



---

Junos<sup>®</sup> OS

# Broadband Subscriber Management Wholesale Feature Guide

Release  
16.2



---

Modified: 2016-11-02

Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://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<sup>®</sup> OS Broadband Subscriber Management Wholesale Feature Guide*  
16.2  
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 . . . . .	xvii
	Documentation and Release Notes . . . . .	xvii
	Supported Platforms . . . . .	xvii
	Using the Examples in This Manual . . . . .	xvii
	Merging a Full Example . . . . .	xviii
	Merging a Snippet . . . . .	xviii
	Documentation Conventions . . . . .	xix
	Documentation Feedback . . . . .	xxi
	Requesting Technical Support . . . . .	xxi
	Self-Help Online Tools and Resources . . . . .	xxi
	Opening a Case with JTAC . . . . .	xxii
<b>Part 1</b>	<b>Configuring DHCP Layer 3 Wholesale Networks</b>	
<b>Chapter 1</b>	<b>Subscriber Management DHCP Layer 3 Wholesale Overview . . . . .</b>	<b>3</b>
	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
<b>Chapter 2</b>	<b>Configuring DHCPv4 Layer 3 Wholesale Networks . . . . .</b>	<b>9</b>
	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
<b>Configuring DHCPv6 Layer 3 Wholesale Networks . . . . .</b>	<b>29</b>
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

<b>Part 2</b>	<b>Configuring PPPoE Layer 3 Wholesale Networks</b>	
<b>Chapter 4</b>	<b>Subscriber Management PPPoE Wholesale Overview</b>	<b>49</b>
	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
<b>Chapter 5</b>	<b>Configuring PPPoE Layer 3 Wholesale Networks</b>	<b>53</b>
	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
<b>Part 3</b>	<b>Configuring Layer 2 Wholesale Networks</b>	
<b>Chapter 6</b>	<b>Subscriber Management Layer 2 Wholesale Overview</b>	<b>67</b>
	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
<b>Chapter 7</b>	<b>Configuring Layer 2 Wholesale Networks</b>	<b>71</b>
	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
<b>Part 4</b>	<b>Configuring ANCP-Triggered Layer 2 Wholesale Services</b>	
<b>Chapter 8</b>	<b>ANCP-Triggered Layer 2 Wholesale Service Overview . . . . .</b>	<b>93</b>
	Layer 2 Wholesale with ANCP-Triggered VLANs Overview . . . . .	93
	RADIUS Authorization for ANCP-Triggered VLANs . . . . .	96
	Instantiation of an ANCP-Triggered, Autosensed, Dynamic VLAN . . . . .	96
	RADIUS Interim Accounting Updates . . . . .	97
	Removal of the Layer 2 Wholesale Service . . . . .	98
	Interactions Between In-Band and Out-of-Band VLAN Autosensing . . . . .	99
	Migration of Subscriber Ownership from Wholesaler to Retailer . . . . .	101
	Migration of Subscriber Ownership from Retailer to Wholesaler . . . . .	102
	Migration of Subscriber Ownership Between Retailers . . . . .	102
	Modification of the Access Line Identifier or Port VLAN Identifier . . . . .	102
	Consequences of a State Transition in the Access-Facing Physical Interface . . . . .	104
	Consequences of a State Transition from Up to Down in the Core-Facing Physical Interface . . . . .	105
	Consequences of a State Transition from Down to Up in the Core-Facing Physical Interface . . . . .	106
	Loss of ANCP TCP Adjacency . . . . .	107
	Junos OS Predefined Variables . . . . .	109
	Juniper Networks VSAs Supported by the AAA Service Framework . . . . .	129
	AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS . . . . .	141
	AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS . . . . .	147
<b>Chapter 9</b>	<b>Configuring ANCP-Triggered Layer 2 Wholesale Services . . . . .</b>	<b>153</b>
	Configuring ANCP Neighbors . . . . .	153
	Configuring the ANCP Agent for ANCP-Triggered, Autosensed Dynamic VLANs . . . . .	154
	Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs . . . . .	155

	Configuring Out-of-Band ANCP Messages to Trigger Dynamic VLAN Instantiation . . . . .	156
	Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages . . .	157
	Configuring the ANCP Agent to Dampen the Effects of Short-Term Adjacency Losses . . . . .	159
	Reestablishing Pending Access Line Sessions for Layer 2 Wholesale . . . . .	160
	Clearing ANCP Access Loops . . . . .	160
<b>Chapter 10</b>	<b>Configuring Flat-File Accounting for Layer 2 Wholesale Services . . . . .</b>	<b>163</b>
	Flat-File Accounting Overview . . . . .	163
	Configuring Flat-File Accounting for Layer 2 Wholesale . . . . .	166
	Configuring Flat-File Accounting for Extensible Subscriber Services Management . . . . .	170
<b>Part 5</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 11</b>	<b>Configuration Statements . . . . .</b>	<b>177</b>
	accept . . . . .	182
	accept-out-of-band . . . . .	183
	access-profile . . . . .	184
	access-profile (Dynamic VLAN) . . . . .	185
	access-profile (Dynamic Stacked VLAN) . . . . .	186
	accounting-server . . . . .	186
	active-server-group . . . . .	187
	address . . . . .	188
	address-assignment (Address-Assignment Pools) . . . . .	191
	adjacency-loss-hold-time (ANCP) . . . . .	192
	ancp . . . . .	193
	authentication . . . . .	195
	authentication (DHCP Local Server) . . . . .	196
	authentication (DHCP Relay Agent) . . . . .	197
	authentication-order . . . . .	198
	authentication-server . . . . .	199
	auto-configure . . . . .	200
	auto-configure-trigger interface (ANCP) . . . . .	201
	backup-on-failure (Accounting Options) . . . . .	202
	circuit-id (VLAN Authentication Username) . . . . .	203
	cleanup-interval (Accounting Options) . . . . .	204
	compress (Accounting Options) . . . . .	204
	connectivity-type . . . . .	205
	core-facing . . . . .	206
	demux0 (Dynamic Interface) . . . . .	207
	demux-options (Dynamic Interface) . . . . .	208
	demux-source (Dynamic IP Demux Interface) . . . . .	209
	demux-source (Dynamic Underlying Interface) . . . . .	210
	demux-source (Underlying Interface) . . . . .	211
	dhcp-attributes (Address-Assignment Pools) . . . . .	212
	dhcp-local-server . . . . .	213
	dhcp-relay . . . . .	219
	dhcpv6 (DHCP Local Server) . . . . .	226

domain-name (Address-Assignment Pools) .....	229
dynamic-profile (DHCP Local Server) .....	230
dynamic-profile (DHCP Relay Agent) .....	231
dynamic-profile (Dynamic PPPoE) .....	232
dynamic-profile (Stacked VLAN) .....	233
dynamic-profile (VLAN) .....	234
dynamic-profiles .....	235
egress-stats (Flat-File Accounting Options) .....	243
encapsulation (Dynamic Interfaces) .....	244
encapsulation (Logical Interface) .....	247
encapsulation (Physical Interface) .....	251
exclude (RADIUS) .....	257
family .....	262
family (Address-Assignment Pools) .....	267
family (Dynamic Demux Interface) .....	268
family (Dynamic PPPoE) .....	269
family (Dynamic Standard Interface) .....	270
fields (Flat-File Accounting Options) .....	272
file (Flat-File Accounting Options) .....	273
flat-file-profile (Accounting Options) .....	274
flat-file-profile (Extensible Subscriber Services) .....	276
flexible-vlan-tagging .....	277
format (Flat-File Accounting Options) .....	278
forwarding-options .....	279
general-param (Flat-File Accounting Options) .....	280
grace-period .....	281
group (DHCP Local Server) .....	282
group (DHCP Relay Agent) .....	285
ingress-stats (Flat-File Accounting Options) .....	288
inner-vlan-id (Dynamic VLANs) .....	289
inner-vlan-id-swap-ranges .....	290
input-vlan-map (Dynamic Interfaces) .....	290
interface (DHCP Local Server) .....	291
interface (DHCP Relay Agent) .....	293
interface (Dynamic Routing Instances) .....	294
interface (Routing Instances) .....	295
interface-mac-limit (VPLS) .....	296
interfaces .....	297
interfaces (Static and Dynamic Subscribers) .....	298
interval (Flat-File Accounting Options) .....	302
instance-role .....	303
instance-type .....	304
ip-address-first .....	306
keepalives (Dynamic Profiles) .....	307
l2-stats (Flat-File Accounting Options) .....	308
mac-validate (Dynamic IP Demux Interface) .....	309
multicast-replication .....	310
neighbor (Define ANCP) .....	311
no-local-switching .....	311

no-tunnel-services	312
maximum-lease-time	313
overall-packet (Flat-File Accounting Options)	314
output-vlan-map (Dynamic Interfaces)	315
pap (Dynamic PPP)	316
pool (Address-Assignment Pools)	317
pool-match-order	318
pop (Dynamic VLANs)	318
pppoe-options (Dynamic PPPoE)	319
pppoe-underlying-options (Static and Dynamic Subscribers)	320
ppp-options (Dynamic PPP)	321
prefix (Address-Assignment Pools)	322
profile (Access)	323
protocols	327
proxy-arp	330
proxy-arp (Dynamic Profiles)	331
push (Dynamic VLANs)	331
push-backup-to-master (Accounting Options)	332
radius (Access Profile)	333
radius-server	335
range (Address-Assignment Pools)	336
ranges (Dynamic VLAN)	337
remote-id (VLAN Authentication Username)	337
route-distinguisher	338
routing-instances (Dynamic Profiles)	340
routing-instances (Multiple Routing Entities)	342
schema-version (Flat-File Accounting Options)	343
secret	344
server (Dynamic PPPoE)	344
server-group	345
services (System Services)	346
site (VPLS Multihoming for FEC 128)	349
site-identifier (VPLS)	349
site-range	350
stacked-vlan-ranges	351
stacked-vlan-tagging	352
system	352
traceoptions (DHCP)	353
underlying-interface (demux0)	355
underlying-interface (Dynamic PPPoE)	356
unit	357
unit (Dynamic Demux Interface)	364
unit (Dynamic Profiles Standard Interface)	366
unnumbered-address (Dynamic PPPoE)	368
unnumbered-address (Dynamic Profiles)	369
unnumbered-address (Ethernet)	371
username-include	372
user-prefix (DHCP Local Server)	373
vlan-id (Dynamic VLANs)	374

	vlan-id (VLAN ID to Be Bound to a Logical Interface) . . . . .	375
	vlan-model . . . . .	376
	vlan-ranges . . . . .	377
	vlan-tags . . . . .	378
	vlan-tags (Stacked VLAN Tags) . . . . .	379
	vpls (Routing Instance) . . . . .	381
	vrf-export . . . . .	383
	vrf-import . . . . .	384
	vrf-target . . . . .	385
<b>Chapter 12</b>	<b>Operational Commands . . . . .</b>	<b>387</b>
	AAA and Address Assignment Pool Commands . . . . .	387
	show network-access aaa statistics . . . . .	388
	show network-access aaa statistics authentication . . . . .	397
	show network-access aaa subscribers . . . . .	400
	show network-access address-assignment pool . . . . .	403
	ANCP and ANCP-Triggered Wholesale Commands . . . . .	404
	clear ancp access-loop . . . . .	405
	clear ancp neighbor . . . . .	406
	clear network-access aaa subscriber . . . . .	408
	request ancp oam port-down . . . . .	409
	request ancp oam port-up . . . . .	411
	request auto-configuration reconnect-pending . . . . .	412
	show ancp neighbor . . . . .	413
	show auto-configuration out-of-band pending . . . . .	421
	DHCPv4 Local Server Commands . . . . .	421
	show dhcp server binding . . . . .	422
	show dhcp server statistics . . . . .	428
	clear dhcp server binding . . . . .	431
	clear dhcp server statistics . . . . .	434
	DHCPv4 Relay Commands . . . . .	435
	show dhcp relay binding . . . . .	436
	show dhcp relay statistics . . . . .	441
	clear dhcp relay binding . . . . .	445
	clear dhcp relay statistics . . . . .	447
	DHCPv6 Local Server Commands . . . . .	449
	show dhcpv6 server binding . . . . .	450
	show dhcpv6 server statistics . . . . .	456
	clear dhcpv6 server binding . . . . .	459
	clear dhcpv6 server statistics . . . . .	461
	Interface Commands . . . . .	461
	show interfaces (Loopback) . . . . .	463
	show interfaces (Aggregated Ethernet) . . . . .	470
	show interfaces (Fast Ethernet) . . . . .	481
	show interfaces (Gigabit Ethernet) . . . . .	497
	show interfaces (PPPoE) . . . . .	521
	show interfaces demux0 (Demux Interfaces) . . . . .	531
	show interfaces filters . . . . .	540
	show interfaces l2-routing-instance . . . . .	542

show interfaces routing .....	544
show interfaces routing-instance .....	550
show ppp interface .....	552
Subscriber Commands .....	561
show subscribers .....	562
show subscribers summary .....	582
VPLS Commands .....	587
show vpls connections .....	588
show vpls flood event-queue .....	599
show vpls flood instance .....	601
show vpls flood route .....	603
show vpls mac-table .....	605
show vpls statistics .....	610

## Part 6

## Index

Index .....	615
-------------	-----



# List of Figures

<b>Part 1</b>	<b>Configuring DHCP Layer 3 Wholesale Networks</b>	
<b>Chapter 2</b>	<b>Configuring DHCPv4 Layer 3 Wholesale Networks . . . . .</b>	<b>9</b>
	Figure 1: Basic Subscriber Management Layer 3 Wholesale Solution Topology . . . . .	10
	Figure 2: DHCPv4 Layer 3 Wholesale Network Reference Topology . . . . .	11
<b>Chapter 3</b>	<b>Configuring DHCPv6 Layer 3 Wholesale Networks . . . . .</b>	<b>29</b>
	Figure 3: Basic Subscriber Management DHCPv6 Layer 3 Wholesale Solution Topology . . . . .	30
	Figure 4: DHCPv6 Layer 3 Wholesale Network Reference Topology . . . . .	31
<b>Part 2</b>	<b>Configuring PPPoE Layer 3 Wholesale Networks</b>	
<b>Chapter 5</b>	<b>Configuring PPPoE Layer 3 Wholesale Networks . . . . .</b>	<b>53</b>
	Figure 5: Basic Subscriber Management PPPoE Layer 3 Wholesale Solution Topology . . . . .	54
	Figure 6: PPPoE Layer 3 Wholesale Network Reference Topology . . . . .	55
<b>Part 3</b>	<b>Configuring Layer 2 Wholesale Networks</b>	
<b>Chapter 7</b>	<b>Configuring Layer 2 Wholesale Networks . . . . .</b>	<b>71</b>
	Figure 7: Basic Subscriber Management Layer 2 Wholesale Solution Topology . . . . .	72
	Figure 8: Layer 2 Wholesale Network Reference Topology . . . . .	74
<b>Part 4</b>	<b>Configuring ANCP-Triggered Layer 2 Wholesale Services</b>	
<b>Chapter 8</b>	<b>ANCP-Triggered Layer 2 Wholesale Service Overview . . . . .</b>	<b>93</b>
	Figure 9: Sample Layer 2 Wholesale Access Topology . . . . .	94



# List of Tables

	<b>About the Documentation</b> . . . . .	<b>xvii</b>
	Table 1: Notice Icons . . . . .	xix
	Table 2: Text and Syntax Conventions . . . . .	xix
<b>Part 1</b>	<b>Configuring DHCP Layer 3 Wholesale Networks</b>	
<b>Chapter 1</b>	<b>Subscriber Management DHCP Layer 3 Wholesale Overview</b> . . . . .	<b>3</b>
	Table 3: Required Juniper Networks VSAs for the Broadband Subscriber Management Wholesale Network Solution . . . . .	6
<b>Part 2</b>	<b>Configuring PPPoE Layer 3 Wholesale Networks</b>	
<b>Chapter 4</b>	<b>Subscriber Management PPPoE Wholesale Overview</b> . . . . .	<b>49</b>
	Table 4: Required Juniper Networks VSAs for the Broadband Subscriber Management Wholesale Network Solution . . . . .	51
<b>Part 3</b>	<b>Configuring Layer 2 Wholesale Networks</b>	
<b>Chapter 6</b>	<b>Subscriber Management Layer 2 Wholesale Overview</b> . . . . .	<b>67</b>
	Table 5: Required Juniper Networks VSAs for the Broadband Subscriber Management Wholesale Network Solution . . . . .	69
<b>Chapter 7</b>	<b>Configuring Layer 2 Wholesale Networks</b> . . . . .	<b>71</b>
	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
<b>Part 4</b>	<b>Configuring ANCP-Triggered Layer 2 Wholesale Services</b>	
<b>Chapter 8</b>	<b>ANCP-Triggered Layer 2 Wholesale Service Overview</b> . . . . .	<b>93</b>
	Table 9: Junos OS Predefined Variables and Definitions . . . . .	109
	Table 10: Supported Juniper Networks VSAs . . . . .	130
	Table 11: AAA Access Messages—Supported RADIUS Attributes and Juniper Networks VSAs . . . . .	141
	Table 12: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs . . . . .	147
<b>Chapter 10</b>	<b>Configuring Flat-File Accounting for Layer 2 Wholesale Services</b> . . . . .	<b>163</b>
	Table 13: Value of Elements in Sample Accounting Flat File XML Header . . . . .	164
<b>Part 5</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 12</b>	<b>Operational Commands</b> . . . . .	<b>387</b>

Table 14: show network-access aaa statistics Output Fields . . . . .	389
Table 15: show network-access aaa statistics authentication Output Fields . . .	397
Table 16: show network-access aaa subscribers Output Fields . . . . .	400
Table 17: show network-access address-assignment pool Output Fields . . . . .	403
Table 18: show ancp neighbor Output Fields . . . . .	413
Table 19: show auto-configuration out-of-band pending Output Fields . . . . .	421
Table 20: show dhcp server binding Output Fields . . . . .	423
Table 21: show dhcp server statistics Output Fields . . . . .	429
Table 22: show dhcp relay binding Output Fields . . . . .	437
Table 23: show dhcp relay statistics Output Fields . . . . .	442
Table 24: clear dhcp relay statistics Output Fields . . . . .	448
Table 25: show dhcpv6 server binding Output Fields . . . . .	451
Table 26: show dhcpv6 server statistics Output Fields . . . . .	457
Table 27: Loopback show interfaces Output Fields . . . . .	463
Table 28: Aggregated Ethernet show interfaces Output Fields . . . . .	470
Table 29: show interfaces Fast Ethernet Output Fields . . . . .	481
Table 30: show interfaces (Gigabit Ethernet) Output Fields . . . . .	498
Table 31: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type . . . . .	512
Table 32: show interfaces (PPPoE) Output Fields . . . . .	521
Table 33: show interfaces demux0 (Demux Interfaces) Output Fields . . . . .	531
Table 34: show interfaces filters Output Fields . . . . .	540
Table 35: show interfaces l2-routing-instance Output Fields . . . . .	542
Table 36: show interfaces routing Output Fields . . . . .	544
Table 37: show ppp interface Output Fields . . . . .	552
Table 38: show subscribers Output Fields . . . . .	565
Table 39: show subscribers summary Output Fields . . . . .	583
Table 40: show vpls connections Output Fields . . . . .	589
Table 41: show vpls flood event-queue Output Fields . . . . .	599
Table 42: show vpls flood instance Output Fields . . . . .	601
Table 43: show vpls flood route Output Fields . . . . .	603
Table 44: show vpls mac-table Output fields . . . . .	606
Table 45: show vpls statistics Output Fields . . . . .	610

# About the Documentation

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

## 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 [CLI Explorer](#).

## Documentation Conventions

[Table 1 on page xix](#) 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 xix](#) defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Bold text like this</b>	Represents text that you type.	To enter configuration mode, type the <b>configure</b> command:  user@host> <b>configure</b>

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> <b>show chassis alarms</b>  No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"><li>Introduces or emphasizes important new terms.</li><li>Identifies guide names.</li><li>Identifies RFC and Internet draft titles.</li></ul>	<ul style="list-style-type: none"><li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li><li><i>Junos OS CLI User Guide</i></li><li>RFC 1997, <i>BGP Communities Attribute</i></li></ul>
<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@# <b>set system domain-name</b> <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"><li>To configure a stub area, include the <b>stub</b> statement at the [edit protocols ospf area area-id] hierarchy level.</li><li>The console port is labeled <b>CONSOLE</b>.</li></ul>
< > (angle brackets)	Encloses optional keywords or variables.	<b>stub &lt;default-metric <i>metric</i>&gt;;</b>
(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.	<b>broadcast   multicast</b>  <b>(<i>string1</i>   <i>string2</i>   <i>string3</i>)</b>
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<b>rsvp { # Required for dynamic MPLS only</b>
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	<b>community name members [</b> <i>community-ids</i> <b>]</b>
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.	
<b>GUI Conventions</b>		
<b>Bold text like this</b>	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"><li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li><li>To cancel the configuration, click <b>Cancel</b>.</li></ul>

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 <b>Protocols&gt;Ospf</b> .

## 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](mailto: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.  <b>NOTE:</b> 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*.

## 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 on page 129](#) 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](#)

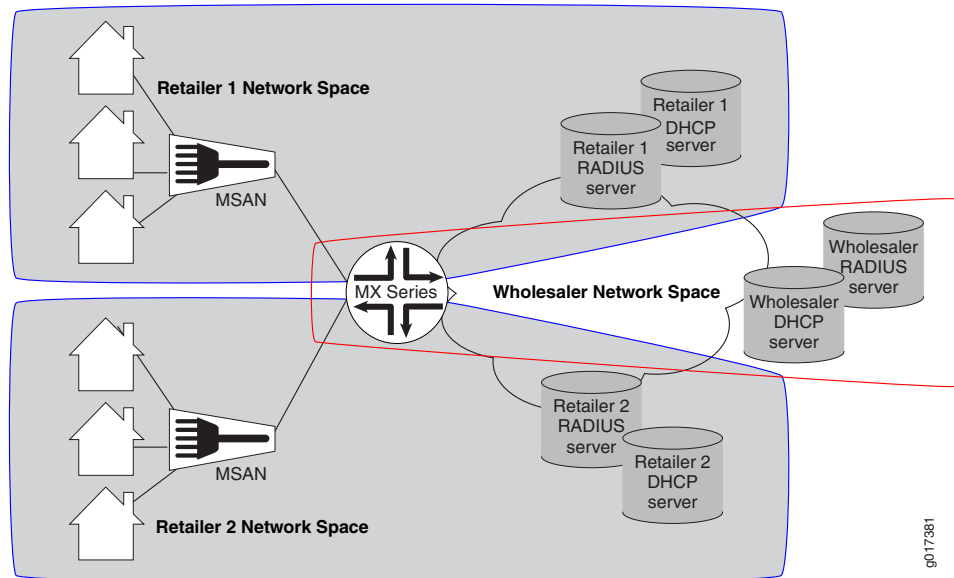
## Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements

---

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**

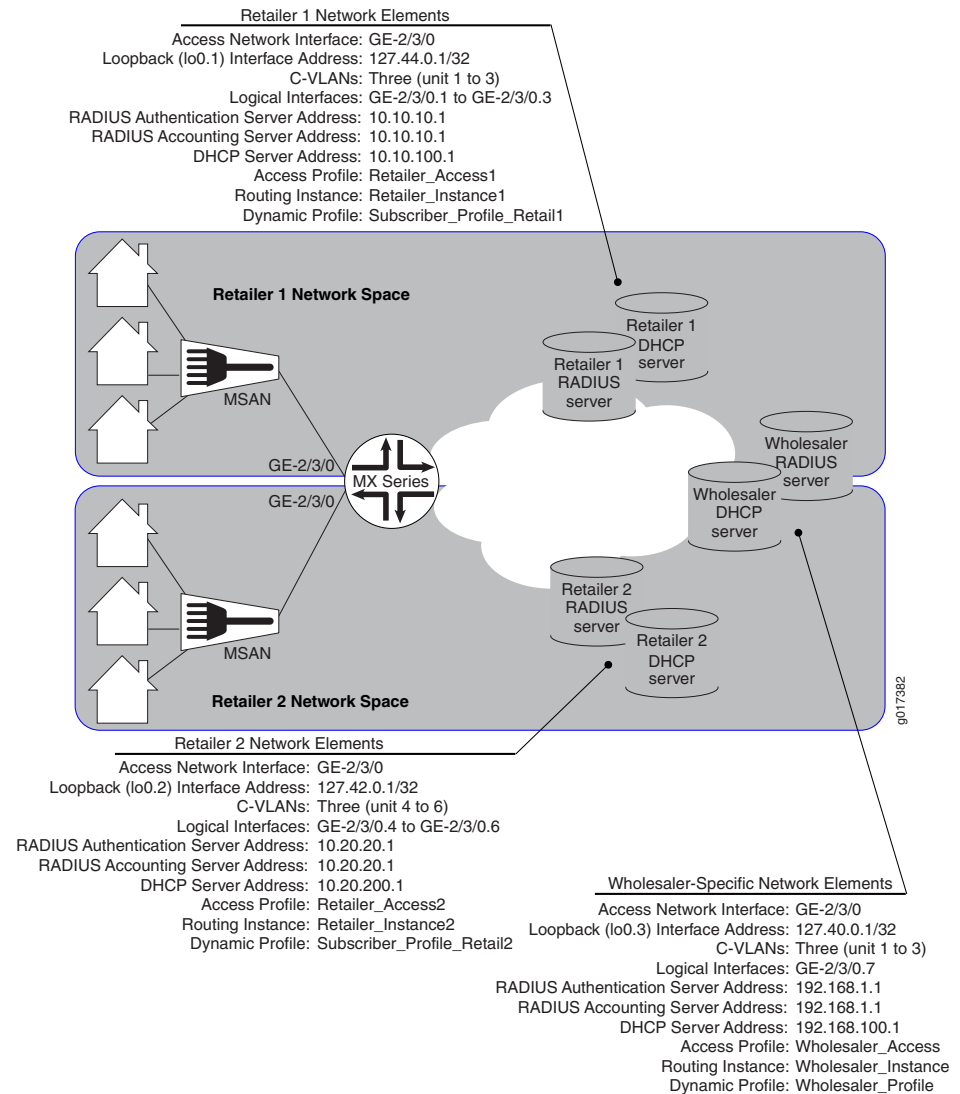
- [Layer 2 and Layer 3 Wholesale Overview on page 3](#)
- [DHCPv4 Layer 3 Wholesale Network Topology Overview on page 10](#)

## 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 127.40.0.1/32
```

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 127.42.0.1/32
```

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 127.44.0.1
```

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 127.33.0.1
```

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 $ABC123$ABC123$ABC123
user@host# set 10.10.10.1 secret $ABC123$ABC123
```

## 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 demux0 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

### 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 $ABC123;
            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 $ABC123;
            username-include {
                user-prefix WholesaleNetwork;
            }
        }
    }
    interface ge-2/3/0.1;
}

```

```
group Retailer1-Group {
  authentication {
    password $ABC123$ABC123;
    username-include {
      user-prefix WholesaleNetwork_Retailer1;
    }
  }
  interface ge-2/3/0.2;
}
group Retailer2-Group {
  authentication {
    password $ABC123$ABC123$ABC123;
    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 $ABC123$ABC123;
          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 $ABC123$ABC123;
            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 $ABC123$ABC123$ABC123;
                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 $ABC123$ABC123$ABC123;
            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

---

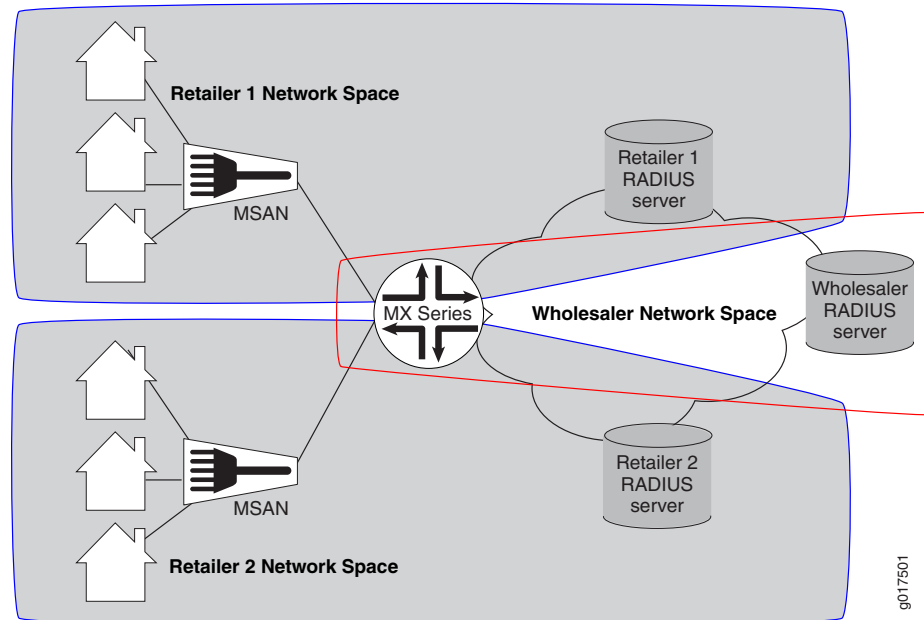
The network topology for the subscriber management DHCPv6 Layer 3 wholesale solution includes configuring separate routing instances for individual retailers that use a portion of the router. This solution uses a DHCPv6 local server configuration.



**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

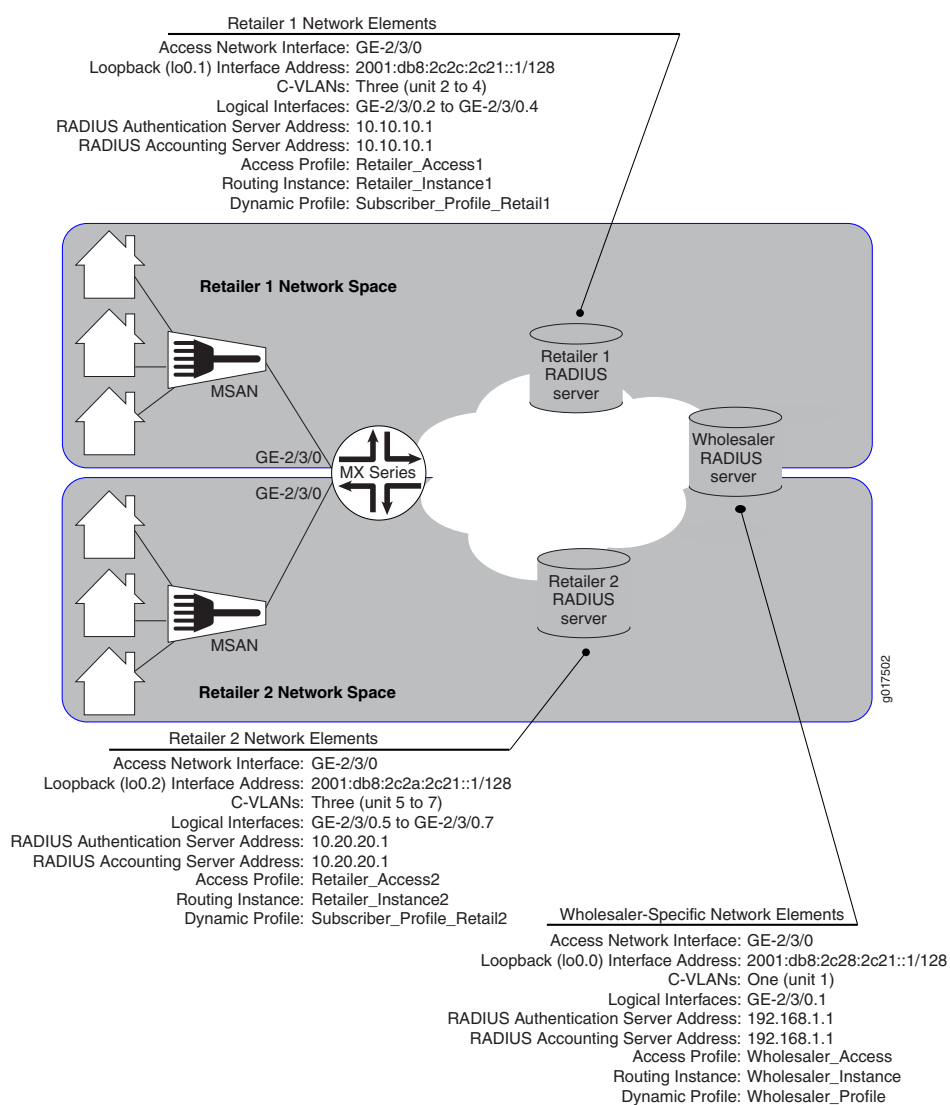
**Related Documentation**

- [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 2001:db8:2c28:2c21::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 2001:db8:2c2c: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
2001:db8:2c28:2c21::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
2001:db8:2c28:2c21::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 $ABC123$ABC123$ABC123
user@host# set 10.10.10.1 secret $ABC123$ABC123
```

