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    inet
ge-0/0/0.0     up    up    iso
ge-5/0/0       up    up
ge-5/0/0.0     up    up    any      f-any
               inet      f-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 (Interface-Name)

```

user@host> show interfaces filters so-2/1/0
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 (PTX Series Packet Transport Routers)

```

user@host > show interfaces filters em0
Interface      Admin Link Proto Input Filter      Output Filter
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 <i>Junos OS Routing Protocols Library for Routing Devices</i> . For information about related operational mode commands for routing instances and protocols, see the CLI Explorer .
Required Privilege Level	view
List of Sample Output	show interfaces routing brief on page 400 show interfaces routing brief (TX Matrix Plus Router) on page 401 show interfaces routing detail on page 401 show interfaces routing detail (TX Matrix Plus Router) on page 402
Output Fields	Table 28 on page 399 lists the output fields for the show interfaces routing command. Output fields are listed in the approximate order in which they appear.

Table 28: 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
RefCount	Number of references to the interface in the routing software.	detail

Table 28: show interfaces routing Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 routing brief

```

user@host> show interfaces 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>
```


show interfaces routing-instance

Syntax	show interfaces routing-instance (<i>instance-name</i> all) <brief detail extensive terse>
Release Information	Command introduced in Junos OS Release 9.1.
Description	Display information about the interfaces configured for either a specific routing instance or for all of the routing instances.
Options	<p>all—Display information about all of the interfaces configured for all of the routing instances on the router.</p> <p><i>instance-name</i>—Display information about the interfaces configured for the specified routing instance.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p>
Required Privilege Level	view
List of Sample Output	show interfaces routing-instance terse on page 405 show interfaces routing-instance all on page 405 show interfaces routing-instance extensive on page 405
Output Fields	The output fields from the show interfaces routing-instance command are identical to those produced by the show interfaces <i>interface-name</i> command. For a description of output fields, see the other chapters in this manual.

Sample Output

show interfaces routing-instance terse

```
user@host> show interfaces routing-instance sample terse
Interface    Admin  Link   Proto  Local          Remote
ge-0/0/0.0   up     up     inet   192.168.4.28/24
```

Sample Output

show interfaces routing-instance all

```
user@host> show interfaces terse routing-instance all
Interface    Admin  Link   Proto  Local          Remote Instance
at-0/0/1     up     up     inet   10.0.0.1/24
ge-0/0/0.0   up     up     inet   192.168.4.28/24      sample-a
at-0/1/0.0   up     up     inet6   fe80::a:0:0:4/64     sample-b
so-0/0/0.0   up     up     inet   10.0.0.1/32
```

show interfaces routing-instance extensive

```
user@host> show interfaces fe-0/1/3 routing-instance instance2 extensive
Logical interface fe-0/1/3.0 (Index 70) (SNMP ifIndex 53) (Generation 211)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
Input bytes : 0
```

```
Output bytes :           42
Input packets:           0
Output packets:          1
IPv6 transit statistics:
  Input bytes :           0
  Output bytes :          0
  Input packets:          0
  Output packets:         0
Local statistics:
  Input bytes :           0
  Output bytes :          42
  Input packets:          0
  Output packets:         1
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: 252, Route table: 4
  Flags: Is-Primary
  Addresses, Flags: Is-Default Is-Preferred Is-Primary
  Destination: 150.1.1/24, Local: 150.1.1.1, Broadcast: 150.1.1.255,
Generation: 263
```

show ppp interface

Syntax	<code>show ppp interface <i>interface-name</i></code> <code><extensive terse></code>
Release Information	Command introduced in Junos OS Release 7.5.
Description	Display information about PPP interfaces.
Options	<i>interface-name</i> —Name of a logical interface. extensive terse —(Optional) Display the specified level of output.
Required Privilege Level	view
List of Sample Output	show ppp interface on page 415 show ppp interface extensive on page 415 show ppp interface terse on page 415
Output Fields	Table 29 on page 407 lists the output fields for the show ppp interface command. Output fields are listed in the approximate order in which they appear.

Table 29: show ppp interface Output Fields

Field Name	Field Description	Level of Output
Session	Name of the logical interface on which the session is running.	All levels
Type	Session type: PPP.	All levels
Phase	PPP process phase: Authenticate , Pending , Establish , LCP , Network , Disabled , and Tunneled .	All levels
Session flags	Special conditions present in the session: Bundled , TCC , No-keepalives , Looped , Monitored , and NCP-only .	All levels
<i>protocol</i> State	Protocol state information. See specific protocol state fields for information.	None specified
AUTHENTICATION	Challenge-Handshake Authentication Protocol (CHAP) authentication state information or Password Authentication Protocol (PAP) state information. See the Authentication field description for further information.	None specified

Table 29: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Keepalive settings	<p>Keepalive settings for the PPP sessions on the L2TP network server (LNS). LNS based PPP sessions are supported only on service interfaces (si).</p> <ul style="list-style-type: none"> • Interval—Time in seconds between successive keepalive requests. Keepalive aging timeout is calculated as a product of the interval and Down-count values. If the keepalive aging timeout is greater than 180 seconds, the keepalive packets are handled by the Routing Engine. If the aging timeout is less than or equal to 180 seconds, the packets are handled by the Packet Forwarding Engine. • Up-count—The number of keepalive packets a destination must receive to change a link's status from down to up. • Down-count—The number of keepalive packets a destination must fail to receive before the network takes down a link. 	extensive
RE Keepalive statistics	<p>Keepalive statistics for the packets handled by the Routing Engine.</p> <ul style="list-style-type: none"> • LCP echo req Tx—LCP echo requests sent from the Routing Engine. • LCP echo req Rx—LCP echo requests received at the Routing Engine. • LCP echo rep Tx—LCP echo responses sent from the Routing Engine. • LCP echo rep Rx—LCP echo responses received at the Routing Engine. • LCP echo req timeout—Number of keepalive packets where the keepalive aging timer has expired. • LCP Rx echo req Magic Num Failures—LCP echo requests where the magic numbers shared between the PPP peers during LCP negotiation did not match. • LCP Rx echo rep Magic Num Failures—LCP echo responses where the magic numbers shared between the PPP peers during LCP negotiation did not match. 	extensive

Table 29: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
LCP	<p>LCP information:</p> <ul style="list-style-type: none"> • State—LCP protocol state (all platforms except M120 and M320 routers): <ul style="list-style-type: none"> • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is not available for traffic. • Opened—Link is administratively available for traffic. • Req-sent—An attempt has been made to configure the connection. • State—LCP protocol state (M120 and M320 routers): <ul style="list-style-type: none"> • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is available (up), but no Open has occurred. • Closing—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Opened—Link is administratively available for traffic. A Configure-Ack has been both sent and received. • Req-sent—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received. • Starting—An administrative Open has been initiated, but the lower layer is still unavailable (Down). • Stopped—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack. • Stopping—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Last started—LCP state start time. • Last completed—LCP state completion time. 	extensive

Table 29: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<ul style="list-style-type: none"> • Negotiated options: <ul style="list-style-type: none"> • ACFC—Address and-Control Field Compression. A configuration option that provides a method to negotiate the compression of the Data Link Layer Address and Control fields. • Asynchronous map—Asynchronous control character map. A configuration option used on asynchronous links such as telephone lines to identify control characters that must be replaced by a two-character sequence to prevent them from being interpreted by equipment used to establish the link. • Authentication protocol—Protocol used for authentication. This option provides a method to negotiate the use of a specific protocol for authentication. It requires a peer to authenticate itself before allowing network-layer protocol packets to be exchanged. By default, authentication is not required. • Authentication algorithm—Type of authentication algorithm. The Message Digest algorithm (MD5) is the only algorithm supported. • Endpoint discriminator class—For multilink PPP (MLPPP), a configuration option that identifies the system transmitting the packet. This option advises a system that the peer on this link could be the same as the peer on another existing link. • Magic number—A configuration option that provides a method to detect looped-back links and other data-link layer anomalies. By default, the magic number is not negotiated. • MRU—Maximum receive unit. A configuration option that may be sent to inform the peer that the implementation can receive larger packets, or to request that the peer send smaller packets. The default value is 1500 octets. • MRRU—For multilink PPP, the maximum receive reconstructed unit. A configuration option that specifies the maximum number of octets in the Information fields of reassembled packets. • Multilink header suspendable classes—For MLPPP, an LCP option that advises the peer that the implementation wishes to receive fragments with a format given by the code number, with the maximum number of suspendable classes given. • Multilink header format classes—For MLPPP, an LCP option that advises the peer that the implementation wishes to receive fragments with a format given by the code number. • PFC—Protocol-Field-Compression. A configuration option that provides a method to negotiate the compression of the PPP Protocol field. • short sequence—For MLPPP, an option that advises the peer that the implementation wishes to receive fragments with short, 12-bit sequence numbers. 	

Table 29: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Authentication	<p>CHAP or PAP authentication state information. For CHAP authentication:</p> <ul style="list-style-type: none"> • Chap-ans-rcvd—Packet was sent from the peer, indicating that the peer received the Chap-resp-sent packet. • Chap-ans-sent—Packet was sent from the authenticator, indicating that the authenticator received the peer's Chap-resp-rcvd packet. • Chap-chal-rcvd—Challenge packet has been received by the peer. • Chap-chal-sent—Challenge packet has been sent by the authenticator to begin the CHAP protocol or has been transmitted at any time during the Network-Layer Protocol (NCP) phase to ensure that the connection has not been altered. • Chap-resp-rcvd—CHAP response packet has been received by the authenticator. • Chap-resp-sent—CHAP response packet has been sent to the authenticator. • Closed—Link is not available for authentication. • Failure—Authenticator compares the response value in the response packet from the peer with its own response value, but the value does not match. Authentication fails. • Success—Authenticator compares the response value in the response packet from the peer with its own response value, and the value matches. Authentication is successful. <p>For PAP authentication:</p> <ul style="list-style-type: none"> • Pap-resp-sent—PAP response sent to peer (ACK/NACK). • Pap-req-rcvd—PAP request packet received from peer. • Pap-resp-rcvd—PAP response received from the peer (ACK/NACK). • Pap-req-sent—PAP request packet sent to the peer. • Closed—Link is not available for authentication. • Failure—Authenticator compares the response value in the response packet from the peer with its own response value, but the value does not match. Authentication fails. • Success—Authenticator compares the response value in the response packet from the peer with its own response value, and the value matches. Authentication is successful. 	None specified

Table 29: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
IPCP	<p>Internet Protocol Control Protocol (IPCP) information.</p> <ul style="list-style-type: none"> • State—(All platforms except M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is not available for traffic. • Opened—Link is administratively available for traffic. • Req-sent—An attempt has been made to configure the connection. • State—(M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is available (up), but no Open has occurred. • Closing—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Opened—Link is administratively available for traffic. A Configure-Ack has been both sent and received. • Req-sent—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received. • Starting—An administrative Open has been initiated, but the lower layer is still unavailable (Down). • Stopped—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack. • Stopping—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Last started—IPCP state start time. • Last completed—IPCP state authentication completion time. • Negotiated options: <ul style="list-style-type: none"> • compression protocol—Negotiate the use of a specific compression protocol. By default, compression is not enabled. • local address—Desired local address of the sender of a Configure-Request. If all four octets are set to zero, the peer provides the IP address. • primary DNS server—Negotiate with the remote peer to select the address of the primary DNS server to be used on the local end of the link. • primary WINS server—Negotiate with the remote peer to select the address of the primary WINS server to be used on the local end of the link. • remote address—IP address of the remote end of the link in dotted quad notation. • secondary DNS server—Negotiate with the remote peer to select the address of the secondary DNS server to be used on the local end of the link. • secondary WINS server—Negotiate with the remote peer to select the address of the secondary WINS server to be used on the local end of the link. • Negotiation mode—PPP Network Control Protocol (NCP) negotiation mode configured for IPCP: Active or Passive 	extensive

Table 29: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
IPV6CP	<p>Internet Protocol version 6 Control Protocol (IPv6CP) information.</p> <ul style="list-style-type: none"> • State—(All platforms except M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is not available for traffic. • Opened—Link is administratively available for traffic. • Req-sent—An attempt has been made to configure the connection. • State—(M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is available (up), but no Open has occurred. • Closing—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Opened—Link is administratively available for traffic. A Configure-Ack has been both sent and received. • Req-sent—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received. • Starting—An administrative Open has been initiated, but the lower layer is still unavailable (Down). • Stopped—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack. • Stopping—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Last started—IPv6CP state start time. • Last completed—IPv6CP state authentication completion time. • Negotiated options: <ul style="list-style-type: none"> • local interface identifier—Desired local address of the sender of a Configure-Request. If all four octets are set to zero, the peer provides the IP address. • remote interface identifier—IP address of the remote end of the link in dotted quad notation. • Negotiation mode—PPP Network Control Protocol (NCP) negotiation mode configured for IPv6CP: Active or Passive 	extensive