## 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**

## 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 2001:db8:2121::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 2001:db8:2121::a/128
user@host# set range Range1 high 2001:db8:2121::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 2001:db8:2222::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 $ABC123
```

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 example.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 2001:db8:2121::0/64;
        range Range1 {
          low 2001:db8:2121::a/128;
          high 2001:db8:2121::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 2001:db8:2222::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 $ABC123;
      username-include {
        domain-name example.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*.

## 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 on page 129](#) 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

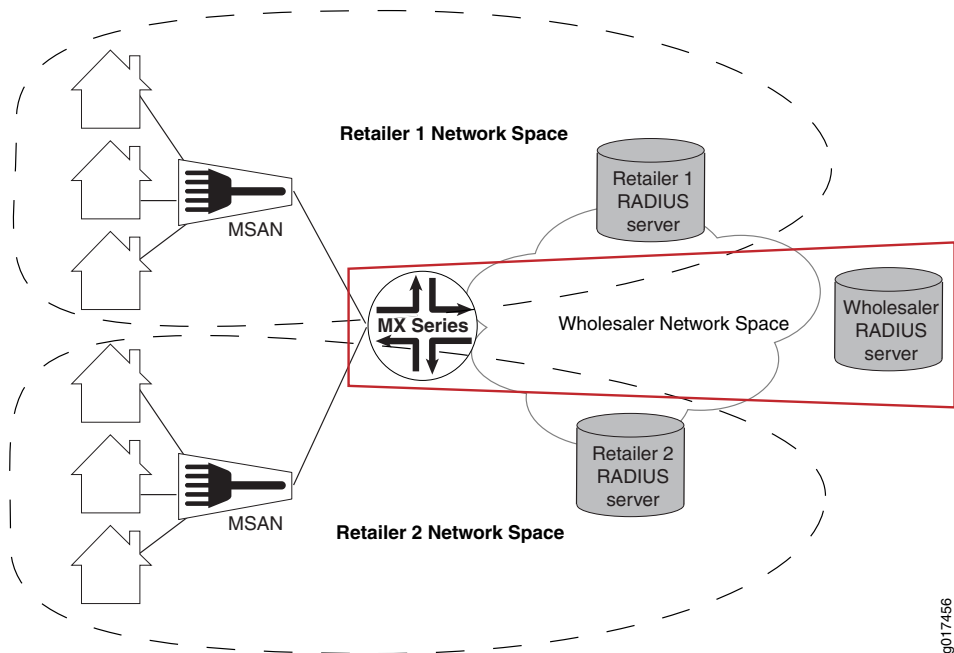
## Broadband Subscriber Management PPPoE Layer 3 Wholesale Topology and Configuration Elements

---

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.

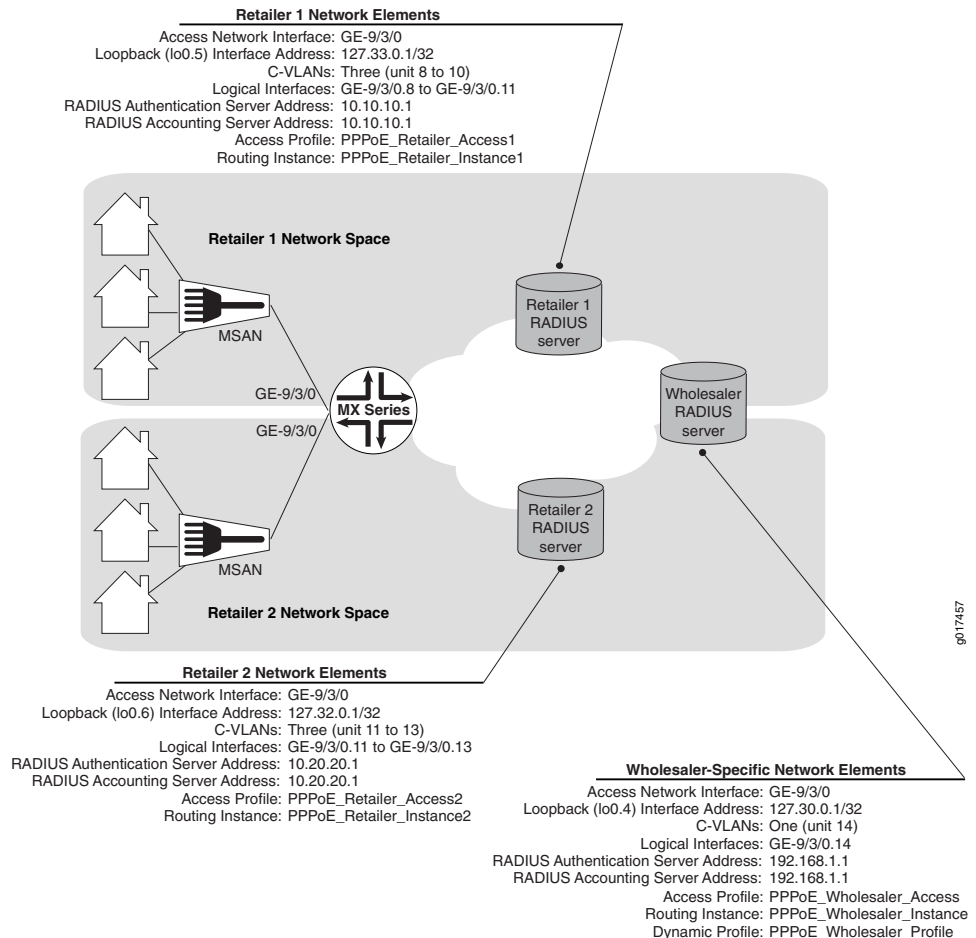
**Related Documentation**

- [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 127.30.0.1/32
```

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

```
[edit interfaces lo0 unit 4 family inet]
user@host# set address 127.30.0.2/32 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 127.33.0.1/32
```

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

```
[edit interfaces lo0 unit 5 family inet]
user@host# set address 127.33.0.2/32 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 $ABC123$ABC123$ABC123
user@host# set 10.10.10.1 secret $ABC123$ABC123
```

## 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

---

### 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.  <b>NOTE:</b> 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 on page 129](#) 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

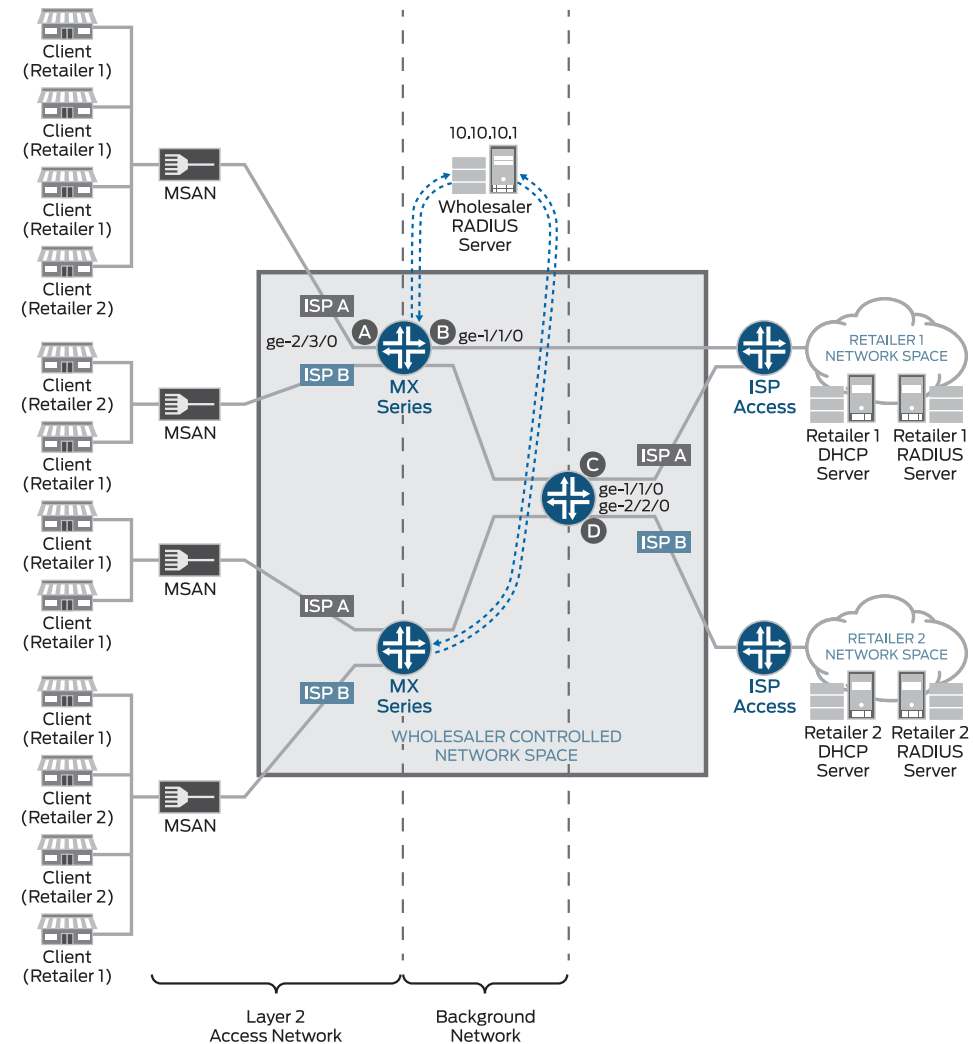
---

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.

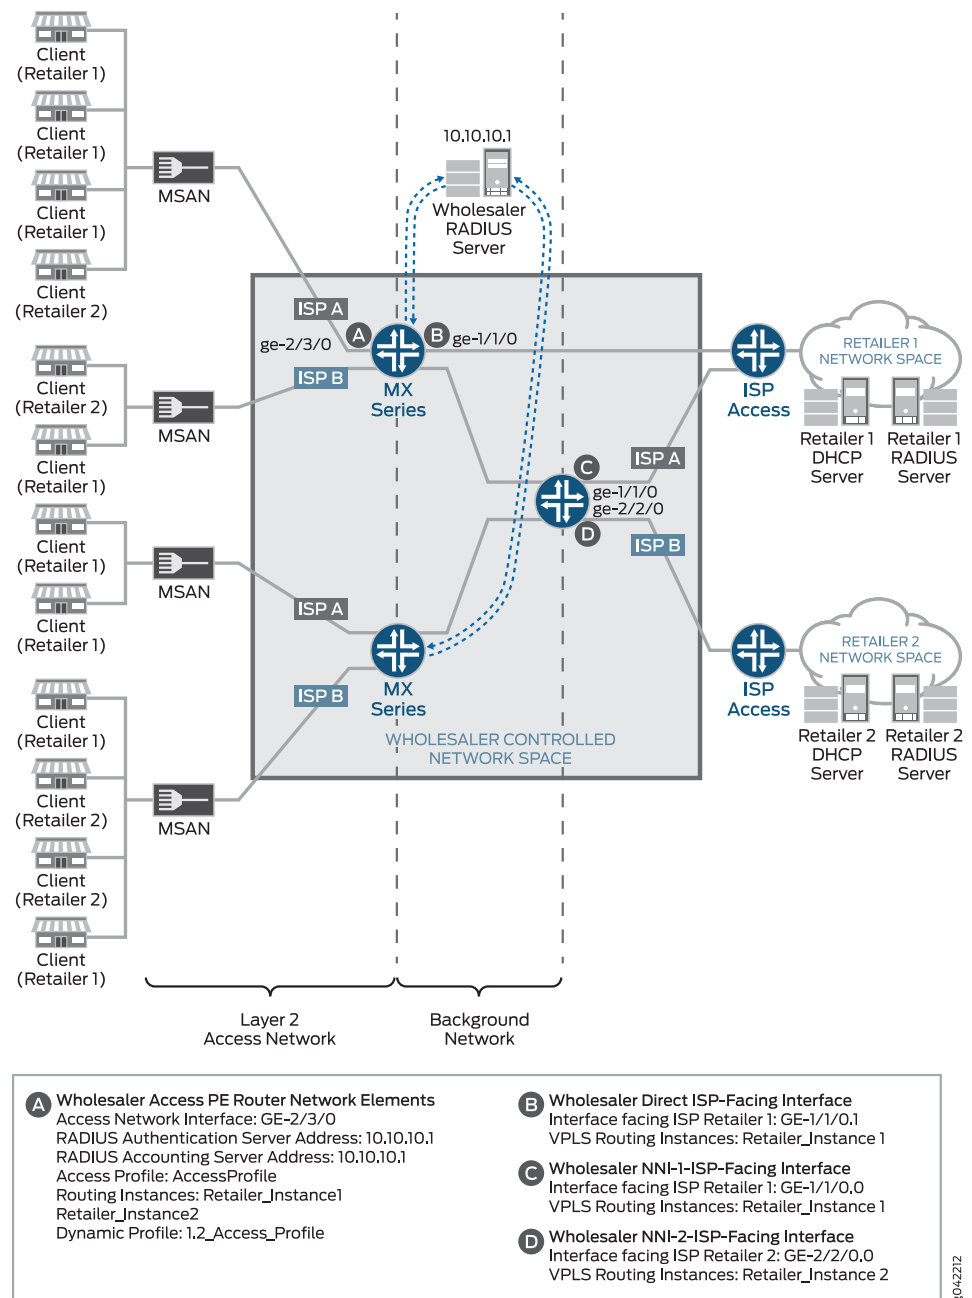
The example also shows two different connection options from one subscriber access router to one of the individual ISP access routers. One connection option uses an interface on the subscriber access router to connect directly to the ISP access router. Another connection option uses two routers: a subscriber 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 information about BGP configuration. See the *Junos OS MPLS Applications Configuration Guide* 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 <b>vlan-vpls</b> 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 <b>flexible-ethernet-services</b> encapsulation type removes any VLAN ID value limitation.
vlan-vpls	extended-vlan-vpls	The <b>extended-vlan-vpls</b> encapsulation type can support multiple TPIDs. Using this encapsulation type removes any VLAN ID value limitation.
No encapsulation type	extended-vlan-vpls	The <b>extended-vlan-vpls</b> 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 10.10.1.1:1
```

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 10.10.10.1
```

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 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 $ABC123$ABC123$ABC123
user@host# set 10.10.10.1 secret $ABC123$ABC123
```

### 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

# Configuring ANCP-Triggered Layer 2 Wholesale Services

- [ANCP-Triggered Layer 2 Wholesale Service Overview on page 93](#)
- [Configuring ANCP-Triggered Layer 2 Wholesale Services on page 153](#)
- [Configuring Flat-File Accounting for Layer 2 Wholesale Services on page 163](#)



## CHAPTER 8

# ANCP-Triggered Layer 2 Wholesale Service Overview

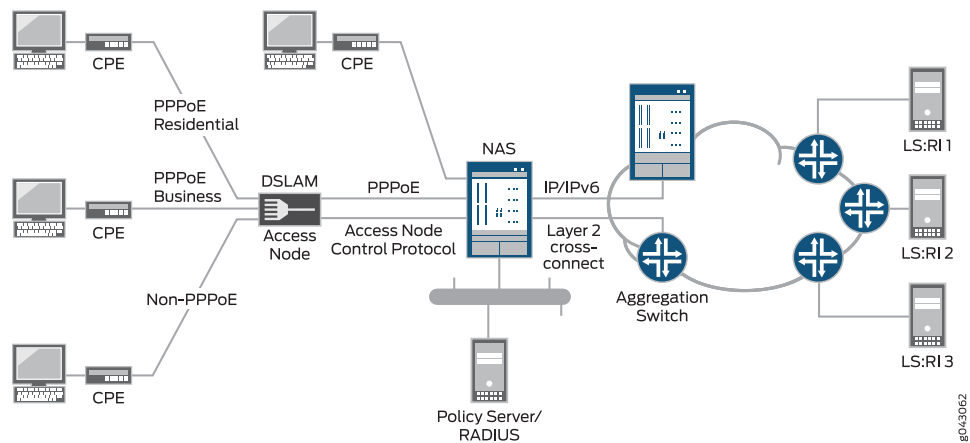
- [Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93](#)
- [Junos OS Predefined Variables on page 109](#)
- [Juniper Networks VSAs Supported by the AAA Service Framework on page 129](#)
- [AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 141](#)
- [AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 147](#)

### Layer 2 Wholesale with ANCP-Triggered VLANs Overview

The conventional mechanism for triggering autosensed dynamic VLANs relies on access line attributes provided by PPPoE or DHCP traffic in upstream control packets. Packets of a specified type are exceptioned and authorization depends on fields extracted from the packet as specified in a dynamic profile assigned to the autosensed VLAN range. However, for some wholesale networks, the traffic might not be PPPoE or DHCP. In this case, a different mechanism is required.

[Figure 9 on page 94](#) shows a sample topology with direct connections between the wholesaler's BNG and the NSP (network service provider) routers for the retailers. Each retailer's network resides in a dedicated routing instance. The wholesaler uses Layer 2 cross-connects to implement the retail networks with 1:1 autosensed, dynamic VLANs and VLAN tag swapping. Core-facing physical interfaces are dedicated to forwarding subscriber connections to the retailer's router. The traffic for an entire outer VLAN can be wholesaled this way. This direct-connect model supports any combination of wholesaler-owned and wholesaled connections for the entire access-facing VLAN range.

Figure 9: Sample Layer 2 Wholesale Access Topology



A wholesaler providing Layer 2 bitstream access to NSP partners might use this model. Bitstream access enables retailers to offer bidirectional transmission of broadband data and other high-speed services directly to customers across the wholesaler's network. In this topology, the PPPoE residential and subscriber customers are retained by the wholesaler (access provider). The non-PPPoE connections (here multiple connections and subscribers are represented by a single line) can be wholesaled to retail NSPs.

In this model, dynamic VLAN detection and creation for the wholesaled connections do not use in-band control packets. Instead, they rely on an out-of-band protocol, ANCP. ANCP Port Up messages both announce to the ANCP agent on the BNG that new access lines are operational and provide updates about previously announced lines. The messages include ANCP DSL attributes that correspond to Juniper Networks DSL VSAs and DSL Forum VSAs.

Starting in Junos OS Release 16.1, you can configure the ANCP agent to trigger the creation of an autosensed VLAN when the ANCP agent receives a Port Up message where the DSL-Line-State attribute has a value of Showtime. The Showtime state indicates that ports are configured, the subscriber is connected, and the DSL modem is online and ready to transfer data. The other possible values of the attribute, Idle and Silent, are ignored for this purpose and are used by the ANCP agent only to update the ANCP session database (SDB).

During VLAN authorization, RADIUS determines which traffic belongs to the access provider's own subscribers and which belongs to the wholesale customer (retail NSP) based on identification of the subscriber's access line by the agent remote identifier.

When the ANCP agent receives the Port Up message, the agent triggers the autoconfiguration daemon, `autoconfd`, to initiate the VLAN detection, authorization, and creation processes. Those processes require the following information:

- Three ANCP subscriber access loop attributes (TLVs) that identify the access line and are conveyed in the Port Up message:
  - Access-Loop-Circuit-ID—Access loop circuit identifier used by the ANCP agent to determine which logical interface or interface set corresponds to the subscriber; corresponds to the Juniper Networks Acc-Loop-Cir-ID VSA (26-110).
  - Access-Loop-Remote-ID—Unique identifier of the access line; corresponds to the Juniper Networks Acc-Loop-Remote-ID VSA (26-182).
  - Access-Aggregation-Circuit-ID-Binary—Identifier that represents the outer VLAN tag that the access node inserts on upstream traffic; corresponds to the Juniper Networks Acc-Aggr-Cir-Id-Bin VSA (26-111).
- The name of the physical interface facing the subscriber. This name derives from the local mapping of an ANCP neighbor to the corresponding subscriber-facing access port.

The Access-Aggregation-Circuit-ID-Binary attribute and the access-facing interface name together provide information equivalent to that used for conventional autosensed VLAN detection.

ANCP Port Down messages indicate that the subscriber access loop is not present or at least is no longer operational. This message triggers the automatic destruction of the dynamic VLAN, regardless of the value of any other ANCP line attribute.

VLAN logical interfaces are created in the default routing-instance unless a nondefault routing instance is provided by local authorization (domain map) or external authorization (RADIUS). Multiple routing instances are required when both access-provider-owned and wholesaled connections are supported at the same time. One routing instance is required for the access provider's own subscribers. An additional routing instance is required for each retail NSP. Consequently, the routing-instance has to be specified when the VLAN is authorized. The RADIUS-based VLAN authorization process determines whether the subscriber access-loop identified by the attributes in the Port Up message is wholesaled to a partner NSP—and therefore maintained as a unique routing-instance—or managed as a subscriber owned by the access provider.

- [RADIUS Authorization for ANCP-Triggered VLANs on page 96](#)
- [Instantiation of an ANCP-Triggered, Autosensed, Dynamic VLAN on page 96](#)
- [RADIUS Interim Accounting Updates on page 97](#)
- [Removal of the Layer 2 Wholesale Service on page 98](#)
- [Interactions Between In-Band and Out-of-Band VLAN Autosensing on page 99](#)
- [Migration of Subscriber Ownership from Wholesaler to Retailer on page 101](#)
- [Migration of Subscriber Ownership from Retailer to Wholesaler on page 102](#)
- [Migration of Subscriber Ownership Between Retailers on page 102](#)

- [Modification of the Access Line Identifier or Port VLAN Identifier on page 102](#)
- [Consequences of a State Transition in the Access-Facing Physical Interface on page 104](#)
- [Consequences of a State Transition from Up to Down in the Core-Facing Physical Interface on page 105](#)
- [Consequences of a State Transition from Down to Up in the Core-Facing Physical Interface on page 106](#)
- [Loss of ANCP TCP Adjacency on page 107](#)

## RADIUS Authorization for ANCP-Triggered VLANs

When a subscriber logs in, the Access-Request message that is sent to the RADIUS server includes a username and optionally a password generated locally on the router. Starting in Junos OS Release 16.1, you can configure the router to create a unique username with the value of ANCP TLVs— Access-Loop-Circuit-ID, Access-Loop-Remote-Id, or both—as received in the ANCP Port Up message from the access node. Alternatively, if you configure the router to convey ANCP-sourced access loop attributes as Juniper Networks VSAs—in this case Acc-Loop-Cir-Id (26-110) and Acc-Loop-Remote-Id (26-182)—then the Access-Request message includes sufficient unique access line information for the RADIUS server to determine whether the access loop is wholesaled to a retailer or retained for the wholesaler.

The RADIUS server responds to the Access-Request with one of the following messages:

- **Access-Accept**—In this case, the VLAN triggered by the Port Up message is wholesaled to a retail NSP. Authorization is similar to that for PPPoE sessions. The Access-Accept includes the Virtual-Router VSA (26-1) with a value that corresponds to the NSP's unique, nondefault routing instance. The message can optionally include client services, such as attributes for parameterized CoS, firewall filters and policies for the logical interface, and Layer 2 service activations.
- **Access-Reject**—In this case, either the VLAN triggered by the Port Up message is one of the wholesaler's own subscribers or RADIUS refuses to grant access to the network. In either case, the VLAN entry is removed from the ANCP SDB. Unless a Port Down message is received first, the router ignores subsequent Port Up messages for this subscriber. However, conventional dynamic stacked VLAN autosensing may be initiated by access protocol negotiation, such as PPPoE.

## Instantiation of an ANCP-Triggered, Autosensed, Dynamic VLAN

When the RADIUS server returns an Access-Accept message, the dynamic profile assigned to the autosensed VLAN range is instantiated with the following results:

1. The dynamic VLAN logical interface that represents the Layer 2 wholesale service within the NSP's unique routing instance is created.
2. A core-facing physical interface is selected by a simple round-robin method from the set of eligible interfaces assigned to the NSP's routing instance. A physical interface is eligible when it is operationally up and has at least one VLAN tag that is available for assignment.

3. The access-facing, autosensed outer VLAN tag is mapped to a unique inner VLAN tag. The outer VLAN tag is derived from the Access-Aggregation-Circuit-ID-Binary TLV carried in the ANCP Port Up message. The inner VLAN tag is allocated from the VLAN range configured for the core-facing physical interface.
4. The inner VLAN tag is swapped with (replaces) the outer VLAN tag when the subscriber traffic is tunneled to the NSP. In the dynamic profile, the inner VLAN tag is provided by the predefined variable, **\$junos-inner-vlan-map-id**.
5. The autosensed outer VLAN tag is swapped with the inner VLAN tag when downstream packets from the NSP (which include the allocated inner VLAN tag) are forwarded to the subscriber

You can configure each core-facing physical interface with a range of up to 4094 VLAN IDs. The inner VLAN swap range is assigned to the physical interface locally. This means that inner VLAN ranges for different physical interfaces can overlap each other completely, partially, or not at all.

6. Optionally, before the subscriber packets are forwarded to an NSP, the outer VLAN tag protocol identifier (TPID) in the subscriber packets can be swapped with a TPID to meet the requirements of an individual NSP. In this case, the original value is swapped with the NSP TPID for packets forwarded to the subscriber.
7. An additional VLAN tag, the Trunk VLAN ID, is used internally to identify the provisioned core-facing physical interface so that the subscriber traffic can be tunneled to the allocated interface. In the dynamic profile, this ID is provided by the predefined variable, **\$junos-vlan-map-id**. This identifier differentiates among multiple core-facing trunk physical interfaces for the same NSP.
8. Any client services, such as CoS or firewall filters, are applied to the subscriber session. These services are optionally specified in the RADIUS configuration and conveyed in the RADIUS message as Juniper Networks VSAs.
9. The VLAN session is activated after the logical interface is created and configured for the dynamic VLAN session. Session activation initiates a RADIUS Accounting-Start message. Any services that were received from RADIUS during authorization are now activated.
10. After the dynamic VLAN has been created, subsequent ANCP Port Up messages do not cause a re-authorization of the dynamic VLAN session. Instead, when the ANCP agent receives another Port Up message for the access loop, it updates the ANCP SDB with any changes in ANCP attributes.

## RADIUS Interim Accounting Updates

Interim accounting reports sent to AAA for out-of-band triggered, autosensed dynamic VLANs are supported in the same manner as for conventional autosensed, dynamic, authorized VLANs or client sessions (such as PPPoE sessions). The ANCP agent sends a notification to AAA when it receives an update from the access node. By default, AAA only reports the update to the RADIUS server at the configured interval.

Starting in Junos OS Release 16.1, you can configure the ANCP agent so that when it notifies AAA, an interim update Accounting-Request message is immediately sent to the

RADIUS server. Immediate interim accounting updates can be sent for an ANCP-triggered dynamic VLAN session only when a change occurs in certain key ANCP attributes for the associated access line that can influence system behavior. To prevent an additional load on the RADIUS server for changes to less critical ANCP attributes, changes to any other ANCP attributes do not trigger immediate accounting-interim-update messages. Instead, those changes are reported in the next scheduled Accounting-Interim-Update message.

Immediate interim accounting updates can be sent for changes to any of the following ANCP attributes for an existing session that corresponds to the access line (based on the Access-Loop-Circuit-ID TLV):

- **Actual-Net-Data-Rate-Upstream**—When the calculated (adjusted) upstream rate results in a change to this attribute, the accounting message reports the attribute in the Juniper Networks Act-Data-Rate-Up VSA (26-113). The calculated speed change is reported in the Upstream-Calculated-QoS-Rate VSA (26-142).
- **Actual-Net-Data-Rate-Downstream**—When the calculated (adjusted) downstream rate results in a change to this attribute, the accounting message reports the attribute in the Juniper Networks Act-Data-Rate-Dn VSA (26-114). The calculated speed change is reported in the Downstream-Calculated-QoS-Rate VSA (26-141).

When the **ancp-speed-change-immediate-update** statement is configured at the **[edit access profile *profile-name* accounting]** hierarchy level, RADIUS immediate interim accounting updates are sent for changes to the Actual-Net-Data-Rate-Upstream and Actual-Net-Data-Rate-Downstream TLVs.

When in addition the **auto-configure-trigger interface *interface-name*** statement is configured at the **[edit protocols ancp *neighbor ip-address*]** hierarchy level, immediate interim accounting updates are also sent for changes to the Access-Loop-Remote-ID and Access-Aggregation-Circuit-ID-Binary TLVs.

For more information about RADIUS immediate interim accounting updates, see *Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications*.

## Removal of the Layer 2 Wholesale Service

Any of the following events can remove the logical interface for the dynamic VLAN that represents the access service provided by the Layer 2 wholesaler:

- The receipt of an ANCP Port Down message for the corresponding access loop. The same ANCP attributes that initiate dynamic VLAN creation also initiate dynamic VLAN destruction.

No action is taken for an ANCP Port Down message for which any of the following is true:

- No corresponding subscriber session exists.
- A corresponding subscriber session is present, but is in the process of being deleted.

- The message refers to a conventional autosensed session (which is removed by normal protocol processing).
- An explicit reset of the connection between the ANCP agent and the access node, which triggers a mass logout of all affected dynamic VLAN sessions that support the wholesaled Layer 2 access connections. Sessions for the wholesaler's own subscribers are not affected.
- The deletion or transition to an operational down state of the subscriber-facing physical interface or the core-facing physical interface.
- The loss of adjacency between the neighbor and the ANCP agent.
- The issuance of the **clear auto-configuration interfaces** command to log out the VLAN or the **clear ancp access-loop** command to force a subscriber reset.
- The receipt of a RADIUS-initiated disconnect message.

Any of these events also deactivates the subscriber session to prevent future service activations and issues a RADIUS Accounting-Stop message for related services and for the dynamic VLAN subscriber session. The dynamic profile is then de-instantiated to remove first the dynamic VLAN logical interface and then the corresponding session entry in the VLAN SDB.

### Interactions Between In-Band and Out-of-Band VLAN Autosensing

The ANCP-triggered Layer 2 wholesale implementation accommodates both subscribers wholesaled to a retailer and subscribers belonging to the wholesaler. Any subscriber session detected on the access-facing physical interface can be one or the other. This means that an overlap exists between the outer tag range for the out-of-band autosensed VLANs and that for in-band, autosensed, stacked VLANs.

Both a PPPoE session and a wholesaled session might be attempted for the same access loop. To avoid the undesirable load on the router and the RADIUS server that can ensue when that happens, the order of session negotiation is determined by the order in which packets (PPPoE PADI or ANCP Port Up message) are received for the same access-facing physical interface and VLAN outer tag.



**NOTE:** The following sequences assume that the **remove-when-no-subscribers** statement is included at the [edit interfaces *interface-name* auto-configure] hierarchy level for the access-facing physical interface.

The following sequence of events occurs when a PPPoE PADI packet is received on an in-band control channel before an ANCP Port Up message is received on an out-of-band control channel, for the same access loop:

1. The PADI receipt triggers creation of a dynamic stacked VLAN logical interface. PPPoE and PPP negotiation are in progress.
2. The ANCP Port-Up message is received for the access loop. Because the corresponding in-band VLAN logical interface already exists for the same access-facing physical interface and outer VLAN tag, the ANCP agent simply stores the ANCP access line

attributes and the name of the physical interface in the session database. The agent takes no other action for the message as long as the PPP session (PPP logical interface and the underlying dynamic VLAN logical interface) is maintained.

3. PPP negotiation terminates due to authentication failure (RADIUS Access-Reject response) or a normal logout, which triggers clean-up of the PPP session and removal of the PPP logical interface.
4. Because the **remove-when-no-subscribers** statement is configured, deletion of the PPP logical interface results in deletion of the dynamic stacked VLAN.
5. The next event depends on whether authorization of the ANCP Port Up message has been attempted before.
  - If authorization was not previously attempted:
    - a. A VLAN-OOB SDB session is created and authorization of the access-loop is initiated.
    - b. All exceptioned PPPoE PADI packets received by in-band VLAN auto-sensing are ignored until RADIUS responds to the authorization request.
    - c. The authorization result determines what happens next:
      - If the authorization succeeds (RADIUS returns an Access-Accept message), then a dynamic Layer 2 wholesale logical interface is created within the retailer NSP's routing-instance.
      - If the authorization fails (RADIUS returns an Access-Reject message), then the VLAN-OOB session is cleaned up. Processing resumes for any exceptioned PPPoE PADI packets that are subsequently received by in-band VLAN autosensing.
  - If authorization was previously attempted, then no action is taken because the failure of the PPP session negotiation is assumed to be a login failure outside the Layer2 wholesale context. This behavior prevents unnecessary authorization in response to the ANCP Port-Up message every time a PPPoE session terminates and cleans up from a normal subscriber logout.

The following sequence of events occurs when an ANCP Port Up message is received on an out-of-band control channel before a PPPoE PADI packet for an access loop is received on an in-band control channel, both for the same access loop:

1. Receipt of the ANCP Port Up message triggers authorization of the access loop.
2. A PPPoE PADI packet is received. The packet is ignored because authorization is already in progress for the same access-facing physical interface and outer VLAN tag.
3. The authorization result determines what happens next:
  - If authorization succeeds (RADIUS returns an Access-Accept message)—represented by a VLAN-OOB session in the SDB—then dynamic creation of the VLAN logical interface is initiated for a Layer 2 wholesale session. When the interface is created, subsequent PPPoE PADI packets detected by in-band VLAN autosensing are ignored and no longer exceptioned.

- If authorization fails (RADIUS returns an Access-Reject message), the VLAN-OOB session is cleaned up.
  - a. Receipt of a subsequent PPPoE PADI packet initiates creation of a dynamic stacked VLAN.
  - b. PPPoE and PPP negotiation takes place and events proceed as usual for a conventional, dynamic autosensed VLAN.

## Migration of Subscriber Ownership from Wholesaler to Retailer

The wholesaler-owned subscribers are connected by means of dynamic PPPoE interfaces over dynamic VLANs. For each subscriber, the PPPoE session must be disconnected and the corresponding PPP logical interface deleted before ANCP Port Up messages for the same underlying physical interface and VLAN outer tag can serve as out-of-band triggers for dynamic VLAN autosensing.

One approach to migrating from wholesale to retail ownership is to do the following:

1. Update the RADIUS server database so that subscriber authentication for the relevant access lines results in a RADIUS Access-Reject response. This causes subsequent attempts to negotiate PPPoE for the access line to fail.
2. Initiate logout of the dynamic PPPoE sessions; for example, by issuing a RADIUS-initiated disconnect. This triggers cleanup of the PPPoE logical interface and associated services, which includes issuing RADIUS Accounting-Stop messages for the session and active services, as well as removing the dynamic PPPoE logical interface.

If the migration requires swapping out the current CPE device for another, and the PPPoE session is not otherwise gracefully logged out, then turning off the CPE results in a PPP keepalive failure on the router and triggers session logout.

3. Remove the underlying dynamic VLAN logical interface. This occurs automatically when the **remove-when-no-subscribers** statement is included at the **[edit interfaces interface-name auto-configure]** hierarchy level for the access-facing physical interface. Otherwise, issue the **clear auto-configuration interfaces interface-name** command to remove the dynamic VLAN logical interface.
4. Trigger a Port Up notification to initiate dynamic VLAN detection, authorization, and creation by one of the following methods:
  - Power cycle the CPE, with a sufficient delay between turning off and turning back on the device so that a Port Down message is sent followed by a Port Up message and the router is given enough time to detect a keepalive failure indicating loss of the session.
  - Issue a **clear ancp access-loop** command.
  - Issue a **request ancp oam port-up** command.

## Migration of Subscriber Ownership from Retailer to Wholesaler

One approach to migrating from retail to wholesale ownership is to do the following:

1. Update the RADIUS server database so that dynamic VLAN authorization for the relevant access lines results in a RADIUS Access-Reject response. Doing this causes subsequent ANCP Port Up messages to be ignored.
2. Initiate logout of the dynamic VLAN sessions supporting the wholesale access service; for example, by issuing a RADIUS-initiated disconnect. Doing this triggers cleanup of the session, which includes issuing RADIUS Accounting-Stop messages for the session, removal of the dynamic VLAN logical interface and active services, and freeing the allocated inner VLAN tag associated with the core-facing physical interface for assignment to a different Layer 2 wholesale subscriber session.

If the migration requires swapping out the current CPE device for another, then turning off the CPE results in an ANCP Port Down message that triggers session logout and cleanup.

3. Allow subscribers to connect to the wholesaler's network using conventional dynamic VLAN autosensing followed by PPPoE and PPP negotiation and creation of PPP logical interfaces.

## Migration of Subscriber Ownership Between Retailers

Typically, you migrate access between NSP retailers by triggering a restart of the existing dynamic VLAN session. The restart initiates a logout from the session followed by immediate dynamic VLAN detection, authorization, and creation within the routing-instance corresponding to the new NSP. A restart is a logical Port Down and Port Up sequence for the VLAN's corresponding access loop. You can use any of the following methods to restart a given dynamic VLAN logical interface:

- Initiate a RADIUS Disconnect-Request message or configure your RADIUS server to send the message. The message must have the Acct-Terminate-Cause RADIUS attribute (49) set to a value of 16 (callback). This cause is processed as a disconnect (logout) followed immediately by a reconnect (login) only for dynamic VLANs created by an ANCP Port Up message. Other clients respond to this value with only a disconnect.
- Trigger a Port Down message followed by a Port UP message by one of the following methods:
  - Power cycle the CPE, with a sufficient delay between turning off and turning back on the device so that a Port Down message is sent followed by a Port Up message and the router is given enough time to detect a keepalive failure indicating loss of the session.
  - Issue a **clear ancp access-loop** command.

## Modification of the Access Line Identifier or Port VLAN Identifier

When the line identifier or port VLAN identifier for an access loop is modified, the access node reports the change in a Port Up message to the ANCP agent. The message conveys

the line ID in the Access-Loop-Remote-ID TLV and the port VLAN ID in the Access-Aggregation-Circuit-ID-Binary TLV.

The access node should send a Port Down message for the access loop before it modifies any of the identification attributes for an existing session. The Port Down message triggers clean up of the corresponding Layer 2 wholesale session. If the access node does not send a Port Down in this case, then the following behavior has the same effect as inserting the Port Down message that the access node failed to send:

- For a line ID change, the ANCP agent treats the reconfiguration as a logical Port Down message for the access line identified by the previous Access-Loop-Remote-Id, followed by a Port Up message for the access line identified by the new Access-Loop-Remote-Id.
- For a port VLAN ID change, the ANCP agent treats the reconfiguration as a logical Port Down message for the access line identified by the previous Access-Aggregation-Circuit-Id-Binary, followed by a Port Up message for the access line identified by the new Access-Aggregation-Circuit-Id-Binary.

In either case, the ANCP agent notifies the autoconfiguration daemon (autoconfd), which takes the following actions:

1. Logs out the corresponding Layer 2 wholesale session.
2. Sends a RADIUS Accounting-Stop message with the final statistics for the associated service sessions and client session.
3. Removes the dynamic VLAN logical interface.
4. Attempts to reestablish the Layer 2 wholesale session by means of a login sequence, including authentication, creation of the dynamic VLAN logical interface, activation of any services, and if successful, sending RADIUS Accounting-Start messages for the service and client sessions.

You must manually log out any PPPoE session corresponding to the access line with the previous identifiers, even if the access node sends the appropriate Port Down message when the values change.



**NOTE:** In the case of a change in the port VLAN ID only, autoconfd takes no action to reinitiate the session when a dynamic stacked VLAN or a Layer 2 wholesale session exists with the same access-facing physical interface and outer VLAN tag. You must manually intervene in this case, such as by issuing a `request ancp oam port-up` command to initiate the creation of the Layer 2 wholesale session.



**BEST PRACTICE:** Because an existing session is not automatically logged out, we recommend that the network operator disconnect the session—for example, by issuing a RADIUS Disconnect-Request message—before modifying any of the access line attributes. Depending on subsequent

subscriber login and successful negotiation, a new session with the new identifier may then be created as usual.

.....

## Consequences of a State Transition in the Access-Facing Physical Interface

The following behavior results when the access-facing physical interface state transitions from Up to Down:

- Conventional in-band VLAN autosensing stops for the interface.
- ANCP-sourced Port Up messages for the interface are ignored. Action on new or unprocessed Port Up messages is deferred until the interface transitions to the Up state. If the ANCP connection is in band with the subscriber traffic on the interface, then all ANCP traffic stops; if the outage lasts long enough, the ANCP adjacency is lost.
- All Layer 2 wholesale sessions that are assigned to the interface are treated as if the ANCP agent received a Port Down message for the corresponding access line. Each session is subject to being logged out. Whether a session is logged out depends on the ANCP adjacency loss hold timer. The timer starts running when the ANCP agent detects the state transition to Down. The subscriber continues using the original session if all three of the following occur before the timer expires:
  1. The physical interface comes back up.
  2. The ANCP adjacency is restored.
  3. A Port Up message is received on the interface.

Otherwise, autoconfd takes the following actions:

1. Logs out the session.
2. Sends a RADIUS Accounting-Stop message with the final statistics for the associated service sessions and client session.
3. Removes the dynamic VLAN logical interface.

These sessions can be reestablished when the physical interface recovers, unless an ANCP Port Down message is received.

- The autoconfiguration daemon does not automatically delete dynamic, autosensed VLAN logical interfaces. The interfaces for the ANCP-triggered Layer 2 wholesale VLANs are maintained because the assumption is that an outage is short-lived. If the outage is not short-lived, then a subsequent Port Down message brings down the session and removes the interface.

For conventional autosensed dynamic VLANs, the interface is removed only when the **remove-when-no-subscribers** statement is configured on the access-facing physical interface and all references to the VLAN logical interface from a higher logical interface or session are removed. This mechanism does not apply to the ANCP-triggered Layer 2 wholesale VLANs because they do not have upper session references.

The following behavior results when the access-facing physical interface state transitions from Down to Up:

1. Conventional in-band VLAN autosensing resumes for the interface. PPPoE sessions owned by the access provider that were logged out due to the transition from Up to Down can renegotiate and undergo a full login sequence.
2. Appropriate actions are taken for all ANCP Port Up messages for the interface, including messages that were deferred because of the previous Down state for the interface. If the ANCP connection is in band with the subscriber traffic, then all ANCP traffic resumes.
3. Forwarding resumes for any dynamic, autosensed VLAN logical interfaces that were not deleted when the interface went down.

Deletion of an access-facing physical interface triggers logout and removal of all upper dynamic VLAN logical interfaces and their corresponding sessions.

### Consequences of a State Transition from Up to Down in the Core-Facing Physical Interface

The following behavior results when the core-facing physical interface state transitions from Up to Down:

- The core-facing physical interface is no longer eligible for assigning new or pending access lines in this routing instance as based on the original RADIUS authorization.
- All Layer 2 wholesale sessions that are assigned to the interface are treated as if the ANCP agent received a Port Down/Port Up message sequence for the corresponding access line. Each session is subject to being logged out. Whether a session is logged out depends on the ANCP adjacency loss hold timer. The timer starts running when the ANCP agent detects the state transition to Down. The subscriber continues using the original session if all three of the following occur before the timer expires:
  1. The physical interface comes back up.
  2. The ANCP adjacency is restored.
  3. A Port Up message is received on the interface.

Otherwise, autoconfd takes the following actions:

1. Logs out the session.
  2. Removes the session entry from the SDB.
  3. Sends a RADIUS Accounting-Stop message with the final statistics for the associated service sessions and client session.
  4. Removes the dynamic VLAN logical interface.
- Next, autoconfd attempts to migrate the sessions to available connections on any remaining eligible core-facing physical interfaces that are assigned to the same routing instance:
    1. The original request is placed on a retry queue.

2. A login sequence is attempted for each session, including authentication, creation of dynamic VLAN logical interfaces, activation of any services, and sending RADIUS Accounting-Start messages for the service and client sessions.
  - If the login sequence is successful, then the request is removed from the retry queue.
  - If the login fails, then the session is logged out, the session entry is removed from the SDB, and the corresponding access line is set to a pending state.

When the available connections are all used—when there are no more available VLAN tags from the configured inner VLAN ID swap range—as a result of successful reconnections, no attempt is made to connect any remaining Layer 2 wholesale sessions. Although authentication can succeed, the creation of dynamic VLAN logical interfaces fails during profile instantiation. In this case, the session is out, the session entry is removed from the SDB, and the corresponding access line is set to a pending state.

- The pending access lines that represent these non-migrated sessions can be reestablished if the interface recovers or if additional VLAN connections become available; for example, by a configuration change that either increases the inner VLAN ID swap range for one or more remaining core-facing physical interfaces or adds new core-facing physical interfaces. However, if the ANCP agent receives a Port Down message for a pending access line, the corresponding session is not reestablished.

You can use the **show auto-configuration out-of-band pending** command to display a count of pending access lines per routing instance.

**NOTE:**

In addition to core-facing physical interface state transitions from Up to Down, these behaviors also apply in the following circumstances:

- A core-facing physical interface is deleted.
- More Layer 2 wholesale sessions are assigned to a routing instance than can be accommodated by the inner VLAN ID swap range configured on the interface assigned to the routing instance.



**BEST PRACTICE:** We recommend that you use aggregated Ethernet for the core-facing physical interfaces to provide link protection, bandwidth aggregation, or both.

---

## Consequences of a State Transition from Down to Up in the Core-Facing Physical Interface

The following behavior results when the core-facing physical interface state transitions from Down to Up:

- The physical interface is now eligible to assign new Layer 2 wholesale subscriber sessions.

- The ANCP agent notifies the autoconfiguration daemon (autoconfd), which attempts to reestablish the Layer 2 wholesale sessions that correspond to pending access line by initiating a conventional login sequence. This sequence includes authentication, creation of dynamic VLAN logical interfaces, activation of any services, and sending RADIUS Accounting-Start messages for the service and client sessions.
- Pending sessions continue to be reestablished until none are left or an error occurs, typically due to exhaustion of inner VLAN tags from the swap range on the interface. In the latter case, the sessions are logged out, the session entry is removed from the SDB, and the access line is set to a pending state.

You can use the **show auto-configuration out-of-band pending** command to display a count of pending access lines per routing instance.

These behaviors also occur in the following cases:

- Additional VLAN connection resources become available, by a configuration change that either increases the inner VLAN ID swap range for one or more remaining core-facing physical interfaces or adds new core-facing physical interfaces. The newly added physical interface must be in the Up state to assume any Layer 2 wholesale sessions.
- A RADIUS-initiated disconnect is received for an existing Layer 2 wholesale session assigned to this routing instance is logged out (disconnect only). For a disconnect with a reconnect qualifier, the affected session is given preference to reconnect over pending access lines.
- You issue the **request auto-configuration reconnect-pending**, **clear ancp access-loop**, or **request ancp oam port-up** command.

## Loss of ANCP TCP Adjacency

The ANCP agent can lose its TCP adjacency with a neighbor in any of the following circumstances:

- The access node renegotiates the connection; for example, as a result of losing synchronization. The renegotiation triggers the local state to change from established to not established. The state transitions back to established when the session is successfully renegotiated.
- An end-of-file (EOF) message is received on the socket indicating the adjacency is closed. This can result when the ANCP configuration is deleted on the access node.
- An adjacency keepalive failure occurs. When no response is received for three consecutive polls, the adjacency is declared to be lost.

The ANCP agent treats the loss of adjacency as if it has received a Port Down message for each access loop represented by the ANCP connection. The agent notifies autoconfd, which takes the following actions:

- Logs out all Layer 2 wholesale sessions that were triggered by this ANCP connection.
- Sends a RADIUS Accounting-Stop message with the final statistics for the associated service sessions and client session.

- Removes the dynamic VLAN logical interface.

If the assigned access-facing or core-facing physical interface is in the Down state, any pending sessions that were triggered by this ANCP connection cannot be reestablished when the interface recovers to the Up state.

Dynamic, conventionally auto-sensed VLAN logical interfaces, such as those supporting PPPoE sessions, are not affected by the TCP adjacency loss.

If the adjacency is reestablished, the expected behavior is a complete replay of Port Down and Port Up messages for all associated configured access lines. The Layer 2 wholesale sessions for which the ANCP agent receives Port Up messages are reestablished.

You can mitigate the effects of short-term adjacency losses by configuring an adjacency loss hold time. The timer starts when adjacency is lost. Even though the adjacency is lost, established sessions are maintained while the timer runs unless a Port Down message is received for the corresponding access line.

Any access line that for which the ANCP agent has not received a Port Up message by the time the timer expires is treated as though the agent has received a Port Down message for the line. The ANCP Agent notifies autoconfd, which takes the following actions:

- Logs out all Layer 2 wholesale sessions that correspond to the access line.
- Sends a RADIUS Accounting-Stop message with the final statistics for the associated service sessions and client session.
- Removes the dynamic VLAN logical interface.

Port Up messages received after the timer expires repopulate the SDB access line table and reestablish the Layer 2 wholesale sessions

The adjacency loss hold timer serves the following purposes:

- Dampens the effect of adjacency loss of short duration thereby maintaining existing Layer 2 wholesale sessions.
- Detects the removal of an access line configuration on the access node. For example, in some circumstances you may want to remove the configuration of an access line on a neighbor. You first disconnect the ANCP session between a neighbor and the BNG, remove the configuration on the neighbor, and then restore the ANCP connection with the BNG. The neighbor does not issue a Port Down message. If the adjacency-loss hold-timer is configured, the ANCP agent detects an access line for which it has not received a Port Up or Port Down message, and consequently triggers logout and removal of the corresponding Layer 2-wholesale session.



**NOTE:** When you deactivate the ANCP protocol, the router does not perform a commit check to determine whether any ANCP or L2-BSA subscribers are present (active or inactive). Any subscribers that are active at the time of deactivation remain active.

#### Related Documentation

- [Configuring Out-of-Band ANCP Messages to Trigger Dynamic VLAN Instantiation on page 156](#)
- [Configuring the ANCP Agent for ANCP-Triggered, Autosensed Dynamic VLANs on page 154](#)
- [Configuring the ANCP Agent](#)
- [ANCP and the ANCP Agent Overview](#)

## Junos OS Predefined Variables

Junos OS contains many predefined variables. The dynamic profile obtains and replaces values for these variables from an incoming client data packet and configuration (local and RADIUS). These variables are predefined—you use them in the body of a dynamic profile without first having to define the variables at the **[dynamic-profiles profile-name variables]** hierarchy level. [Table 9 on page 109](#) provides a list of predefined variables, their descriptions, and where in the Junos OS hierarchy you can configure them.

**Table 9: Junos OS Predefined Variables and Definitions**

Variable	Definition
<b>Access and Access-Internal Routes</b>	
\$junos-framed-route-cost	Cost metric of an IPv4 access route. You specify this variable at the <b>[edit dynamic-profiles profile-name routing-options access route address]</b> hierarchy level for the <b>metric</b> statement.
\$junos-framed-route-distance	Distance of an IPv4 access route. You specify this variable at the <b>[edit dynamic-profiles profile-name routing-options access route address]</b> hierarchy level for the <b>preference</b> statement.
\$junos-framed-route-ip-address-prefix	Route prefix of an IPv4 access route. You specify this variable at the <b>[edit dynamic-profiles profile-name routing-options access]</b> hierarchy level for the <b>route</b> statement.
\$junos-framed-route-ipv6-address-prefix	Route prefix of an IPv6 access route. You specify this variable at the <b>[edit dynamic-profiles profile-name routing-options access]</b> hierarchy level for the <b>route</b> statement.

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
\$junos-framed-route-ipv6-cost	Cost metric of an IPv6 access route. You specify this variable with the <b>metric</b> statement at the [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib \$junos-ipv6-rib access route \$junos-framed-route-ipv6-address-prefix] hierarchy level.
\$junos-framed-route-ipv6-distance	Distance of an IPv6 access route. You specify this variable with the <b>preference</b> statement at the [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib \$junos-ipv6-rib access route \$junos-framed-route-ipv6-address-prefix] hierarchy level.
\$junos-framed-route-ipv6-nexthop	IPv6 next-hop address of an access route. You specify this variable at the [edit dynamic-profiles <i>profile-name</i> routing-options access route <i>address</i> ] hierarchy level for the <b>next-hop</b> statement.
\$junos-framed-route-ipv6-tag	Tag value of an IPv6 access route. You specify this variable with the <b>tag</b> statement at the [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib \$junos-ipv6-rib access route \$junos-framed-route-ipv6-address-prefix] hierarchy level.
\$junos-framed-route-nexthop	IPv4 next-hop address of an access route. You specify this variable at the [edit dynamic-profiles <i>profile-name</i> routing-options access route <i>address</i> ] hierarchy level for the <b>next-hop</b> statement.
\$junos-framed-route-tag	Tag value of an IPv4 access route. You specify this variable at the [edit dynamic-profiles <i>profile-name</i> routing-options access route <i>address</i> ] hierarchy level for the <b>tag</b> statement.
\$junos-interface-name	<p>Logical interface of an access-internal route. DHCP or PPP supplies this information when the subscriber logs in. You specify this variable at the [edit dynamic-profiles <i>profile-name</i> routing-options access-internal route <i>address</i>] hierarchy level for the <b>qualified-next-hop</b> statement.</p> <p>This variable is also used for creating dynamic IP demux interfaces.</p>

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
\$junos-ipv6-rib	<p>Routing table for an IPv6 access route. You specify this variable with the <b>rib</b> statement at the <b>[edit dynamic-profiles profile-name routing-instances \$junos-routing-instance routing-options]</b> hierarchy level.</p> <p>You can use this variable to specify a nondefault routing instance for the route.</p>
\$junos-subscriber-ip-address	<p>IP address of a subscriber identified in an access-internal route. You specify this variable at the <b>[edit dynamic-profiles profile-name routing-options access-internal]</b> hierarchy level for the <b>route</b> statement.</p> <p>This variable is also used for creating dynamic IP demux interfaces.</p>
\$junos-subscriber-mac-address	<p>MAC address for a subscriber identified in an access-internal route. You specify this variable at the <b>[edit dynamic-profiles profile-name routing-options access-internal route address qualified-next hop underlying-interface]</b> hierarchy level for the <b>mac-address</b> statement.</p>
<b>Dynamic Protocols</b>	
\$junos-igmp-access-group-name	<p>Specifies the access list to use for the source (S) filter.</p>
\$junos-igmp-access-source-group-name	<p>Specifies the access list to use for the source-group (S,G) filter.</p>
\$junos-igmp-enable	<p>Ensures that IGMP is not disabled on the interface by an AAA-based authentication and management method (for example, RADIUS). You specify this variable at the <b>[dynamic-profiles profile-name protocols igmp]</b> hierarchy level for the <b>interface</b> statement.</p>
\$junos-igmp-immediate-leave	<p>Enables IGMP immediate leave on the interface. You specify this variable at the <b>[dynamic-profiles profile-name protocols igmp]</b> hierarchy level for the <b>interface</b> statement.</p>
\$junos-igmp-version	<p>IGMP version configured in a client access profile. Junos OS obtains this information from the RADIUS server when a subscriber accesses the router. The version is applied to the accessing subscriber when the profile is instantiated. You specify this variable at the <b>[dynamic-profiles profile-name protocols igmp]</b> hierarchy level for the <b>interface</b> statement.</p>

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
\$junos-interface-name	<p>Name of the dynamic interface to which the subscriber access client connects. Its use is in dynamically enabling IGMP on the subscriber interface. You specify this variable at the <b>[dynamic-profiles <i>profile-name</i> protocols igmp]</b> hierarchy level for the <b>interface</b> statement.</p> <p>The interface name is derived from concatenating the <b>\$junos-interface-ifs-name</b> and the <b>\$junos-underlying-interface-unit</b> variables obtained when a subscriber is created dynamically at the <b>[dynamic-profiles <i>profile-name</i> interfaces]</b> hierarchy level.</p>
\$junos-ipv6-ndra-prefix	<p>Prefix value for the router advertisement interface. Junos OS obtains this information from the RADIUS server when a subscriber accesses the router. The prefix value is applied to the accessing subscriber when the profile is instantiated. You specify this variable at the <b>[dynamic-profiles <i>profile-name</i> protocols router-advertisement interface <i>\$junos-interface-name</i>]</b> hierarchy level.</p>
\$junos-mld-access-group-name	Specifies the access list to use for the group (G) filter.
\$junos-mld-access-source-group-name	Specifies the access list to use for the source-group (S,G) filter.
\$junos-mld-enable	<p>Ensures that MLD is not disabled on the interface by an AAA-based authentication and management method (for example, RADIUS). You specify this variable at the <b>[dynamic-profiles <i>profile-name</i> protocols mld]</b> hierarchy level for the <b>interface</b> statement.</p>
\$junos-mld-immediate-leave	<p>Enables MLD immediate leave on the interface. You specify this variable at the <b>[dynamic-profiles <i>profile-name</i> protocols mld]</b> hierarchy level for the <b>interface</b> statement.</p>
\$junos-mld-version	<p>MLD version configured in a client access profile. Junos OS obtains this information from the RADIUS server when a subscriber accesses the router. The version is applied to the accessing subscriber when the profile is instantiated. You specify this variable at the <b>[dynamic-profiles <i>profile-name</i> protocols mld]</b> hierarchy level for the <b>interface</b> statement.</p>

#### Dynamic CoS — Traffic-Control Profile Parameters

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
\$junos-cos-adjust-minimum	<p>Minimum adjusted shaping rate configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>adjust-minimum</b> statement at the <b>[edit dynamic-profiles <i>profile-name</i> class-of-service traffic-control-profiles <i>profile-name</i>]</b> hierarchy level.</p>
\$junos-cos-byte-adjust	<p>Byte adjustment value configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>bytes</b> option with the <b>overhead-accounting</b> statement at the <b>[edit dynamic-profiles <i>profile-name</i> class-of-service traffic-control-profiles <i>profile-name</i>]</b> hierarchy level.</p>
\$junos-cos-byte-adjust-cell	<p>Overhead bytes when downstream ATM traffic is in cell-mode.</p> <p><b>NOTE:</b> Do not configure the \$junos-cos-byte-adjust-cell variable when the \$junos-cos-byte-adjust variable is configured.</p>
\$junos-cos-byte-adjust-frame	<p>Overhead bytes when downstream ATM traffic is in frame-mode.</p> <p><b>NOTE:</b> Do not configure the \$junos-cos-byte-adjust-frame variable when the \$junos-cos-byte-adjust variable is configured.</p>
\$junos-cos-delay-buffer-rate	<p>Delay-buffer rate configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>delay-buffer-rate</b> statement at the <b>[edit dynamic-profiles <i>profile-name</i> class-of-service traffic-control-profiles <i>profile-name</i>]</b> hierarchy level.</p>

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
\$junos-cos-excess-rate	<p>Excess rate configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>excess-rate</b> statement at the <b>[edit dynamic-profiles <i>profile-name</i> class-of-service traffic-control-profiles <i>profile-name</i>]</b> hierarchy level.</p>
\$junos-cos-excess-rate-high	<p>Rate configured for excess high-priority traffic in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>excess-rate-high</b> statement at the <b>[edit dynamic-profiles <i>profile-name</i> class-of-service traffic-control-profiles <i>profile-name</i>]</b> hierarchy level.</p>
\$junos-cos-excess-rate-low	<p>Rate configured for excess low-priority traffic in a traffic-control profile in a dynamic profile for subscriber access. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>excess-rate-low</b> statement at the <b>[edit dynamic-profiles <i>profile-name</i> class-of-service traffic-control-profiles <i>profile-name</i>]</b> hierarchy level.</p>
\$junos-cos-guaranteed-rate	<p>Guaranteed rate configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>guaranteed-rate</b> statement at the <b>[edit dynamic-profiles <i>profile-name</i> class-of-service traffic-control-profiles <i>profile-name</i>]</b> hierarchy level.</p>

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
\$junos-cos-guaranteed-rate-burst	<p>Burst size for the guaranteed rate that is configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable with the <b>burst-size</b> option in the <b>guaranteed-rate</b> statement at the <b>[edit dynamic-profiles <i>profile-name</i> class-of-service traffic-control-profiles <i>profile-name</i>]</b> hierarchy level.</p>
\$junos-cos-scheduler-map	<p>Scheduler-map name configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>scheduler-map</b> statement at the <b>[edit dynamic-profiles <i>profile-name</i> class-of-service traffic-control-profiles <i>profile-name</i>]</b> hierarchy level.</p> <p><b>NOTE:</b> The scheduler map can be defined dynamically (at the <b>[edit dynamic-profiles <i>profile-name</i> class-of-service scheduler-maps]</b> hierarchy level) or statically (at the <b>[edit class-of-service scheduler-maps]</b> hierarchy level).</p>
\$junos-cos-shaping-mode	<p>Shaping mode configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>overhead-accounting</b> statement at the <b>[edit dynamic-profiles <i>profile-name</i> class-of-service traffic-control-profiles <i>profile-name</i>]</b> hierarchy level.</p>

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
\$junos-cos-shaping-rate	<p>Shaping rate configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>shaping-rate</b> statement at the <b>[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]</b> hierarchy level.</p>
\$junos-cos-shaping-rate-burst	<p>Burst size for the shaping rate configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable with the <b>burst-size</b> option in the <b>shaping-rate</b> statement at the <b>[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]</b> hierarchy level.</p>
\$junos-cos-shaping-rate-excess-high	<p>Shaping rate configured for excess high-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.</p>