Table 29: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
OSINLCP State	<p>OSI Network Layer Control Protocol (OSINLCP) protocol state information (all platforms except M120 and M320 routers):</p> <ul style="list-style-type: none"> • State: <ul style="list-style-type: none"> • Ack-rcvd—Configure-Request has been sent and Configure-Ack has been received. • Ack-sent—Configure-Request and Configure-Ack have both been sent, but Configure-Ack has not yet been received. • Closed—Link is not available for traffic. • Opened—Link is administratively available for traffic. • Req-sent—Attempt has been made to configure the connection. • Last started—OSINLCP state start time. • Last completed—OSINLCP state completion time. 	extensive
TAGCP	<p>TAGCP information.</p> <ul style="list-style-type: none"> • State—(All platforms except M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is not available for traffic. • Opened—Link is administratively available for traffic. • Req-sent—An attempt has been made to configure the connection. • State—(M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is available (up), but no Open has occurred. • Closing—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Opened—Link is administratively available for traffic. A Configure-Ack has been both sent and received. • Req-sent—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received. • Starting—An administrative Open has been initiated, but the lower layer is still unavailable (Down). • Stopped—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack. • Stopping—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Last started—TAGCP state start time. • Last completed—TAGCP state authentication completion time. 	extensive none

Sample Output

show ppp interface

```
user@host> show ppp interface si-1/3/0.0
Session si-1/3/0.0, Type: PPP, Phase: Authenticate
Session flags: Monitored
LCP State: Opened
AUTHENTICATION: CHAP State: Chap-resp-sent, Chap-ans-sent
IPCP State: Closed, OSINLCP State: Closed
```

show ppp interface extensive

```
user@host> show ppp interface si-0/0/3.0 extensive

Session si-0/0/3.0, Type: PPP, Phase: Network
Keepalive settings: Interval 30 seconds, Up-count 1, Down-count 3
RE Keepalive statistics:
LCP echo req Tx      : 657 (last sent 00:50:10 ago)
LCP echo req Rx      : 0 (last seen: never)
LCP echo rep Tx      : 0
LCP echo rep Rx      : 657
LCP echo req timeout : 0
LCP Rx echo req Magic Num Failures : 0
LCP Rx echo rep Magic Num Failures : 0
LCP
State: Opened
Last started: 2007-01-29 10:43:50 PST
Last completed: 2007-01-29 10:43:50 PST
Negotiated options:
Authentication protocol: PAP, Magic number: 2341124815, MRU: 4470
Authentication: PAP
State: Success
Last started: 2007-01-29 10:43:50 PST
Last completed: 2007-01-29 10:43:50 PST
IPCP
State: Opened
Last started: 2007-01-29 10:43:50 PST
Last completed: 2007-01-29 10:43:50 PST
Negotiated options:
Local address: 10.10.10.1, Remote address: 10.10.10.2
Negotiation mode: Active
IPV6CP
State: Opened
Last started: 2007-01-29 10:43:50 PST
Last completed: 2007-01-29 10:43:50 PST
Negotiated options:
Local interface identifier: 2a0:a522:64:d319, Remote interface identifier: 0:0:0:c
Negotiation mode: Passive
```

show ppp interface terse

```
user@host> show ppp interface si-1/3/0 terse
Session name  Session type  Session phase  Session flags
si-1/3/0.0    PPP           Authenticate   Monitored
```

Subscriber Commands

- `show subscribers`
- `show subscribers summary`

show subscribers

Syntax show subscribers
 <detail | extensive | terse>
 <aci-interface-set-name *aci-interface-set-name*>
 <address *address*>
 <agent-circuit-identifier *agent-circuit-identifier-substring*>
 <client-type *client-type*>
 <count>
 <id>
 <interface *interface*>
 <logical-system *logical-system*>
 <mac-address *mac-address*>
 <physical-interface *physical-interface-name*>
 <profile-name *profile-name*>
 <routing-instance *routing-instance*>
 <stacked-vlan-id *stacked-vlan-id*>
 <subscriber-state *subscriber-state*>
 <user-name *user-name*>
 <vci *vci-identifier*>
 <vpi *vpi-identifier*>
 <vlan-id *vlan-id*>

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**, and **extensive** 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.
 Options **aci-interface-set-name** and **agent-circuit-identifier** introduced in Junos OS Release 12.2.
 The **physical-interface** and **user-name** options introduced in Junos OS Release 12.3.
 Options **vci** and **vpi** introduced in Junos OS Release 12.3R3 and supported in later 12.3Rx releases.
 Options **vci** and **vpi** supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
 Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description Display information for active subscribers.

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

aci-interface-set-name—(Optional) Display all dynamic subscriber sessions that use the specified agent circuit identifier (ACI) interface set. Use the ACI interface set name generated by the router, such as aci-1003-ge-1/0/0.4001, and not the actual ACI value found in the DHCP or PPPoE control packets.

address—(Optional) Display subscribers whose IP address matches the specified address. You must specify the IPv4 or IPv6 address prefix without a netmask (for example,

192.168.17.1). If you specify the IP address as a prefix with a netmask (for example, 192.168.17.1/32), the router displays a message that the IP address is invalid, and rejects the command.

agent-circuit-identifier-substring—(Optional) Display all dynamic subscriber sessions whose ACI value matches the specified substring.

client-type—(Optional) Display subscribers whose client type matches one of the following client types:

- ***dhcp***—DHCP clients only.
- ***dot1x***—Dot1x clients only.
- ***essm***—ESSM clients only.
- ***fwauth***—FwAuth (authenticated across a firewall) clients only.
- ***l2tp***—L2TP clients only.
- ***mlppp***—MLPPP clients only.
- ***ppp***—PPP clients only.
- ***pppoe***—PPPoE clients only.
- ***static***—Static clients only.
- ***vlan***—VLAN clients only.
- ***vlan-oob***—VLAN out-of-band (ANCP-triggered) clients only.
- ***vpls-pw***—VPLS pseudowire clients only.
- ***xauth***—Xauth clients only.

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.

physical-interface-name—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose physical interface matches the specified physical interface.

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.

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

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

user-name—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose username matches the specified subscriber name.

vci-identifier—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual circuit identifier (VCI) matches the specified VCI identifier. The range of values is 0 through 255.

vpi-identifier—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual path identifier (VPI) matches the specified VPI identifier. The range of values is 0 through 65535.

vlan-id—(Optional) Display subscribers whose VLAN ID matches the specified VLAN ID, regardless of whether the subscriber uses a single-tagged or double-tagged VLAN. For subscribers using a double-tagged VLAN, this option displays subscribers where the inner VLAN tag matches the specified VLAN ID. To display only subscribers where the specified value matches only double-tagged VLANs, use the **stacked-vlan-id** option to match the outer VLAN tag.



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

Required Privilege Level

view

Related Documentation

- [show subscribers summary on page 437](#)
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*
- *Verifying and Managing Junos OS Enhanced Subscriber Management*

List of Sample Output

[show subscribers \(IPv4\) on page 424](#)
[show subscribers \(IPv6\) on page 424](#)
[show subscribers \(IPv4 and IPv6 Dual Stack\) on page 424](#)
[show subscribers \(LNS on MX Series Routers\) on page 425](#)
[show subscribers \(L2TP Switched Tunnels\) on page 425](#)
[show subscribers client-type dhcp detail on page 425](#)
[show subscribers client-type vlan-oob detail on page 425](#)
[show subscribers count on page 426](#)
[show subscribers address detail \(IPv6\) on page 426](#)

[show subscribers detail \(IPv4\) on page 426](#)
[show subscribers detail \(IPv6\) on page 427](#)
[show subscribers detail \(IPv6 Static Demux Interface\) on page 427](#)
[show subscribers detail \(L2TP LNS Subscribers on MX Series Routers\) on page 427](#)
[show subscribers detail \(L2TP Switched Tunnels\) on page 427](#)
[show subscribers detail \(Tunneled Subscriber\) on page 428](#)
[show subscribers detail \(IPv4 and IPv6 Dual Stack\) on page 428](#)
[show subscribers detail \(ACI Interface Set Session\) on page 429](#)
[show subscribers detail \(PPPoE Subscriber Session with ACI Interface Set\) on page 429](#)
[show subscribers extensive on page 430](#)
[show subscribers extensive \(RPF Check Fail Filter\) on page 430](#)
[show subscribers extensive \(L2TP LNS Subscribers on MX Series Routers\) on page 430](#)
[show subscribers extensive \(IPv4 and IPv6 Dual Stack\) on page 431](#)
[show subscribers extensive \(ADF Rules \) on page 432](#)
[show subscribers extensive \(Effective Shaping-Rate\) on page 432](#)
[show subscribers aci-interface-set-name detail \(Subscriber Sessions Using Specified ACI Interface Set\) on page 432](#)
[show subscribers agent-circuit-identifier detail \(Subscriber Sessions Using Specified ACI Substring\) on page 433](#)
[show subscribers interface extensive on page 433](#)
[show subscribers logical-system terse on page 434](#)
[show subscribers physical-interface count on page 434](#)
[show subscribers routing-instance inst1 count on page 434](#)
[show subscribers stacked-vlan-id detail on page 434](#)
[show subscribers stacked-vlan-id vlan-id detail \(Combined Output\) on page 434](#)
[show subscribers stacked-vlan-id vlan-id interface detail \(Combined Output for a Specific Interface\) on page 435](#)
[show subscribers user-name detail on page 435](#)
[show subscribers vlan-id on page 435](#)
[show subscribers vlan-id detail on page 435](#)
[show subscribers vpi vci extensive \(PPPoE-over-ATM Subscriber Session\) on page 436](#)
[show subscribers address detail \(Enhanced Subscriber Management\) on page 436](#)

Output Fields [Table 30 on page 420](#) lists the output fields for the **show subscribers** command. Output fields are listed in the approximate order in which they appear.

Table 30: show subscribers Output Fields

Field Name	Field Description
Interface	<p>Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface.</p> <p>The * character indicates a continuation of addresses for the same session.</p>
IP Address/VLAN ID	<p>Subscriber IP address or VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i></p> <p>No IP address or VLAN ID is assigned to an L2TP tunnel-switched session. For these subscriber sessions the value is Tunnel-switched.</p>
User Name	Name of subscriber.

Table 30: show subscribers Output Fields (*continued*)

Field Name	Field Description
LS:RI	Logical system and routing instance associated with the subscriber.
Type	Subscriber client type (DHCP, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN).
IP Address	Subscriber IPv4 address.
IP Netmask	Subscriber IP netmask.
Primary DNS Address	IP address of primary DNS server.
Secondary DNS Address	IP address of secondary DNS server.
Primary WINS Address	IP address of primary WINS server.
Secondary WINS Address	IP address of secondary WINS server.
IPv6 Address	Subscriber IPv6 address, or multiple addresses.
IPv6 Prefix	Subscriber IPv6 prefix. If you are using DHCPv6 prefix delegation, this is the delegated prefix.
IPv6 User Prefix	IPv6 prefix obtained through ND/RA.
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	(Enhanced subscriber management for MX Series routers) Name of the enhanced subscriber management logical interface, in the form demux0.nnnn (for example, demux0.3221225472), to which access-internal and framed subscriber routes are mapped.
Interface Type	Whether the subscriber interface is Static or Dynamic .
Interface Set	Internally generated name of the dynamic ACI interface set used by the subscriber session.
Interface Set Type	Interface type of the ACI interface set: Dynamic . This is the only ACI interface set type currently supported.
Interface Set Session ID	Identifier of the dynamic ACI interface set entry in the session database.