\$junos-cos-shaping-rate-excess-high-burst	<p>Shaping rate burst size configured for excess high-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.</p>
\$junos-cos-shaping-rate-excess-low	<p>Shaping rate configured for excess low-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.</p>
\$junos-cos-shaping-rate-excess-low-burst	<p>Shaping rate burst size configured for excess low-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.</p>

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
\$junos-cos-shaping-rate-priority-high	Shaping rate configured for high-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
\$junos-cos-shaping-rate-priority-high-burst	Shaping rate burst size configured for high-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
\$junos-cos-shaping-rate-priority-low	Shaping rate configured for low-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
\$junos-cos-shaping-rate-priority-low-burst	Shaping rate burst size configured for low-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
\$junos-cos-shaping-rate-priority-medium	Shaping rate configured for medium-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
\$junos-cos-shaping-rate-priority-medium-burst	Shaping rate burst size configured for medium-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
\$junos-cos-traffic-control-profile	<p>Traffic-control profile configured in a dynamic profile for subscriber access. The Junos OS obtains the profile information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>traffic-control-profiles</b> statement at the <b>[edit dynamic-profiles <i>profile-name</i> class-of-service]</b> hierarchy level.</p>

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
<b>Dynamic CoS — Scheduler Parameters</b>	
\$junos-cos-scheduler	<p>Name of a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable at the <b>[edit dynamic-profiles <i>profile-name</i> class-of-service schedulers]</b> hierarchy level.</p>
\$junos-cos-scheduler-bs	<p>Buffer size as a percentage of total buffer, specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>buffer-size</b> statement with the <b>percent</b> option at the <b>[edit dynamic-profiles <i>profile-name</i> class-of-service schedulers <i>scheduler-name</i>]</b> hierarchy level.</p>
\$junos-cos-scheduler-pri	<p>Packet-scheduling priority value specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>priority</b> statement at the <b>[edit dynamic-profiles <i>profile-name</i> class-of-service schedulers <i>scheduler-name</i>]</b> hierarchy level.</p>
\$junos-cos-scheduler-dropfile-any	<p>Name of the drop profile for random early detection (RED) for loss-priority level <b>any</b> specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>drop-profile</b> statement at the <b>[edit dynamic-profiles <i>profile-name</i> class-of-service schedulers <i>scheduler-name</i> drop-profile-map loss-priority <i>any</i> protocol <i>any</i>]</b> hierarchy level.</p> <p><b>NOTE:</b> The drop profile must be configured statically (at the <b>[edit class-of-service drop-profiles]</b> hierarchy level).</p>

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
\$junos-cos-scheduler-dropfile-high	<p>Name of the drop profile for random early detection (RED) for loss-priority level <b>high</b> specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>drop-profile</b> statement at the <b>[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name drop-profile-map loss-priority high protocol any]</b> hierarchy level.</p> <p><b>NOTE:</b> The drop profile must be configured statically (at the <b>[edit class-of-service drop-profiles]</b> hierarchy level).</p>
\$junos-cos-scheduler-dropfile-low	<p>Name of the drop profile for random early detection (RED) for loss-priority level <b>low</b> specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>drop-profile</b> statement at the <b>[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name drop-profile-map loss-priority low protocol any]</b> hierarchy level.</p> <p><b>NOTE:</b> The drop profile must be configured statically (at the <b>[edit class-of-service drop-profiles]</b> hierarchy level) for loss-priority <b>low</b>.</p>

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
\$junos-cos-scheduler-dropfile-medium-high	<p>Name of the drop profile for random early detection (RED) for loss-priority level <b>medium-high</b> specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>drop-profile</b> statement at the <b>[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name drop-profile-map loss-priority medium-high protocol any]</b> hierarchy level.</p> <p><b>NOTE:</b> The drop profile must be configured statically (at the <b>[edit class-of-service drop-profiles]</b> hierarchy level).</p>
\$junos-cos-scheduler-dropfile-medium-low	<p>Name of the drop profile for random early detection (RED) for loss-priority level <b>medium-low</b> specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>drop-profile</b> statement at the <b>[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name drop-profile-map loss-priority medium-low protocol any]</b> hierarchy level.</p> <p><b>NOTE:</b> The drop profile must be configured statically (at the <b>[edit class-of-service drop-profiles]</b> hierarchy level).</p>
\$junos-cos-scheduler-excess-priority	<p>Priority value of the excess rate specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>excess-priority</b> statement at the <b>[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]</b> hierarchy level.</p>

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
\$junos-cos-scheduler-excess-rate	<p>Value of the excess rate specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>excess-rate</b> statement at the <b>[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]</b> hierarchy level.</p>
\$junos-cos-scheduler-shaping-rate	<p>Value of the shaping rate specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>shaping-rate</b> statement at the <b>[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]</b> hierarchy level.</p>
\$junos-cos-scheduler-tx	<p>Transmit rate specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached.</p> <p>You reference this variable in the <b>transmit-rate</b> statement at the <b>[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]</b> hierarchy level.</p>
<b>Dynamic Connectivity Fault Management Parameters</b>	
\$junos-action-profile	Name of the action profile configured in a dynamic profile.
\$junos-ccm-interval	Continuity check interval time configured in a dynamic profile.
\$junos-loss-threshold	The number of continuity check messages lost before marking the remote MEP as down, configured in a dynamic profile.
\$junos-ma-name-format	Name of the maintenance association name format configured in a dynamic profile.
\$junos-md-name-format	Name of the maintenance domain format configured in a dynamic profile.

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
\$junos-ma-name	Name of the maintenance association configured in a dynamic profile.
\$junos-md-level	Value of 'Level', configured in a dynamic profile.
\$junos-md-name	Name of the maintenance domain configured in a dynamic profile.
\$junos-mep-id	The 'MEP' value configured in the dynamic profile.
\$junos-remote-mep-id	The 'Remote MEP' value configured in the dynamic profile.
<b>Filters — RADIUS-obtained Policies</b>	
\$junos-input-filter	Name of an input filter to be attached; filter name is derived from RADIUS VSA 26-10 (Ingress-Policy-Name) or RADIUS attribute 11 (Filter-ID) to the interface.
\$junos-input-interface-filter	<p>Name of an input filter to be attached to a <b>family any</b> interface; filter name is derived from RADIUS VSA 26-191 (Input-Interface-Filter) to the interface.</p> <p>You can also specify the filter name with the <code>\$junos-input-interface-filter</code> statement at the <code>[edit dynamic-profiles profile-name interfaces interface-name unit logical-interface-number filter input]</code> hierarchy level.</p>
\$junos-input-ipv6-filter	Name of an IPv6 input filter to be attached; filter name is derived from RADIUS VSA 26-106 (IPv6-Ingress-Policy-Name) to the interface.
\$junos-output-filter	Name of an output filter to be attached; filter name is derived from RADIUS VSA 26-11 (Egress-Policy-Name) to the interface.
\$junos-output-interface-filter	<p>Name of an output filter to be attached to a <b>family any</b> interface; filter name is derived from RADIUS VSA 26-191 (Output-Interface-Filter) to the interface.</p> <p>You can also specify the filter name with the <code>\$junos-output-interface-filter</code> statement at the <code>[edit dynamic-profiles profile-name interfaces interface-name unit logical-interface-number filter output]</code> hierarchy level.</p>

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
\$junos-output-ipv6-filter	Name of an IPv6 output filter to be attached; filter name is derived from RADIUS VSA 26-107 (IPv6-Egress-Policy-Name) to the interface.
<b>Subscriber Interfaces — Dynamic Demux Interfaces</b>	
\$junos-interface-ifd-name	<p>Name of the device to which the subscriber access client connects. All interfaces are created on this device. Its primary use is in creating single or multiple subscribers on a statically created interface. You specify this variable at the <b>[dynamic-profiles <i>profile-name</i> interfaces]</b> hierarchy level.</p> <p>When creating a logical underlying interface for a dynamic VLAN demux interface, you must also specify this variable at the <b>[dynamic-profiles <i>profile-name</i> interfaces <i>demux0</i> unit <i>\$junos-interface-unit</i> demux-options underlying-interface]</b> hierarchy level.</p>
\$junos-interface-unit	Creates a unit number assigned to the logical interface. The router supplies this information when the subscriber accesses the network. You specify this variable at the <b>[dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i>]</b> hierarchy level for the <b>unit</b> statement.
\$junos-ipv6-address	Selects the IPv6 address of the interface the subscriber uses. You specify this variable at the <b>[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]</b> , <b>[edit dynamic-profiles <i>profile-name</i> interfaces <i>demux0</i> unit <i>logical-unit-number</i> family <i>family</i>]</b> , <b>[edit dynamic-profiles <i>profile-name</i> interfaces <i>pp0</i> unit "<i>\$junos-interface-unit</i>" family <i>family</i>]</b> , and <b>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]</b> hierarchy level for the <b>address</b> statement.
\$junos-loopback-interface	Selects the loopback interface the subscriber uses. You specify this variable at the <b>[dynamic-profiles <i>profile-name</i> interfaces <i>demux0</i> unit "<i>\$junos-interface-unit</i>" family inet]</b> hierarchy level for the <b>unnumbered-address</b> statement.

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
\$junos-preferred-source-address	<p>Selects the preferred IPv4 source address (family inet) associated with the loopback address used for the subscriber. You specify this variable at the [dynamic profiles <i>profile-name</i> interfaces demux0 unit "\$junos-interface-unit" family inet unnumbered-address "\$junos-loopback-interface"] hierarchy level for the preferred-source-address statement.</p> <p><b>NOTE:</b> Starting in Junos OS Release 16.1, when you specify a static logical interface for the unnumbered interface in a dynamic profile that includes the \$junos-routing-instance predefined variable, you must not configure an IPv4 preferred source address, whether with the \$junos-preferred-source-address predefined variable or the preferred-source-address statement. Configuring the preferred source address in this circumstance causes a commit failure.</p>
\$junos-preferred-source-ipv6-address	<p>Selects the preferred IPv6 source address (family inet6) associated with the loopback address used for the subscriber. You specify this variable at the [dynamic profiles <i>profile-name</i> interfaces demux0 unit "\$junos-interface-unit" family inet6 unnumbered-address "\$junos-loopback-interface"] hierarchy level for the preferred-source-address statement.</p> <p><b>NOTE:</b> Starting in Junos OS Release 16.1, when you specify a static logical interface for the unnumbered interface in a dynamic profile that includes the \$junos-routing-instance predefined variable, you must not configure an IPv6 preferred source address, whether with the \$junos-preferred-source-ipv6-address predefined variable or the preferred-source-address statement. Configuring the preferred source address in this circumstance causes a commit failure.</p>
\$junos-subscriber-ip-address	<p>IP address of the subscriber. You specify this variable at the [dynamic-profiles <i>profile-name</i> interfaces demux0 unit family inet demux-source] hierarchy level.</p> <p>This variable is also used for creating access-internal routes.</p>

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
\$junos-subscriber-ipv6-address	IPv6 address for subscriber. You specify this variable at the <b>[dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 demux-source]</b> hierarchy level.
\$junos-subscriber-ipv6-multi-address	<p>Expands the <b>demux-source</b> into multiple addresses; for example, the IPv6 prefix and /128 address for the subscriber.</p> <p>You specify this variable at the <b>[dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 demux-source]</b> hierarchy level.</p>
\$junos-underlying-interface	<p>Creates a logical underlying interface for a dynamic IP demux interface. The client logs in on this interface. You specify this variable at the <b>[dynamic profiles <i>profile-name</i> interfaces demux0 unit "\$junos-interface-unit" demux-options]</b> hierarchy level for the <b>underlying-interface</b> statement.</p> <p>When configured, the underlying interface is used to determine the <b>\$junos-underlying-interface</b>, <b>\$junos-underlying-interface-unit</b>, and <b>\$junos-ifd-name</b> variables. For example, if the receiving logical interface is ge-0/0/0.1, the <b>\$junos-underlying-interface</b> variable is set to ge-0/0/0 and the <b>\$junos-underlying-interface-unit</b> variable is set to 1.</p> <p>This variable is also used for creating access-internal routes.</p>
<b>Subscriber Interfaces — Static VLAN Interfaces</b>	
\$junos-interface-ifd-name	Name of the device to which the subscriber access client connects. All interfaces are created on this device. Its primary use is in creating single or multiple subscribers on a statically created interface. You specify this variable at the <b>[dynamic-profiles <i>profile-name</i> interfaces]</b> hierarchy level.
\$junos-underlying-interface-unit	Obtains the unit number for the underlying interface. It specifies the use of the underlying interface for the subscriber. You specify this variable at the <b>[dynamic-profiles <i>profile-name</i> interfaces <i>\$junos-interface-ifd-name</i>]</b> hierarchy level for the <b>unit</b> statement.

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
<b>Subscriber Interfaces — Dynamic PPPoE Interfaces</b>	
<code>\$junos-interface-unit</code>	Specifies the logical unit number when the router dynamically creates a PPPoE logical interface. The <code>\$junos-interface-unit</code> predefined variable is dynamically replaced with the unit number supplied by the network when the PPPoE subscriber logs in. You specify this variable at the <code>[edit dynamic-profiles profile-name interfaces pp0]</code> hierarchy level for the <code>unit</code> statement.
<code>\$junos-underlying-interface</code>	Specifies the name of the underlying Ethernet interface on which the router dynamically creates the PPPoE logical interface. The <code>\$junos-underlying-interface</code> predefined variable is dynamically replaced with the name of the underlying interface supplied by the network when the PPPoE subscriber logs in. You specify this variable at the <code>[edit dynamic-profiles profile-name interfaces pp0 unit "\$junos-interface-unit" pppoe-options]</code> hierarchy level for the <code>underlying-interface</code> statement.
<b>Subscriber Interfaces — Dynamic Interface Sets</b>	
<code>\$junos-interface-set-name</code>	Name of an interface set configured in a dynamic profile. To represent the name of a dynamically created agent circuit identifier (ACI) interface set, use the <code>\$junos-interface-set-name</code> predefined variable in the <code>interface-set</code> statement at the <code>[edit dynamic-profiles profile-name interfaces]</code> hierarchy level.

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
<code>\$junos-phy-ifd-interface-set-name</code>	<p>Name of an interface set associated with the underlying physical interface in a dynamic profile.</p> <p>In a heterogeneous topology where residential and business subscribers share the same physical interface, although only two levels of CoS are required for residential access, business access requires three levels. Because they share the same physical interface, three levels are configured for both, causing an unnecessary level 2 node to be consumed for each residential connection.</p> <p>Starting in Junos OS Release 16.1, you can reduce the CoS resources wasted on residential access by collecting the residential subscribers into an interface set associated with the physical interface, so that a level 2 node is used for the interface set rather than for each residential interface. To do so, specify the <code>\$junos-phy-ifd-interface-set-name</code> predefined variable with the <code>interface-set</code> statement at the <code>[edit dynamic-profiles profile-name interfaces]</code> hierarchy level to create the interface set based on the underlying physical interface.</p>
<code>\$junos-svlan-interface-set-name</code>	<p>Locally generated interface set name for use by dual-tagged VLAN interfaces based on the outer tag of the dual-tagged VLAN. The format of the generated variable is <i>physical_interface_name - outer_VLAN_tag</i>.</p>
<b>Wholesale Networking</b>	
<code>\$junos-interface-name</code>	<p>Name of the dynamic interface to which the subscriber access client connects. Its use is in identifying the subscriber interface. You specify this variable at the <code>[dynamic-profiles profile-name routing-instance \$junos-routing-instance]</code> hierarchy level for the <code>interface</code> statement.</p> <p>The interface name is derived from concatenating the <code>\$junos-interface-ifd-name</code> and the <code>\$junos-underlying-interface-unit</code> variables obtained when a subscriber is created dynamically at the <code>[dynamic-profiles profile-name routing-instance \$junos-routing-instance interface]</code> hierarchy level.</p>

Table 9: Junos OS Predefined Variables and Definitions (*continued*)

Variable	Definition
\$junos-routing-instance	<p>Name of the routing instance to which the subscriber is assigned. This variable triggers a return value from the RADIUS server for Virtual-Router (VSA 26-1).</p> <p>You reference this variable in the statement at the <b>[dynamic-profiles <i>profile-name</i>]</b> hierarchy level for the <b>routing-instance</b> statement.</p> <p><b>NOTE:</b> Starting in Junos OS Release 16.1, when you specify a static logical interface for the unnumbered interface in a dynamic profile that includes the <b>\$junos-routing-instance</b> predefined variable, you must not configure a preferred source address, whether with the <b>\$junos-preferred-source-address</b> predefined variable, the <b>\$junos-preferred-source-ipv6-address</b> predefined variable, or the <b>preferred-source-address</b> statement. Configuring the preferred source address in this circumstance causes a commit failure.</p>
\$junos-inner-vlan-map-id	<p>Identifier for the inner VLAN tag for Layer 2 wholesale, ANCP-triggered, autosensed dynamic VLANs. The VLAN tag is allocated from the inner VLAN ID swap ranges that are provisioned on the core-facing physical interface. The inner VLAN tag is swapped with (replaces) the outer VLAN tag when the subscriber traffic is tunneled to the NSP.</p> <p>You specify this variable with the <b>inner-vlan-id</b> statement at the <b>[edit dynamic-profiles <i>profile-name</i> interfaces <i>profile-name</i> interface-ifd-name unit \$junos-interface-unit input-vlan-map]</b> hierarchy level.</p>
\$junos-vlan-map-id	<p>Identifier for a VLAN that is rewritten at the input or output interface as specified by a VLAN map.</p> <p>You specify this variable with the <b>vlan-id</b> statement at the <b>[edit dynamic-profiles <i>profile-name</i> interfaces <i>profile-name</i> interface-ifd-name unit \$junos-interface-unit input-vlan-map]</b> or <b>[edit dynamic-profiles <i>profile-name</i> interfaces <i>profile-name</i> interface-ifd-name unit \$junos-interface-unit input-vlan-map]</b> hierarchy levels.</p>

Release History Table

Release	Description
16.1	Starting in Junos OS Release 16.1, when you specify a static logical interface for the unnumbered interface in a dynamic profile that includes the <b>\$junos-routing-instance</b> predefined variable, you must not configure an IPv4 preferred source address, whether with the <b>\$junos-preferred-source-address</b> predefined variable or the <b>preferred-source-address</b> statement. Configuring the preferred source address in this circumstance causes a commit failure.
16.1	Starting in Junos OS Release 16.1, when you specify a static logical interface for the unnumbered interface in a dynamic profile that includes the <b>\$junos-routing-instance</b> predefined variable, you must not configure an IPv6 preferred source address, whether with the <b>\$junos-preferred-source-ipv6-address</b> predefined variable or the <b>preferred-source-address</b> statement. Configuring the preferred source address in this circumstance causes a commit failure.
16.1	Starting in Junos OS Release 16.1, you can reduce the CoS resources wasted on residential access by collecting the residential subscribers into an interface set associated with the physical interface, so that a level 2 node is used for the interface set rather than for each residential interface.
16.1	Starting in Junos OS Release 16.1, when you specify a static logical interface for the unnumbered interface in a dynamic profile that includes the <b>\$junos-routing-instance</b> predefined variable, you must not configure a preferred source address, whether with the <b>\$junos-preferred-source-address</b> predefined variable, the <b>\$junos-preferred-source-ipv6-address</b> predefined variable, or the <b>preferred-source-address</b> statement. Configuring the preferred source address in this circumstance causes a commit failure.

**Related  
Documentation**

- *Dynamic Variables Overview*
- *Configuring Predefined Dynamic Variables in Dynamic Profiles*
- *Junos OS Predefined Variables That Correspond to RADIUS Attributes and VSAs*
- *User-Defined Variables*

## Juniper Networks VSAs Supported by the AAA Service Framework

Table 10 on page 130 describes Juniper Networks VSAs supported by the Junos OS AAA Service Framework. The AAA Service Framework uses vendor ID 4874, which is assigned to Juniper Networks by the Internet Assigned Numbers Authority (IANA).



**NOTE:** A “Yes” entry in the Dynamic CoA Support column indicates that the attribute can be dynamically configured by Access-Accept messages and dynamically modified by CoA-Request messages.

Table 10: Supported Juniper Networks VSAs

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-1	Virtual-Router	<p>Client logical system:routing instance name. Allowed only from AAA server for default logical system:routing instance.</p> <p>When this VSA is not included in the subscriber profile, the routing instance assigned to the subscriber—the one in which the subscriber session comes up—varies by subscriber type.</p> <p>For DHCP and PPPoE subscribers, it is the default routing instance.</p> <p>For L2TP tunnel subscribers, it is the routing instance in which the tunnel resides, whether default or non-default. If the tunnel routing instance is not default and you want the L2TP session to be in the default routing instance, you must use the Virtual-Router VSA to set the desired routing instance.</p>	string: <i>logical system:routing instance</i>	No
26-4	Primary-DNS	Client DNS address negotiated during IPCP.	integer: 4-byte <i>primary-dns-address</i>	No
26-5	Secondary-DNS	Client DNS address negotiated during IPCP	integer: 4-byte <i>secondary-dns-address</i>	No
26-6	Primary-WINS	Client WINS (NBNS) address negotiated during IPCP.	integer: 4-byte <i>primary-wins-address</i>	No
26-7	Secondary-WINS	Client WINS (NBNS) address negotiated during IPCP.	integer: 4-byte <i>secondary-wins-address</i>	No
26-8	Tunnel-Virtual-Router	Virtual router name for tunnel connection.	string: <i>tunnel-virtual-router</i>	No
26-9	Tunnel-Password	<p>Tunnel password in cleartext.</p> <p>Do not use both this VSA and the standard RADIUS attribute Tunnel-Password [69]. We recommend that you use the standard attribute because the password is encrypted when that attribute is used.</p>	string: <i>tunnel-password</i>	No
26-10	Ingress-Policy-Name	Input policy name to apply to client interface.	string: <i>input-policy-name</i>	Yes

Table 10: Supported Juniper Networks VSAs (*continued*)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-11	Egress-Policy-Name	Output policy name to apply to client interface.	string: <i>output-policy-name</i>	Yes
26-23	IGMP-Enable	Whether IGMP is enabled or disabled on a client interface.	integer: <ul style="list-style-type: none"> <li>0=disable</li> <li>1=enable</li> </ul>	Yes
26-24	PPPoE-Description	Client MAC address.	string: <i>pppoe client-mac-address</i>	No
26-25	Redirect-VRouter-Name	Client logical system:routing instance name indicating to which logical system:routing instance the request is redirected for user authentication.	string: <i>logical-system:routing-instance</i>	No
26-30	Tunnel-Nas-Port-Method	<p>Method that determines whether the RADIUS server conveys to the LNS the physical NAS port number identifier and the type of the physical port, such as Ethernet or ATM. This information is conveyed only when the VSA value is 1.</p> <p>The VSA is formatted such that the first octet indicates the tunnel and the remaining three bytes are the attribute value.</p>	4-octet integer: <ul style="list-style-type: none"> <li>0 = none</li> <li>1 = Cisco CLID</li> </ul>	Yes
26-31	Service-Bundle	SSC service bundle.	string <i>bundle-name</i>	No
26-33	Tunnel-Max-Sessions	Maximum number of sessions allowed in a tunnel.	integer: 4-octet	No
26-34	Framed-IP-Route-Tag	Route tag to apply to returned framed-ip-address.	integer: 4-octet	No
26-42	Input-Gigapackets	Number of times the input-packets attribute rolls over its 4-octet field.	integer	No
26-43	Output-Gigapackets	Number of times the output-packets attribute rolls over its 4-octet field.	integer	No
26-47	Ipv6-Primary-DNS	Client primary IPv6 DNS address negotiated by DHCP.	hexadecimal string: <i>ipv6-primary-dns-address</i>	No
26-48	Ipv6-Secondary-DNS	Client secondary IPv6 DNS address negotiated by DHCP.	hexadecimal string: <i>ipv6-secondary-dns-address</i>	No

Table 10: Supported Juniper Networks VSAs (*continued*)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-51	Disconnect-Cause	Disconnect cause when a tunneled subscriber is disconnected, and L2TP layer of the LNS initiates the termination. The PPP Disconnect Cause Code (L2TP AVP 46) is included in VSA 26-51 in the Accounting-Stop message that the router sends to the RADIUS server.	hexadecimal string: <i>disconnect-cause</i>	No
26-55	DHCP-Options	Client DHCP options.	string: <i>dhcp-options</i>	No
26-56	DHCP-MAC-Address	Client MAC address.	string: <i>mac-address</i>	No
26-57	DHCP-GI-Address	DHCP relay agent IP address.	integer: 4-octet	No
26-58	LI-Action	<p>Traffic mirroring action.</p> <p>For dynamic CoA, VSA 26-58 changes the action on the mirrored traffic identified by VSA 26-59.</p> <p>CoA-Request messages that include any of the RADIUS-based mirroring attributes (VSAs 26-58, 26-59, 26-60, or 26-61) must always include all four VSAs.</p> <p>If the CoA action is to stop mirroring (VSA 26-58 value is 0), then the values of the other three attributes in the CoA message must match the existing attribute values, or the action fails.</p>	<p>salt-encrypted integer</p> <p>0=stop mirroring 1=start mirroring 2=no action</p>	Yes
26-59	Med-Dev-Handle	<p>Identifier that associates mirrored traffic to a specific subscriber.</p> <p>For dynamic CoA, VSA 26-58 changes the action on the mirrored traffic identified by VSA 26-59.</p> <p>CoA-Request messages that include any of the RADIUS-based mirroring attributes (VSAs 26-58, 26-59, 26-60, or 26-61) must always include all four VSAs.</p>	salt-encrypted string	No

Table 10: Supported Juniper Networks VSAs (*continued*)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-60	Med-Ip-Address	IP address of content destination device to which mirrored traffic is forwarded.  CoA-Request messages that include any of the RADIUS-based mirroring attributes (VSAs 26-58, 26-59, 26-60, or 26-61) must always include all four VSAs.	salt-encrypted IP address	No
26-61	Med-Port-Number	UDP port in the content destination device to which mirrored traffic is forwarded.  CoA-Request messages that include any of the RADIUS-based mirroring attributes (VSAs 26-58, 26-59, 26-60, or 26-61) must always include all four VSAs.	salt-encrypted integer	No
26-63	Interface-Desc	Text string that identifies the subscriber's access interface.	string: <i>interface-description</i>	No
26-64	Tunnel-Group	Name of the tunnel group (profile) assigned to a domain map.	string: <i>tunnel-group-name</i>	No
26-65	Activate-Service	Service to activate for the subscriber. Tagged VSA, which supports 8 tags (1-8).	string: <i>service-name</i>	Yes
26-66	Deactivate-Service	Service to deactivate for the subscriber.	string: <i>service-name</i>	Yes
26-67	Service-Volume	Amount of traffic, in MB, that can use the service; service is deactivated when the volume is exceeded. Tagged VSA, which supports 8 tags (1-8).	<ul style="list-style-type: none"> <li>range = 0 through 16777215 MB</li> <li>0 = no limit</li> </ul>	Yes
26-68	Service-Timeout	Number of seconds that the service can be active; service is deactivated when the timeout expires. Tagged VSA, which supports 8 tags (1-8).	<ul style="list-style-type: none"> <li>range = 0 through 16777215 seconds</li> <li>0 = no timeout</li> </ul>	Yes
26-69	Service-Statistics	Whether statistics for the service is enabled or disabled. Tagged VSA, which supports 8 tags (1-8).	<ul style="list-style-type: none"> <li>0 = disable</li> <li>1 = enable time statistics</li> <li>2 = enable time and volume statistics</li> </ul>	Yes
26-71	IGMP-Access-Name	Access list to use for the group (G) filter.	string: 32-octet	Yes

Table 10: Supported Juniper Networks VSAs (*continued*)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-72	IGMP-Access-Src-Name	Access list to use for the source-group (S,G) filter.	string: 32-octet	Yes
26-74	MLD-Access-Name	Access list to use for the group (G) filter.	string: 32-octet	Yes
26-75	MLD-Access-Src-Name	Access list to use for the source-group (S,G) filter.	string: 32-octet	Yes
26-77	MLD-Version	MLD protocol version.	integer: 1-octet <ul style="list-style-type: none"> <li>1=MLD version 1</li> <li>2=MLD version 2</li> </ul>	Yes
26-78	IGMP-Version	IGMP protocol version.	integer: 1-octet <ul style="list-style-type: none"> <li>1=IGMP version 1</li> <li>2=IGMP version 2</li> <li>3=IGMP version 3</li> </ul>	Yes
26-83	Service-Session	Name of the service.	string: <i>service-name</i>	No
26-84	Mobile-IP-Algorithm	Authentication algorithm used for Mobile IP registration.	integer: 4-octet	No
26-85	Mobile-IP-SPI	Security parameter index number for Mobile IP registration.	integer: 4-octet	No
26-86	Mobile-IP-Key	Security association MD5 key for Mobile IP registration.	string: key	No
26-87	Mobile-IP-Replay	Replay timestamp for Mobile IP registration.	integer: 4-octet	No
26-89	Mobile-IP-Lifetime	Registration lifetime for Mobile IP registration.	integer: 4-octet	No
26-91	Tunnel-Switch-Profile	Tunnel switch profile that determines whether a subscriber session is switched to a second session to a remote LNS. Takes precedence over tunnel switch profiles applied in any other manner.	string: <i>profile-name</i>	No
26-92	L2C-Up-Stream-Data	Actual upstream rate access loop parameter (ASCII encoded) as defined in GSMP extensions for layer2 control (L2C) Topology Discovery and Line Configuration.	string: actual upstream rate access loop parameter (ASCII encoded)	No

Table 10: Supported Juniper Networks VSAs (*continued*)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-93	L2C-Down-Stream-Data	Actual downstream rate access loop parameter (ASCII encoded) as defined in GSMP extensions for layer2 control (L2C) Topology Discovery and Line Configuration.	string: actual downstream rate access loop parameter (ASCII encoded)	No
26-94	Tunnel-Tx-Speed-Method	Method that determines the source from which the transmit speed is derived. Overrides global configuration in the CLI.	integer: 4-octet <ul style="list-style-type: none"> <li>• 0 = none</li> <li>• 1 = static Layer 2</li> <li>• 2 = dynamic layer 2. This method is not supported; the static Layer 2 method is used instead.</li> <li>• 3 = CoS. This method is not supported; the actual method is used instead.</li> <li>• 4 = actual</li> <li>• 5 = ANCP</li> <li>• 6 = PPPoE IA tags</li> </ul>	No
26-97	IGMP-Immediate-Leave	IGMP Immediate Leave.	integer: 4-octet <ul style="list-style-type: none"> <li>• 0=disable</li> <li>• 1=enable</li> </ul>	Yes
26-100	MLD-Immediate-Leave	MLD Immediate Leave.	integer: 4-octet <ul style="list-style-type: none"> <li>• 0=disable</li> <li>• 1=enable</li> </ul>	Yes
26-106	IPv6-Ingress-Policy-Name	Input policy name to apply to a user IPv6 interface.	string: <i>policy-name</i>	Yes
26-107	IPv6-Egress-Policy-Name	Output policy name to apply to a user IPv6 interface.	string: <i>policy-name</i>	Yes

Table 10: Supported Juniper Networks VSAs (*continued*)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-108	CoS-Traffic-Control-Profile-Parameter-Type	CoS traffic-shaping parameter type and description: <ul style="list-style-type: none"> <li><b>T01:</b> Scheduler-map name</li> <li><b>T02:</b> Shaping rate</li> <li><b>T03:</b> Guaranteed rate</li> <li><b>T04:</b> Delay-buffer rate</li> <li><b>T05:</b> Excess rate</li> <li><b>T06:</b> Traffic-control profile</li> <li><b>T07:</b> Shaping mode</li> <li><b>T08:</b> Byte adjust</li> <li><b>T09:</b> Adjust minimum</li> <li><b>T10:</b> Excess-rate high</li> <li><b>T11:</b> Excess-rate low</li> <li><b>T12:</b> Shaping rate burst</li> <li><b>T13:</b> Guaranteed rate burst</li> </ul>	Two parts, delimited by white space: <ul style="list-style-type: none"> <li>Parameter type</li> <li>Parameter value</li> </ul> Examples: <ul style="list-style-type: none"> <li><b>T01 smap_basic</b></li> <li><b>T02 50m</b></li> <li><b>T03 1m</b></li> <li><b>T04 2000</b></li> <li><b>T05 200</b></li> <li><b>T06 tcp-gold</b></li> <li><b>T07 frame-mode</b></li> <li><b>T08 50</b></li> </ul>	Yes
26-109	DHCP-Guided-Relay-Server	IP address of DHCP server that DHCP relay agent uses to forward the discover PDUs.	integer: 4-byte <i>ip-address</i>	No
26-110	Acc-Loop-Cir-Id	Identification of the subscriber node connection to the access node.	string: up to 63 ASCII characters	No
26-111	Acc-Aggr-Cir-Id-Bin	Unique identification of the DSL line.	integer: 8-octet	No
26-112	Acc-Aggr-Cir-Id-Asc	Identification of the uplink on the access node, as in the following examples: <ul style="list-style-type: none"> <li>Ethernet access aggregation—<b>ethernet slot/port [:inner-vlan-id] [:outer-vlan-id]</b></li> <li>ATM aggregation—<b>atm slot/port:vpi.vci</b></li> </ul>	string: up to 63 ASCII characters	No
26-113	Act-Data-Rate-Up	Actual upstream data rate of the subscriber's synchronized DSL link.	integer: 4-octet	No
26-114	Act-Data-Rate-Dn	Actual downstream data rate of the subscriber's synchronized DSL link.	integer: 4-octet	No
26-115	Min-Data-Rate-Up	Minimum upstream data rate configured for the subscriber.	integer: 4-octet	No
26-116	Min-Data-Rate-Dn	Minimum downstream data rate configured for the subscriber.	integer: 4-octet	No

Table 10: Supported Juniper Networks VSAs (*continued*)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-117	Att-Data-Rate-Up	Maximum upstream data rate that the subscriber can attain.	integer: 4-octet	No
26-118	Att-Data-Rate-Dn	Maximum downstream data rate that the subscriber can attain.	integer: 4-octet	No
26-119	Max-Data-Rate-Up	Maximum upstream data rate configured for the subscriber.	integer: 4-octet	No
26-120	Max-Data-Rate-Dn	Maximum downstream data rate configured for the subscriber.	integer: 4-octet	No
26-121	Min-LP-Data-Rate-Up	Minimum upstream data rate in low power state configured for the subscriber.	integer: 4-octet	No
26-122	Min-LP-Data-Rate-Dn	Minimum downstream data rate in low power state configured for the subscriber.	integer: 4-octet	No
26-123	Max-Interlv-Delay-Up	Maximum one-way upstream interleaving delay configured for the subscriber.	integer: 4-octet	No
26-124	Act-Interlv-Delay-Up	Subscriber's actual one-way upstream interleaving delay..	integer: 4-octet	No
26-125	Max-Interlv-Delay-Dn	Maximum one-way downstream interleaving delay configured for the subscriber.	integer: 4-octet	No
26-126	Act-Interlv-Delay-Dn	Subscriber's actual one-way downstream interleaving delay.	integer: 4-octet	No
26-127	DSL-Line-State	State of the DSL line.	integer: 4-octet <ul style="list-style-type: none"> <li>• 1 = Show uptime</li> <li>• 2 = Idle</li> <li>• 3 = Silent</li> </ul>	No
26-128	DSL-Type	Encapsulation used by the subscriber associated with the DSLAM interface from which requests are initiated.	integer: 4-octet	No
26-130	Qos-Set-Name	Interface set to apply to the dynamic profile.	string: <i>interface-set-name</i>	No

Table 10: Supported Juniper Networks VSAs (*continued*)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-140	Service-Interim-Acct-Interval	Amount of time between interim accounting updates for this service. Tagged VSA, which supports 8 tags (1-8).	<ul style="list-style-type: none"> <li>range = 600 through 86400 seconds</li> <li>0 = disabled</li> </ul> <p><b>NOTE:</b> Values are rounded up to the next higher multiple of 10 minutes. For example, a setting of 900 seconds (15 minutes) is rounded up to 20 minutes (1200 seconds).</p>	Yes
26-141	Downstream-Calculated-QoS-Rate	Calculated (adjusted) downstream QoS rate in Kbps as set by the ANCP configuration.	range = 1000 through 4,294,967,295	No
26-142	Upstream-Calculated-QoS-Rate	Calculated (adjusted) upstream QoS rate in Kbps as set by the ANCP configuration.	range = 1000 through 4,294,967,295	No
26-143	Max-Clients-Per-Interface	Maximum allowable client sessions per interface. For DHCP clients, this value is the maximum sessions per logical interface. For PPPoE clients, this value is the maximum sessions (PPPoE interfaces) per PPPoE underlying interface.	integer: 4-octet	No
26-146	CoS-Scheduler-Pmt-Type	<p>CoS scheduler parameter type and description:</p> <ul style="list-style-type: none"> <li>Null: CoS scheduler name</li> <li><b>T01:</b> CoS scheduler transmit rate</li> <li><b>T02:</b> CoS scheduler buffer size</li> <li><b>T03:</b> CoS scheduler priority</li> <li><b>T04:</b> CoS scheduler drop-profile low</li> <li><b>T05:</b> CoS scheduler drop-profile medium-low</li> <li><b>T06:</b> CoS scheduler drop-profile medium-high</li> <li><b>T07:</b> CoS scheduler drop-profile high</li> <li><b>T08:</b> CoS scheduler drop-profile any</li> </ul>	<p>Three parts, delimited by white space:</p> <ul style="list-style-type: none"> <li>Scheduler name</li> <li>Parameter type</li> <li>Parameter value</li> </ul> <p>Examples:</p> <ul style="list-style-type: none"> <li><b>be_sched</b></li> <li><b>be_sched T01 12m</b></li> <li><b>be_sched T02 26</b></li> </ul>	Yes
26-151	IPv6-Acct-Input-Octets	IPv6 receive octets.	integer	No
26-152	IPv6-Acct-Output-Octets	IPv6 transmit octets.	integer	No

Table 10: Supported Juniper Networks VSAs (*continued*)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-153	IPv6-Acct-Input-Packets	IPv6 receive packets.	integer	No
26-154	IPv6-Acct-Output-Packets	IPv6 transmit packets.	integer	No
26-155	IPv6-Acct-Input-Gigawords	IPv6 receive gigawords.	integer	No
26-156	IPv6-Acct-Output-Gigawords	IPv6 transmit gigawords.	integer	No
26-158	PPPoE-Padn	Route add for PPPoE sessions	string	No
26-160	Vlan-Map-Id	Trunk VLAN tag corresponding to the core-facing trunk physical interface.	integer	No
26-161	IPv6-Delegated-Pool-Name	Address pool used to locally allocate a delegated prefix (IA_PD).	string	No
26-162	Tx-Connect-Speed	Indication of transmit speed of the user's connection.	string	No
26-163	Rx-Connect-Speed	Indication of receive speed of the user's connection.	string	No
26-173	Service-Activate-Type	Indication of service activation type. This is a tagged attribute.	integer: 4-octet <ul style="list-style-type: none"> <li>1 = dynamic-profile for residential services</li> <li>2 = op-script for business services</li> </ul>	No
26-174	Client-Profile-Name	Enables RADIUS to override an assigned client dynamic profile with the included profile.	string	No
26-177	Cos-Shaping-Rate	Effective downstream shaping rate for subscriber.	string	No
26-179	Service-Volume-Gigawords	Amount of traffic, in 4GB units, that can use the service; service is deactivated when the volume is exceeded. Tagged VSA, which supports 8 tags (1-8).	<ul style="list-style-type: none"> <li>range = 0 through 16777215 4GB units</li> <li>0 = no limit</li> </ul>	Yes
26-180	Update-Service	New values of service and time quotas for existing service. Tagged VSA, which supports 8 tags (1-8).	string: <i>service-name</i>	Yes

Table 10: Supported Juniper Networks VSAs (*continued*)

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-181	DHCPv6-Guided-Relay-Server	IPv6 addresses of DHCPv6 servers to which DHCPv6 relay agent forwards the Solicit and subsequent PDUs. Use multiple instances of the VSA to specify a list of servers.	hexadecimal string: <i>ipv6-address</i>	No
26-182	Acc-Loop-Remote-Id	Reports the ANCP Access-Loop-Remote-ID attribute.	string	No
26-183	Acc-Loop-Encap	Reports the ANCP Access-Loop-Encapsulation attribute.	hexadecimal string	No
26-184	Inner-Vlan-Map-Id	Inner VLAN tag allocated from the ranges provisioned on the core-facing physical interface, used to swap (replace) the autosensed VLAN tag on the access interface.	integer	No
26-185	Core-Facing-Interface	Name of the core-facing physical interface that forwards the Layer 2 wholesale session's downstream and upstream traffic relative to the network service provider (NSP) router.	string	No
26-187	DHCP-First-Relay-IPv4-Address	IPv4 address of the first relay link of a client/server binding.	integer: 4-byte <i>ip-address</i>	No
26-188	DHCP-First-Relay-IPv6-Address	IPv6 address of the first relay link of a client/server binding.	hexadecimal string: <i>ipv6-address</i>	No
26-191	Input-Interface-Filter	Name of an input filter to be attached to a <b>family any</b> interface.	string	No
26-192	Output-Interface-Filter	Name of an output filter to be attached to a <b>family any</b> interface.	string	No

- Related Documentation**
- [AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 141](#)
  - [AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 147](#)

## AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS

Table 11 on page 141 shows the RADIUS attributes and Juniper Networks VSAs support in AAA access messages. A checkmark in a column indicates that the message type supports that attribute.

**Table 11: AAA Access Messages—Supported RADIUS Attributes and Juniper Networks VSAs**

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
1	User-Name	✓	✓	–	–	✓	✓
2	User-Password	✓	–	–	–	–	–
3	CHAP-Password	✓	–	–	–	–	–
4	NAS-IP-Address	✓	–	–	–	–	–
5	NAS-Port	✓	–	–	–	–	–
6	Service-Type	✓	✓	–	–	–	–
7	Framed-Protocol	✓	✓	–	–	–	–
8	Framed-IP-Address	✓	✓	–	–	✓	–
9	Framed-IP-Netmask	–	✓	–	–	–	–
11	Filter-Id	–	✓	–	–	–	–
12	Framed-MTU	✓	–	–	–	–	–
18	Reply-Message	–	✓	✓	✓	–	–
22	Framed-Route	–	✓	–	–	–	–
25	Class	–	✓	–	–	–	–
26-1	Virtual-Router	✓	✓	–	–	–	–
26-4	Primary-DNS	–	✓	–	–	–	–
26-5	Secondary-DNS	–	✓	–	–	–	–
26-6	Primary-WINS	–	✓	–	–	–	–
26-7	Secondary-WINS	–	✓	–	–	–	–

**Table 11: AAA Access Messages—Supported RADIUS Attributes and Juniper Networks VSAs (continued)**

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
26-8	Tunnel-Virtual-Router	–	✓	–	–	–	–
26-9	Tunnel-Password	–	✓	–	–	–	–
26-10	Ingress-Policy-Name	–	✓	–	–	–	–
26-11	Egress-Policy-Name	–	✓	–	–	–	–
26-23	IGMP-Enable	–	✓	–	–	–	–
26-24	PPPoE-Description	✓	–	–	–	–	–
26-25	Redirect-VR-Name	–	✓	–	–	–	–
26-31	Service-Bundle	–	✓	–	–	–	–
26-33	Tunnel-Maximum-Sessions	–	✓	–	–	–	–
26-34	Framed-IP-Route-Tag	–	✓	–	–	–	–
26-47	Ipv6-Primary-DNS	–	✓	–	–	–	–
26-48	Ipv6-Secondary-DNS	–	✓	–	–	–	–
26-55	DHCP-Options	✓	–	–	–	–	–
26-56	DHCP-MAC-Address	✓	✓	–	–	–	–
26-57	DHCP-GI-Address	✓	–	–	–	–	–
26-58	LI-Action	–	✓	–	–	✓	–
26-59	Med-Dev-Handle	–	✓	–	–	✓	–
26-60	Med-Ip-Address	–	✓	–	–	✓	–
26-61	Med-Port-Number	–	✓	–	–	✓	–
26-63	Interface-Desc	✓	–	–	–	–	–
26-64	Tunnel-Group	–	✓	–	–	–	–
26-65	Activate-Service	–	✓	–	–	✓	–
26-66	Deactivate-Service	–	✓	–	–	✓	–

Table 11: AAA Access Messages—Supported RADIUS Attributes and Juniper Networks VSAs (*continued*)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
26-67	Service-Volume	–	✓	–	–	✓	–
26-68	Service-Timeout	–	✓	–	–	✓	–
26-69	Service-Statistics	–	✓	–	–	✓	–
26-71	IGMP-Access-Name	–	✓	–	–	–	–
26-72	IGMP-Access-Src-Name	–	✓	–	–	–	–
26-74	MLD-Access-Name	–	✓	–	–	–	–
26-75	MLD-Access-Src-Name	–	✓	–	–	–	–
26-77	MLD-Version	–	✓	–	–	–	–
26-78	IGMP-Version	–	✓	–	–	–	–
26-91	Tunnel-Switch-Profile	–	✓	–	–	–	–
26-94	Tunnel-Tx-Speed-Method	–	✓	–	–	–	–
26-97	IGMP-Immediate-Leave	–	✓	–	–	–	–
26-100	MLD-Immediate-Leave	–	✓	–	–	–	–
26-106	IPv6-Ingress-Policy-Name	–	✓	–	–	–	–
26-107	IPv6-Egress-Policy-Name	–	✓	–	–	–	–
26-108	CoS-Parameter-Type	–	✓	–	–	✓	–
26-109	DHCP-Guided-Relay-Server	–	✓	–	–	–	–
26-110	Acc-Loop-Cir-Id	✓	–	–	–	–	–
26-111	Acc-Aggr-Cir-Id-Bin	✓	–	–	–	–	–
26-112	Acc-Aggr-Cir-Id-Asc	✓	–	–	–	–	–
26-113	Act-Data-Rate-Up	✓	–	–	–	–	–
26-114	Act-Data-Rate-Dn	✓	–	–	–	–	–
26-115	Min-Data-Rate-Up	✓	–	–	–	–	–

Table 11: AAA Access Messages—Supported RADIUS Attributes and Juniper Networks VSAs (continued)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
26-116	Min-Data-Rate-Dn	✓	–	–	–	–	–
26-117	Att-Data-Rate-Up	✓	–	–	–	–	–
26-118	Att-Data-Rate-Dn	✓	–	–	–	–	–
26-119	Max-Data-Rate-Up	✓	–	–	–	–	–
26-120	Max-Data-Rate-Dn	✓	–	–	–	–	–
26-121	Min-LP-Data-Rate-Up	✓	–	–	–	–	–
26-122	Min-LP-Data-Rate-Dn	✓	–	–	–	–	–
26-123	Max-Interlv-Delay-Up	✓	–	–	–	–	–
26-124	Act-Interlv-Delay-Up	✓	–	–	–	–	–
26-125	Max-Interlv-Delay-Dn	✓	–	–	–	–	–
26-126	Act-Interlv-Delay-Dn	✓	–	–	–	–	–
26-127	DSL-Line-State	✓	–	–	–	–	–
26-128	DSL-Type	✓	–	–	–	–	–
26-130	QoS-Set-Name	–	✓	–	–	–	–
26-140	Service-Interim-Account-Interval	–	✓	–	–	✓	–
26-141	Downstream-Calculated-QoS-Rate	✓	–	–	–	–	–
26-142	Upstream-Calculated-QoS-Rate	✓	–	–	–	–	–
26-143	Max-Clients-Per-Interface	–	✓	–	–	–	–
26-146	Cos-Scheduler-Pmt-Type	–	✓	–	–	✓	–
26-158	PPPoE-Padn	–	✓	–	–	–	–
26-160	Vlan-Map-ID	–	✓	–	–	–	–
26-161	IPv6-Delegated-Pool-Name	–	✓	–	–	–	–
26-162	Tx-Connect-Speed	✓	–	–	–	–	–

**Table 11: AAA Access Messages—Supported RADIUS Attributes and Juniper Networks VSAs (continued)**

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
26-163	Rx-Connect-Speed	✓	–	–	–	–	–
26-173	Service-Activate-Type	–	✓	–	–	✓	–
26-174	Client-Profile-Name	–	✓	–	–	–	–
26-179	Service-Volume-Gigawords	–	✓	–	–	✓	–
26-180	Update-Service	–	–	–	–	✓	–
26-181	DHCPv6-Guided-Relay-Server	–	✓	–	–	–	–
26-182	Acc-Loop-Remote-Id	✓	–	–	–	–	–
26-183	Acc-Loop-Encap	✓	–	–	–	–	–
26-184	Inner-Vlan-Map-Id	–	✓	–	–	–	–
26-185	Core-Facing-Interface	–	–	–	–	–	–
26-187	DHCP-First-Relay-IPv4-Address	✓	–	–	–	–	–
26-188	DHCP-First-Relay-IPv6-Address	✓	–	–	–	–	–
26-191	Input-Interface-Filter	✓	–	–	–	–	–
26-192	Output-Interface-Filter	✓	–	–	–	–	–
27	Session-Timeout	–	✓	–	✓	–	–
31	Calling-Station-ID	✓	–	–	–	✓	–
32	NAS-Identifier	✓	–	–	–	–	–
44	Acct-Session-ID	✓	–	–	–	✓	✓
61	NAS-Port-Type	✓	–	–	–	–	–
64	Tunnel-Type	–	✓	–	–	–	–
65	Tunnel-Medium-Type	–	✓	–	–	–	–
66	Tunnel-Client-Endpoint	–	✓	–	–	–	–
67	Tunnel-Server-Endpoint	–	✓	–	–	–	–

Table 11: AAA Access Messages—Supported RADIUS Attributes and Juniper Networks VSAs (*continued*)

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
68	Acct-Tunnel-Connection	–	✓	–	–	–	–
69	Tunnel-Password	–	✓	–	–	–	–
82	Tunnel-Assignment-Id	–	✓	–	–	–	–
83	Tunnel-Preference	–	✓	–	–	–	–
85	Acct-Interim-Interval	–	✓	–	–	–	–
87	NAS-Port-Id	✓	–	–	–	✓	–
88	Framed-Pool	–	✓	–	–	–	–
90	Tunnel-Client-Auth-Id	–	✓	–	–	–	–
91	Tunnel-Server-Auth-Id	–	✓	–	–	–	–
96	Framed-Interface-ID	–	✓	–	–	–	–
97	Framed-IPv6-Prefix	–	✓	–	–	–	–
99	Framed-IPv6-Route	–	✓	–	–	–	–
100	Framed-IPv6-Pool	–	✓	–	–	–	–
101	Error-Cause	–	–	–	–	✓	✓
123	Delegated-IPv6-Prefix	–	✓	–	–	–	–
242	Ascend-Data-Filter	–	✓	–	–	✓	–

**Related Documentation**

- [AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 147](#)
- *RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework*
- *RADIUS IETF Attributes Supported by the AAA Service Framework*
- [Juniper Networks VSAs Supported by the AAA Service Framework on page 129](#)

## AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS

Table 12 on page 147 shows the RADIUS attributes and Juniper Networks VSAs support in AAA accounting messages. A checkmark in a column indicates that the message type supports that attribute.

**Table 12: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs**

Attribute Number	Attribute Name	Acct Start	Acct Stop	Interim Acct	Acct On	Acct Off
1	User-Name	✓	✓	✓	–	–
3	CHAP-Password	✓	–	–	–	–
4	NAS-IP-Address	✓	✓	✓	✓	✓
5	NAS-Port	✓	✓	✓	–	–
6	Service-Type	✓	✓	✓	–	–
7	Framed-Protocol	✓	✓	✓	–	–
8	Framed-IP-Address	✓	✓	✓	–	–
9	Framed-IP-Netmask	✓	✓	✓	–	–
11	Filter-Id	–	✓	✓	–	–
22	Framed-Route	✓	✓	✓	–	–
25	Class	✓	✓	✓	–	–
26-1	Virtual-Router	✓	✓	✓	–	–
26-10	Ingress-Policy-Name	✓	✓	✓	–	–
26-11	Egress-Policy-Name	✓	✓	✓	–	–
26-24	PPPoE-Description	✓	✓	✓	–	–
26-42	Input-Gigapackets	–	✓	✓	–	–
26-43	Output-Gigapackets	–	✓	✓	–	–
26-47	Ipv6-Primary-DNS	✓	✓	✓	–	–
26-48	Ipv6-Secondary-DNS	✓	✓	✓	–	–

**Table 12: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs (continued)**

Attribute Number	Attribute Name	Acct Start	Acct Stop	Interim Acct	Acct On	Acct Off
26-51	Disconnect-Cause	–	✓	–	–	–
26-55	DHCP-Options	✓	✓	✓	–	–
26-56	DHCP-MAC-Address	✓	✓	✓	–	–
26-57	DHCP-GI-Address	✓	✓	✓	–	–
26-63	Interface-Desc	✓	✓	✓	–	–
26-83	Service-Session	–	✓	✓	–	–
26-110	Acc-Loop-Cir-Id	✓	✓	✓	–	–
26-111	Acc-Aggr-Cir-Id-Bin	✓	✓	✓	–	–
26-112	Acc-Aggr-Cir-Id-Asc	✓	✓	✓	–	–
26-113	Act-Data-Rate-Up	✓	✓	✓	–	–
26-114	Act-Data-Rate-Dn	✓	✓	✓	–	–
26-115	Min-Data-Rate-Up	✓	✓	✓	–	–
26-116	Min-Data-Rate-Dn	✓	✓	✓	–	–
26-117	Att-Data-Rate-Up	✓	✓	✓	–	–
26-118	Att-Data-Rate-Dn	✓	✓	✓	–	–
26-119	Max-Data-Rate-Up	✓	✓	✓	–	–
26-120	Max-Data-Rate-Dn	✓	✓	✓	–	–
26-121	Min-LP-Data-Rate-Up	✓	✓	✓	–	–
26-122	Min-LP-Data-Rate-Dn	✓	✓	✓	–	–
26-123	Max-Interlv-Delay-Up	✓	✓	✓	–	–
26-124	Act-Interlv-Delay-Up	✓	✓	✓	–	–
26-125	Max-Interlv-Delay-Dn	✓	✓	✓	–	–
26-126	Act-Interlv-Delay-Dn	✓	✓	✓	–	–

**Table 12: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs (continued)**

Attribute Number	Attribute Name	Acct Start	Acct Stop	Interim Acct	Acct On	Acct Off
26-127	DSL-Line-State	✓	✓	✓	–	–
26-128	DSL-Type	✓	✓	✓	–	–
26-141	Downstream-Calculated-QoS-Rate	✓	✓	✓	–	–
26-142	Upstream-Calculated-QoS-Rate	✓	✓	✓	–	–
26-151	IPv6-Acct-Input-Octets	–	✓	✓	–	–
26-152	IPv6-Acct-Output-Octets	–	✓	✓	–	–
26-153	IPv6-Acct-Input-Packets	–	✓	✓	–	–
26-154	IPv6-Acct-Output-Packets	–	✓	✓	–	–
26-155	IPv6-Acct-Input-Gigawords	–	✓	✓	–	–
26-156	IPv6-Acct-Output-Gigawords	–	✓	✓	–	–
26-160	Vlan-Map-Id	✓	✓	✓	–	–
26-162	Tx-Connect-Speed	✓	✓	✓	–	–
26-163	Rx-Connect-Speed	✓	✓	✓	–	–
26-177	Cos-Shaping-Rate	✓	✓	✓	–	–
26-182	Acc-Loop-Remote-Id	✓	✓	–	–	–
26-183	Acc-Loop-Encap	✓	✓	–	–	–
26-184	Inner-Vlan-Map-Id	✓	✓	–	–	–
26-185	Core-Facing-Interface	✓	✓	–	–	–
26-187	DHCP-First-Relay-IPv4-Address	✓	✓	✓	–	–
26-188	DHCP-First-Relay-IPv6-Address	✓	✓	✓	–	–
26-191	Input-Interface-Filter	✓	✓	✓	–	–
26-192	Output-Interface-Filter	✓	✓	✓	–	–
31	Calling-Station-ID	✓	✓	✓	–	–

**Table 12: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs (continued)**

Attribute Number	Attribute Name	Acct Start	Acct Stop	Interim Acct	Acct On	Acct Off
32	NAS-Identifier	✓	✓	✓	–	–
40	Acct-Status-Type	✓	✓	✓	✓	✓
41	Acct-Delay-Time	✓	✓	✓	✓	✓
42	Acct-Input-Octets	–	✓	✓	–	–
43	Acct-Output-Octets	–	✓	✓	–	–
44	Acct-Session-ID	✓	✓	✓	✓	✓
45	Acct-Authentic	✓	✓	✓	✓	✓
46	Acct-Session-Time	–	✓	✓	–	–
47	Acct-Input-Packets	–	✓	✓	–	–
48	Acct-Output-Packets	–	✓	✓	–	–
49	Acct-Terminate-Cause	–	✓	✓	–	–
52	Acct-Input-Gigawords	–	✓	✓	–	–
53	Acct-Output-Gigawords	–	✓	✓	–	–
55	Event-Timestamp	✓	✓	✓	✓	✓
61	NAS-Port-Type	✓	✓	✓	–	–
64	Tunnel-Type	✓	✓	✓	–	–
65	Tunnel-Medium-Type	✓	✓	✓	–	–
66	Tunnel-Client-Endpoint	✓	✓	✓	–	–
67	Tunnel-Server-Endpoint	✓	✓	✓	–	–
68	Acct-Tunnel-Connection	✓	✓	✓	–	–
82	Tunnel-Assignment-Id	✓	✓	✓	–	–
87	NAS-Port-Id	✓	✓	✓	–	–
90	Tunnel-Client-Auth-Id	✓	✓	✓	–	–

**Table 12: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs (continued)**

Attribute Number	Attribute Name	Acct Start	Acct Stop	Interim Acct	Acct On	Acct Off
91	Tunnel-Server-Auth-Id	✓	✓	✓	–	–
99	Framed-IPv6-Route	✓	✓	✓	–	–
100	Framed-IPv6-Pool	✓	✓	✓	–	–
123	Delegated-IPv6-Prefix	✓	✓	✓	–	–

**Related Documentation**

- [AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 141](#)
- *RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework*
- *RADIUS IETF Attributes Supported by the AAA Service Framework*
- [Juniper Networks VSAs Supported by the AAA Service Framework on page 129](#)



## CHAPTER 9

# Configuring ANCP-Triggered Layer 2 Wholesale Services

- [Configuring ANCP Neighbors on page 153](#)
- [Configuring the ANCP Agent for ANCP-Triggered, Autosensed Dynamic VLANs on page 154](#)
- [Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs on page 155](#)
- [Configuring Out-of-Band ANCP Messages to Trigger Dynamic VLAN Instantiation on page 156](#)
- [Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages on page 157](#)
- [Configuring the ANCP Agent to Dampen the Effects of Short-Term Adjacency Losses on page 159](#)
- [Reestablishing Pending Access Line Sessions for Layer 2 Wholesale on page 160](#)
- [Clearing ANCP Access Loops on page 160](#)

## Configuring ANCP Neighbors

---

Starting in Junos OS Release 16.1, you must configure each neighboring access node that you want the ANCP agent to monitor and potentially shape traffic for. Some neighbor settings override globally configured values.

To configure an ANCP neighbor:

1. Specify the IP address of the neighbor.

```
[edit protocols ancp]  
user@host# set neighbor 203.0.113.234
```

2. (Optional) Configure the neighbor to operate in a backward-compatible mode when it does not support the current IETF standard and the backward-compatible mode is not configured globally.

```
[edit protocols ancp neighbor 203.0.113.234]  
user@host# set pre-ietf-mode
```

3. (Optional) Override the globally configured backward-compatible mode when the neighbor supports the current IETF standard.

```
[edit protocols ancp neighbor 203.0.113.234]
user@host# set ietf-mode
```

4. (Optional) Configure the interval in seconds between ANCP adjacency messages exchanged with this neighbor.

```
[edit protocols ancp neighbor 203.0.113.234]
user@host# set adjacency-timer 20
```

5. (Optional) Specify the maximum number of discovery table entries that are accepted from this neighbor.

```
[edit protocols ancp neighbor 203.0.113.234]
user@host# set maximum-discovery-table-entries 10000
```

6. (Optional) Enable out-of-band ANCP triggering of autosensed, dynamic VLANs on the physical interface.

```
[edit protocols ancp neighbor 203.0.113.234]
user@host# set auto-configure-trigger interface ge-1/0/0
```

7. (Optional) Configure how long the ANCP agent maintains a Layer 2 wholesale session when an adjacency loss occurs.

```
[edit protocols ancp neighbor 203.0.113.234]
user@host# set adjacency-loss-hold-time 10
```

**Related  
Documentation**

- *Configuring the ANCP Agent*
- *Configuring the ANCP Agent for Backward Compatibility*
- *Specifying the Interval Between ANCP Adjacency Messages*
- *Specifying the Maximum Number of Discovery Table Entries*
- [Configuring the ANCP Agent to Dampen the Effects of Short-Term Adjacency Losses on page 159](#)

---

## Configuring the ANCP Agent for ANCP-Triggered, Autosensed Dynamic VLANs

Starting in Junos OS Release 16.1, you can configure the ANCP agent to associate a neighbor with an access-facing physical interface for the creation of autosensed dynamic VLANs on the interface. When the ANCP agent receives a Port Up message from the neighbor, it triggers notification to the autoconfd daemon to initiate the detection, authorization, and creation of dynamic VLANs. Receipt of an out-of-band ANCP Port Down message triggers notification to the autoconfd daemon to initiate the destruction of an existing VLAN on the interface.

This configuration assumes the following:

- The dynamic profile is configured to instantiate a dynamic VLAN when notified by the ANCP agent that it has received an out-of-band ANCP Port Up message.
- The RADIUS authentication server is properly configured to authorize the VLANs and apply services as needed.

- The ANCP agent is configured to initiate interim accounting updates (which also enables immediate interim accounting updates) in response to information received in Port Up messages.

To map a neighbor to a physical interface for autosensed dynamic VLANs:

- Specify the physical interface name.

```
[edit protocols ancp]
user@host# set auto-configure-trigger interface physical-interface-name
```

#### Related Documentation

- *Configuring the ANCP Agent*
- [Configuring ANCP Neighbors on page 153](#)
- *ANCP and the ANCP Agent Overview*
- [Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93](#)
- [Configuring Out-of-Band ANCP Messages to Trigger Dynamic VLAN Instantiation on page 156](#)

## Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs

When a subscriber logs in, the Access-Request message that is sent to the RADIUS server includes a username and optionally a password generated locally on the router to authenticate the subscriber during the VLAN authorization process. Starting in Junos OS Release 16.1, for a Layer 2 network that is wholesaled to a retailer where the dynamic VLANs are instantiated by out-of-band ANCP Port Up messages, you can configure the router to create a unique username with the value of the ANCP TLVs—Access-Loop-Circuit-ID, Access-Loop-Remote-ID, or both—as received in the ANCP Port Up message from the access node.

This configuration assumes the following:

- The ANCP agent is configured to notify AAA when it receives ANCP Port Up and Port Down messages.
- The dynamic profile is configured to instantiate a dynamic VLAN when notified by the ANCP agent that it has received an out-of-band ANCP Port Up message.
- The RADIUS authentication server is properly configured.

To include ANCP TLVs in the authentication username

1. (Optional) Specify inclusion of the Access-Loop-Circuit-ID TLV value.

```
[edit interfaces ge-0/0/0 auto-configure vlan-ranges username-include]
user@host# set circuit-id
```

2. (Optional) Specify inclusion of the Access-Loop-Remote-ID TLV value.

```
[edit interfaces ge-0/0/0 auto-configure vlan-ranges username-include]
user@host# set remote-id
```



**NOTE:** This ANCP information is not supported in stacked VLANs.



**NOTE:** You can use any of the attributes available to the `username-include` statement, except: `mac-address`, `option-18`, `option-37`, and `option-82`.

You can include other information in the username as for conventional autosensed dynamic VLANs. Alternatively, if you configure the router to convey ANCP-sourced access loop attributes as Juniper Networks VSAs—in this case `Acc-Loop-Cir-Id` (26-110) and `Acc-Loop-Remote-Id` (26-182)—the Access-Request message includes sufficient unique access line information for the RADIUS server to determine whether the access loop is wholesaled to a retailer or retained for the wholesaler.

**Related  
Documentation**

- [Configuring VLAN Interface Username Information for AAA Authentication](#)
- [Configuring Out-of-Band ANCP Messages to Trigger Dynamic VLAN Instantiation on page 156](#)
- [Configuring the ANCP Agent for ANCP-Triggered, Autosensed Dynamic VLANs on page 154](#)
- [Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93](#)

---

## Configuring Out-of-Band ANCP Messages to Trigger Dynamic VLAN Instantiation

The instantiation of conventional autosensed dynamic VLANs is triggered by in-band PPPoE or DHCP control packets that the Packet Forwarding Engine exceptions to the Routing Engine. A VLAN is authorized based on information extracted from specific fields and created according to a dynamic profile assigned to the VLAN range or stacked VLAN range.

Another way to instantiate an autosensed dynamic VLAN is with the processing of packets from an out-of-band protocol, ANCP. The out-of-band protocol method is useful where the traffic received might not be PPPoE or DHCP, such as in a Layer 2 wholesale scenario, where the traffic for an entire outer VLAN is wholesaled to a retailer and the VLANs are based on access line identifiers.

For this method, you configure the dynamic profile to accept packets from the out-of-band protocol. The dynamic profile is on an access-facing physical interface and is associated with a VLAN range available for the autosensed VLANs.

This configuration assumes the following:

- The dynamic profile is configured to instantiate a dynamic VLAN when notified by the ANCP agent that it has received an out-of-band ANCP Port Up message.
- The RADIUS authentication server is properly configured to authorize the VLANs and apply services as needed.

- The ANCP agent is configured to notify AAA when it receives ANCP Port Up and Port Down messages.
- The ANCP agent is configured to initiate interim accounting updates (which also enables immediate interim accounting updates) in response to information received in Port Up messages.



**NOTE:** Out-of-band triggering is supported only for single-tag VLANs; it is not supported for stacked VLANs.

To configure the instantiation of autosensed dynamic VLANs by out-of-band ANCP packets:

- Specify that ANCP packets are accepted.

```
[edit interfaces interface-name auto-configure vlan-ranges dynamic-profile profile-name]
user@host# set accept-out-of-band ancp
```

#### Related Documentation

- [Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93](#)
- [Configuring the ANCP Agent for ANCP-Triggered, Autosensed Dynamic VLANs on page 154](#)
- [Configuring VLAN Interface Username Information for AAA Authentication](#)
- [Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs](#)

## Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages

Starting in Junos OS Release 16.1, you can trigger ANCP OAM to simulate the sending of an ANCP Port-Down or Port-Up message. Typically you use this feature only when troubleshooting an ANCP issue or to mitigate an error condition when ANCP is not operating normally.

When you issue either the **request ancp oam port-down** command or the **request ancp oam port-up** command from operational mode, you must specify either an IP address for an ANCP neighbor or the physical interface used for subscriber access. You must also specify all of the following; all three are required together to identify the access node:

- circuit-id *aci*—ANCP Access-Loop-Circuit-ID TLV
- remote-id *ari*—ANCP Access-Loop-Remote-ID TLV
- outer-vlan-id *vlan-id*—ANCP Access-Aggregation-Circuit-ID-Binary TLV

You can use the **request ancp oam port-up** command to trigger reauthorization and re-creation of the dynamic VLAN session and logical interface that is supporting Layer 2 wholesale after they have been removed by any of the following:

- Issuance of the **clear network-access aaa subscriber** command.

- Receipt of a RADIUS disconnect message that does not include the RADIUS Acct-Terminate-Cause attribute (49).
- Action by the ANCP agent.

The previous instance of the VLAN can be either ANCP-triggered (a wholesaled VLAN) or a conventionally autosensed dynamic VLAN (an access-provider-owned VLAN).

If no access line parameters are available from ANCP for a given access line, you can use the **request ancp oam port-up** command as a test mechanism to trigger authorization of a dynamic VLAN session and logical interface. The session and interface are created when a RADIUS Access-Accept message is subsequently received.

These commands have no effect on conventionally autosensed dynamic VLANs (for the access provider's own subscriber sessions) that have matching access loop attributes.



**NOTE:** Genuine ANCP Port-Down and Port-Up messages take precedence over these simulated messages. This means that when a Port-Down message has already been received, you cannot use the **request ancp oam port-up** command to initiate the Port-Up condition. When a Port-Up message has already been received, you cannot use the **request ancp oam port-down** command to initiate the Port-Down condition.

You can use the **request ancp oam port-down** command to trigger removal of the ANCP-triggered, autosensed, dynamic VLAN that corresponds to the specified attributes. The typical use for this command is to remove the VLAN created by sending a **request ancp oam port-up** command.

To simulate an ANCP Port Up message:

- Identify the loop by the neighbor's IP address or the access-facing physical interface, and the ACI, ARI, and outer VLAN ID.

```
user@host> request ancp oam port-up neighbor 192.168.32.5 circuit-id line-aci-1
remote-id line-ari-1 outer-vlan-id 126
user@host> request ancp oam port-up subscriber-interface ge-1/0/1 circuit-id line-aci-1
remote-id line-ari-1 outer-vlan-id 126
```

To simulate an ANCP Port Down message:

- Identify the loop by the neighbor's IP address or the access-facing physical interface, and the ACI, ARI, and outer VLAN ID.

```
user@host> request ancp oam port-down neighbor 192.168.32.5 circuit-id line-aci-1
remote-id line-ari-1 outer-vlan-id 126
user@host> request ancp oam port-down subscriber-interface ge-1/0/1 circuit-id
line-aci-1 remote-id line-ari-1 outer-vlan-id 126
```

To verify the operation of either request, you can enter the following commands before and after initiating the Port Down or Port Up message:

- **show subscribers client-type vlan-oob detail**—Subscriber information is displayed for the VLAN on Port Up, or disappears on Port Down.
- **show subscribers summary**—The VLAN-OOB counter reflects the creation or removal of the VLAN-OOB session by incrementing (Port Up) or decrementing (Port Down).
- **show l2-routing-instance routing-instance-name**—The VLAN counters reflect the creation or removal of the VLAN-OOB session by incrementing (Port Up) or decrementing (Port Down).

**Related Documentation**

- [Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93](#)

## Configuring the ANCP Agent to Dampen the Effects of Short-Term Adjacency Losses

By default, the ANCP agent treats a loss of adjacency as if it has received a Port Down message for every access loop that is represented by the adjacency. All Layer 2 wholesale sessions are logged out and cleaned up. If the associated physical interface is in the Down state, then pending sessions cannot be reestablished when the interface transitions back to the Up state.

Starting in Junos OS Release 16.1, you can configure the ANCP agent to maintain the corresponding ANCP-triggered Layer 2 wholesale sessions for a configurable period in the event that an ANCP adjacency is lost. If the adjacency is restored before the timer expires, the session continues. If the timer expires before the adjacency is restored, then the session is logged out and cleaned up. This behavior dampens the effect of unstable ANCP connections. The hold timer can also detect when an access line is unconfigured on a neighbor and trigger logout and cleanup of the related sessions.



**NOTE:** The default value of the timer is 0, which means that the loss of neighbor adjacency immediately triggers a logout of all corresponding Layer 2 wholesale sessions.

To configure how long the ANCP agent maintains sessions in the event of an adjacency loss for any neighbor:

- Specify the hold timer duration in seconds.

```
[edit protocols ancp]
user@host# set adjacency-loss-hold-time seconds
```

To configure how long the ANCP agent maintains sessions in the event of an adjacency loss for a specific neighbor:

- Specify the hold timer duration in seconds.

```
[edit protocols ancp neighbor ip-address]
user@host# set adjacency-loss-hold-time seconds
```

- Related Documentation**
- [Configuring the ANCP Agent](#)
  - [Configuring ANCP Neighbors on page 153](#)
  - [ANCP and the ANCP Agent Overview](#)
  - [Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93](#)

---

## Reestablishing Pending Access Line Sessions for Layer 2 Wholesale

The access lines for ANCP-triggered, Layer 2 wholesale sessions can transition to a pending state after an ANCP adjacency loss when the inner VLAN ID swap range has been exhausted of tags and no other eligible core-facing physical interfaces are available. Typically, the sessions are reestablished when more VLAN IDs are made available, such as by extending the swap range, or more interfaces are available, such as by reconfiguration. When that does not happen, starting in Junos OS Release 16.1, you can manually initiate the reestablishment process by issuing the **request auto-configuration reconnect-pending** command.

To manually reestablish sessions for which the corresponding access lines are in the pending state:

- Specify the routing instance with the reconnection request.  
`user@host> request auto-configuration reconnect-pending`

- Related Documentation**
- [Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93](#)

---

## Clearing ANCP Access Loops

Starting in Junos OS Release 16.1, you can force a reset of a particular Layer 2 wholesale connection while the access loop is operationally up by issuing the **clear ancp access-loop** command. The command initiates logout of an ANCP-triggered, dynamic VLAN session, which includes issuing RADIUS Accounting-Stop messages for the session, and removal of the dynamic VLAN logical interface and active services. After the session is cleaned up, the command initiates re-authorization of the dynamic VLAN session, simulating receipt of an ANCP Port Up message. The session may then be recreated.

You must identify the access loop by either the IP address of the ANCP neighbor or the name of the subscriber-facing physical interface. You must also specify one or more of the following additional identifiers for the access loop:

- **circuit-id**—The value of the ANCP Access-Loop-Circuit-ID TLV.
- **remote-id**—The value of the ANCP Access-Loop-Remote-ID TLV.
- **outer-vlan-id**—The value of the ANCP Access-Aggregation-Circuit-ID-binary TLV.

**NOTE:**

The `clear ancp access-loop` command has no effect in the following circumstances:

- The access line is reported to be down, as indicated by an ANCP Port Down message, when the command is issued.
- An ANCP Port Down message is received for the access line while the dynamic VLAN logical interface and the services are being removed. In this case, re-authorization of the dynamic VLAN cannot take place until an ANCP Port Up message is received for that access line.
- A conventionally autosensed dynamic VLAN (for the access provider's own subscriber sessions) has matching access loop attributes. In this case, the Layer 2 wholesale access line for which the command is intended is cleared, but the other VLAN, for sessions owned by the access-provider, is cleared as expected.

To clear an ANCP access loop:

- Identify the loop by the neighbor's IP address or the access-facing physical interface, and one or more of the ACI, ARI, and outer VLAN ID.