Table 30: show subscribers Output Fields (*continued*)

Field Name	Field Description
Underlying Interface	Name of the underlying interface for the subscriber session.
Dynamic Profile Name	Dynamic profile used for the subscriber.
Dynamic Profile Version	Version number of the 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).
L2TP State	Current state of the L2TP session, Tunneled or Tunnel-switched . When the value is Tunnel-switched , two entries are displayed for the subscriber; the first entry is at the LNS interface on the LTS and the second entry is at the LAC interface on the LTS.
Tunnel switch Profile Name	Name of the L2TP tunnel switch profile that initiates tunnel switching.
Local IP Address	IP address of the local gateway (LAC).
Remote IP Address	IP address of the remote peer (LNS).
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	<p>For the dhcp client type, 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.</p> <p>For the vlan-oob client type, the agent circuit ID or access-loop circuit identifier that identifies the subscriber line based on the subscriber-facing DSLAM interface on which the subscriber request originates.</p>
Agent Remote ID	<p>For the dhcp client type, 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.</p> <p>For the vlan-oob client type, the agent remote ID or access-loop remote identifier that identifies the subscriber line based on the NAS-facing DSLAM interface on which the subscriber request originates.</p>
DHCP Relay IP Address	IP address used by the DHCP relay agent.
ATM VPI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual path identifier (VPI) on the subscriber's physical interface.
ATM VCI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual circuit identifier (VCI) for each VPI configured on the subscriber interface.

Table 30: show subscribers Output Fields (*continued*)

Field Name	Field Description
Login Time	Date and time at which the subscriber logged in.
Effective shaping-rate	Actual downstream traffic shaping rate for the subscriber, in kilobits per second.
IPv4 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv4 packets that fail the RPF check.
IPv6 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv6 packets that fail the RPF check.
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.
Underlying Session ID	For DHCPv6 subscribers on a PPPoE network, displays the session ID of the underlying PPPoE interface.
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.
IPv6 Delegated Address Pool	Name of the pool used for DHCPv6 prefix delegation.
IPv6 Delegated Network Prefix Length	Length of the prefix configured for the IPv6 delegated address pool.
IPv6 Interface Address	Address assigned by the Framed-Ipv6-Prefix AAA attribute.
IPv6 Framed Interface Id	Interface ID assigned by the Framed-Interface-Id AAA attribute.
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.

Table 30: show subscribers Output Fields (*continued*)

Field Name	Field Description
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).

Sample Output

show subscribers (IPv4)

```

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   10.0.0.10           RETAILER1-CLIENT  test1:retailer1
demux0.1073741825   192.3.0.3           RETAILER1-CLIENT  test1:retailer1
demux0.1073741826   198.53.102.3        RETAILER2-CLIENT  test1:retailer2

```

show subscribers (IPv6)

```

user@host> show subscribers
Interface          IP Address/VLAN ID  User Name          LS:RI
ge-1/0/0.0         2001:db8::c0:0:0:0/74 WHOLESALER-CLIENT default:default
*                  2001:db8::1/128     subscriber-25      default:default

```

show subscribers (IPv4 and IPv6 Dual Stack)

```

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     192.168.1.1        dualstackuser1@EXAMPLE1.com
default:ASP-1
*                  2001:db8:1::/48
*                  2001:db8:1:1::/64
pp0.1073741837     192.168.1.3        dualstackuser2@EXAMPLE1.com
default:ASP-1
*                  2001:db8:1:2:5::/64

```

show subscribers (LNS on MX Series Routers)

```

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

```

show subscribers (L2TP Switched Tunnels)

```

user@host> show subscribers
Interface          IP Address/VLAN ID  User Name      LS:RI
si-2/1/0.1073741842 Tunnel-switched    ap@example.com  default:default

si-2/1/0.1073741843 Tunnel-switched    ap@example.com  default:default

```

show subscribers client-type dhcp detail

```

user@host> show subscribers client-type dhcp detail
Type: DHCP
IP Address: 192.20.9.7
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: 10.20.10.7
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 client-type vlan-oob detail

```

user@host> show subscribers client-type vlan-oob detail
Type: VLAN-OOB
User Name: L2WS.line-aci-1.line-ari-1
Logical System: default
Routing Instance: ISP1
Interface: demux0.1073744127
Interface type: Dynamic
Underlying Interface: ge-1/0/0
Dynamic Profile Name: Prof_L2WS
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 2304
Session ID: 77
VLAN Id: 126

```

Core-Facing Interface: ge-2/1/1
VLAN Map Id: 6
Inner VLAN Map Id: 2001
Agent Circuit ID: line-aci-1
Agent Remote ID: line-ari-1
Login Time: 2013-10-29 14:43:52 EDT

show subscribers count

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

show subscribers address detail (IPv6)

```
user@host> show subscribers address 10.16.12.137 detail
Type: PPPoE
User Name: pppoeTerV6User1Svc
IP Address: 10.16.12.137
IP Netmask: 255.0.0.0
IPv6 User Prefix: 2001:db8:0:c88::/32
Logical System: default
Routing Instance: default
Interface: pp0.1073745151
Interface type: Dynamic
Underlying Interface: demux0.8201
Dynamic Profile Name: pppoe-client-profile
MAC Address: 00:0d:02:01:00:01
Session Timeout (seconds): 31622400
Idle Timeout (seconds): 86400
State: Active
Radius Accounting ID: jnpr demux0.8201:6544
Session ID: 6544
Agent Circuit ID: if13720
Agent Remote ID: if13720
Login Time: 2012-05-21 13:37:27 PDT
Service Sessions: 1
```

show subscribers detail (IPv4)

```
user@host> show subscribers detail
Type: DHCP
IP Address: 10.20.9.7
IP Netmask: 255.255.0.0
Primary DNS Address: 192.168.17.1
Secondary DNS Address: 192.168.17.2
Primary WINS Address: 192.168.22.1
Secondary WINS Address: 192.168.22.2
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
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 detail (IPv6)

```
user@host> show subscribers detail
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2001:db8: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 detail (IPv6 Static Demux Interface)

```
user@host> show subscribers detail
Type: STATIC-INTERFACE
User Name: demux0.1@example.net
IPv6 Prefix: 2001:db8:3:4:5:6:7:aa/32
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 detail (L2TP LNS Subscribers on MX Series Routers)

```
user@host> show subscribers detail
Type: L2TP
User Name: user1@example.net
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 (L2TP Switched Tunnels)

```
user@host> show subscribers detail
Type: L2TP
User Name: ap@example.com
Logical System: default
```

```
Routing Instance: default
Interface: si-2/1/0.1073741842
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 10.50.1.1
Remote IP Address: 192.168.20.3
Radius Accounting ID: 21
Session ID: 21
Login Time: 2013-01-18 03:01:11 PST
```

```
Type: L2TP
User Name: ap@example.com
Logical System: default
Routing Instance: default
Interface: si-2/1/0.1073741843
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 10.30.1.1
Remote IP Address: 192.20.1.10
Session ID: 22
Login Time: 2013-01-18 03:01:14 PST
```

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 detail (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@EXAMPLE1.com
IP Address: 10.1.1.1
IPv6 Prefix: 2001:db8:1::/32
IPv6 User Prefix: 2001:db8:1:1::/32
Logical System: default
Routing Instance: ASP-1
```



```

Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST

Type: DHCP
IPv6 Prefix: 2001:db8:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: jnpr :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00

```

show subscribers detail (ACI Interface Set Session)

```

user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0
Interface Set: aci-1001-ge-1/0/0.2800
Interface Set Session ID: 0
Underlying Interface: ge-1/0/0.2800
Dynamic Profile Name: aci-vlan-set-profile-2
Dynamic Profile Version: 1
State: Active
Session ID: 1
Agent Circuit ID: aci-ppp-dhcp-20
Login Time: 2012-05-26 01:54:08 PDT

```

show subscribers detail (PPPoE Subscriber Session with ACI Interface Set)

```

user@host> show subscribers detail
Type: PPPoE
User Name: ppphint2
IP Address: 10.10.1.5
Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Dynamic
Interface Set: aci-1001-demux0.1073741824
Interface Set Type: Dynamic
Interface Set Session ID: 2
Underlying Interface: demux0.1073741824
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:64:39:01:02

```

```
State: Active
Radius Accounting ID: 3
Session ID: 3
Agent Circuit ID: aci-ppp-dhcp-dvlan-50
Login Time: 2012-03-07 13:46:53 PST
```

show subscribers extensive

```
user@host> show subscribers extensive
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2001:db8:db2:ffff:1::/32
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
IPv6 Address Pool: pd_pool
IPv6 Network Prefix Length: 48
```

show subscribers extensive (RPF Check Fail Filter)

```
user@host> show subscribers extensive
...
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ae0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof
State: Active
Session ID: 9
VLAN Id: 100
Login Time: 2011-08-26 08:17:00 PDT
IPv4 rpf-check Fail Filter Name: rpf-allow-dhcp
IPv6 rpf-check Fail Filter Name: rpf-allow-dhcpv6
...
```

show subscribers extensive (L2TP LNS Subscribers on MX Series Routers)

```
user@host> show subscribers extensive
Type: L2TP
User Name: user1@example.net
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 extensive (IPv4 and IPv6 Dual Stack)

```

user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@EXAMPLE1.com
IP Address: 192.1.1.1
IPv6 Prefix: 2001:db8:1::/32
IPv6 User Prefix: 2001:db8:1:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST
IPv6 Delegated Network Prefix Length: 48
IPv6 Interface Address: 2001:db8:1:1::1/32
IPv6 Framed Interface Id: 10:1:2:2
IPv4 Input Filter Name: FILTER-IN-pp0.1073741825-in
IPv4 Output Filter Name: FILTER-OUT-pp0.1073741825-out
IPv6 Input Filter Name: FILTER-IN6-pp0.1073741825-in
IPv6 Output Filter Name: FILTER-OUT6-pp0.1073741825-out

Type: DHCP
IPv6 Prefix: 2001:db8:1:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: jnpr :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
IPv6 Delegated Network Prefix Length: 48

```

show subscribers extensive (ADF Rules)

```
user@host> show subscribers extensive
...
Service Session ID: 12
Service Session Name: SERVICE-PROFILE
State: Active
Family: inet
  ADF IPv4 Input Filter Name: __junos_adf_12-demux0.3221225474-inet-in
    Rule 0: 010101000b0101020b020200201811
      from {
        source-address 10.1.1.2/32;
        destination-address 10.2.2.0/24;
        protocol 17;
      }
      then {
        accept;
      }
    }
```

show subscribers extensive (Effective Shaping-Rate)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741837
Interface type: Dynamic
Interface Set: ifset-1
Underlying Interface: ae1
Dynamic Profile Name: svlan-dhcp-test
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.201
VLAN Id: 0x8100.201
Login Time: 2011-11-30 00:18:04 PST
Effective shaping-rate: 31000000k
...
```

show subscribers aci-interface-set-name detail (Subscriber Sessions Using Specified ACI Interface Set)

```
user@host> show subscribers aci-interface-set-name aci-1003-ge-1/0/0.4001 detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
User Name: ppphint2
IP Address: 10.10.1.7
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
```

```

Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:65:26:01:02
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT

```

show subscribers agent-circuit-identifier detail (Subscriber Sessions Using Specified ACI Substring)

```

user@host> show subscribers agent-circuit-identifier aci-ppp-vlan detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
User Name: ppphint2
IP Address: 10.10.1.7
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:65:26:01:02
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT

```

show subscribers interface extensive

```

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: 192.168.200.6
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 terse

```
user@host> show subscribers logical-system test1 terse
Interface          IP Address/VLAN ID  User Name          LS:RI
demux0.1073741825  10.0.0.3            RETAILER1-CLIENT  test1:retailer1
demux0.1073741826  10.0.0.6            RETAILER2-CLIENT  test1:retailer2
```

show subscribers physical-interface count

```
user@host> show subscribers physical-interface ge-1/0/0 count
Total subscribers: 3998, Active Subscribers: 3998
```

show subscribers routing-instance inst1 count

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

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-name detail

```
user@host> show subscribers user-name larry1 detail
Type: DHCP
User Name: larry1
IP Address: 10.0.0.37
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.1
Interface type: Static
Dynamic Profile Name: foo
MAC Address: 00:10:94:00:00:01
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-11-07 08:25:59 PST
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 01 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 32 2f
37 2d 30 2d 30 37 05 01 06 0f 21 2c
```

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 vpi vci extensive (PPPoE-over-ATM Subscriber Session)

```
user@host> show subscribers vpi 40 vci 50 extensive
Type: PPPoE
User Name: testuser
IP Address: 10.0.0.2
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: pp0.0
Interface type: Static
MAC Address: 00:00:65:23:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
ATM VPI: 40
ATM VCI: 50
Login Time: 2012-12-03 07:49:26 PST
IP Address Pool: pool_1
IPv6 Framed Interface Id: 200:65ff:fe23:102
```

show subscribers address detail (Enhanced Subscriber Management)