```
user@host> clear ancp access-loop neighbor 192.168.32.5 circuit-id line-aci-1 remote-id
line-ari-1 outer-vlan-id 126
user@host> clear ancp access-loop subscriber-interface ge-1/0/1 circuit-id line-aci-1
remote-id line-ari-1 outer-vlan-id 126
```

**Related  
Documentation**

- [Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93](#)



## CHAPTER 10

# Configuring Flat-File Accounting for Layer 2 Wholesale Services

- [Flat-File Accounting Overview on page 163](#)
- [Configuring Flat-File Accounting for Layer 2 Wholesale on page 166](#)
- [Configuring Flat-File Accounting for Extensible Subscriber Services Management on page 170](#)

### Flat-File Accounting Overview

---

Accounting statistics can be collected from the Packet Forwarding Engine and reported in an XML flat file, which both contains and describes the data. Starting in Junos OS Release 16.1, you can use a flat-file profile that acts as a template to define attributes for accounting flat files.

When the accounting file is created, a file header is also created if the file format is IP Detail Record (IPDR). The header is not created if the format is comma-separated variable (CSV). The file header includes the following information:

- XML namespace—Static link to the World Wide Web Consortium (W3C) organization's XML Schema Instance (XSI) definition.
- Schema version—Configurable name of the schema that defines the information conveyed in the accounting file. The schema version is associated with a specific XML format and output based on the flat-file profile configuration that is used for the business purpose. This structure enables the XML-formatted contents of the file to be correctly interpreted by the service provider's external file processor.
- NAS ID—Name of the BNG host (network access server) where the accounting statistics are collected.
- File creation timestamp—UTC time zone date and time when the accounting file was created.
- File ID—Number identifying the file. The ID is incremented when a new accounting file is created and can range from 1 through 2,147,483,647.

For example, consider the following sample header for an accounting file for Extensible Subscriber Services Manager (ESSM) business subscribers:

```

<BNGFile xmlns:xsi=http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="BNG_IPDR_20130423.xsd" NAS-ID="host-mx480-x5"

  FileCreationTimeStamp="2015-10-09T08:25:50" FileID="29">
<IPDR>
.....
.....
</IPDR>
</BNGFile>

```

Table 13 on page 164 lists the elements and their values in the sample header.

**Table 13: Value of Elements in Sample Accounting Flat File XML Header**

Description	Header Element	Value
XML namespace	xmlns	:xsi=http://www.w3.org/2001/XMLSchema-instance"
schema version	xsi:noNamespaceSchemaLocation	BNG_IPDR_20130423.xsd
NAS ID	NAS-ID	host-mx480-x5
file creation timestamp	FileCreationTimeStamp	2015-10-09T08:25:50
File ID	FileID	29

You can configure the following options for flat-file accounting at the **[edit accounting-options file filename]** hierarchy level:

- Maximum size of the accounting file.
- Number of files that are saved before overwriting.
- One or more sites where the files are sent for archiving.
- Frequency at which the files are transferred to an archive site.
- Start time for file transfer.
- Compression for the transferred files.
- Local backup on the router for files when transfer fails.
- Whether accounting files are saved when a change in mastership occurs for both the new master Routing Engine and the new backup Routing Engine or for only the new master Routing Engine.
- How long files are kept before being deleted from the local backup directory.

You can also create one or more flat-file profiles at the **[edit accounting-options flat-file-profile *profile-name*]** hierarchy level that act as templates to specify the following attributes for new accounting files when they are created:

- Statistics fields that you want to collect, such as egress statistics or ingress statistics fields.
- Name and format of the accounting file.
- Frequency at which the Packet Forwarding Engine is polled for the statistics.
- Schema version.

Archive sites provide security and storage for the accounting files, which are transferred at regular intervals. When more than one archive site is configured, the router attempts to transfer the files to the first site on the list. If that fails, the router tries each of the other sites in turn until the transfer either succeeds for one site or fails for all sites. If you configure the last site in the list to be a local directory on the router rather than another remote site, then the files are backed up locally if all remote sites fail. The failed files are simply stored in the designated site. They are not automatically resubmitted to the archival sites. You must use an event script or some other means to have these files resubmitted. Any files remaining in the local directory are deleted when the **cleanup-interval** expires.

Alternatively, you can use the **backup-on-failure** statement at the **[edit accounting-options file *filename*]** hierarchy level to back up the files locally if all the remote attempts fail. If that occurs, the router compresses the accounting files and backs them up to the **/var/log/pfedBackup/** directory. Whenever any of the archive sites is reachable, the router attempts to transfer the data from **/var/log/pfedBackup/** to that site in compressed format. If the transfer of the backed-up files to the reachable site fails, the system tries to transfer the files to any other site that becomes reachable during the transfer interval. Any files that fail to transfer are compressed and kept in **/var/log/pfedBackup/** until an archival site is reachable and the files are successfully transferred. Any files that remain in that directory are deleted when the **cleanup-interval** expires.



**BEST PRACTICE:** Use the **backup-on-failure** feature to reliably and automatically back up files and retransmit them to archives rather than relying on a local site listed as the last archive site.

If the backup Routing Engine does not have access to the archive site—for example, when the site is not connected by means of an out-of-band interface or when the path to the site is routed through a line card—you can ensure that the backup Routing Engine's accounting files are backed up by using the **push-backup-to-master** statement at the **[edit accounting-options file *filename*]** hierarchy level. When a change in mastership occurs, the new backup Routing Engine saves its files to the **/var/log/pfedBackup/** directory. The master Routing Engine subsequently includes these files when it sends its own accounting files to the archive site at every transfer interval.

To conserve resources during transfer of accounting files and at the archive site, use the **compress** statement at the **[edit accounting-options file *filename*]** hierarchy level to compress the files when they are transferred. This option is disabled by default.

A system logging message is generated when a transfer succeeds (**transfer-file: Transferred *filename***) or fails (**transfer-file failed to transfer**). In the event of a failure, an error message is logged to indicate the nature of the failure.

**Related  
Documentation**

- [Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170](#)
- [Configuring Flat-File Accounting for Layer 2 Wholesale on page 166](#)
- [Configuring Accounting-Data Log Files](#)

---

## Configuring Flat-File Accounting for Layer 2 Wholesale

---

Although flat-file accounting is typically used for recording accounting statistics on logical interfaces for Extensible Subscriber Services Manager (ESSM) business subscribers, starting in Junos OS Release 16.1, you can also use flat-file accounting to collect and archive various accounting statistics for your Layer 2 wholesale environment. You do this by creating a flat-file profile and applying it to a core-facing physical interface.

You can also configure a flat-file profile to monitor and report Layer 2 multicast statistics; you assign this profile to the logical interface configured on the core-facing physical interface. This approach enables you to have separate accounting files that overlap in content only in the non-statistical, general parameters. The Layer 2 multicast statistics are available only when the encapsulation on the logical interface is **ethernet-vpls**.

You can configure multiple accounting profiles with different combinations of fields for specific accounting requirements, and then assign the profiles as needed to provisioned interfaces to satisfy the accounting requirements for each interface depending on how the interface is used.

A given flat-file profile can be assigned to both use cases; for example, by specifying **all-fields** for a global or group level. In this case, the fields you configure appear in the accounting record only if they make sense in the context.



---

**BEST PRACTICE:** We recommend you use separate flat-file profiles for Layer 2 wholesale core-facing physical interfaces and ESSM business subscriber logical interfaces.

---

Some statistics and general parameter fields are available either only for logical interfaces or only for physical interfaces. The **accounting-type**, **line-id**, **nas-port-id**, and **vlan-id** general parameters are not available for core-facing physical interfaces. Because the core-facing physical interfaces carry Layer 2 cross-connected sessions, no useful IPv6 statistics are available. Accordingly, do not configure the **input-v6-bytes**, **input-v6-packets**, **output-v6-bytes**, or **output-v6-packets** overall packet fields.

To configure flat-file accounting for a Layer 2 wholesale network:

1. Create a flat-file profile.

```
[edit accounting-options]
user@host# edit flat-file-profile profile-name
```

2. (Optional) Configure the name of the XML schema for the accounting flat file.

```
[edit accounting-options flat-file-profile profile-name]
user@host# set schema-version schema-name
```

3. Specify the filename for the accounting file.

```
[edit accounting-options flat-file-profile profile-name]
user@host# set file accounting-filename
```

4. Specify the general, nonstatistical parameters for the accounting file that are displayed as part of the accounting record header.

```
[edit accounting-options flat-file-profile profile-name fields]
user@host# set general-param option
```



**BEST PRACTICE:** We recommend that you include the general parameter `all-fields` option for both core-facing physical interfaces and, when you are collecting Layer 2 multicast statistics, on the logical interface that represents the physical interface.

5. Specify the accounting statistics that are collected and recorded in the accounting file for the core-facing physical interface.

```
[edit accounting-options flat-file-profile profile-name fields]
user@host# set egress-stats option
user@host# set ingress-stats option
user@host# set overall-packet option
```



**BEST PRACTICE:** We recommend that you include the following statistics fields in flat-file profiles for core-facing physical interfaces:

- Egress statistics fields: `all-fields`
- Ingress statistics fields: `all-fields`
- Overall packet fields: `input-bytes`, `input-discards`, `input-errors`, `input-packets`, `output-bytes`, `output-errors`, `output-packets`

6. (Optional) For Layer 2 multicast statistics, specify the accounting statistics that are collected and recorded in the accounting file for the logical interface representing the core-facing physical interface.

```
[edit accounting-options flat-file-profile profile-name fields]
user@host# set l2-stats option
```



**BEST PRACTICE:** We recommend that you include the following statistics fields in flat-file profiles for logical interfaces on the core-facing physical interfaces:

- Layer 2 statistics fields: all-fields

7. (Optional) Specify the format of the accounting file.

```
[edit accounting-options flat-file-profile profile-name]
user@host# set format (csv | ipdr)
```

8. (Optional) Specify the interval at which the Packet Forwarding Engine associated with the interface is polled for the statistics specified in the profile.

```
[edit accounting-options flat-file-profile profile-name]
user@host# set interval minutes
```



**NOTE:** When you do not configure this option, the polling interval is 15 minutes.

9. Configure the maximum size of the accounting file.

```
[edit accounting-options file filename]
user@host# set size bytes
```

10. Configure one or more archive sites for the files.

```
[edit accounting-options file filename]
user@host# set archive-sites site-name
```

The **site-name** is any valid FTP or Secure FTP URL. When the file is archived, the router attempts to transfer the file to the archive site. If you have specified more than one site, it tries the first site in the list. If that fails, it tries each site in turn until a transfer succeeds. The log file is stored at the archive site with a filename of the format **router-name\_log-filename\_timestamp**. The last site in a list is often a local directory, in case no remote site is reachable.

11. (Optional) Configure the start time for transferring files.

```
[edit accounting-options file filename]
user@host# set start-time YYYY-MM-DD.hh:mm
```

12. (Optional) Configure how frequently the file is transferred.

```
[edit accounting-options file filename]
user@host# set transfer-interval minutes
```



**NOTE:** When you do not configure this option, the file is transferred every 30 minutes.

13. (Optional) Configure the maximum number of files (3 through 1000) to save.

```
[edit accounting-options file filename]
user@host# set files number
```



**NOTE:** When you do not configure this option, a maximum of 10 files are saved.

14. (Optional) Configure the router to save a backup copy of the accounting file to the `/var/log/pfedBackup` directory if the normal transfer of the files to the archive sites fails.

```
[edit accounting-options file filename]
user@host# set backup-on-failure
```



**NOTE:** When you do not configure this option, the file is saved on failure into the local directory specified as the last site in the list of archive sites.

15. (Optional) Configure the accounting file to be compressed during transfer to an archive site.

```
[edit accounting-options file filename]
user@host# set compress
```

16. (Optional) Configure the router's new backup Routing Engine to send its accounting file to the `/var/log/pfedBackup` directory on the new master Routing Engine when a change in mastership occurs.

```
[edit accounting-options file filename]
user@host# set push-backup-to-master
```

17. (Optional) Configure the number of days after which accounting files are deleted from the local backup directory.

```
[edit accounting-options]
user@host# set cleanup-interval days
```



**NOTE:** Files are retained for 1 day if you do not configure this option.

18. Assign the profile to the relevant core-facing physical interface.

```
[edit interfaces physical-interface-name]
user@host# set accounting-profile flat-file-profile-name
```

19. Assign the profile to the relevant interface.

For the core-facing physical interface:

```
[edit interfaces physical-interface-name]
user@host# set accounting-profile flat-file-profile-name
```

For the logical interface representing the core-facing physical interface:

```
[edit interfaces physical-interface-name unit logical-unit-number]
user@host# set accounting-profile flat-file-profile-name
```

#### Related Documentation

- [Configuring Accounting-Data Log Files](#)

- [Flat-File Accounting Overview on page 163](#)
- [Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170](#)
- [Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93](#)

## Configuring Flat-File Accounting for Extensible Subscriber Services Management

---

Flat-file accounting is typically used to collect and archive various accounting statistics on logical interfaces for Extensible Subscriber Services Manager (ESSM) business subscribers. Other applications include accounting for wholesaler and retailer subscriber activity in a Layer 2 wholesale environment. Starting in Junos OS Release 16.1, you can create a flat-file profile to use as a template to define attributes for accounting flat files. The profile specifies the following:

- The statistics fields that are collected.
- The filename where the statistics are logged.
- The format of the file, the interval at which the statistics are collected.
- The name of the XML schema file that specifies the contents of the accounting file.

You can configure multiple accounting profiles with different combinations of fields for specific accounting requirements, and then assign the profiles as needed to provisioned interfaces to satisfy the accounting requirements for each interface depending on how it is used.

A given flat-file profile can be assigned to both use cases; for example, by specifying **all-fields** for a global or group level. In this case, the fields you configure appear in the accounting record only if they make sense in the context.



**BEST PRACTICE:** We recommend you use separate flat-file profiles for ESSM business subscriber logical interfaces and Layer 2 wholesale core-facing physical interfaces.

To configure flat-file accounting for ESSM business services:

1. Create a flat-file profile.

```
[edit accounting-options]  
user@host# edit flat-file-profile profile-name
```

2. (Optional) Configure the name of the XML schema for the accounting flat file.

```
[edit accounting-options flat-file-profile profile-name]  
user@host# set schema-version schema-name
```

3. Specify the filename for the accounting file.

```
[edit accounting-options flat-file-profile profile-name]  
user@host# set file accounting-filename
```

4. Specify the general, nonstatistical parameters for the accounting file that are displayed as part of the accounting record header.

```
[edit accounting-options flat-file-profile profile-name fields]
user@host# set general-param option
```



**BEST PRACTICE:** We recommend that you include the following general parameter fields in flat-file profiles for ESSM subscribers:

- General parameter fields: `accounting-type`, `descr`, `line-id`, `logical-interface`, `nas-port-id`, `timestamp`, and `vlan-id`

5. Specify the accounting statistics that are collected and recorded in the accounting file.

```
[edit accounting-options flat-file-profile profile-name fields]
user@host# set egress-stats option
user@host# set ingress-stats option;
user@host# set overall-packet option;
```



**BEST PRACTICE:** We recommend that you include the following statistics fields in flat-file profiles for core-facing physical interfaces:

- Egress statistics fields: `all-fields`
- Ingress statistics fields: `all-fields`
- Overall packet fields: `all-fields`

6. (Optional) Specify the format of the accounting file.

```
[edit accounting-options flat-file-profile profile-name]
user@host# set format (csv | ipdr)
```

7. (Optional) Specify the interval at which the Packet Forwarding Engine associated with the interface is polled for the statistics specified in the profile.

```
[edit accounting-options flat-file-profile profile-name]
user@host# set interval minutes
```



**NOTE:** When you do not configure this option, the polling interval is 15 minutes.

8. Configure the maximum size of the accounting file.

```
[edit accounting-options file filename]
user@host# set size bytes
```

9. Configure one or more archive sites for the files.

```
[edit accounting-options file filename]
user@host# set archive-sites site-name
```

The **site-name** is any valid FTP or Secure FTP URL. When the file is archived, the router attempts to transfer the file to the archive site. If you have specified more than one site, it tries the first site in the list. If that fails, it tries each site in turn until a transfer succeeds. The log file is stored at the archive site with a filename of the format **router-name\_log-filename\_timestamp**. The last site in a list is often a local directory, in case no remote site is reachable.

10. (Optional) Configure the start time for transferring the file.

```
[edit accounting-options file filename]  
user@host# set start-time YYYY-MM-DD.hh:mm
```

11. (Optional) Configure how frequently the file is transferred.

```
[edit accounting-options file filename]  
user@host# set transfer-interval minutes
```



**NOTE:** When you do not configure this option, the file is transferred every 30 minutes.

12. (Optional) Configure the maximum number of files (3 through 1000) to save.

```
[edit accounting-options file filename]  
user@host# set files number
```



**NOTE:** When you do not configure this option, a maximum of 10 files are saved.

13. (Optional) Configure the router to save a backup copy of the accounting file to the **/var/log/pfedBackup** directory if the normal transfer of the files to the archive sites fails. Specify whether only the current file from the master Routing Engine is saved or both that file and the file from the backup Routing Engine.

```
[edit accounting-options file filename]  
user@host# set backup-on-failure (master-and-slave | master-only)
```



**NOTE:** When you do not configure this option, the file is saved on failure into the local directory specified as the last site in the list of archive sites.

14. (Optional) Configure the accounting file to be compressed during transfer to an archive site.

```
[edit accounting-options file filename]  
user@host# set compress
```

15. (Optional) Configure the router's new backup Routing Engine to send its accounting file to the **/var/log/pfedBackup** directory on the new master Routing Engine when a change in mastership occurs.

```
[edit accounting-options file filename]  
user@host# set push-backup-to-master
```

16. (Optional) Configure the number of days after which accounting files are deleted from the local backup directory.

```
[edit accounting-options]  
user@host# set cleanup-interval days
```



**NOTE:** Files are retained for 1 day if you do not configure this option.

17. Assign the profile to an ESSM subscriber.

```
[edit system services extensible-subscriber-services]  
user@host# set flat-file-profile flat-file-profile-name
```

**Related  
Documentation**

- *Configuring Accounting-Data Log Files*
- [Configuring Flat-File Accounting for Layer 2 Wholesale on page 166](#)



## PART 5

# Configuration Statements and Operational Commands

- [Configuration Statements on page 177](#)
- [Operational Commands on page 387](#)



## CHAPTER 11

# Configuration Statements

- [accept](#) on page 182
- [accept-out-of-band](#) on page 183
- [access-profile](#) on page 184
- [access-profile \(Dynamic VLAN\)](#) on page 185
- [access-profile \(Dynamic Stacked VLAN\)](#) on page 186
- [accounting-server](#) on page 186
- [active-server-group](#) on page 187
- [address](#) on page 188
- [address-assignment \(Address-Assignment Pools\)](#) on page 191
- [adjacency-loss-hold-time \(ANCP\)](#) on page 192
- [ancp](#) on page 193
- [authentication](#) on page 195
- [authentication \(DHCP Local Server\)](#) on page 196
- [authentication \(DHCP Relay Agent\)](#) on page 197
- [authentication-order](#) on page 198
- [authentication-server](#) on page 199
- [auto-configure](#) on page 200
- [auto-configure-trigger interface \(ANCP\)](#) on page 201
- [backup-on-failure \(Accounting Options\)](#) on page 202
- [circuit-id \(VLAN Authentication Username\)](#) on page 203
- [cleanup-interval \(Accounting Options\)](#) on page 204
- [compress \(Accounting Options\)](#) on page 204
- [connectivity-type](#) on page 205
- [core-facing](#) on page 206
- [demux0 \(Dynamic Interface\)](#) on page 207
- [demux-options \(Dynamic Interface\)](#) on page 208
- [demux-source \(Dynamic IP Demux Interface\)](#) on page 209
- [demux-source \(Dynamic Underlying Interface\)](#) on page 210



- [demux-source \(Underlying Interface\)](#) on page 211
- [dhcp-attributes \(Address-Assignment Pools\)](#) on page 212
- [dhcp-local-server](#) on page 213
- [dhcp-relay](#) on page 219
- [dhcpv6 \(DHCP Local Server\)](#) on page 226
- [domain-name \(Address-Assignment Pools\)](#) on page 229
- [dynamic-profile \(DHCP Local Server\)](#) on page 230
- [dynamic-profile \(DHCP Relay Agent\)](#) on page 231
- [dynamic-profile \(Dynamic PPPoE\)](#) on page 232
- [dynamic-profile \(Stacked VLAN\)](#) on page 233
- [dynamic-profile \(VLAN\)](#) on page 234
- [dynamic-profiles](#) on page 235
- [egress-stats \(Flat-File Accounting Options\)](#) on page 243
- [encapsulation \(Dynamic Interfaces\)](#) on page 244
- [encapsulation \(Logical Interface\)](#) on page 247
- [encapsulation \(Physical Interface\)](#) on page 251
- [exclude \(RADIUS\)](#) on page 257
- [family](#) on page 262
- [family \(Address-Assignment Pools\)](#) on page 267
- [family \(Dynamic Demux Interface\)](#) on page 268
- [family \(Dynamic PPPoE\)](#) on page 269
- [family \(Dynamic Standard Interface\)](#) on page 270
- [fields \(Flat-File Accounting Options\)](#) on page 272
- [file \(Flat-File Accounting Options\)](#) on page 273
- [flat-file-profile \(Accounting Options\)](#) on page 274
- [flat-file-profile \(Extensible Subscriber Services\)](#) on page 276
- [flexible-vlan-tagging](#) on page 277
- [format \(Flat-File Accounting Options\)](#) on page 278
- [forwarding-options](#) on page 279
- [general-param \(Flat-File Accounting Options\)](#) on page 280
- [grace-period](#) on page 281
- [group \(DHCP Local Server\)](#) on page 282
- [group \(DHCP Relay Agent\)](#) on page 285
- [ingress-stats \(Flat-File Accounting Options\)](#) on page 288
- [inner-vlan-id \(Dynamic VLANs\)](#) on page 289
- [inner-vlan-id-swap-ranges](#) on page 290
- [input-vlan-map \(Dynamic Interfaces\)](#) on page 290

- [interface \(DHCP Local Server\)](#) on page 291
- [interface \(DHCP Relay Agent\)](#) on page 293
- [interface \(Dynamic Routing Instances\)](#) on page 294
- [interface \(Routing Instances\)](#) on page 295
- [interface-mac-limit \(VPLS\)](#) on page 296
- [interfaces](#) on page 297
- [interfaces \(Static and Dynamic Subscribers\)](#) on page 298
- [interval \(Flat-File Accounting Options\)](#) on page 302
- [instance-role](#) on page 303
- [instance-type](#) on page 304
- [ip-address-first](#) on page 306
- [keepalives \(Dynamic Profiles\)](#) on page 307
- [l2-stats \(Flat-File Accounting Options\)](#) on page 308
- [mac-validate \(Dynamic IP Demux Interface\)](#) on page 309
- [multicast-replication](#) on page 310
- [neighbor \(Define ANCP\)](#) on page 311
- [no-local-switching](#) on page 311
- [no-tunnel-services](#) on page 312
- [maximum-lease-time](#) on page 313
- [overall-packet \(Flat-File Accounting Options\)](#) on page 314
- [output-vlan-map \(Dynamic Interfaces\)](#) on page 315
- [pap \(Dynamic PPP\)](#) on page 316
- [pool \(Address-Assignment Pools\)](#) on page 317
- [pool-match-order](#) on page 318
- [pop \(Dynamic VLANs\)](#) on page 318
- [pppoe-options \(Dynamic PPPoE\)](#) on page 319
- [pppoe-underlying-options \(Static and Dynamic Subscribers\)](#) on page 320
- [ppp-options \(Dynamic PPP\)](#) on page 321
- [prefix \(Address-Assignment Pools\)](#) on page 322
- [profile \(Access\)](#) on page 323
- [protocols](#) on page 327
- [proxy-arp](#) on page 330
- [proxy-arp \(Dynamic Profiles\)](#) on page 331
- [push \(Dynamic VLANs\)](#) on page 331
- [push-backup-to-master \(Accounting Options\)](#) on page 332
- [radius \(Access Profile\)](#) on page 333
- [radius-server](#) on page 335

- [range \(Address-Assignment Pools\)](#) on page 336
- [ranges \(Dynamic VLAN\)](#) on page 337
- [remote-id \(VLAN Authentication Username\)](#) on page 337
- [route-distinguisher](#) on page 338
- [routing-instances \(Dynamic Profiles\)](#) on page 340
- [routing-instances \(Multiple Routing Entities\)](#) on page 342
- [schema-version \(Flat-File Accounting Options\)](#) on page 343
- [secret](#) on page 344
- [server \(Dynamic PPPoE\)](#) on page 344
- [server-group](#) on page 345
- [services \(System Services\)](#) on page 346
- [site \(VPLS Multihoming for FEC 128\)](#) on page 349
- [site-identifier \(VPLS\)](#) on page 349
- [site-range](#) on page 350
- [stacked-vlan-ranges](#) on page 351
- [stacked-vlan-tagging](#) on page 352
- [system](#) on page 352
- [traceoptions \(DHCP\)](#) on page 353
- [underlying-interface \(demux0\)](#) on page 355
- [underlying-interface \(Dynamic PPPoE\)](#) on page 356
- [unit](#) on page 357
- [unit \(Dynamic Demux Interface\)](#) on page 364
- [unit \(Dynamic Profiles Standard Interface\)](#) on page 366
- [unnumbered-address \(Dynamic PPPoE\)](#) on page 368
- [unnumbered-address \(Dynamic Profiles\)](#) on page 369
- [unnumbered-address \(Ethernet\)](#) on page 371
- [username-include](#) on page 372
- [user-prefix \(DHCP Local Server\)](#) on page 373
- [vlan-id \(Dynamic VLANs\)](#) on page 374
- [vlan-id \(VLAN ID to Be Bound to a Logical Interface\)](#) on page 375
- [vlan-model](#) on page 376
- [vlan-ranges](#) on page 377
- [vlan-tags](#) on page 378
- [vlan-tags \(Stacked VLAN Tags\)](#) on page 379
- [vpls \(Routing Instance\)](#) on page 381
- [vrf-export](#) on page 383


- [vrf-import on page 384](#)
- [vrf-target on page 385](#)

## accept

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

- 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](#)

## accept-out-of-band

<b>Syntax</b>	<code>accept-out-of-band protocol;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> auto-configure vlan-ranges <b>dynamic-profile</b> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Configure the protocol for which packets are accepted as out-of-band traffic to trigger instantiation or deletion of autosensed dynamic VLANs.
<div style="display: flex; align-items: center;">  <p><b>NOTE:</b> A given physical interface can support VLANs created by either conventional packet-triggering or out-of-band triggering, but not both at the same time.</p> </div>	
<b>Options</b>	<p><b>protocol</b>—Out-of-band protocol. The following out-of-band protocol is supported:</p> <ul style="list-style-type: none"> <li>• <b>ancp</b>—ANCP Port Up and Port Down messages trigger instantiation and deletion, respectively, of autosensed, dynamic VLAN.</li> </ul>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Out-of-Band ANCP Messages to Trigger Dynamic VLAN Instantiation on page 156</a></li> <li>• <a href="#">Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93</a></li> </ul>

## access-profile

<b>Syntax</b>	<code>access-profile <i>profile-name</i>;</code>
<b>Hierarchy Level</b>	<p>[edit],</p> <p>[edit forwarding-options <a href="#">dhcp-relay</a>]</p> <p>[edit forwarding-options <a href="#">dhcp-relay</a> dual-stack-group <i>dual-stack-group</i>],</p> <p>[edit forwarding-options <a href="#">dhcp-relay</a> group <i>group-name</i>]</p> <p>[edit forwarding-options <a href="#">dhcp-relay</a> dhcpv6]</p> <p>[edit forwarding-options <a href="#">dhcp-relay</a> dhcpv6 group <i>group-name</i>]</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>]</p> <p>[edit interfaces <i>interface-name</i> auto-configure <a href="#">vlan-ranges</a>],</p> <p>[edit interfaces <i>interface-name</i> auto-configure <a href="#">stacked-vlan-ranges</a>],</p> <p>[edit routing-instances <i>routing-instances-name</i>]</p> <p>[edit system services <a href="#">dhcp-local-server</a>]</p> <p>[edit system services <a href="#">dhcp-local-server</a> group <i>group-name</i>]</p> <p>[edit system services <a href="#">dhcp-local-server</a> <a href="#">dhcpv6</a>]</p> <p>[edit system services <a href="#">dhcp-local-server</a> <a href="#">dhcpv6</a> group <i>group-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
<b>Description</b>	<p>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 <b>[edit]</b> hierarchy level, or you can apply them to DHCP clients or subscribers, VLANs, or to a routing instance.</p>
<b>Options</b>	<p><b><i>profile-name</i></b>—Name of the access profile that you configured at the <b>[edit access profile name]</b> hierarchy level.</p>
<b>Required Privilege Level</b>	<p>access—To view this statement in the configuration.</p> <p>access-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Attaching Access Profiles</a></li> <li>• <a href="#">Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces</a></li> <li>• <a href="#">Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution on page 15</a></li> <li>• <a href="#">Configuring Access Components for the PPPoE Wholesale Network Solution on page 57</a></li> </ul>

## access-profile (Dynamic VLAN)

<b>Syntax</b>	<code>access-profile <i>vlan-access-profile-name</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> auto-configure vlan-ranges <b>dynamic-profile</b> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.2.
<b>Description</b>	<p>Access profiles contain subscriber access authentication, authorization and accounting configuration parameters. You can create an access profiles and then attach it at various configuration levels. When you attach an access profile to an interface configured for dynamic VLAN/SVLAN, all the VLANs/SVLANS use the same set of authentication, authorization and accounting parameters configured in that access profile. When an access profile is assigned at the new level, it takes precedence over any other access profile assignments.</p> <p>You can assign different access profiles to different dynamic profiles on the same interface. If you assign an access profile at the global level, but a different access profile is assigned at the interface level, the access profile at the interface level authenticates all dynamic VLANs and SVLANs on the interface.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs</i></li> <li>• <a href="#">show subscribers on page 562</a></li> <li>• <i>shmlog (Shared Memory Log)</i></li> </ul>

## access-profile (Dynamic Stacked VLAN)

---

<b>Syntax</b>	<code>access-profile <i>svlan-access-profile-name</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> auto-configure stacked-vlan-ranges dynamic-profile <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.2.
<b>Description</b>	<p>Access profiles contain subscriber access authentication, authorization and accounting configuration parameters. You can create an access profiles and then attach it at various configuration levels. When you attache an access profile to an interface configured for dynamic VLAN/SVLAN, all the VLANs/SVLANS use the same set of authentication, authorization and accounting parameters configured in that access profile. When an access profile is assigned at the new level, it takes precedence over any other access profile assignments.</p> <p>You can assign different access profiles to different dynamic profiles on the same interface. If you assign an access profile at the global level, but a different access profile is assigned at the interface level, the access profile at the interface level authenticates all dynamic VLANs and SVLANs on the interface.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs</i></li><li>• <a href="#">show subscribers on page 562</a></li><li>• <i>shmlog (Shared Memory Log)</i></li></ul>

## accounting-server

---

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

## active-server-group

<b>Syntax</b>	<code>active-server-group <i>server-group-name</i>;</code>
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay],          [edit forwarding-options dhcp-relay dhcpv6],          [edit forwarding-options dhcp-relay <b>group</b> <i>group-name</i>],          [edit forwarding-options dhcp-relay <b>group</b> <i>group-name</i> dhcpv6],          [edit logical-systems <i>logical-system-name</i> forwarding-options <b>dhcp-relay</b>],          [edit logical-systems <i>logical-system-name</i> forwarding-options <b>dhcp-relay</b> dhcpv6],          [edit logical-systems <i>logical-system-name</i> forwarding-options <b>group</b> <i>group-name</i>],          [edit logical-systems <i>logical-system-name</i> forwarding-options <b>group</b> <i>group-name</i> dhcpv6],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options <b>dhcp-relay</b>],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options <b>dhcp-relay</b> dhcpv6],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <b>group</b> <i>group-name</i>],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <b>group</b> <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 <b>dhcp-relay</b> dhcpv6],          [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <b>group</b> <i>group-name</i>],          [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 <b>group</b> <i>group-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.          Support at the [edit ... <b>dhcpv6</b>] hierarchy levels introduced in Junos OS Release 11.4.          Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
<b>Description</b>	<p>Apply a DHCP relay agent configuration to the named group of DHCP server addresses. Use the statement at the [edit ... <b>dhcpv6</b>] hierarchy levels to configure DHCPv6 support.</p> <p>A group-specific configuration overrides a global option.</p>
<b>Options</b>	<b><i>server-group-name</i></b> —Name of the group of DHCP or DHCPv6 server addresses to which the DHCP or DHCPv6 relay agent configuration applies.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.          interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Extended DHCP Relay Agent Overview</i></li> <li>• <i>Configuring Active Server Groups</i></li> <li>• <i>Configuring Group-Specific DHCP Relay Options</i></li> <li>• <b>dhcp-relay on page 219</b></li> </ul>

## address

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

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

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

**Description** Configure the interface address.

**Options** *address*—Address of the interface.

- 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, and 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 192.168.1.1/8;
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.1.1/8;
      }
    }
  }
}
```

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

- Related Documentation**
- *Configuring the Protocol Family*
  - *Junos OS System Basics Configuration Guide*
  - *family*
  - *negotiate-address*
  - [unnumbered-address \(Ethernet\) on page 371](#)

## 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;
      }
    }
    hold-down;
    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*

## adjacency-loss-hold-time (ANCP)

---

<b>Syntax</b>	adjacency-loss-hold-time <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit protocols <a href="#">ancp</a> ], [edit protocols ancp <a href="#">neighbor ip-address</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Configure the ANCP agent to monitor, either globally or for a specified neighbor, how long an ANCP adjacency is down and to trigger a state change for the subscriber access line if the hold timer expires before the adjacency comes back up, as indicated by a Port Up message on the access line. By default, there is no delay between detecting an adjacency loss and triggering the state change.
<b>Options</b>	<b>seconds</b> —Duration of period that the ANCP agent monitors loss of adjacency. <b>Default:</b> 0 seconds <b>Range:</b> 0 through 1800 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the ANCP Agent to Dampen the Effects of Short-Term Adjacency Losses on page 159</a></li><li>• <a href="#">Configuring ANCP Neighbors on page 153</a></li><li>• <i>Configuring the ANCP Agent</i></li><li>• <a href="#">Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93</a></li></ul>

## ancp

```

Syntax  ancp {
    adjacency-loss-hold-time seconds;
    adjacency-timer seconds;
    gsmp-syn-timeout seconds;
    gsmp-syn-wait;
    interfaces {
        interface-set interface-set-name {
            access-identifier identifier-string;
            underlying-interface underlying-interface-name;
        }
        interface-name {
            access-identifier identifier-string;
        }
    }
    maximum-discovery-table-entries entry-number;
    maximum-helper-restart-time;
    neighbor ip-address {
        adjacency-loss-hold-time seconds;
        adjacency-timer;
        auto-configure-trigger interface interface-name;
        ietf-mode;
        maximum-discovery-table-entries entry-number;
        pre-ietf-mode;
    }
    pre-ietf-mode;
    qos-adjust {
        other-bytes bytes;
        other-overhead-adjust percentage;
        sdsl-bytes bytes;
        sdsl-overhead-adjust percentage;
        vdsl-bytes bytes;
        vdsl-overhead-adjust percentage;
        vdsl2-bytes bytes;
        vdsl2-overhead-adjust percentage;
    }
    qos-adjust-adsl adjustment-factor;
    qos-adjust-adsl2 adjustment-factor;
    qos-adjust-adsl2-plus adjustment-factor;
    qos-adjust-other adjustment-factor;
    qos-adjust-sdsl adjustment-factor;
    qos-adjust-vdsl adjustment-factor;
    qos-adjust-vdsl2 adjustment-factor;
    traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size>
        <world-readable | no-world-readable>;
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}

```

Hierarchy Level [edit protocols]

**Release Information** Statement introduced in Junos OS Release 9.4.

**Description** Configure Junos OS ANCP agent features.

The remaining statements are explained separately.



**NOTE:** When you deactivate the ANCP protocol, the router does not perform a commit check to determine whether any ANCP or L2-BSA subscribers are present (active or inactive). Any subscribers that are active at the time of deactivation remain active.

---

**Required Privilege** routing—To view this statement in the configuration.

**Level** routing-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring the ANCP Agent*

## authentication

<b>Syntax</b>	<pre> authentication {   packet-types [<i>packet-types</i>];   password <i>password-string</i>;   username-include {     circuit-id;     circuit-type;     delimiter <i>delimiter-character</i>;     domain-name <i>domain-name-string</i>;     interface-name;     mac-address;     option-18;     option-37;     option-82 &lt;circuit-id&gt; &lt;remote-id&gt;;     radius-realm <i>radius-realm-string</i>;     remote-id;     user-prefix <i>user-prefix-string</i>;   } }</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> auto-configure <a href="#">vlan-ranges</a> ], [edit interfaces <i>interface-name</i> auto-configure <a href="#">stacked-vlan-ranges</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0.
<b>Description</b>	<p>Specify the authentication parameters that trigger the Access-Request message to AAA for the interface.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Subscribers over Static Interfaces</a></li> <li>• <a href="#">Configuring the Static Subscriber Global Authentication Password</a></li> <li>• <a href="#">Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs on page 155</a></li> <li>• <a href="#">Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93</a></li> </ul>

## authentication (DHCP Local Server)

---


<b>Syntax</b>	<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 &lt;circuit-id&gt; &lt;remote-id&gt;;     relay-agent-interface-id;     relay-agent-remote-id;     relay-agent-subscriber-id;     routing-instance-name;     user-prefix <i>user-prefix-string</i>;   } }</pre>
<b>Hierarchy Level</b>	<pre>[edit system services <a href="#">dhcp-local-server</a>], [edit system services dhcp-local-server <a href="#">dhcpv6</a>], [edit system services dhcp-local-server dhcpv6 <a href="#">group group-name</a>], [edit system services dhcp-local-server <a href="#">group group-name</a>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system   services <a href="#">dhcp-local-server</a> ...], [edit logical-systems <i>logical-system-name</i> system services <a href="#">dhcp-local-server</a> ...], [edit routing-instances <i>routing-instance-name</i> system services <a href="#">dhcp-local-server</a> ...]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>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.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Using External AAA Authentication Services with DHCP</i></li></ul>

## authentication (DHCP Relay Agent)

<b>Syntax</b>	<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 &lt;circuit-id&gt; &lt;remote-id&gt;;     relay-agent-interface-id;     relay-agent-remote-id;     relay-agent-subscriber-id;     routing-instance-name;     user-prefix <i>user-prefix-string</i>;   } }</pre>
<b>Hierarchy Level</b>	<pre> [edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit forwarding-options dhcp-relay dhcpv6 <b>group</b> <i>group-name</i>], [edit forwarding-options dhcp-relay dual-stack-group <i>dual-stack-group-name</i>], [edit forwarding-options dhcp-relay <b>group</b> <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>
<b>Release Information</b>	<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 <b>[edit ... dhcpv6]</b> hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Support at the <b>[edit ... dual-stack-group <i>dual-stack-group-name</i>]</b> hierarchy level introduced in Junos OS Release 15.1.</p>
<b>Description</b>	<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 <b>[edit...dhcpv6]</b> hierarchy levels to configure DHCPv6 support.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">dhcp-relay on page 219</a></li> <li>• <i>Using External AAA Authentication Services with DHCP</i></li> </ul>

## authentication-order

---

<b>Syntax</b>	<code>authentication-order [ <i>authentication-methods</i> ];</code>
<b>Hierarchy Level</b>	<code>[edit access <i>profile</i> <i>profile-name</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. <b>none</b> option added in Junos OS Release 11.2. <b>nasreq</b> option added in Junos OS Release 16.1.
<b>Description</b>	Set the order in which AAA tries different authentication methods when verifying that a client can access the router or switch. For each login attempt, AAA tries the authentication methods in order, from first to last.
<b>Options</b>	<b><i>authentication-methods</i></b> —Ordered list of methods to use for authentication attempts. The list includes one or more of the following methods in any combination: <ul style="list-style-type: none"><li>• <b>nasreq</b>—Verify the client using NASREQ authentication services.</li><li>• <b>none</b>—No authentication is performed. 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.</li><li>• <b>password</b>—Verify the client using the information configured at the <code>[edit access <i>profile</i> <i>profile-name</i> client <i>client-name</i>]</code> hierarchy level.</li><li>• <b>radius</b>—Verify the client using RADIUS authentication services.</li></ul>
<div> <b>NOTE:</b> Subscriber access management does not support the <b>password</b> option, and authentication fails when no method (<b>none</b>) is specified.</div>	
<b>Default:</b> <code>password</code>	
<b>Required Privilege Level</b>	<code>admin</code> —To view this statement in the configuration. <code>admin-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Configuring CHAP Authentication with RADIUS</i></li><li>• <i>Specifying the Authentication and Accounting Methods for Subscriber Access</i></li><li>• <i>Configuring Access Profiles for L2TP or PPP Parameters</i></li></ul>

---

## authentication-server

---

<b>Syntax</b>	authentication-server [ <i>ip-address</i> ];
<b>Hierarchy Level</b>	[edit access <a href="#">profile</a> <i>profile-name</i> <a href="#">radius</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1.
<b>Description</b>	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.
<b>Options</b>	<i>ip-address</i> —IPv4 address.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring RADIUS Server Parameters for Subscriber Access</i></li></ul>

## auto-configure

```
Syntax auto-configure {
  vlan-ranges {
    access-profile profile-name;
    authentication {
      packet-types [packet-types];
      password password-string;
      username-include {
        circuit-id;
        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;
        remote-id;
        user-prefix user-prefix-string;
      }
    }
    dynamic-profile profile-name {
      accept (any | dhcp-v4 | dhcp-v6 | inet | inet6 | pppoe);
      accept-out-of-band protocol;
      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;
    }

```

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

## auto-configure-trigger interface (ANCP)

<b>Syntax</b>	auto-configure-trigger interface <i>interface-name</i> ;
<b>Hierarchy Level</b>	[edit protocols ancp <a href="#">neighbor</a> <i>ip-address</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Map an ANCP neighbor to a subscriber-facing physical interface on the router, so that ANCP Port Up and Port Down messages trigger notifications to the auto-configuration daemon (autoconfd) to initiate VLAN creation (Port Up) or removal (Port Down).
<b>Options</b>	<i>interface-name</i> —Name of the physical interface.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring ANCP Neighbors on page 153</a></li> <li>• <i>Configuring the ANCP Agent</i></li> </ul>

## backup-on-failure (Accounting Options)

---

<b>Syntax</b>	backup-on-failure (master-and-slave   master-only);
<b>Hierarchy Level</b>	[edit accounting-options <i>file filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Configure the router to save a copy of the accounting file locally, to the <b>/var/log/pfedBackup</b> directory of the relevant Routing Engine, in the event that file transfer to the remote archive sites cannot be completed.
<b>Options</b>	<b>master-and-slave</b> —Back up accounting files from both the master Routing Engine and the backup Routing Engine.  <b>master-only</b> —Back up accounting files from only the master Routing Engine.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Accounting-Data Log Files</a></li><li>• <a href="#">Configuring Flat-File Accounting for Layer 2 Wholesale on page 166</a></li><li>• <a href="#">Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170</a></li><li>• <a href="#">Flat-File Accounting Overview on page 163</a></li></ul>

## circuit-id (VLAN Authentication Username)

<b>Syntax</b>	circuit-id;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> auto-configure vlan-ranges authentication <b>username-include</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Include the agent circuit identifier (ACI) in the username sent to RADIUS for authentication of the dynamic VLAN. The ACI is conveyed by the Access-Loop-Circuit-ID TLV in an out-of-band ANCP Port Up message.



**NOTE:** This statement is not supported for stacked VLANs.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs on page 155</a></li> <li>• <a href="#">Configuring VLAN Interface Username Information for AAA Authentication</a></li> <li>• <a href="#">Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93</a></li> </ul>

## cleanup-interval (Accounting Options)

---

<b>Syntax</b>	cleanup-interval <i>days</i> ;
<b>Hierarchy Level</b>	[edit accounting-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Configure the interval to delete files from the local backup directory.
<b>Options</b>	<b>days</b> —Number of days after which accounting-options files are to be deleted from the backup directory. <b>Range:</b> 1 through 31 days <b>Default:</b> 1 day
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Accounting-Data Log Files</a></li><li>• <a href="#">Configuring Flat-File Accounting for Layer 2 Wholesale on page 166</a></li><li>• <a href="#">Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170</a></li><li>• <a href="#">Flat-File Accounting Overview on page 163</a></li></ul>

## compress (Accounting Options)

---

<b>Syntax</b>	compress;
<b>Hierarchy Level</b>	[edit accounting-options file <i>filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Compress the accounting file during file transfer to the backup site.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Accounting-Data Log Files</a></li><li>• <a href="#">Configuring Flat-File Accounting for Layer 2 Wholesale on page 166</a></li><li>• <a href="#">Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170</a></li><li>• <a href="#">Flat-File Accounting Overview on page 163</a></li></ul>

## connectivity-type

<b>Syntax</b>	<code>connectivity-type (ce   irb   permanent);</code>
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. <b>irb</b> option introduced in Junos OS Release 9.3. <b>permanent</b> option introduced in Junos OS Release 10.4.
<b>Description</b>	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).



**NOTE:** The **connectivity-type** statement is not supported for FEC 129 VPLS (also known as LDP VPLS with BGP-based autodiscovery).

<b>Default</b>	<b>ce</b>
<b>Options</b>	<p><b>ce</b>—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><b>irb</b>—Allow a VPLS connection to remain up so long as an IRB interface is configured for the VPLS routing instance.</p> <p><b>permanent</b>—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>



**NOTE:** To specifically take down a VPLS routing instance that is using the **permanent** option, all associated static logical interfaces must also be down.

<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Configuring VPLS Routing Instances</li> <li>Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers on page 84</li> </ul>

## core-facing

---

<b>Syntax</b>	core-facing;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>family</b> <i>family</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>family</b> <i>family</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2.
<b>Description</b>	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.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Junos OS Broadband Subscriber Management and Services Library</i></li></ul>

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

## demux-source (Dynamic IP Demux Interface)

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

## demux-source (Dynamic Underlying Interface)

---

<b>Syntax</b>	<code>demux-source <i>family</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">dynamic-profiles</a> <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6.
<b>Description</b>	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.

---

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

## demux-source (Underlying Interface)

<b>Syntax</b>	<code>demux-source <i>family</i>;</code>
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0. Support for aggregated Ethernet added in Junos OS Release 9.4.
<b>Description</b>	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.

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

## dhcp-attributes (Address-Assignment Pools)

---

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[edit access address-assignment <b>pool</b> <i>pool-name</i> <b>family</b> <i>family</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.3 for EX Series switches.
<b>Description</b>	Configure address pools that can be used by different client applications.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Address-Assignment Pools Overview</i></li><li>• <i>DHCP Attributes for Address-Assignment Pools</i></li><li>• <i>Configuring Address-Assignment Pools</i></li><li>• <i>Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address</i></li></ul>

## 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 {
            ...
        }
        duplicate-clients incoming-interface;
        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;
  server-duid-type type;
  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 **dhcpcv6** 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;
        }
    }
}
dual-stack-group dual-stack-group-name {
    access-profile profile-name;
    authentication {
        ... authentication-configuration
    }
    dynamic-profile profile-name {
        ... dynamic-profile-configuration
    }
    relay-agent-interface-id {
        ... relay-agent-interface-id-configuration
    }
    relay-agent-remote-id {
        ... relay-agent-remote-id-configuration
    }
    service-profile dynamic-profile-name;
}
```

```
}
duplicate-clients incoming-interface;
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;
  dual-stack dual-stack-group-name;
  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;
dual-stack dual-stack-group-name;
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-id-override
}
}
server-group {
  server-group-name {
    server-ip-address;
  }
}
route-suppression:
server-response-time seconds;
service-profile dynamic-profile-name;
}
```

<b>Hierarchy Level</b>	<p>[edit forwarding-options],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	<p>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.</p> <p>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.</p> <p>The extended DHCP and DHCPv6 relay agent options configured with the <b>dhcp-relay</b> and <b>dhcpv6</b> statements are incompatible with the DHCP/BOOTP relay agent options configured with the <b>bootp</b> 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.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Extended DHCP Relay Agent Overview</i></li> <li>• <i>DHCPv6 Relay Agent Overview</i></li> <li>• <i>DHCP Relay Proxy Overview</i></li> <li>• <i>Using External AAA Authentication Services with DHCP</i></li> </ul>

## dhcipv6 (DHCP Local Server)

---

```
Syntax  dhcipv6 {
    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;
        }
    }
    duplicate-clients incoming-interface;
    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 {
        client-negotiation-match incoming-interface;
        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 {
    client-negotiation-match incoming-interface;
    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 {
    client-negotiation-match incoming-interface;
    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;
server-duid-type type;
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)

---

<b>Syntax</b>	<code>domain-name <i>domain-name</i>;</code>
<b>Hierarchy Level</b>	[edit access address-assignment pool <i>pool-name</i> family inet <a href="#">dhcp-attributes</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0.
<b>Description</b>	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.
<b>Options</b>	<i>domain-name</i> —Name of the domain.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Address-Assignment Pools</i></li></ul>

## dynamic-profile (DHCP Local Server)


---

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

## dynamic-profile (DHCP Relay Agent)

<b>Syntax</b>	<pre>dynamic-profile <i>profile-name</i> {     aggregate-clients (merge   replace);     use-primary <i>primary-profile-name</i>; }</pre>
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay],          [edit forwarding-options <b>dhcp-relay</b> dhcpv6],          [edit forwarding-options dhcp-relay dhcpv6 <b>group</b> <i>group-name</i>],          [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <b>interface</b> <i>interface-name</i>],          [edit forwarding-options dhcp-relay dual-stack-group <i>dual-stack-group-name</i>],          [edit forwarding-options dhcp-relay <b>group</b> <i>group-name</i>],          [edit forwarding-options dhcp-relay group <i>group-name</i> <b>interface</b> <i>interface-name</i>],          [edit logical-systems <i>logical-system-name</i> forwarding-options <b>dhcp-relay</b> ...],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>          forwarding-options <b>dhcp-relay</b> ...],          [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.2.          Support at the <b>[edit ... dhcpv6]</b> hierarchy levels introduced in Junos OS Release 11.4.          Statement introduced in Junos OS Release 12.1 for EX Series switches.          Support at the <b>[edit ... dual-stack-group dual-stack-group-name]</b> hierarchy level introduced in Junos OS Release 15.1.</p>
<b>Description</b>	<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>
<b>Options</b>	<p><b><i>profile-name</i></b>—Name of the dynamic profile.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.          interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">dhcp-relay on page 219</a></li> <li>• <i>Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces</i></li> <li>• <i>Grouping Interfaces with Common DHCP Configurations</i></li> <li>• <i>Configuring a Default Subscriber Service</i></li> </ul>

## dynamic-profile (Dynamic PPPoE)

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

## dynamic-profile (Stacked VLAN)

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

## dynamic-profile (VLAN)

---

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

## 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;
    }
    only-at-create;
  }
  filter filter-name {
    interface-specific;
    term term-name {
      from {
        match-conditions;
      }
      then {
        action;
        action-modifiers;
      }
    }
  }
}
policer policer-name {
  filter-specific;
  if-exceeding {
    (bandwidth-limit bps | bandwidth-percent percentage);
    burst-size-limit bytes;
  }
  logical-bandwidth-policer;
  logical-interface-policer;
  physical-interface-policer;
  then {
    policer-action;
  }
}
}
hierarchical-policer 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 {...}
}
policy-options {
    prefix-list uid {
        ip-addresses;
        dynamic-db;
    }
}
predefined-variable-defaults predefined-variable <variable-option> default-value;
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 Level** routing—To view this statement in the configuration.  
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*

## egress-stats (Flat-File Accounting Options)

<b>Syntax</b>	<pre>egress-stats {   all-fields;   input-bytes;   input-packets;   output-bytes;   output-packets;   queue-id;   red-drop-bytes;   red-drop-packets;   tail-drop-packets; }</pre>
<b>Hierarchy Level</b>	[edit accounting-options flat-file-profile <i>profile-name</i> <b>fields</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Specify egress queue statistics to be collected for the interface.
<b>Options</b>	<p><b>all-fields</b>—Collect all egress queue statistics available for the interface context, logical or physical.</p> <p><b>input-bytes</b>—Collect the number of octets queued including traffic dropped because of congestion.</p> <p><b>input-packets</b>—Collect the number of packets queued including traffic dropped because of congestion.</p> <p><b>output-bytes</b>—Collect the number of octets transmitted by the egress queue.</p> <p><b>output-packets</b>—Collect the number of packets transmitted by the egress queue.</p> <p><b>queue-id</b>—Collect the logical identifier for the egress queue; identifies the traffic class.</p> <p><b>red-drop-bytes</b>—Collect the number of octets dropped on the egress queue because of random early detection.</p> <p><b>red-drop-packets</b>—Collect the number of packets dropped on the egress queue because of random early detection.</p> <p><b>tail-drop-packets</b>—Collect the number of packets dropped in the egress queue because of tail drop.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Flat-File Accounting for Layer 2 Wholesale on page 166</a></li> <li>• <a href="#">Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170</a></li> <li>• <a href="#">Flat-File Accounting Overview on page 163</a></li> </ul>

## encapsulation (Dynamic Interfaces)

---

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[edit <a href="#">dynamic-profiles</a> <i>profile-name</i> <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	Dynamic interface configuration of the logical link-layer encapsulation type.
<b>Options</b>	<p><b>atm-ccc-cell-relay</b>—Use ATM cell-relay encapsulation.</p> <p><b>atm-ccc-vc-mux</b>—Use ATM virtual circuit (VC) multiplex encapsulation on circuit cross-connect (CCC) circuits. When you use this encapsulation type, you can configure the <b>ccc</b> family only.</p> <p><b>atm-cisco-nlpid</b>—Use Cisco ATM network layer protocol ID (NLPID) encapsulation. When you use this encapsulation type, you can configure the <b>inet</b> family only.</p> <p><b>atm-mlppp-llc</b>—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><b>atm-nlpid</b>—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the <b>inet</b> family only.</p> <p><b>atm-ppp-llc</b>—For ATM2 IQ interfaces only, use PPP over AAL5 LLC encapsulation.</p> <p><b>atm-ppp-vc-mux</b>—For ATM2 IQ interfaces only, use PPP over ATM AAL5 multiplex encapsulation.</p> <p><b>atm-snap</b>—Use ATM subnetwork attachment point (SNAP) encapsulation.</p> <p><b>atm-tcc-snap</b>—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.</p> <p><b>atm-tcc-vc-mux</b>—Use ATM VC multiplex encapsulation on TCC circuits. When you use this encapsulation type, you can configure the <b>tcc</b> family only.</p> <p><b>atm-vc-mux</b>—Use ATM VC multiplex encapsulation. When you use this encapsulation type, you can configure the <b>inet</b> family only.</p> <p><b>ether-over-atm-llc</b>—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.

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

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

## encapsulation (Logical Interface)

<b>Syntax</b>	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);
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ], [edit interfaces rlsq <i>number</i> <b>unit</b> <i>logical-unit-number</i> ] [edit protocols evpn]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers ( <b>ethernet</b> , <b>vlan-ccc</b> , and <b>vlan-tcc</b> options only). Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Access Routers. Only the <b>atm-ccc-cell-relay</b> and <b>atm-ccc-vc-mux</b> options are supported on ACX Series routers.
<b>Description</b>	Configure a logical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.
<b>Options</b>	<p><b>atm-ccc-cell-relay</b>—Use ATM cell-relay encapsulation.</p> <p><b>atm-ccc-vc-mux</b>—Use ATM virtual circuit (VC) multiplex encapsulation on CCC circuits. When you use this encapsulation type, you can configure the <b>ccc</b> family only.</p> <p><b>atm-cisco-nlpid</b>—Use Cisco ATM network layer protocol identifier (NLPID) encapsulation. When you use this encapsulation type, you can configure the <b>inet</b> family only.</p> <p><b>atm-mlppp-llc</b>—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><b>atm-nlpid</b>—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the <b>inet</b> family only.</p> <p><b>atm-ppp-llc</b>—(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><b>atm-ppp-vc-mux</b>—(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><b>atm-snap</b>—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM subnetwork attachment point (SNAP) encapsulation.</p>

**atm-tcc-snap**—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.

**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 router over a time-division multiplexing (TDM) link. This encapsulation type enables the PE router 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**—Use PPP over Ethernet encapsulation to configure an underlying Ethernet interface for a dynamic PPPoE logical interface on M120 and M320 routers with Intelligent Queuing 2 (IQ2) PICs, and on MX Series routers with MPCs.

**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.

**vxlan**—Use VXLAN data plane encapsulation for EVPN.

**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)

<b>Syntax</b>	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);
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> ], [edit interfaces rlsq <i>number:number</i> ]
<b>Release Information</b>	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 ( <b>flexible-ethernet-services</b> , <b>ethernet-ccc</b> , and <b>ethernet-tcc</b> options only).
<b>Description</b>	Specify the physical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.
<b>Default</b>	<b>ppp</b> —Use serial PPP encapsulation.

### Options



**NOTE:** Frame Relay, ATM, PPP, SONET, and SATSOP options are not supported on the EX Series switches.

**atm-ccc-cell-relay**—Use ATM cell-relay encapsulation.

**atm-pvc**—Defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*. When you configure physical ATM interfaces with ATM PVC encapsulation, an RFC 2684-compliant ATM Adaptation Layer 5 (AAL5) tunnel is set up to route the ATM cells over a Multiprotocol Label Switching (MPLS) path that is typically established between two MPLS-capable routers using the Label Distribution Protocol (LDP).

**cisco-hdlc**—Use Cisco-compatible High-Level Data Link Control (HDLC) framing. E1, E3, SONET/SDH, T1, and T3 interfaces can use Cisco HDLC encapsulation. Two related versions are supported:

- CCC version (**cisco-hdlc-ccc**)—The logical interface does not require an encapsulation statement. When you use this encapsulation type, you can configure the **ccc** family only.
- TCC version (**cisco-hdlc-tcc**)—Similar to CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.

**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 BPDUs 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. When you use this encapsulation type, you can configure the **ccc** family only.

**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. Extended VLAN CCC encapsulation supports TPIDs 0x8100, 0x9100, and 0x9901. When you use this encapsulation type, you can configure the **ccc** family only. 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 is defined in RFC 1490, *Multiprotocol Interconnect over Frame Relay*. E1, E3, link services, SONET/SDH, T1, T3, and voice services interfaces can use Frame Relay encapsulation.

**frame-relay-ccc**—Use Frame Relay encapsulation on CCC circuits. This encapsulation is same as standard Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to CCC. The logical interface must also have **frame-relay-ccc** encapsulation. 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 the Cisco Frame Relay. IETF frame relay encapsulation identifies the payload format using NLPID and SNAP formats. Cisco-compatible Frame Relay encapsulation uses the Ethernet type to identify the type of payload.



**NOTE:** When the encapsulation type is set to Cisco-compatible Frame Relay encapsulation, ensure that the LMI type is set to ANSI or Q933-A.

---

**frame-relay-ether-type-tcc**—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media. This encapsulation is Cisco-compatible Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to TCC.

**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. The connection between the two CE routers can be either user-to-network interface (UNI) or network-to-network interface (NNI); this is completely transparent to the PE routers. When you use this encapsulation type, you can configure the **ccc** family only.

**frame-relay-tcc**—This encapsulation is similar to Frame Relay CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.

**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. This encapsulation is defined in RFC 1661, *The Point-to-Point Protocol (PPP) for the Transmission of Multiprotocol Datagrams over Point-to-Point Links*. PPP is the default encapsulation type for physical interfaces. E1, E3, SONET/SDH, T1, and T3 interfaces can use 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 CCC encapsulation supports TPID 0x8100 only. 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. 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.

<b>Required Privilege Level</b>	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*
- *Configuring VLAN and 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*

## exclude (RADIUS)

```
Syntax  exclude {
    acc-aggr-cir-id-asc [ access-request | accounting-start | accounting-stop ];
    acc-aggr-cir-id-bin [ access-request | accounting-start | accounting-stop ];
    acc-loop-cir-id [ access-request | accounting-start | accounting-stop ];
    acc-loop-encap [ access-request | accounting-on | accounting-off | accounting-start |
        accounting-stop ];
    acc-loop-remote-id [ access-request | accounting-on | accounting-off | accounting-start
        | accounting-stop ];
    accounting-authentic [ accounting-on | accounting-off ];
    accounting-delay-time [ accounting-on | accounting-off ];
    accounting-session-id [ access-request | accounting-on | accounting-off | accounting-stop
        ];
    accounting-terminate-cause [ accounting-off ];
    acct-tunnel-connection [ accounting-start | accounting-stop ];
    act-data-rate-dn [ access-request | accounting-start | accounting-stop ];
    act-data-rate-up [ access-request | accounting-start | accounting-stop ];
    act-interlv-delay-dn [ access-request | accounting-start | accounting-stop ];
    act-interlv-delay-up [ access-request | accounting-start | accounting-stop ];
    att-data-rate-dn [ access-request | accounting-start | accounting-stop ];
    att-data-rate-up [ access-request | accounting-start | accounting-stop ];
    called-station-id [ access-request | accounting-start | accounting-stop ];
    calling-station-id [ access-request | accounting-start | accounting-stop ];
    chap-challenge [ access-request ];
    chargeable-user-identity [ access-request ];
    class [ accounting-start | accounting-stop ];
    cos-shaping-rate [ accounting-start | accounting-stop ];
    dhcp-gi-address [ access-request | accounting-start | accounting-stop ];
    dhcp-mac-address [ access-request | accounting-start | accounting-stop ];
    dhcp-options [ access-request | accounting-start | accounting-stop ];
    downstream-calculated-qos-rate [ access-request | accounting-start | accounting-stop
        ];
    dsl-forum-attributes [ access-request | accounting-start | accounting-stop ];
    dsl-line-state [ access-request | accounting-start | accounting-stop ];
    dsl-type [ access-request | accounting-start | accounting-stop ];
    event-timestamp [ accounting-on | accounting-off | accounting-start | accounting-stop
        ];
    filter-id [ accounting-start | accounting-stop ];
    first-relay-ipv4-address [ access-request | accounting-start | accounting-stop ];
    first-relay-ipv6-address [ access-request | accounting-start | accounting-stop ];
    framed-ip-address [ accounting-start | accounting-stop ];
    framed-ip-netmask [ accounting-start | accounting-stop ];
    input-filter [ accounting-start | accounting-stop ];
    input-gigapackets [ accounting-stop ];
    input-gigawords [ accounting-stop ];
    interface-description [ access-request | accounting-start | accounting-stop ];
    l2tp-rx-connect-speed [ access-request | accounting-start | accounting-stop ];
    l2tp-tx-connect-speed [ access-request | accounting-start | accounting-stop ];
    max-data-rate-dn [ access-request | accounting-start | accounting-stop ];
    max-data-rate-up [ access-request | accounting-start | accounting-stop ];
    max-interlv-delay-dn [ access-request | accounting-start | accounting-stop ];
    max-interlv-delay-up [ access-request | accounting-start | accounting-stop ];
    min-data-rate-dn [ access-request | accounting-start | accounting-stop ];
}
```

```

min-data-rate-up [ access-request | accounting-start | accounting-stop ];
min-lp-data-rate-dn [ access-request | accounting-start | accounting-stop ];
min-lp-data-rate-up [ access-request | accounting-start | accounting-stop ];
nas-identifier [ access-request | accounting-on | accounting-off | accounting-start |
  accounting-stop ];
nas-port [ access-request | accounting-start | accounting-stop ];
nas-port-id [ access-request | accounting-start | accounting-stop ];
nas-port-type [ access-request | accounting-start | accounting-stop ];
output-filter [ accounting-start | accounting-stop ];
output-gigapackets [ accounting-stop ];
output-gigawords [ accounting-stop ];
pppoe-description [ access-request | accounting-start | accounting-stop ];
tunnel-assignment-id [ accounting-start | accounting-stop ];
tunnel-client-auth-id [ accounting-start | accounting-stop ];
tunnel-client-endpoint [ accounting-start | accounting-stop ];
tunnel-medium-type [ accounting-start | accounting-stop ];
tunnel-server-auth-id [ accounting-start | accounting-stop ];
tunnel-server-endpoint [ accounting-start | accounting-stop ];
tunnel-type [ accounting-start | accounting-stop ];
upstream-calculated-qos-rate [ access-request | accounting-start | accounting-stop ];
virtual-router [ access-request | accounting-start | accounting-stop ]
}

```

**Hierarchy Level** [edit access profile *profile-name* radius attributes]

**Release Information** Statement introduced in Junos OS Release 9.1.  
Statement introduced in Junos OS Release 9.1 for EX Series switches.  
Options **downstream-calculated-qos-rate**, **dsl-forum-attributes**, and **upstream-calculated-qos-rate** added in Junos OS Release 11.4.  
Options **cos-shaping-rate** and **filter-id** added in Junos OS Release 13.2.  
Option **pppoe-description** added in Junos OS Release 14.2.  
Option **virtual-router** added in Junos OS Release 15.1.  
Options **first-relay-ipv4-address** and **first-relay-ipv6-address** added in Junos OS Release 16.1.

**Description** Configure the router or switch to exclude the specified attributes from the specified type of RADIUS message.

Not all attributes are available in all types of RADIUS messages. By default, the router or switch includes the specified attributes in RADIUS Access-Request, Acct-On, Acct-Off, Acct-Start, and Acct-Stop messages.



**NOTE:** If you exclude an attribute from Acct-Off messages, the attributes are then excluded from Interim-Acct messages.

**Options** RADIUS attribute type—RADIUS attribute, Juniper Networks (vendor ID 4874) VSA number and name, or DSL Forum (vendor ID 3561) VSA number and name:

- **acc-aggr-cir-id-asc**—Juniper Networks VSA 26-112, Acc-Aggr-Cir-Id-Asc.
- **acc-aggr-cir-id-bin**—Juniper Networks VSA 26-111, Acc-Aggr-Cir-Id-Bin.
- **acc-loop-cir-id**—Juniper Networks VSA 26-110, Acc-Loop-Cir-Id.
- **acc-loop-encap**—Juniper Networks VSA 26-183, Acc-Loop-Encap.
- **acc-loop-remote-id**—Juniper Networks VSA 26-182, Acc-Loop-Remote-Id.
- **accounting-authentic**—RADIUS attribute 45, Acct-Authentic.
- **accounting-delay-time**—RADIUS attribute 41, Acct-Delay-Time.
- **accounting-session-id**—RADIUS attribute 44, Acct-Session-Id.
- **accounting-terminate-cause**—RADIUS attribute 49, Acct-Terminate-Cause.
- **acct-tunnel-connection**—RADIUS attribute 68, Acct-Tunnel-Connection.
- **act-data-rate-dn**—Juniper Networks VSA 26-114, Act-Data-Rate-Dn
- **act-data-rate-up**—Juniper Networks VSA 26-113, Act-Data-Rate-Up
- **act-interlv-delay-dn**—Juniper Networks VSA 26-126, Act-Interlv-Delay-Dn
- **act-interlv-delay-up**—Juniper Networks VSA 26-124, Act-Interlv-Delay-Up
- **att-data-rate-dn**—Juniper Networks VSA 26-118, Att-Data-Rate-Dn
- **att-data-rate-up**—Juniper Networks VSA 26-117, Att-Data-Rate-Up
- **called-station-id**—RADIUS attribute 30, Called-Station-Id.
- **calling-station-id**—RADIUS attribute 31, Calling-Station-Id.
- **chargeable-user-identity**—RADIUS attribute 89, Chargeable-User-Identity.
- **class**—RADIUS attribute 25, Class.
- **cos-shaping-rate**—Juniper Networks VSA 26-177, Cos-Shaping-Rate.
- **dhcp-gi-address**—Juniper Networks VSA 26-57, DHCP-GI-Address.
- **dhcp-mac-address**—Juniper Networks VSA 26-56, DHCP-MAC-Address.
- **dhcp-options**—Juniper Networks VSA 26-55, DHCP-Options.
- **downstream-calculated-qos-rate**—Juniper Networks VSA 26-141
- **dsl-forum-attributes**—DSL Forum VSA (vendor ID 3561) as described in RFC 4679, *DSL Forum Vendor-Specific RADIUS Attributes*
- **dsl-line-state**—Juniper Networks VSA 26-127, DSL-Line-State
- **dsl-type**—Juniper Networks VSA 26-128, DSL-Type
- **event-timestamp**—RADIUS attribute 55, Event-Timestamp.
- **filter-id**—RADIUS attribute 11, Filter-Id.