```
user@host> show subscribers address 100.20.0.111 detail
Type: DHCP
User Name: simple_filters_service
IP Address: 10.0.0.2
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: demux0.3221225482
Interface type: Dynamic
Underlying Interface: demux0.3221225472
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:01:02:03:04:0f
State: Active
Radius Accounting ID: 11
Session ID: 11
PFE Flow ID: 15
Stacked VLAN Id: 210
VLAN Id: 209
Login Time: 2014-03-24 12:53:48 PDT
Service Sessions: 1
DHCP Options: len 3
35 01 01
```


show subscribers summary

Syntax	<pre>show subscribers summary <all> <detail extensive terse> <count> <physical-interface <i>physical-interface-name</i>> <logical-system <i>logical-system</i> pic port routing-instance <i>routing-instance</i> slot></pre>
Release Information	Command introduced in Junos OS Release 10.2.
Description	Display summary information for subscribers.
Options	<p>all—(Optional) Display full subscriber summary.</p> <p>detail extensive terse—(Optional) Display the specified level of output.</p> <p>count—(Optional) Display the count of total subscribers and active subscribers for any specified option.</p> <p>logical-system—(Optional) Display subscribers whose logical system matches the specified logical system.</p> <p>physical-interface-name—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers whose physical interface matches the specified physical interface, by subscriber state, client type and LS:RI.</p> <p>pic—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by PIC number and the total number of subscribers.</p> <p>port—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by port number and the total number of subscribers.</p> <p>routing-instance—(Optional) Display subscribers whose routing instance matches the specified routing instance.</p> <p>slot—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by FPC slot number and the total number of subscribers.</p>
	<div>  <p>NOTE: Due to display limitations, logical system and routing instance output values are truncated when necessary.</p> </div>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show subscribers on page 417
List of Sample Output	show subscribers summary on page 439

[show subscribers summary all on page 439](#)
[show subscribers summary physical-interface on page 440](#)
[show subscribers summary physical-interface pic on page 440](#)
[show subscribers summary physical-interface port on page 440](#)
[show subscribers summary physical-interface slot on page 440](#)
[show subscribers summary pic on page 441](#)
[show subscribers summary pic \(Aggregated Ethernet Interfaces\) on page 441](#)
[show subscribers summary port on page 441](#)
[show subscribers summary port extensive on page 441](#)
[show subscribers summary slot on page 441](#)
[show subscribers summary terse on page 442](#)

Output Fields [Table 31 on page 438](#) lists the output fields for the **show subscribers summary** command. Output fields are listed in the approximate order in which they appear.

Table 31: show subscribers summary Output Fields

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. • Total—Total number of subscribers for all states.
Subscribers by Client Type	<p>Number of subscribers summarized by client type. Client types can include DHCP, L2TP, PPP, PPPOE, STATIC-INTERFACE, VLAN, and VLAN-OOB. Also displays the total number of subscribers for all client types (Total).</p>
Subscribers by LS:RI	<p>Number of subscribers summarized by logical system:routing instance (LS:RI) combination. Also displays the total number of subscribers for all LS:RI combinations (Total).</p>
Subscribers by Connection Type	<p>Number of subscribers summarized by connection type, Cross-connected or Terminated.</p>
Interface	<p>Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface.</p> <p>The * character indicates a continuation of addresses for the same session.</p> <p>For aggregated Ethernet interfaces, the output of the summary (pic port slot) options prefixes the interface name with ae0:.</p>
Count	<p>Count of subscribers displayed for each PIC, port, or slot when those options are specified with the summary option. For an aggregated Ethernet configuration, the total subscriber count does not equal the sum of the individual PIC, port, or slot counts, because each subscriber can be in more than one aggregated Ethernet link.</p>
Total Subscribers	<p>Total number of subscribers for all physical interfaces, all PICS, all ports, or all LS:RI slots.</p>
IP Address/VLAN ID	<p>Subscriber IP address or VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i></p>

Table 31: show subscribers summary Output Fields (*continued*)

Field Name	Field Description
User Name	Name of subscriber.
LS:RI	Logical system and routing instance associated with the subscriber.

Sample Output

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
VLAN-OOB   2
TOTAL     193
```

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 summary physical-interface

```
user@host> show subscribers summary physical-interface ge-1/0/0
Subscribers by State
  Active: 3998
  Total: 3998

Subscribers by Client Type
  DHCP: 3998
  Total: 3998

Subscribers by LS:RI
  default:default: 3998
  Total: 3998
```

show subscribers summary physical-interface pic

```
user@host> show subscribers summary physical-interface ge-0/2/0 pic
Subscribers by State
  Active: 4825
  Total: 4825

Subscribers by Client Type
  DHCP: 4825
  Total: 4825

Subscribers by LS:RI
  default:default: 4825
  Total: 4825
```

show subscribers summary physical-interface port

```
user@host> show subscribers summary physical-interface ge-0/3/0 port
Subscribers by State
  Active: 4825
  Total: 4825

Subscribers by Client Type
  DHCP: 4825
  Total: 4825

Subscribers by LS:RI
  default:default: 4825
  Total: 4825
```

show subscribers summary physical-interface slot

```
user@host> show subscribers summary physical-interface ge-2/0/0 slot
Subscribers by State
  Active: 4825
  Total: 4825

Subscribers by Client Type
  DHCP: 4825
  Total: 4825

Subscribers by LS:RI
  default:default: 4825
  Total: 4825
```

show subscribers summary pic

```

user@host> show subscribers summary pic
Interface          Count
ge-1/0             1000
ge-1/3             1000

Total Subscribers: 2000

```

show subscribers summary pic (Aggregated Ethernet Interfaces)

```

user@host> show subscribers summary pic
Interface          Count
ae0: ge-1/0        801
ae0: ge-1/3        801

Total Subscribers: 801

```

show subscribers summary port

```

user@host> show subscribers summary port
Interface          Count
ge-1               2000

Total Subscribers: 2000

```

show subscribers summary port extensive

```

user@host> show subscribers summary port extensive
Interface: ge-5/0/1
Count: 201
Detail:
Subscribers by Client Type
  DHCP: 100
  PPPoE: 100
  VLAN-00B: 1
Subscribers by Connection Type
  Terminated: 200
  Cross-connected: 1

Interface: ge-5/0/2
Count: 301
Detail:
Subscribers by Client Type
  DHCP: 200
  PPPoE: 100
  VLAN-00B: 1
Subscribers by Connection Type
  Terminated: 300
  Cross-connected: 1

Total Subscribers: 502

```

show subscribers summary slot

```

user@host> show subscribers summary slot
Interface          Count
ge-1               2000

Total Subscribers: 2000

```

show subscribers summary terse

```
user@host> show subscribers summary terse
Interface      IP Address/VLAN ID  User Name      LS:RI
ge-1/3/0.1073741824  100                WHOLESALE-CLIENT default:default
demux0.1073741824    100.0.0.10         RETAILER1-CLIENT test1:retailer1
demux0.1073741825    101.0.0.3          RETAILER2-CLIENT test1:retailer2
demux0.1073741826    102.0.0.3          RETAILER2-CLIENT test1:retailer2
```

VPLS Commands

- [show vpls connections](#)
- [show vpls flood event-queue](#)
- [show vpls flood instance](#)
- [show vpls flood route](#)
- [show vpls mac-table](#)
- [show vpls statistics](#)

show vpls connections

Syntax	<pre>show vpls connections <brief extensive> <down up up-down> <history> <instance <i>instance-name</i> local-site <i>local-site-name</i> remote-site <i>remote-site-name</i>> <instance-history> <logical-system (all <i>logical-system-name</i>)> <status> <summary></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>instance-history option introduced in Junos OS Release 12.3R2.</p>
Description	(T Series and M Series routers, except for the M160 router) Display virtual private LAN service (VPLS) connection information.
Options	<p>none—Display information about all VPLS connections for all routing instances.</p> <p>brief extensive—(Optional) Display the specified level of output.</p> <p>down up up-down—(Optional) Display nonoperational, operational, or both types of connections.</p> <p>history—(Optional) Display information about connection history.</p> <p>instance <i>instance-name</i>—(Optional) Display the VPLS connections for the specified routing instance only.</p> <p>instance-history—(Optional) Display information about connection history for a particular instance.</p> <p>local-site <i>local-site-name</i>—(Optional) Display the VPLS connections for the specified local site name or ID only.</p> <p>remote-site <i>remote-site-name</i>—(Optional) Display the VPLS connections for the specified remote site name or ID only. Label block size information is always shown as 0 when using this option.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>status—(Optional) Display information about the connection and interface status.</p> <p>summary—(Optional) Display summary of all VPLS connections information.</p>
Required Privilege Level	view
List of Sample Output	<p>show vpls connections on page 448</p> <p>show vpls connections (with multiple pseudowires) on page 450</p>

[show vpls connections extensive \(Static VPLS Neighbors\) on page 451](#)

Output Fields [Table 32 on page 444](#) lists the output fields for the **show vpls connections** command. Output fields are listed in the approximate order in which they appear.

Table 32: show vpls connections Output Fields

Field Name	Field Description
Instance	Name of the VPLS instance.
Local site	Name of the local site.
VPLS-id	Identifier for the VPLS site.
Number of local interfaces	Number of interfaces configured for the local site.
Number of local interfaces up	Number of interfaces configured for the local site that are currently up.
IRB interface present	Indicates whether or not an integrated routing and bridging (IRB) interface is present (yes or no).
Intf	<p>List of all of the interfaces configured for the local site. The types of interfaces can include VPLS virtual loopback tunnel interfaces and label-switched interfaces. Any interface that supports VPLS could be listed here.</p> <p>Virtual loopback tunnel interfaces are displayed using the vt-fpc/pic/port.nnnnn format. Label-switched interfaces are displayed using the lsi.nnnnn format. In both cases, nnnnn is a dynamically generated virtual port used to transport and receive packets from other provider edge (PE) routers in the VPLS domain.</p> <p>Each interface might include the following information:</p> <ul style="list-style-type: none"> • Identification as a VPLS interface • Name of the associated VPLS routing instance • Local site number • Remote site number • VPLS neighbor address • VPLS identifier
Interface flags	<p>Flag associated with the interface. Can include the following:</p> <ul style="list-style-type: none"> • VC-Down—The virtual circuit associated with this interface is down.
Label-base	First label in a block of labels. A remote PE router uses this first label when sending traffic toward the advertising PE router.
Offset	Displays the VPLS Edge (VE) block offset in the Layer 2 VPN NLRI. The VE block offset is used to identify a label block from which a particular label value is selected to setup a pseduowire for a remote site. The block offset value itself indicates the starting VE ID that maps to the label base contained in the VPLS NLRI advertisement.

Table 32: show vpls connections Output Fields (*continued*)

Field Name	Field Description
Size	Label block size. A configurable value that represents the number of label blocks required to cover all the pseudowires for the remote peer. Acceptable configuration values are: 2 , 4 , 8 and 16 . The default value is 2 . A value of 0 will be displayed when using the remote-site option.
Range	Label block range. A value that keeps track of the numbers of remote sites discovered within each label block.
Preference	Preference value advertised for a VPLS site. When multiple PE routers are assigned the same VE ID for multihoming, you might need to specify that a particular PE router acts as the designated forwarder by configuring the site preference value. The site preference indicates the degree of preference for a particular customer site. The site preference is one of the tie-breaking criteria used in a designated forwarder election.
status-vector	Bit vector advertising the state of local PE-CE circuits to remote PE routers. A bit value of 0 indicates that the local circuit and LSP tunnel to the remote PE router are up, whereas a value of 1 indicates either one or both are down.
connection-site	Name of the connection site.
Neighbor	IP address and VPLS identifier for the VPLS neighbor. If multiple pseudowires have been configured, the IP address will also show the PW-specific <i>vpls-id-list</i> , for example, 10.255.144.4 (vpls-id 200).
Type	Type of connection: loc (local) or rmt (remote).

Table 32: show vpls connections Output Fields (*continued*)

Field Name	Field Description
St	<p>Status of the VPLS connection (corresponds with Legend for Connection Status):</p> <ul style="list-style-type: none"> • EI—The local VPLS interface is configured with an encapsulation that is not supported. • EM—The encapsulation type received on this VPLS connection from the neighbor does not match the local VPLS connection interface encapsulation type. • VC-Dn—The virtual circuit is currently down. • CM—The two routers do not agree on a control word, which causes a control word mismatch. • CN—The virtual circuit is not provisioned properly. • OR—The label associated with the virtual circuit is out of range. • OL—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit. • LD—All of the CE-facing interfaces to the local site are down. Therefore, the connection to the local site is signaled as down to the other PE routers. No pseudowires can be established. • RD—All the interfaces to the remote neighbor are down. Therefore, the remote site has been signaled as down to the other PE routers. No pseudowires can be established. • LN—The local site has lost path selection to the remote site and therefore no pseudowires can be established from this local site. • RN—The remote site has lost path selection to a local site or other remote site and therefore no pseudowires are established to this remote site. In a multihoming configuration, one multihomed PE site displays the state LN, and the other multihomed PE site displays the state RN in the following circumstances: <ul style="list-style-type: none"> • The multihomed links are both configured to be the backup site. • The two multihomed PE routers have the same site ID, but have a peering relationship with a route reflector (RR) that has a different site ID. • XX—The VPLS connection is down for an unknown reason. This is a programming error. • MM—The MTU for the local site and the remote site do not match. • BK—The router is using a backup connection. • PF—Profile parse failure. • RS—The remote site is in a standby state. • NC—The interface encapsulation is not configured as an appropriate CCC, TCC, or VPLS encapsulation. • WE—The encapsulation configured for the interface does not match the encapsulation configured for the associated connection within the VPLS routing instance.

Table 32: show vpls connections Output Fields (*continued*)

Field Name	Field Description
	<ul style="list-style-type: none"> • NP—The router detects that interface hardware is not present. The hardware might be offline, a PIC might not be of the desired type, or the interface might be configured in a different routing instance. • -->—Only the outbound connection is up. • <--—Only the inbound connection is up. • Up—The VPLS connection is operational. • Dn—The VPLS connection is down. • CF—The router cannot find enough bandwidth to the remote router to satisfy the VPLS connection bandwidth requirement. • SC—The local site identifier matches the remote site identifier. No pseudowire can be established between these two sites. You should configure different values for the local and remote site identifiers. • LM—The local site identifier is not the minimum designated, meaning it is not the lowest. There is another local site with a lower site identifier. Pseudowires are not being established to this local site, and the associated local site identifier is not being used to distribute VPLS label blocks. However, this is not an error state. Traffic continues to be forwarded to the PE router interfaces connected to the local sites when the local sites are in this state. • RM—The remote site identifier is not the minimum designated, meaning it is not the lowest. There is another remote site connected to the same PE router which has lower site identifier. The PE router cannot establish a pseudowire to this remote site and the associated remote site identifier cannot be used to distribute VPLS label blocks. However, this is not an error state. Traffic can continue to be forwarded to the PE router interface connected to this remote site when the remote site is in this state. • IL—The incoming packets for the VPLS connection have no MPLS label. • MI—The configured mesh group identifier is in use by another system in the network. • ST—The router has switched to a standby connection. • PB—Profile busy. • SN—The VPLS neighbor is static.
Time last up	Time connection was last in the Up condition.
# Up trans	Number of transitions from Down to Up condition.
Status	Status of the (local or remote circuit) local interface: <ul style="list-style-type: none"> • Up—Operational • Dn—Down • NP—Not present • DS—Disabled • WE—Wrong encapsulation • UN—Uninitialized
Encapsulation	Type of encapsulation: VPLS .
Remote PE	Address of the remote provider edge router.

Table 32: show vpls connections Output Fields (*continued*)

Field Name	Field Description
Negotiated control-word	Whether a control word has been negotiated: Yes or No .
Incoming label	Name of the incoming label.
Outgoing label	Name of the outgoing label.
Negotiated PW status TLV	Indicates whether or not the pseudowire status TLV has been negotiated for the VPLS connection.
Local interface	Provides the following information about the local interface configured for the VPLS neighbor: <ul style="list-style-type: none"> • Name of the local interface • Status—Interface status (Up or Down) • Encapsulation—Interface encapsulation (for example, ETHERNET) • Description—Includes the VPLS instance name, the VPLS neighbor address, and the VPLS identifier
Time	Date and time of VPLS connection event.
Event	Type of event.
Interface/Lbl/PE	Interface, label, or PE router.
Connection History	Each entry can include the date, time, year, and the connection event. Connection events include any of a variety of events related to VPLS connections, such as route changes, label updates, and interfaces going down or coming up.

Sample Output

show vpls connections

```
user@host> show vpls connections
Layer-2 VPN connections:
```

Legend for connection status (St)

EI -- encapsulation invalid	NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch	WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down	NP -- interface hardware not present
CM -- control-word mismatch	-< -- only outbound connection is up
CN -- circuit not provisioned	>- -- only inbound connection is up
OR -- out of range	Up -- operational
OL -- no outgoing label	Dn -- down
LD -- local site signaled down	CF -- call admission control failure
RD -- remote site signaled down	SC -- local and remote site ID collision
LN -- local site not designated	LM -- local site ID not minimum designated
RN -- remote site not designated	RM -- remote site ID not minimum designated
XX -- unn connection status	IL -- no incoming label
MM -- MTU mismatch	MI -- Mesh-Group ID not availble
BK -- Backup connection	ST -- Standby connection

PF -- Profile parse failure PB -- Profile busy

Legend for interface status
 Up -- operational
 Dn -- down

Instance: vpls-1
 Local site: 1 (11)
 Number of local interfaces: 1
 Number of local interfaces up: 1
 IRB interface present: no
 lt-1/3/0.10496

vt-1/3/0.1048588	1	Intf - vpls vpls-1 local site 11 remote site 1
vt-1/2/0.1048591	2	Intf - vpls vpls-1 local site 11 remote site 2
vt-1/2/0.1048585	3	Intf - vpls vpls-1 local site 11 remote site 3
vt-1/2/0.1048587	4	Intf - vpls vpls-1 local site 11 remote site 4
vt-1/2/0.1048589	5	Intf - vpls vpls-1 local site 11 remote site 5
vt-1/3/0.1048586	6	Intf - vpls vpls-1 local site 11 remote site 6
vt-1/3/0.1048590	7	Intf - vpls vpls-1 local site 11 remote site 7
vt-1/3/0.1048584	8	Intf - vpls vpls-1 local site 11 remote site 8

Label-base	Offset	Size	Range	Preference
+ 800256	1	16	16	100

Timer Values:
 Startup wait time: 120 seconds
 New site wait-time: 20 seconds
 Collision detect time: 30 seconds
 Reclaim wait time: 748 milliseconds

connection-site	Type	St	Time last up	# Up trans
1	rmt	Up	Apr 28 13:28:24 2009	2

Remote PE: 124.1.2.1, Negotiated control-word: No
 Incoming label: 800256, Outgoing label: 800026
 Local interface: vt-1/3/0.1048588, Status: Up, Encapsulation: VPLS
 Description: Intf - vpls vpls-1 local site 11 remote site 1

Connection History:

Apr 28 13:28:24 2009	status update timer
Apr 28 13:28:24 2009	PE route down
Apr 28 13:24:27 2009	status update timer
Apr 28 13:24:27 2009	loc intf up vt-1/3/0.1048588
Apr 28 13:24:27 2009	PE route changed
Apr 28 13:24:27 2009	Out lbl Update 800026
Apr 28 13:24:27 2009	In lbl Update 800256
Apr 28 13:24:27 2009	loc intf down

2	rmt	Up	Apr 28 13:28:24 2009	2
---	-----	----	----------------------	---

Remote PE: 124.1.7.1, Negotiated control-word: No
 Incoming label: 800257, Outgoing label: 800034
 Local interface: vt-1/2/0.1048591, Status: Up, Encapsulation: VPLS
 Description: Intf - vpls vpls-1 local site 11 remote site 2

Connection History:

Apr 28 13:28:24 2009	status update timer
Apr 28 13:28:24 2009	PE route down
Apr 28 13:24:28 2009	status update timer
Apr 28 13:24:28 2009	loc intf up vt-1/2/0.1048591
Apr 28 13:24:28 2009	PE route changed

```

Apr 28 13:24:28 2009 Out lbl Update      800034
Apr 28 13:24:28 2009 In lbl Update      800257
Apr 28 13:24:28 2009 loc intf down
3      rmt Up      Apr 28 13:28:24 2009      2
Remote PE: 124.1.4.1, Negotiated control-word: No
Incoming label: 800258, Outgoing label: 800026
Local interface: vt-1/2/0.1048585, Status: Up, Encapsulation: VPLS
Description: Intf - vpls vpls-1 local site 11 remote site 3
Connection History:
Apr 28 13:28:24 2009 status update timer
Apr 28 13:28:24 2009 PE route down
Apr 28 13:24:26 2009 status update timer
Apr 28 13:24:26 2009 loc intf up      vt-1/2/0.1048585
Apr 28 13:24:26 2009 PE route changed
Apr 28 13:24:26 2009 Out lbl Update      800026
Apr 28 13:24:26 2009 In lbl Update      800258
Apr 28 13:24:26 2009 loc intf down
4      rmt Up      Apr 28 13:28:24 2009      2
Remote PE: 124.1.6.1, Negotiated control-word: No
Incoming label: 800259, Outgoing label: 800026
Local interface: vt-1/2/0.1048587, Status: Up, Encapsulation: VPLS
Description: Intf - vpls vpls-1 local site 11 remote site 4
Connection History:
Apr 28 13:28:24 2009 status update timer
Apr 28 13:28:24 2009 PE route down
Apr 28 13:24:27 2009 status update timer
Apr 28 13:24:27 2009 loc intf up      vt-1/2/0.1048587
Apr 28 13:24:27 2009 PE route changed
Apr 28 13:24:27 2009 Out lbl Update      800026
Apr 28 13:24:27 2009 In lbl Update      800259
Apr 28 13:24:27 2009 loc intf down
5      rmt Up      Apr 28 13:28:24 2009      2
Remote PE: 124.1.3.1, Negotiated control-word: No
Incoming label: 800260, Outgoing label: 800034
Local interface: vt-1/2/0.1048589, Status: Up, Encapsulation: VPLS
Description: Intf - vpls vpls-1 local site 11 remote site 5
Connection History:
Apr 28 13:28:24 2009 status update timer
Apr 28 13:28:24 2009 PE route down
Apr 28 13:24:28 2009 status update timer
Apr 28 13:24:28 2009 loc intf up      vt-1/2/0.1048589
Apr 28 13:24:28 2009 PE route changed
Apr 28 13:24:28 2009 Out lbl Update      800034
Apr 28 13:24:27 2009 In lbl Update      800260
Apr 28 13:24:27 2009 loc intf down

```

show vpls connections (with multiple pseudowires)

```
user@host> show vpls connections
```

```
Layer-2 VPN connections:
```

```
Legend for connection status (St)
```

```

EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down   NP -- interface hardware not present
CM -- control-word mismatch     -> -- only outbound connection is up
CN -- circuit not provisioned   <- -- only inbound connection is up
OR -- out of range             Up -- operational
OL -- no outgoing label        Dn -- down
LD -- local site signaled down  CF -- call admission control failure
RD -- remote site signaled down SC -- local and remote site ID collision

```

```

LN -- local site not designated   LM -- local site ID not minimum designated
RN -- remote site not designated  RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure       PB -- Profile busy
RS -- remote site standby         SN -- Static Neighbor
LB -- Local site not best-site    RB -- Remote site not best-site
VM -- VLAN ID mismatch

```

Legend for interface status

```

Up -- operational
Dn -- down

```

Instance: vpls

VPLS-id: 100

Mesh-group connections: __ves__

```

Neighbor      Type  St      Time last up      # Up trans
10.255.114.3 (vpls-id 100) rmt  Up      Apr 11 23:38:38 2013      1
Remote PE: 10.255.114.3, Negotiated control-word: No
Incoming label: 262145, Outgoing label: 262145
Negotiated PW status TLV: No
Local interface: lsi.1049090, Status: Up, Encapsulation: ETHERNET
Description: Intf - vpls h-vpls neighbor 10.255.114.3 vpls-id 100

```

Mesh-group connections: spokes

```

Neighbor      Type  St      Time last up      # Up trans
10.255.114.4 (vpls-id 200) rmt  Up      Apr 11 23:39:25 2013      1
Remote PE: 10.255.114.4, Negotiated control-word: No
Incoming label: 262148, Outgoing label: 304224
Negotiated PW status TLV: Yes
Local PW status code: 0x00000000, Neighbor PW status code: 0x00000000
Local interface: lsi.1049091, Status: Up, Encapsulation: ETHERNET
Description: Intf - vpls h-vpls neighbor 10.255.114.4 vpls-id 200
10.255.114.4 (vpls-id 201) rmt  Up      Apr 11 23:39:25 2013      1
Remote PE: 10.255.114.4, Negotiated control-word: No
Incoming label: 262149, Outgoing label: 304225
Negotiated PW status TLV: Yes
Local PW status code: 0x00000000, Neighbor PW status code: 0x00000000
Local interface: lsi.1049096, Status: Up, Encapsulation: ETHERNET
Description: Intf - vpls h-vpls neighbor 10.255.114.4 vpls-id 201

```

show vpls connections extensive (Static VPLS Neighbors)

```
user@host> show vpls connections extensive instance red
```

Layer-2 VPN connections:

Legend for connection status (St)

```

EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down   NP -- interface hardware not present
CM -- control-word mismatch     -> -- only outbound connection is up
CN -- circuit not provisioned   <- -- only inbound connection is up
OR -- out of range              Up -- operational
OL -- no outgoing label         Dn -- down
LD -- local site signaled down  CF -- call admission control failure
RD -- remote site signaled down SC -- local and remote site ID collision
LN -- local site not designated LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unn connection status     IL -- no incoming label
MM -- MTU mismatch              MI -- Mesh-Group ID not availble

```

BK -- Backup connection ST -- Standby connection
 PF -- Profile parse failure PB -- Profile busy
 RS -- remote site standby SN -- Static Neighbor

Legend for interface status

Up -- operational
 Dn -- down

Instance: static

VPLS-id: 1

Number of local interfaces: 1

Number of local interfaces up: 1

ge-0/0/5.0

lsi.1049344

Intf - vpls static neighbor 10.255.114.3 vpls-id

1

Neighbor	Type	St	Time last up	# Up trans
10.255.114.3(vpls-id 1)(SN)	rmt	Up	Mar 4 08:48:41 2010	1

Remote PE: 10.255.114.3, Negotiated control-word: No

Incoming label: 29696, Outgoing label: 29697

Negotiated PW status TLV: No

Local interface: lsi.1049344, Status: Up, Encapsulation: ETHERNET

Description: Intf - vpls static neighbor 10.255.114.3 vpls-id 1

Connection History:

Mar 4 08:48:41 2010	status update timer	
Mar 4 08:48:41 2010	PE route changed	
Mar 4 08:48:41 2010	Out lbl Update	29697
Mar 4 08:48:41 2010	In lbl Update	29696
Mar 4 08:48:41 2010	loc intf up	lsi.1049344

user@PE1> show vpls connections extensive (Multihoming with FEC 129)

Layer-2 VPN connections:

Legend for connection status (St)

EI -- encapsulation invalid	NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch	WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down	NP -- interface hardware not present
CM -- control-word mismatch	-> -- only outbound connection is up
CN -- circuit not provisioned	<- -- only inbound connection is up
OR -- out of range	Up -- operational
OL -- no outgoing label	Dn -- down
LD -- local site signaled down	CF -- call admission control failure
RD -- remote site signaled down	SC -- local and remote site ID collision
LN -- local site not designated	LM -- local site ID not minimum designated
RN -- remote site not designated	RM -- remote site ID not minimum designated
XX -- unknown connection status	IL -- no incoming label
MM -- MTU mismatch	MI -- Mesh-Group ID not available
BK -- Backup connection	ST -- Standby connection
PF -- Profile parse failure	PB -- Profile busy
RS -- remote site standby	SN -- Static Neighbor
LB -- Local site not best-site	RB -- Remote site not best-site
VM -- VLAN ID mismatch	

Legend for interface status

Up -- operational
 Dn -- down

Instance: green

L2vpn-id: 100:100

Local-id: 1.1.1.2

Number of local interfaces: 2

Number of local interfaces up: 2


```

ge-0/3/1.0
ge-0/3/3.0
lsi.101711873          Intf - vpls green local-id 1.1.1.2 remote-id
1.1.1.4 neighbor 1.1.1.4
Remote-id              Type  St    Time last up          # Up trans
1.1.1.4                rmt  Up    Jan 31 13:49:52 2012      1
Remote PE: 1.1.1.4, Negotiated control-word: No
Incoming label: 262146, Outgoing label: 262146
Local interface: lsi.101711873, Status: Up, Encapsulation: ETHERNET
Description: Intf - vpls green local-id 1.1.1.2 remote-id 1.1.1.4 neighbor
1.1.1.4
Connection History:
Jan 31 13:49:52 2012  status update timer
Jan 31 13:49:52 2012  PE route changed
Jan 31 13:49:52 2012  Out lbl Update                262146
Jan 31 13:49:52 2012  In lbl Update                  262146
Jan 31 13:49:52 2012  loc intf up                    lsi.101711873
Multi-home:
Local-site              Id    Pref  State
test                    1     100   Up
Number of interfaces: 1
Number of interfaces up: 1
ge-0/3/1.0
Received multi-homing advertisements:
Remote-PE              Pref  flag  Description
1.1.1.4                100   0x0

```

show vpls flood event-queue

Syntax	show vpls flood event-queue
Release Information	Command introduced in Junos OS Release 8.0.
Description	Display the pending events in the VPLS flood queue.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show vpls flood event-queue on page 454
Output Fields	Table 33 on page 454 lists the output fields for the show vpls flood event-queue command. Output fields are listed in the approximate order in which they appear.

Table 33: show vpls flood event-queue Output Fields

Field Name	Field Description
Current Pending Event	Provides information on the current event in the VPLS flood event queue.
Name	Name of the event.
Owner Name	Name of the interface associated with the flood event.
Pending Op	Pending operation for the event.
Last Error	Name of the last error encountered.
Number of Retries	Number of attempts made to update the event queue.
Pending Event List	List of the events awaiting processing.
Event Name	Name of the event.
Pending Op	Pending operation for the event.
Event Identifier	Name of the interface associated with the flood event.

Sample Output

show vpls flood event-queue

```

user@host> show vpls flood event-queue
Current Pending Event
  Name:          Flood Nexthop
  Owner Name: ge-4/3/0.0
  Pending Op: ADD

```

```
Last Error:ENOMEM
Number of Retries:3
Pending Event List:
Event Name      Pending Op      Event Identifier
Flood Nexthop   ADD             ge-4/3/0.0
Flood Route     ADD             ge-4/3/0.0
```

show vpls flood instance

Syntax	show vpls flood instance <brief detail extensive> <instance-name> <logical-system <i>logical-system-name</i> >
Release Information	Command introduced in Junos OS Release 8.0.
Description	Display VPLS information related to the flood process.
Options	<p>none—Display VPLS information related to the flood process for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance-name—(Optional) Display VPLS information related to the flood process for the specified routing instance.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Display VPLS information related to the flood process for the specified logical system.</p>
Required Privilege Level	view
List of Sample Output	show vpls flood instance on page 457 show vpls flood instance logical-system-name on page 457 show vpls flood instance detail on page 457
Output Fields	Table 34 on page 456 lists the output fields for the show vpls flood instance command. Output fields are listed in the approximate order in which they appear.

Table 34: show vpls flood instance Output Fields

Field Name	Field Description
Logical system	Name of the logical system.
Name	Name of the VPLS routing instance.
CEs	Number of CE routers connected to the VPLS instance.
VEs	Number of VE routers connected to the VPLS instance.
Flood routes	List of all flood routes associated with the VPLS instance.
Prefix	Prefix for the route.
Type	Type of route.
Owner	VPLS routing instance or interface associated with the route.
Nhype	Next-hop type. For example, flood for a flood route.

Table 34: show vpls flood instance Output Fields (*continued*)

Field Name	Field Description
Nhindex	Next-hop index number for the route.

Sample Output

show vpls flood instance

```
user@host> show vpls flood instance

Logical system: __juniper_ls1__
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix   Type      Owner      NhType      NhIndex
  default  ALL_CE_FLOOD green      flood       383
  0x47/16  CE_FLOOD  fe-1/2/1.0 flood       388
```

show vpls flood instance logical-system-name

```
user@host: __juniper_ls1__> show vpls flood instance juniper_ls1

Logical system: __juniper_ls1__
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix   Type      Owner      NhType      NhIndex
  default  ALL_CE_FLOOD green      flood       383
  0x47/16  CE_FLOOD  fe-1/2/1.0 flood       388
```

show vpls flood instance detail

```
user@host: __juniper_ls1__> show vpls flood instance detail

Logical system: __juniper_ls1__
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix   Type      Owner      NhType      NhIndex
  default  ALL_CE_FLOOD green      flood       383
  0x47/16  CE_FLOOD  fe-1/2/1.0 flood       388
```

show vpls flood route

Syntax	show vpls flood route (all-ce-flood instance-name <i>instance-name</i> <logical-system-name <i>logical-system-name</i> > ce-flood interface <i>interface-name</i>)
Release Information	Command introduced in Junos OS Release 8.0.
Description	Display VPLS route information related to the flood process for either the specified routing instance or the specified interface.
Options	<p>all-ce-flood—Display the flood next-hop route for all customer edge routers for traffic coming from the core of the network.</p> <p>ce-flood interface <i>interface-name</i>—Display the flood next-hop route for traffic coming from the specified customer edge interface.</p> <p>instance-name <i>instance-name</i>—Display the flood routes for the specified instance.</p> <p>logical-system-name <i>logical-system-name</i>—(Optional) Specify the logical system whose flood routes you want to display. You can only specify the default logical system name for VPLS. The default logical system name is __juniper_ls1__ (the name must be entered in the command with the underscore characters).</p>
Required Privilege Level	view
List of Sample Output	show vpls flood route all-ce-flood on page 459 show vpls flood route ce-flood on page 459
Output Fields	Table 35 on page 458 lists the output for the show vpls flood route command. Output fields are listed in the approximate order in which they appear.

Table 35: show vpls flood route Output Fields

Field Name	Field Description
Flood route prefix	Prefix for the flood route.
Flood route type	Type of flood route (either CE_FLOOD or ALL_CE_FLOOD).
Flood route owner	VPLS routing instance or interface associated with the flood route.
Nexthop type	Next-hop type. For example, flood for a flood route.
Nexthop index	Next-hop index number for the route.
Interfaces flooding to	Interfaces to which VPLS routes are being flooded.
Name	Name of the interface.

Table 35: show vpls flood route Output Fields (*continued*)

Field Name	Field Description
Type	Type of VPLS router (CE or VE).
Nh type	Next-hop type.
Index	Index number for the flood route.

Sample Output

show vpls flood route all-ce-flood

```
user@host:~juniper_ls1~> show vpls flood route all-ce-flood logical-system-name
~juniper_ls1_instance-name green
```

```
Flood route prefix: default
Flood route type: ALL_CE_FLOOD
Flood route owner: green
Nexthop type: flood
Nexthop index: 383
  Interfaces Flooding to:
    Name      Type      NhType      Index
    fe-1/2/1.0 CE
```

show vpls flood route ce-flood

```
user@host:~juniper_ls1~> show vpls flood route ce-flood interface fe-1/2/1.0
```

```
Flood route prefix: 0x47/16
Flood route type: CE_FLOOD
Flood route owner: fe-1/2/1.0
Nexthop type: flood
Nexthop index: 388
  Interfaces Flooding to:
    Name      Type      NhType      Index
    lsi.49152 VE      indr      262142
```

show vpls mac-table

Syntax	show vpls mac-table <brief detail extensive summary> <bridge-domain <i>bridge-domain-name</i> > <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> <mac-address> <vlan-id <i>vlan-id-number</i> >
Release Information	Command introduced in Junos OS Release 8.5. Command introduced in Junos OS Release 15.1
Description	Display learned virtual private LAN service (VPLS) media access control (MAC) address information.
Options	none —Display all learned VPLS MAC address information. brief detail extensive summary —(Optional) Display the specified level of output. bridge-domain <i>bridge-domain-name</i> —(Optional) Display learned VPLS MAC addresses for the specified bridge domain. instance <i>instance-name</i> —(Optional) Display learned VPLS MAC addresses for the specified instance. interface <i>interface-name</i> —(Optional) Display learned VPLS MAC addresses for the specified instance. logical-system (all <i>logical-system-name</i>) —(Optional) Display learned VPLS MAC addresses for all logical systems or for the specified logical system. mac-address —(Optional) Display the specified learned VPLS MAC address information.. vlan-id <i>vlan-id-number</i> —(Optional) Display learned VPLS MAC addresses for the specified VLAN.
Required Privilege Level	view
List of Sample Output	show vpls mac-table on page 461 show vpls mac-table (with Layer 2 Services over GRE Interfaces) on page 462 show vpls mac-table (with VXLAN enabled) on page 462 show vpls mac-table count on page 462 show vpls mac-table detail on page 463 show vpls mac-table extensive on page 463
Output Fields	Table 36 on page 461 describes the output fields for the show vpls mac-table command. Output fields are listed in the approximate order in which they appear.

Table 36: show vpls mac-table Output fields

Field Name	Field Description
Routing instance	Name of the routing instance.
Bridging domain	Name of the bridging domain.
MAC address	MAC address or addresses learned on a logical interface.
MAC flags	Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> • S—Static MAC address configured. • D—Dynamic MAC address learned. • SE—MAC accounting is enabled. • NM—Nonconfigured MAC.
Logical interface	Name of the logical interface.
MAC count	Number of MAC addresses learned on a specific routing instance or interface.
Learning interface	Logical interface or logical Label Switched Interface (LSI) the address is learned on.
Base learning interface	Base learning interface of the MAC address. This field is introduced in Junos OS Release 14.2.
Learn VLAN ID/VLAN	VLAN ID of the routing instance or bridge domain in which the MAC address was learned.
VXLAN ID/VXLAN	VXLAN Network Identifier (VNI)
Layer 2 flags	Debugging flags signifying that the MAC address is present in various lists.
Epoch	Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging.
Sequence number	Sequence number assigned to this MAC address. Used for debugging.
Learning mask	Mask of Packet Forwarding Engines where this MAC address was learned. Used for debugging.
IPC generation	Creation time of the logical interface when this MAC address was learned. Used for debugging.

Sample Output

show vpls mac-table

```

user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC,
          SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vpls_ldp1
VLAN : 223
  MAC          MAC          Logical
  address      flags        interface
00:90:69:9c:1c:5d  D          ge-0/2/5.400

```

MAC flags (S -static MAC, D -dynamic MAC,
SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vpls_red
VLAN : 401

MAC address	MAC flags	Logical interface
00:00:aa:12:12:12	D	lsi.1051138
00:05:85:74:9f:f0	D	lsi.1051138

show vpls mac-table (with Layer 2 Services over GRE Interfaces)

user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : vpls_4site:1000
Bridging domain : __vpls_4site:1000__, MAC

MAC address	MAC flags	Logical interface
00:01:01:00:01:f4	D,SE	ge-4/2/0.1000
00:02:01:33:01:f4	D,SE	lsi.1052004
00:03:00:32:01:f4	D,SE	lsi.1048840
00:04:00:14:01:f4	D,SE	lsi.1052005
00:02:01:33:02:f7	D,SE	gr-1/2/10.10

show vpls mac-table (with VXLAN enabled)

user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : vpls_4site:1000
Bridging domain : __vpls_4site:1000__, VLAN : 4094,4093
VXLAN: Id : 300, Multicast group: 226.1.1.3

MAC address	MAC flags	Logical interface
00:01:01:00:01:f4	D,SE	ge-4/2/0.1000
00:02:01:33:01:f4	D,SE	lsi.1052004
00:03:00:32:01:f4	D,SE	lsi.1048840
00:04:00:14:01:f4	D,SE	lsi.1052005
00:02:01:33:02:f7	D,SE	vtep.1052010
00:04:00:14:02:f7	D,SE	vtep.1052011

show vpls mac-table count

user@host> show vpls mac-table count
0 MAC address learned in routing instance __juniper_private1__

MAC address count per interface within routing instance:

Logical interface	MAC count
lc-0/0/0.32769	0
lc-0/1/0.32769	0
lc-0/2/0.32769	0
lc-2/0/0.32769	0
lc-0/3/0.32769	0
lc-2/1/0.32769	0
lc-9/0/0.32769	0
lc-11/0/0.32769	0
lc-2/2/0.32769	0
lc-9/1/0.32769	0
lc-11/1/0.32769	0

1c-2/3/0.32769	0
1c-9/2/0.32769	0
1c-11/2/0.32769	0
1c-11/3/0.32769	0
1c-9/3/0.32769	0

MAC address count per learn VLAN within routing instance:

Learn VLAN ID	MAC count
0	0

1 MAC address learned in routing instance vpls_ldp1

MAC address count per interface within routing instance:

Logical interface	MAC count
lsi.1051137	0
ge-0/2/5.400	1

MAC address count per learn VLAN within routing instance:

Learn VLAN ID	MAC count
0	1

1 MAC address learned in routing instance vpls_red

MAC address count per interface within routing instance:

Logical interface	MAC count
ge-0/2/5.300	1

MAC address count per learn VLAN within routing instance:

Learn VLAN ID	MAC count
0	1

show vpls mac-table detail

```

user@host> show vpls mac-table detail
MAC address: 00:90:69:9c:1c:5d
Routing instance: vpls_ldp1
Learning interface: ge-0/2/5.400
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                       IPC generation: 0

MAC address: 00:90:69:9c:1c:5d
Routing instance: vpls_red
Learning interface: ge-0/2/5.300
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                       IPC generation: 0

```

show vpls mac-table extensive

```

user@host> show vpls mac-table extensive

MAC address: 00:10:00:01:00:00
Routing instance: vpls_1
Bridging domain: __vpls_1__, VLAN : NA
Learning interface: lsi.1049165
Base learning interface: lsi.1049165
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 1
Learning mask: 0x00000001

```

MAC address: 00:10:00:01:00:01
Routing instance: vpls_1
Bridging domain: __vpls_1__, VLAN : NA
Learning interface: lsi.1049165
Base learning interface: lsi.1049165
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0 Sequence number: 1
Learning mask: 0x00000001

MAC address: 00:10:00:01:00:02
Routing instance: vpls_1
Bridging domain: __vpls_1__, VLAN : NA
Learning interface: lsi.1049165
Base learning interface: lsi.1049165
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0 Sequence number: 1
Learning mask: 0x00000001

MAC address: 00:10:00:01:00:03
Routing instance: vpls_1
Bridging domain: __vpls_1__, VLAN : NA
Learning interface: lsi.1049165
Base learning interface: lsi.1049165
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0 Sequence number: 1
Learning mask: 0x00000001

show vpls statistics

Syntax	show vpls statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	(T Series and M Series routers, except for the M160 router) Display virtual private LAN service (VPLS) statistics.
Options	<p>none—Display VPLS statistics for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display VPLS statistics for a specific VPLS routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show vpls statistics on page 466 show vpls statistics instance on page 466
Output Fields	Table 37 on page 465 lists the output fields for the show vpls statistics command. Output fields are listed in the approximate order in which they appear.

Table 37: show vpls statistics Output Fields

Field Name	Field Description
Instance	Name of the VPLS instance.
Local interface	Name of the local VPLS virtual loopback tunnel interface, <i>vt-fpc/pic/port.nnnnn</i> , where <i>nnnnn</i> is a dynamically generated virtual port used to transport and receive packets from other provider edge (PE) routers in the VPLS domain.
Index	Number associated with the next hop.
Remote provider edge router	Address of the remote provider edge router.
Multicast packets	Number of multicast packets received.
Multicast bytes	Number of multicast bytes received.
Flood packets	Number of VPLS flood packets received.
Flood bytes	Number of VPLS flood bytes received.

Table 37: show vpls statistics Output Fields (*continued*)

Field Name	Field Description
Current MAC count	Number of MAC addresses learned by the interface and the configured maximum limit on the number of MAC addresses that can be learned.

Sample Output

show vpls statistics

```

user@host> show vpls statistics

VPLS statistics:

Instance: green

  Local interface: fe-2/2/1.0, Index: 69
    Multicast packets:      1
    Multicast bytes   :      60
    Flooded packets   :      18
    Flooded bytes    :    2556
    Current MAC count:      1

  Local interface: lt-0/3/0.2, Index: 72
    Multicast packets:      3
    Multicast bytes   :    153
    Flooded packets   :      1
    Flooded bytes    :     51
    Current MAC count:      1

  Local interface: lsi.32769, Index: 75
    Current MAC count:      0

  Local interface: lsi.32771, Index: 77
  Remote PE: 10.255.14.222
    Current MAC count:      2

Instance: red

  Local interface: vt-0/3/0.32768, Index: 74
    Multicast packets:      0
    Multicast bytes   :      0
    Flooded packets   :      0
    Flooded bytes    :      0
    Current MAC count:      0

  Local interface: vt-0/3/0.32770, Index: 76
    Multicast packets:      0
    Multicast bytes   :      0
    Flooded packets   :      0
    Flooded bytes    :      0
    Current MAC count:      0

```

show vpls statistics instance

```

user@host> show vpls statistics instance red

```

Layer-2 VPN Statistics:

Instance: red

Local interface: vt-3/2/0.32768, Index: 73

Remote provider edge router: 10.255.17.35

Multicast packets: 0

Multicast bytes : 0

Flood packets : 0

Flood bytes : 0

Current MAC count: 1 (Limit 20)

PART 5

Index

- [Index on page 471](#)

Index

Symbols

#, comments in configuration statements.....	xviii
(), in syntax descriptions.....	xviii
802.1Q VLANs	
dynamic.....	110
mixed VLAN tagging.....	171
VLAN IDs.....	255
< >, in syntax descriptions.....	xviii
[], in configuration statements.....	xviii
{ }, in configuration statements.....	xviii
(pipe), in syntax descriptions.....	xviii

A

AAA	
subscriber statistics	
displaying.....	268, 273
subscribers	
displaying.....	276
accept statement.....	97
access	
Layer 2 wholesale.....	86
Layer 3 wholesale	
DHCP.....	15, 35
PPPoE.....	57
access profile	
retailer	
DHCP.....	16, 36
PPPoE.....	59
wholesaler.....	86
DHCP.....	16, 36
PPPoE.....	58
access-profile statement	
routing instances.....	98
accounting-server statement.....	99
active-server-group statement	
DHCP relay agent.....	100

address assignment pool	
128-bit address	
DHCPv6 example.....	44
74-bit prefix	
DHCPv6 example.....	45
configuring.....	40
address assignment pools	
displaying.....	279
address server	
configuring.....	40
address statement.....	101
address-assignment pools	
DHCPv6 attributes	
specifying the preferred-lifetime.....	119
specifying the t1-percentage.....	119
specifying the t2-percentage.....	119
specifying the valid-lifetime.....	119
address-assignment statement	
address-assignment pools.....	105
aggregated Ethernet interfaces	
status information, displaying.....	327
authentication statement	
DHCP local server.....	106
DHCP relay agent.....	107
authentication-order statement	
access.....	108
authentication-server statement.....	109
auto-configure statement.....	110

B

binding state of DHCP client	
clearing.....	290, 303
displaying.....	281, 295
binding state of DHCPv6 client	
clearing.....	317
displaying.....	308
braces, in configuration statements.....	xviii
brackets	
angle, in syntax descriptions.....	xviii
square, in configuration statements.....	xviii
broadband subscriber management	
Layer 2 wholesale topology.....	71
Layer 3 wholesale topology	
DHCPv4.....	9
DHCPv6.....	29
PPPoE.....	53

C

clear dhcp relay binding command.....	303
---------------------------------------	-----

clear dhcp relay statistics command.....	305
clear dhcp server binding command.....	290
clear dhcp server statistics command.....	293
clear dhcpv6 server binding command.....	317
clear dhcpv6 server statistics command.....	319
comments, in configuration statements.....	xviii
connectivity-type statement.....	112
conventions	
text and syntax.....	xvii
core-facing statement.....	113
curly braces, in configuration statements.....	xviii
customer support.....	xix
contacting JTAC.....	xix

D

demux interfaces	
unit statement.....	245
demux-options statement	
dynamic IP demux interface.....	115
demux-source statement.....	118
dynamic IP demux interfaces.....	116
dynamic underlying interface.....	117
demux0 statement	
dynamic IP demux interface.....	114
demux0 statements	
underlying-interface.....	237
DHCP	
extended server binding	
clearing.....	290
displaying.....	281
extended server statistics	
clearing.....	293
displaying.....	287
relay binding	
clearing.....	303
relay binding state	
displaying.....	295
relay statistics	
clearing.....	305
displaying.....	300
DHCP client	
binding state	
clearing.....	290, 303
displaying.....	281, 295
statistics	
clearing.....	293, 305
displaying.....	287
DHCP local server	
lease timers.....	200

DHCP local server statements	
dhcp-local-server.....	120
dhcpv6.....	133
dynamic-profile.....	137
group.....	174
interface.....	181
ip-address-first.....	195
pool-match-order.....	203
traceoptions.....	235
DHCP relay agent statements	
active-server-group.....	100
dhcp-relay.....	126
dynamic-profile.....	138
group.....	177
interface.....	183
server-group.....	228
traceoptions.....	235
user-prefix.....	253
username-include.....	107
dhcp-attributes statement	
address-assignment pools.....	119
dhcp-local-server statement.....	120
dhcp-relay statement.....	126
DHCPv6	
extended server binding	
clearing.....	317
displaying.....	308
extended server statistics	
clearing.....	319
displaying.....	314
DHCPv6 client	
binding state	
clearing.....	317
displaying.....	308
statistics	
clearing.....	319
displaying.....	314
DHCPv6 local server	
lease timers.....	200
dhcpv6 statement.....	133
documentation	
comments on.....	xix
domain-name statement	
address-assignment pools.....	136
dynamic IP demux interface statements	
demux-source.....	116
family.....	167
dynamic IP demux statements	
mac-validate.....	197

dynamic PPP statements	
pap.....	201
ppp-options.....	206
dynamic PPPoE statements	
dynamic-profile.....	139
family.....	168
pppoe-options.....	204
pppoe-underlying-options.....	205
server.....	227
underlying-interface.....	238
unnumbered-address.....	249
dynamic profiles	
DHCPv4 retailer example.....	24
DHCPv4 wholesaler example.....	24
DHCPv6 retailer example.....	43
Layer 2 wholesale retailer example.....	87
PPPoE wholesaler example.....	62
retailer	
DHCPv4.....	19
DHCPv6.....	38
Layer 2 wholesale.....	75
wholesale network	
DHCPv4.....	17
DHCPv6.....	37
PPPoE.....	59
wholesaler	
DHCP.....	37
DHCPv4.....	18
PPPoE.....	60
dynamic profiles statements	
dynamic-profiles.....	142
interface.....	184
interfaces.....	188
keepalives.....	196
routing-instances.....	224
unnumbered-address.....	250
vlan-tags.....	258
dynamic subscribers	
interfaces statement.....	188
pppoe-underlying-options statement.....	205
dynamic underlying interface statements	
demux-source.....	117
dynamic VLANs.....	110
dynamic-profile statement	
DHCP local server.....	137
DHCP relay agent.....	138
dynamic PPPoE.....	139
stacked VLAN ranges.....	140
VLAN ranges.....	141
dynamic-profiles	
interfaces statement.....	188
dynamic IP demux.....	188
dynamic-profiles statement.....	142
E	
encapsulation statement	
dynamic profiles.....	150
logical interfaces.....	153
physical interface.....	157
Ethernet interfaces	
dynamic VLANs.....	110
mixed VLAN tagging.....	171
status information, displaying	
aggregated.....	327
Fast Ethernet.....	338
Gigabit Ethernet.....	354
VLAN IDs.....	255
extended DHCP	
configuring	
local server.....	42
F	
family statement.....	162
address-assignment pools.....	166
dynamic IP demux interface.....	167
dynamic PPPoE.....	168
dynamic profiles.....	169
Fast Ethernet interfaces	
dynamic VLANs.....	110
status information, displaying.....	338
VLAN IDs.....	255
firewall filters	
displaying.....	397
flexible-vlan-tagging statement.....	171
font conventions.....	xvii
forwarding options	
configuring.....	22
DHCP relay agent.....	295, 300, 303, 305
DHCPv4 example.....	25
forwarding-options statement.....	172
G	
Gigabit Ethernet interfaces	
demultiplexing interface information,	
displaying.....	388
dynamic VLANs.....	110
status information, displaying.....	354
VLAN IDs.....	255

grace-period statement.....	173
group statement	
DHCP local server.....	174
DHCP relay agent.....	177

I

input-vlan-map statement	
dynamic interfaces.....	180
instance-role statement.....	192
instance-type statement.....	193
interface	
direct ISP-facing.....	89
Layer 2 wholesale example.....	87
NNI ISP-facing.....	89
interface statement	
DHCP local server.....	181
DHCP relay agent.....	183
dynamic profiles.....	184
IS-IS.....	185
OSPF.....	185
VPNs.....	185
interface-mac-limit statement.....	186
interfaces	
loopback	
DHCPv4 Layer 3 wholesale.....	11
DHCPv6 Layer 3 wholesale.....	32
PPPoE Layer 3 wholesale.....	55
mixed VLAN tagging.....	171
unit statement.....	247
interfaces statement.....	187
dynamic profiles.....	188
IP demultiplexing interface statements	
unit.....	245
ip-address-first statement.....	195

K

keepalives statement	
dynamic profiles.....	196

L

Layer 2 VPNs	
VPLS connections, displaying.....	443
VPLS statistics, displaying.....	465
Layer 2 wholesale	
access.....	86
access profile.....	86
basic topology.....	71
configuration elements.....	71

configuration options	
DHCPv4.....	4, 68
DHCPv6.....	4, 68
direct interface.....	89
dynamic profiles.....	75
retailer example.....	87
interface.....	87
NNI interface.....	89
overview.....	3, 49, 67
reference topology.....	73
routing instances.....	81, 84, 88

Layer 3 wholesale

access	
DHCP.....	15, 35
PPPoE.....	57
RADIUS server.....	15, 35, 58, 86
access profile	
DHCP.....	16, 36
PPPoE.....	58, 59
address assignment pool	
DHCPv6 example.....	44, 45
basic topology	
DHCPv4.....	9
DHCPv6.....	29
PPPoE.....	53
configuration elements	
DHCPv4.....	9
DHCPv6.....	29
PPPoE.....	53
configuration options	
DHCPv4.....	4, 68
DHCPv6.....	4, 68
PPPoE.....	4, 68
DHCP support.....	5
DHCPv4	
forwarding options example.....	25
dynamic profiles	
DHCP.....	37
DHCPv4.....	17, 18, 19
DHCPv4 retailer example.....	24
DHCPv4 wholesaler example.....	24
DHCPv6.....	37, 38
DHCPv6 retailer example.....	43
PPPoE.....	59, 60
PPPoE wholesaler example.....	62
dynamic VLAN	
DHCPv4.....	13
DHCPv6.....	33

forwarding options	
configuring.....	22
interface support.....	5, 50
local server	
DHCPv6 example.....	45
overview.....	3, 49, 67
RADIUS VSAs.....	6, 51, 69
reference topology	
DHCPv4.....	10
DHCPv6.....	31
PPPoE.....	55
routing instances	
DHCP.....	20
DHCPv4 example.....	26
DHCPv6.....	40
DHCPv6 example.....	44
PPPoE.....	61
PPPoE example.....	63
static VLAN	
DHCPv4.....	12
DHCPv6.....	33
PPPoE.....	57
vans	
DHCPv4.....	12
DHCPv6.....	32
local server	
configuring DHCP.....	42
DHCPv6 example.....	45
logical interface statements	
family.....	169
logical interfaces	
VLAN IDs.....	255
logical systems	
subscriber relationship with.....	6, 50
loopback interface	
Layer 3 wholesale	
DHCPv4.....	11
DHCPv6.....	32
PPPoE.....	55
status information, displaying.....	320
M	
mac-validate statement	
dynamic IP demux interface.....	197
manuals	
comments on.....	xix
maximum-lease-time statement.....	200
mixed VLAN tagging.....	171
multicast-replication statement.....	198
N	
no-tunnel-services statement.....	199
O	
OSPF	
enabling.....	185
output-vlan-map statement	
dynamic interfaces.....	201
P	
pap statement	
dynamic PPP.....	201
parentheses, in syntax descriptions.....	xviii
physical interfaces	
dynamic VLANs.....	110
mixed VLAN tagging.....	171
pool statement	
address-assignment pools.....	202
pool-match-order statement.....	203
pop statement	
dynamic VLAN interfaces.....	203
PPP	
interfaces, displaying.....	407
ppp-options statement	
dynamic PPP.....	206
PPPoE	
interfaces, displaying.....	378
pppoe-options statement	
dynamic PPPoE.....	204
pppoe-underlying-options statement	
static and dynamic PPPoE.....	205
prefix statement	
address-assignment pools.....	207
profile statement	
subscriber access.....	208
protocols statement.....	212
proxy-arp statement.....	215
subscriber interfaces.....	216
pseudowires	
mixed VLAN tagging.....	171
push statement	
dynamic VLAN interfaces.....	216
R	
RADIUS	
wholesale VSA support.....	6, 51, 69
RADIUS server	
access configuration	
configuring.....	15, 35, 58, 86

radius statement	
subscriber access.....	217
radius-server statement.....	219
range statement	
address-assignment pools.....	220
ranges statement	
VLAN.....	221
route-distinguisher statement.....	222
routing information	
interfaces	
state, displaying.....	399
routing instances	
retailer	
DHCPv4.....	20
DHCPv4 example.....	26
DHCPv6.....	40
DHCPv6 example.....	44
Layer 2 wholesale.....	81, 84
Layer 2 wholesale example.....	88
PPPoE.....	61
PPPoE example.....	63
subscriber relationship with.....	6, 50
routing-instances statement.....	226
dynamic profiles.....	224

S

secret statement	
access.....	227
server statement	
dynamic PPPoE.....	227
server-group statement	
DHCP relay agent.....	228
services statement	
remote router access.....	229
show dhcp relay binding command.....	295
show dhcp relay statistics command.....	300
show dhcp server binding command.....	281
show dhcp server statistics command.....	287
show dhcpv6 server binding command.....	308
show dhcpv6 server statistics command.....	314
show interfaces (Aggregated Ethernet)	
command.....	327
show interfaces (Fast Ethernet) command.....	338
show interfaces (Gigabit Ethernet) command.....	354
show interfaces (Loopback) command.....	320
show interfaces (PPPoE) command.....	378
show interfaces demux0 (Demux Interfaces)	
command.....	388
show interfaces filters command.....	397

show interfaces routing command.....	399
show interfaces routing-instance command.....	405
show network-access aaa statistics authentication	
command.....	273
show network-access aaa statistics	
command.....	268
show network-access aaa subscribers	
command.....	276
show network-access address-assignment pool	
command.....	279
show ppp interface command.....	407
show subscribers command.....	417
show subscribers summary command.....	437
show vpls connections command.....	443
show vpls flood event-queue command.....	454
show vpls flood instance command.....	456
show vpls flood route command.....	458
show vpls mac-table command.....	460
show vpls statistics command.....	465
site statement.....	231
site-identifier statement.....	232
site-range statement.....	232
stacked VLAN-tag framing	
VLAN ID list.....	259
stacked-vlan-ranges statement.....	233
stacked-vlan-tagging statement.....	234
static PPPoE statements	
pppoe-underlying-options.....	205
static subscribers	
interfaces statement.....	188
pppoe-underlying-options statement.....	205
subscriber access	
subscriber information, displaying.....	417
subscriber summary information,	
displaying.....	437
subscriber interface statements	
demux-options.....	115
demux-source.....	116, 117
demux0.....	114
dynamic-profile.....	139
family.....	167, 168, 169
interfaces.....	188
mac-validate.....	197
pap.....	201
ppp-options.....	206
pppoe-options.....	204
pppoe-underlying-options.....	205
proxy-arp.....	216
server.....	227

- underlying-interface.....237, 238
 - unit.....245, 247
 - unnumbered-address.....249, 250
 - subscribers
 - displaying.....417
 - displaying summary.....437
 - support, technical See technical support
 - syntax conventions.....xvii
 - system statement.....234
- T**
- technical support
 - contacting JTAC.....xix
 - topology
 - Layer 2 wholesale network.....71
 - Layer 3 wholesale network
 - DHCPv4.....9
 - DHCPv6.....29
 - PPPoE.....53
 - traceoptions statement
 - DHCP local server.....235
 - DHCP relay agent.....235
- U**
- underlying-interface statement
 - dynamic PPPoE.....238
 - dynamic profiles.....237
 - unit statement.....239
 - demux interfaces.....245
 - interfaces.....247
 - unnumbered-address statement
 - dynamic PPPoE.....249
 - dynamic profiles.....250
 - Ethernet.....252
 - user-prefix statement
 - DHCP local server.....253
- V**
- virtual private LAN service See VPLS
 - VLAN
 - DHCPv4 wholesale
 - dynamic configuration.....13
 - static configuration.....12
 - DHCPv6 wholesale
 - dynamic configuration.....33
 - static configuration.....33
 - PPPoE wholesale
 - static configuration.....57
 - vlan-id statement
 - 802.1Q VLANs.....255
 - dynamic VLAN map
 - rewriting at ingress or egress.....254
 - Ethernet interfaces.....255
 - vlan-model statement.....256
 - vlan-ranges statement.....257
 - vlan-tags statement
 - dynamic profiles.....258
 - stacked VLAN tags.....259
 - VLANs
 - configuring VLAN ranges.....233, 257
 - VPLS
 - connections, displaying.....443
 - statistics, displaying.....465
 - vpls statement
 - routing-instances.....261
 - vrf-export statement.....263
 - vrf-import statement.....264
 - vrf-target statement.....265
- W**
- wholesale See Layer 2 wholesale See Layer 3 wholesale