- **first-relay-ipv4-address**—Juniper Networks VSA 26-187, DHCP-First-Relay-IPv4-Address.
- **first-relay-ipv6-address**—Juniper Networks VSA 26-188, DHCP-First-Relay-IPv6-Address.
- **framed-ip-address**—RADIUS attribute 8, Framed-IP-Address.
- **framed-ip-netmask**—RADIUS attribute 9, Framed-IP-Netmask.
- **input-filter**—Juniper Networks VSA 26-10, Ingress-Policy-Name.
- **input-gigapackets**—Juniper Networks VSA 26-42, Acct-Input-Gigapackets.
- **input-gigawords**—RADIUS attribute 52, Acct-Input-Gigawords.
- **interface-description**—Juniper Networks VSA 26-53, Interface-Desc.
- **l2tp-rx-connect-speed**—Juniper Networks VSA 26-163, Rx-Connect-Speed
- **l2tp-tx-connect-speed**—Juniper Networks VSA 26-162, Tx-Connect-Speed
- **max-data-rate-dn**—Juniper Networks VSA 26-120, Max-Data-Rate-Dn
- **max-data-rate-up**—Juniper Networks VSA 26-119, Max-Data-Rate-Up
- **max-interlv-delay-dn**—Juniper Networks VSA 26-125, Max-Interlv-Delay-Dn
- **max-interlv-delay-up**—Juniper Networks VSA 26-123, Max-Interlv-Delay-Up
- **min-data-rate-dn**—Juniper Networks VSA 26-116, Min-Data-Rate-Dn
- **min-data-rate-up**—Juniper Networks VSA 26-115, Min-Data-Rate-Up
- **min-lp-data-rate-dn**—Juniper Networks VSA 26-122, Min-Lp-Data-Rate-Dn
- **min-lp-data-rate-up**—Juniper Networks VSA 26-121, Min-Lp-Data-Rate-Up
- **nas-identifier**—RADIUS attribute 32, NAS-Identifier.
- **nas-port**—RADIUS attribute 5, NAS-Port.
- **nas-port-id**—RADIUS attribute 87, NAS-Port-Id.
- **nas-port-type**—RADIUS attribute 61, NAS-Port-Type.
- **output-filter**—Juniper Networks VSA 26-11, Egress-Policy-Name.
- **output-gigapackets**—Juniper Networks VSA 26-43, Acct-Output-Gigapackets.
- **output-gigawords**—RADIUS attribute 53, Acct-Output-Gigawords.
- **pppoe-description**—Juniper Networks VSA 26-24, PPPoE-Description.
- **tunnel-assignment-id**—RADIUS attribute 82, Tunnel-Assignment-ID.
- **tunnel-client-auth-id**—RADIUS attribute 90, Tunnel-Client-Auth-ID.
- **tunnel-client-endpoint**—RADIUS attribute 66, Tunnel-Client-Endpoint.
- **tunnel-medium-type**—RADIUS attribute 65, Tunnel-Medium-Type.
- **tunnel-server-auth-id**—RADIUS attribute 91, Tunnel-Server-Auth-ID.
- **tunnel-server-endpoint**—RADIUS attribute 67, Tunnel-Server-Endpoint.

- **tunnel-type**—RADIUS attribute 64, Tunnel-Type.
- **upstream-calculated-qos-rate**—Juniper Networks VSA 26-142
- **virtual-router**—Juniper Networks VSA 26-1

RADIUS message type:

- **access-request**—RADIUS Access-Request messages.
- **accounting-off**—RADIUS Accounting-Off messages.
- **accounting-on**—RADIUS Accounting-On messages.
- **accounting-start**—RADIUS Accounting-Start messages.
- **accounting-stop**—RADIUS Accounting-Stop messages.

<b>Required Privilege Level</b>	admin—To view this statement in the configuration.
	admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	• <i>Configuring How RADIUS Attributes Are Used for Subscriber Access</i>
	• <i>Configuring RADIUS Server Parameters for Subscriber Access</i>

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



**NOTE:** Not all subordinate statements 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. You can optionally configure this protocol family for the logical interface on which you configure VPLS.
- **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. You can configure this protocol family for the logical interface of CCC physical interfaces. When you use this encapsulation type, you can configure the **ccc** family only.
- **inet**—Internet Protocol version 4 suite. You must configure this protocol family for the logical interface to support IP protocol traffic, including Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), Internet Control Message Protocol (ICMP), and Internet Protocol Control Protocol (IPCP).
- **inet6**—Internet Protocol version 6 suite. You must configure this protocol family for the logical interface to support IPv6 protocol traffic, including Routing Information Protocol for IPv6 (RIPng), Intermediate System-to-Intermediate System (IS-IS), BGP, and Virtual Router Redundancy Protocol for IPv6 (VRRP).
- **iso**—International Organization for Standardization Open Systems Interconnection (ISO OSI) protocol suite. You must configure this protocol family for the logical interface to support IS-IS traffic.
- **mlfr-end-to-end**—Multilink Frame Relay FRF.15. You must configure this protocol or multilink Point-to-Point Protocol (MLPPP) for the logical interface to support multilink bundling.
- **mlfr-uni-nni**—Multilink Frame Relay FRF.16. You must configure this protocol or **mlfr-end-to-end** for the logical interface to support link services and voice services bundling.
- **multilink-ppp**—Multilink Point-to-Point Protocol. You must configure this protocol (or **mlfr-end-to-end**) for the logical interface to support multilink bundling.
- **mpls**—Multiprotocol Label Switching (MPLS). You must configure this protocol family for the logical interface to participate in an MPLS path.
- **pppoe**—Point-to-Point Protocol over Ethernet
- **tcc**—Translational cross-connect protocol suite. You can configure this protocol family for the logical interface of TCC physical interfaces.

- **tnp**—Trivial Network Protocol. This protocol is used to communicate between the Routing Engine and the router's packet forwarding components. The Junos OS automatically configures this protocol family on the router's internal interfaces only, as discussed in *Understanding Internal Ethernet Interfaces*.
- **vpls**—(M Series and T Series routers only) Virtual private LAN service. You can optionally configure this protocol family for the logical interface on which you configure VPLS. VPLS provides an Ethernet-based point-to-multipoint Layer 2 VPN to connect customer edge (CE) routers across an MPLS backbone. When you configure a VPLS encapsulation type, the **family vpls** statement is assumed by default.

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

For more information about VPLS, see the *Junos OS VPNs Library for Routing Devices* and the *Junos OS, Release 16.1*.

The remaining statements are explained separately.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration.
	interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Protocol Family</i></li></ul>

## 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> <filter [aci]>;  
     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-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;
            }
            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> <filter [aci]>;
        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*

## fields (Flat-File Accounting Options)

---

```
Syntax  fields {
        all-fields;
        egress-stats {
            all-fields;
            input-bytes;
            input-packets;
            output-bytes;
            output-packets;
            queue-id;
            red-drop-bytes;
            red-drop-packets;
            tail-drop-packets;
        }
        general-param {
            all-fields;
            accounting-type;
            descr;
            line-id;
            logical-interface;
            nas-port-id;
            physical-interface;
            routing-instance;
            timestamp;
            vlan-id;
        }
        ingress-stats {
            all-fields;
            drop-packets;
            input-bytes;
            input-packets;
            output-bytes;
            output-packets;
            queue-id;
        }
        l2-stats {
            all-fields;
            input-mcast-bytes;
            input-mcast-packets;
        }
        overall-packet {
            all-fields;
            input-bytes;
            input-discards;
            input-errors;
            input-packets;
            inputv6-bytes;
            inputv6-packets;
            output-bytes;
            output-errors;
            output-packets;
            outputv6-bytes;
            outputv6-packets;
        }
    }
```

```
}
}
```

<b>Hierarchy Level</b>	[edit accounting-options <b>flat-file-profile</b> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Specify the accounting statistics and nonstatistical information to be collected for an interface and recorded in the accounting flat file created by the profile.
<b>Options</b>	<p><b>all-fields</b>—Include all available statistics fields in the accounting file. Many fields are available for both logical interfaces and physical interfaces, but some fields are available only for one or the other interface type.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Flat-File Accounting for Layer 2 Wholesale on page 166</a></li> <li>• <a href="#">Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170</a></li> <li>• <a href="#">Flat-File Accounting Overview on page 163</a></li> </ul>

## file (Flat-File Accounting Options)

<b>Syntax</b>	file <i>filename</i> ;
<b>Hierarchy Level</b>	[edit accounting-options <b>flat-file-profile</b> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Specify the name of the accounting file created by a flat-file profile. By default, the filename becomes the name of the local directory where the accounting file is backed up: <code>/var/log/pfedBackup/<i>filename</i></code> .
<b>Options</b>	<b>filename</b> —Name of the accounting file. The complete output filename is in the format <i>filename.hostname.file-number_timestamp.gz</i> .
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Flat-File Accounting for Layer 2 Wholesale on page 166</a></li> <li>• <a href="#">Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170</a></li> <li>• <a href="#">Flat-File Accounting Overview on page 163</a></li> </ul>

## flat-file-profile (Accounting Options)

---

**Syntax**    flat-file-profile *profile-name*{

```
    fields {
        all-fields;
        egress-stats {
            all-fields;
            input-bytes;
            input-packets;
            output-bytes;
            output-packets;
            queue-id;
            red-drop-bytes;
            red-drop-packets;
            tail-drop-packets;
        }
        general-param {
            all-fields;
            accounting-type;
            descr;
            line-id;
            logical-interface;
            nas-port-id;
            physical-interface;
            routing-instance;
            timestamp;
            vlan-id;
        }
        ingress-stats {
            all-fields;
            drop-packets;
            input-bytes;
            input-packets;
            output-bytes;
            output-packets;
            queue-id;
        }
        l2-stats {
            all-fields;
            input-mcast-bytes;
            input-mcast-packets;
        }
        overall-packet {
            all-fields;
            input-bytes;
            input-discards;
            input-errors;
            input-packets;
            inputv6-bytes;
            inputv6-packets;
            output-bytes;
            output-errors;
            output-packets;
            outputv6-bytes;
```

```

        outputv6-packets;
    }
}
file filename;
format (csv | ipdr)
interval minutes;
schema-version schema-name;
}

```

**Hierarchy Level** [edit accounting-options]

**Release Information** Statement introduced in Junos OS Release 16.1.

**Description** Configure a flat-file accounting profile that defines the contents of a flat file that records accounting statistics collected from the Packet Forwarding Engine for an interface at regular intervals. To be used, the profile is associated with a subscriber interface. The accounting flat file is archived by the accounting-options archiving mechanism.

**Options** *profile-name*—Name of the flat-file profile.

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 Accounting-Data Log Files](#)
- [Configuring Flat-File Accounting for Layer 2 Wholesale on page 166](#)
- [Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170](#)
- [Flat-File Accounting Overview on page 163](#)
- *show extensible-subscriber-services accounting*

## flat-file-profile (Extensible Subscriber Services)

---

<b>Syntax</b>	<code>flat-file-profile <i>profile-name</i></code>
<b>Hierarchy Level</b>	<code>[edit system services extensible-subscriber-services]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Specify the name of an accounting flat file profile that applies to an ESSM subscriber.
<b>Options</b>	<i>profile name</i> —Name of an accounting flat file profile configured at the <code>[edit accounting-options]</code> hierarchy level.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170</a></li></ul>

## flexible-vlan-tagging

<b>Syntax</b>	flexible-vlan-tagging;
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	<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>
<b>Description</b>	<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>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Mixed Tagging</i></li> <li>• <i>Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers</i></li> </ul>

## format (Flat-File Accounting Options)

---

<b>Syntax</b>	format (csv   ipdr);
<b>Hierarchy Level</b>	[edit accounting-options flat-file-profile <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Specify the format for logging the flat-file accounting statistics.
<b>Options</b>	<p>csv—Comma-separated values (CSV) format.</p> <p>ipdr—IP Detail Record (IPDR) format.</p> <p><b>Default:</b> ipdr</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Flat-File Accounting for Layer 2 Wholesale on page 166</a></li><li>• <a href="#">Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170</a></li><li>• <a href="#">Flat-File Accounting Overview on page 163</a></li></ul>

## forwarding-options

<b>Syntax</b>	forwarding-options { ... }
<b>Hierarchy Level</b>	[edit] [edit routing-instance <i>routing-instance-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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</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.

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

## general-param (Flat-File Accounting Options)

---

<b>Syntax</b>	<pre>general-param {     all-fields;     accounting-type;     descr;     line-id;     logical-interface;     nas-port-id;     physical-interface;     routing-instance;     timestamp;     vlan-id; }</pre>
<b>Hierarchy Level</b>	[edit accounting-options flat-file-profile <i>profile-name</i> <b>fields</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Specify general, nonstatistical interface parameters that are displayed as part of the header for the accounting file.
<b>Options</b>	<p><b>all-fields</b>—Use all available nonstatistical fields. Many fields are available for both logical interfaces and physical interfaces, but some fields are available for only one interface type.</p> <p><b>accounting-type</b>—(Logical interfaces only ) Use the accounting status type.</p> <p><b>descr</b>—Use the description of the interface as configured.</p> <p><b>line-id</b>—(Logical interfaces only ) Use the access line identifier.</p> <p><b>logical-interface</b>—(Logical interfaces only ) Use the name of the logical interface.</p> <p><b>nas-port-id</b>—(Logical interfaces only ) Use the NAS port ID.</p> <p><b>physical-interface</b>—(Physical interfaces only ) Use the name of the physical interface.</p> <p><b>routing-instance</b>—Use the name of the routing instance to which the interface belongs.</p> <p><b>timestamp</b>—Use the timestamp of the accounting record.</p> <p><b>vlan-id</b>—(Logical interfaces only ) Use the VLAN identifier.</p>
<b>Required Privilege Level</b>	<p><b>system</b>—To view this statement in the configuration.</p> <p><b>system-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Flat-File Accounting for Layer 2 Wholesale on page 166</a></li><li>• <a href="#">Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170</a></li><li>• <a href="#">Flat-File Accounting Overview on page 163</a></li></ul>

---

## grace-period

---

<b>Syntax</b>	<code>grace-period <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit access address-assignment pool <i>pool-name</i> family (inet   inet6) <a href="#">dhcp-attributes</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0.
<b>Description</b>	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.
<b>Options</b>	<b><i>seconds</i></b> —Number of seconds the lease is retained. <b>Range:</b> 0 through 4,294,967,295 seconds <b>Default:</b> 0 (no grace period)
<b>Required Privilege Level</b>	<b>admin</b> —To view this statement in the configuration. <b>admin-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Address-Assignment Pools</i></li></ul>

## 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>;
                client-negotiation-match incoming-interface;
                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>;
    client-negotiation-match incoming-interface;
    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;
}

```

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

- 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>;
    client-negotiation-match incoming-interface;
    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-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>;
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        local-server-group local-server-group;
        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);
  use-option-82;
}
remote-id {
  prefix prefix;
  use-interface-description (logical | device);
}
server-id-override
}
route-suppression;
service-profile dynamic-profile-name;
}

```

<b>Hierarchy Level</b>	<p>[edit forwarding-options <a href="#">dhcp-relay</a>],  [edit forwarding-options dhcp-relay dhcpv6],  [edit logical-systems <i>logical-system-name</i> forwarding-options <a href="#">dhcp-relay</a> ...],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay</a> ...],  [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.  Support at the [edit ... <a href="#">dhcpv6</a>] hierarchy levels introduced in Junos OS Release 11.4.  Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
<b>Description</b>	<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 ... <a href="#">dhcpv6</a>] hierarchy levels to configure DHCPv6 support.</p>
<b>Options</b>	<p><b><i>group-name</i></b>—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>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.  interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">dhcp-relay on page 219</a></li> <li>• <i>Extended DHCP Relay Agent Overview</i></li> <li>• <i>Understanding the Extended DHCP Relay Agent for EX Series Switches</i></li> <li>• <i>Configuring an Extended DHCP Relay Server on EX Series Switches (CLI Procedure)</i></li> <li>• <i>Configuring Group-Specific DHCP Relay Options</i></li> <li>• <i>Grouping Interfaces with Common DHCP Configurations</i></li> <li>• <i>Using External AAA Authentication Services with DHCP</i></li> <li>• <i>Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces</i></li> </ul>

## ingress-stats (Flat-File Accounting Options)

---

<b>Syntax</b>	<pre>ingress-stats {   all-fields;   drop-packets;   input-bytes;   input-packets;   output-bytes;   output-packets;   queue-id; }</pre>
<b>Hierarchy Level</b>	[edit accounting-options flat-file-profile <i>profile-name</i> <b>fields</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Specify ingress queue statistics to be collected for the interface.
<b>Options</b>	<p><b>all-fields</b>—Use all ingress queue statistics available for the interface context, logical or physical.</p> <p><b>drop-packets</b>—Use the number of received packets dropped on the Ingress queue.</p> <p><b>input-bytes</b>—Use the number of octets received on the queue for the traffic class indicated by the queue identifier.</p> <p><b>input-packets</b>—Use the number of packets received on the queue for the traffic class indicated by the queue identifier.</p> <p><b>output-bytes</b>—Use the number of octets forwarded for the traffic class indicated by the queue identifier. Same value as <b>input-bytes</b> unless oversubscription is present at the ingress.</p> <p><b>output-packets</b>—Use the number of packets forwarded for the traffic class indicated by the queue identifier. Same value as <b>input-packets</b> unless oversubscription is present at the ingress.</p> <p><b>queue-id</b>—Use the logical identifier for the ingress queue; identifies the traffic class.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Flat-File Accounting for Layer 2 Wholesale on page 166</a></li><li>• <a href="#">Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170</a></li><li>• <a href="#">Flat-File Accounting Overview on page 163</a></li></ul>

## inner-vlan-id (Dynamic VLANs)

<b>Syntax</b>	<code>inner-vlan-id <i>number</i>;</code>
<b>Hierarchy Level</b>	<p>[edit <b>dynamic-profiles</b> <i>profile-name</i> <b>interfaces</b> <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>input-vlan-map</b>],</p> <p>[edit <b>dynamic-profiles</b> <i>profile-name</i> <b>interfaces</b> <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>output-vlan-map</b>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	<p>For dynamic VLAN interfaces, specify the VLAN ID to rewrite for the inner tag of the final packet.</p> <p>You cannot include the <b>inner-vlan-id</b> statement with the <b>swap</b> statement, <b>swap-push</b> statement, <b>push-push</b> statement, or <b>push-swap</b> statement and the <b>inner-vlan-id</b> statement at the [edit <b>logical-systems</b> <i>logical-system-name</i> <b>interfaces</b> <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>output-vlan-map</b>] 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 <b>inner-vlan-id</b> statement you include at the [edit <b>interfaces</b> <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i>] hierarchy level.</p>
<b>Options</b>	<p><b>number</b>—VLAN ID number. When used for input VLAN maps, you can specify the <code>\$junos-inner-vlan-map-id</code> predefined variable to dynamically obtain the VLAN identifier.</p> <p><b>Range:</b> 0 through 4094</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Inner and Outer TPIDs and VLAN IDs</i></li> </ul>

## inner-vlan-id-swap-ranges

---

<b>Syntax</b>	<code>inner-vlan-id-swap-ranges low-inner-tag–high-inner-tag;</code>
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Define a core-facing VLAN ID ranges from which an inner VLAN ID tag can be allocated to replace the outer VLAN tag that was appended by the access node on the upstream packets to the BNG. You can configure multiple ranges.
<b>Options</b>	<p><b>low-inner-tag</b>—Inner (core-facing) VLAN ID tag representing the lower limit of the swap range. <b>Range:</b> 1 through 4094.</p> <p><b>high-inner-tag</b>—Inner (core-facing) VLAN ID tag representing the upper limit of the swap range. <b>Range:</b> 1 through 4094.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Inner and Outer TPIDs and VLAN IDs</a></li></ul>

## input-vlan-map (Dynamic Interfaces)

---

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

## interface (DHCP Local Server)

**Syntax** interface *interface-name* {  
     access-profile *profile-name*;  
     exclude;  
     overrides {  
         client-discover-match <option60-and-option82 | incoming-interface>;  
         client-negotiation-match incoming-interface;  
         interface-client-limit *number*;  
         rapid-commit;  
     }  
     service-profile *dynamic-profile-name*;  
     trace;  
     upto *upto-interface-name*;  
 }

**Hierarchy Level** [edit system services dhcp-local-server *group group-name*],  
 [edit system services dhcp-local-server *dhcpv6 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 9.0.  
 Statement introduced in Junos OS Release 12.3R2 for EX Series switches.  
 Options **upto** and **exclude** introduced in Junos OS Release 9.1.

**Description** 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 *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 relay agent.



**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.

**Options** **exclude**—Exclude an interface or a range of interfaces from the group. This option and the **overrides** option are mutually exclusive.

***interface-name***—Name of the interface. You can repeat this option multiple times.

***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.

<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Extended DHCP Local Server Overview</i></li><li>• <i>Grouping Interfaces with Common DHCP Configurations</i></li><li>• <i>Using External AAA Authentication Services with DHCP</i></li></ul>

## interface (DHCP Relay Agent)

**Syntax** `interface dhcpinterface-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>;  
         client-negotiation-match 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. .

.....

**Options**    **exclude**—Exclude an interface or a range of interfaces from the group. This option and the **overrides** option are mutually exclusive.

**interface-name**—Name of the interface. You can repeat this option multiple times.

**overrides**—Override the specified default configuration settings for the interface. The **overrides** statement is described separately.

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

**Related Documentation**

- *Extended DHCP Relay Agent Overview*
- [dhcp-relay on page 219](#)
- *Grouping Interfaces with Common DHCP Configurations*
- *Using External AAA Authentication Services with DHCP*

---

## interface (Dynamic Routing Instances)

---

**Syntax**    interface *interface-name*;

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

**Release Information**    Statement introduced in Junos OS Release 9.6.

**Description**    Assign the specified interface to the dynamically created routing instance.

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

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

**Related Documentation**

## interface (Routing Instances)

<b>Syntax</b>	<pre>interface <i>interface-name</i> {     description <i>text</i>; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ], [edit routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	<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>
<b>Description</b>	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 <b>vrf</b> is specified for the <b>instance-type</b> statement included in the routing instance configuration, this statement is required.
<b>Options</b>	<p><b><i>interface-name</i></b>—Name of the interface.</p> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Routing Instances on PE Routers in VPNs</i></li> <li>• <i>Configuring EVPN Routing Instances</i></li> <li>• <i>Configuring EVPN Routing Instances on EX9200 Switches</i></li> <li>• <i>interface (VPLS Routing Instances)</i></li> </ul>

## interface-mac-limit (VPLS)

---

<b>Syntax</b>	<code>interface-mac-limit <i>limit</i> {     packet-action drop; }</code>
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	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.
<b>Description</b>	<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 <b>interface-mac-limit</b> 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 <b>interface-mac-limit</b> option with a valid value for it to be saved in the configuration.</p>
<b>Options</b>	<p><b>limit</b>—Number of MAC addresses that can be learned from each interface.</p> <p><b>Range:</b> 16 through 65,536 MAC addresses</p> <p><b>Default:</b> 1024 addresses</p> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring EVPN Routing Instances</i></li><li>• <i>Configuring EVPN Routing Instances on EX9200 Switches</i></li><li>• <i>Configuring VPLS Routing Instances</i></li><li>• <i>interface</i></li><li>• <i>mac-table-size</i></li></ul>

## interfaces

---

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

## 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 {
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                    output filter-name {
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                }
            }
            max-sessions number;
            max-sessions-vsa-ignore;
            rpf-check {
                mode loose;
            }
            service {
                input {
                    service-set service-set-name {
                        service-filter filter-name;
                    }
                    post-service-filter filter-name;
                }
                output {
                    service-set service-set-name {
                        service-filter filter-name;
                    }
                }
            }
            service-name-table table-name
            short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
                maximum-seconds>;
        }
    }
```

```

    unnumbered-address interface-name <preferred-source-address address>;
}
filter {
    input filter-name (
        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 {
            aaa-options aaa-options-name;
            authentication [ authentication-protocols ];
            chap {
                challenge-length minimum minimum-length maximum maximum-length;
            }
            initiate-ncp (ip | ipv6 | dual-stack-passive)
            ipcp-suggest-dns-option;
            mru size;
            mtu (size | use-lower-layer);
            on-demand-ip-address;
            pap;
            peer-ip-address-optional;
        }
        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*

## interval (Flat-File Accounting Options)

---

<b>Syntax</b>	<code>interval <i>minutes</i>;</code>
<b>Hierarchy Level</b>	[edit accounting-options <a href="#">flat-file-profile</a> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Specify the interval in minutes at which the Packet Forwarding Engine associated with the interface is polled to collect the statistics specified in the flat-file accounting profile. The results are recorded in the flat file.
<b>Options</b>	<p><i>minutes</i>—Polling interval.</p> <p><b>Range:</b> 1 through 2880 minutes</p> <p><b>Default:</b> 15 minutes</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Flat-File Accounting for Layer 2 Wholesale on page 166</a></li><li>• <a href="#">Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170</a></li><li>• <a href="#">Flat-File Accounting Overview on page 163</a></li></ul>

## instance-role

<b>Syntax</b>	<code>instance-role (access   nni);</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ], [edit routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2.
<b>Description</b>	Define the role of the routing instance in a Layer 2 Wholesale network.
<b>Options</b>	<p><b>access</b>—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><b>nni</b>—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>  <p><b>NOTE:</b> 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.</p> </div>	
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers on page 81</a></li> <li>• <a href="#">Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers on page 84</a></li> <li>• <a href="#">Subscriber Management Overview</a></li> </ul>

## instance-type

<b>Syntax</b>	<code>instance-type type;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ], [edit routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>virtual-switch</b> and <b>layer2-control</b> 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. <b>mpls-internet-multicast</b> option introduced in Junos OS Release 11.1 for the EX Series, M Series, MX Series, and T Series. <b>evpn</b> option introduced in Junos OS Release 13.2 for MX 3D Series routers. <b>forwarding</b> option introduced in Junos OS Release 14.2 for the PTX Series. <b>mpls-forwarding</b> option introduced in Junos OS Release 16.1 for the MX Series.
<b>Description</b>	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 3D 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.

- **mpls-forwarding**—(MX Series routers only) Allow filtering and translation of route distinguisher (RD) values in IPv4 and IPv6 VPN address families on both routes received and routes sent for selected BGP sessions. In particular, for Inter-AS VPN Option-B networks, this option can prevent the malicious injection of VPN labels from one peer AS boundary router to another.
- **mpls-internet-multicast**—(EX Series, M Series, MX Series, and T Series routers only) Provide support for ingress replication provider tunnels to carry IP multicast data between routers through an MPLS cloud, using MBGP or next-generation MVPN.
- **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.

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

- Related Documentation**
- *Configuring Routing Instances on PE Routers in VPNs*
  - *Configuring EVPN Routing Instances*
  - *Configuring EVPN Routing Instances on EX9200 Switches*
  - *Configuring Virtual Router Routing Instances*
  - *Example: Configuring Filter-Based Forwarding on the Source Address*
  - *Example: Configuring Filter-Based Forwarding on Logical Systems*
  - *Layer 2 Routing Instance Types*

---

## ip-address-first

---

- Syntax**    ip-address-first;
- Hierarchy Level**    [edit logical-systems *logical-system-name* system services dhcp-local-server [pool-match-order](#)],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server [pool-match-order](#)],  
[edit routing-instances *routing-instance-name* 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**
- *Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use*
  - *Extended DHCP Local Server Overview*
  - *Address-Assignment Pools Overview*

## keepalives (Dynamic Profiles)

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

## l2-stats (Flat-File Accounting Options)

---


<b>Syntax</b>	<pre>l2-stats {     all-fields;     input-mcast-bytes;     input-mcast-packets; }</pre>
<b>Hierarchy Level</b>	[edit accounting-options flat-file-profile <i>profile-name</i> <b>fields</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Specify the statistics to collect for the named flat-file-profile field.
<b>Options</b>	<p><b>all-fields</b>—Collect all Layer 2 statistics for the named flat-file profile.</p> <p><b>input-mcast-bytes</b>—Count multicast bytes from the input side for the named flat-file profile.</p> <p><b>input-mcast-packets</b>—Count multicast packets from the input side for the named flat-file profile.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Flat-File Accounting for Layer 2 Wholesale on page 166</a></li><li>• <a href="#">Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170</a></li><li>• <a href="#">Flat-File Accounting Overview on page 163</a></li></ul>

## mac-validate (Dynamic IP Demux Interface)

<b>Syntax</b>	mac-validate (loose   strict);
<b>Hierarchy Level</b>	[edit dynamic-profiles <i>profile-name</i> interfaces demux0 unit <i>logical-unit-number</i> <b>family</b> inet]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3.
<b>Description</b>	Enable IP and MAC address validation for dynamic IP demux interfaces in a dynamic profile.
<b>Options</b>	<p><b>loose</b>—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><b>strict</b>—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>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring MAC Address Validation for Subscriber Interfaces</i></li> </ul>

## multicast-replication

---

<b>Syntax</b>	<pre>multicast-replication {     ingress;     local-latency-fairness; }</pre>
<b>Hierarchy Level</b>	[edit forwarding-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 15.1 for MX Series routers.
<b>Description</b>	Configure the mode of multicast replication that helps to optimize multicast latency.
<div> <b>NOTE:</b> The <code>multicast-replication</code> statement is supported only on platforms with the <code>enhanced-ip</code> mode enabled.</div>	
<b>Default</b>	This statement is disabled by default.
<b>Options</b>	<p><b>ingress</b>—Complete ingress replication of the multicast data packets where all the egress Packet Forwarding Engines receive packets from the ingress Packet Forwarding Engines directly.</p> <p><b>local-latency-fairness</b>—Complete parallel replication of the multicast data packets.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">forwarding-options on page 279</a></li></ul>

## neighbor (Define ANCP)

<b>Syntax</b>	<pre>neighbor <i>ip-address</i> {     adjacency-loss-hold-time <i>seconds</i>;     adjacency-timer;     auto-configure-trigger interface <i>interface-name</i>;     ietf-mode;     maximum-discovery-table-entries <i>entry-number</i>;     pre-ietf-mode; }</pre>
<b>Hierarchy Level</b>	[edit protocols <a href="#">ancp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4.
<b>Description</b>	Configure an ANCP neighbor with which the ANCP agent on the router forms an adjacency for reporting and shaping traffic.
<b>Options</b>	<p><i>ip-address</i>—IP address of the ANCP neighbor.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the ANCP Agent</a></li> <li>• <a href="#">Configuring ANCP Neighbors on page 153</a></li> </ul>

## no-local-switching

<b>Syntax</b>	no-local-switching
<b>Hierarchy Level</b>	[edit vlans <i>vlan-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3 for EX Series switches.
<b>Description</b>	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.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

## no-tunnel-services

---

<b>Syntax</b>	no-tunnel-services;
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Support for static VPLS added in Junos OS Release 10.2.
<b>Description</b>	Configure VPLS on a router without a Tunnel Services PIC. Configuring the <b>no-tunnel-services</b> 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.

---

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

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

## maximum-lease-time

---

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

## overall-packet (Flat-File Accounting Options)

---

<b>Syntax</b>	<pre>overall-packet {     all-fields;     input-bytes;     input-discards;     input-errors;     input-packets;     inputv6-bytes;     inputv6-packets;     output-bytes;     output-errors;     output-packets;     outputv6-bytes;     outputv6-packets; }</pre>
<b>Hierarchy Level</b>	[edit accounting-options flat-file-profile <i>profile-name</i> <b>fields</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Specify overall packet statistics to be collected for the interface.
<b>Options</b>	<p><b>all-fields</b>—Collect all overall packet statistics available for the interface context, logical or physical.</p> <p><b>input-bytes</b>—Collect the number of octets received on the interface.</p> <p><b>input-discards</b>—(Physical interfaces only ) Collect the number of received packets that were discarded on the interface.</p> <p><b>input-errors</b>—(Physical interfaces only ) Collect the number of frames with errors received on the interface.</p> <p><b>input-packets</b>—Collect the number of packets received on the interface.</p> <p><b>input-v6-bytes</b>—Collect the number of IPv6 octets received on the interface.</p> <p><b>input-v6-packets</b>—Number of IPv6 packets received on the interface.</p> <p><b>output-bytes</b>—Collect the number of octets transmitted on the interface.</p> <p><b>output-errors</b>—(Physical interfaces only ) Collect the number of frames that could not be transmitted on the interface because of errors.</p> <p><b>output-packets</b>—Collect the number of packets transmitted on the interface.</p> <p><b>output-v6-bytes</b>—Collect the number of IPv6 octets transmitted on the interface.</p> <p><b>output-v6-packets</b>—Collect the number of IPv6 packets transmitted on the interface.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Configuring Flat-File Accounting for Layer 2 Wholesale on page 166](#)
  - [Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170](#)
  - [Flat-File Accounting Overview on page 163](#)

## output-vlan-map (Dynamic Interfaces)

---

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[edit <b>dynamic-profiles</b> <i>profile-name</i> <b>interfaces</b> <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	<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>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution on page 76</a></li> </ul>

## pap (Dynamic PPP)

---

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

## 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;
    }
  }
  hold-down;
  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

---

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

## pop (Dynamic VLANs)

---

<b>Syntax</b>	<pre>pop;</pre>
<b>Hierarchy Level</b>	[edit <a href="#">dynamic-profiles</a> <i>profile-name</i> <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">output-vlan-map</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Removing a VLAN Tag</i></li><li>• <a href="#">Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution on page 76</a></li></ul>

---

## pppoe-options (Dynamic PPPoE)

---

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

## pppoe-underlying-options (Static and Dynamic Subscribers)

---

<b>Syntax</b>	<pre>pppoe-underlying-options {   access-concentrator <i>name</i>;   <b>dynamic-profile</b> <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 &lt;lockout-time-min <i>minimum-seconds</i>&gt; &lt;lockout-time-max     <i>maximum-seconds</i>&gt; &lt;filter [<i>aci</i>]&gt;; }</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0.
<b>Description</b>	<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 (<b>ppp-over-ether</b>) encapsulation.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring PPPoE</i> (for static interfaces)</li><li>• <i>Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces</i></li><li>• <i>Assigning a Service Name Table to a PPPoE Underlying Interface</i></li></ul>

## ppp-options (Dynamic PPP)

<b>Syntax</b>	<pre> ppp-options {   aaa-options <i>aaa-options-name</i>;   authentication [ <i>authentication-protocols</i> ];   mru <i>size</i>;   mtu (<i>size</i>   use-lower-layer);   chap {     challenge-length minimum <i>minimum-length</i> maximum <i>maximum-length</i>;   }   initiate-ncp (ip   ipv6   dual-stack-passive)   ipcp-suggest-dns-option;   mru <i>size</i>;   mtu (<i>size</i>   use-lower-layer);   on-demand-ip-address;   pap;   peer-ip-address-optional; } </pre>
<b>Hierarchy Level</b>	[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"]
<b>Release Information</b>	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.
<b>Description</b>	Configure PPP-specific interface properties in a dynamic profile.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Dynamic Profiles Overview</i></li> <li>• <i>Configuring Dynamic Authentication for PPP Subscribers</i></li> <li>• <i>Attaching Dynamic Profiles to Static PPP Subscriber Interfaces</i></li> <li>• <i>Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface</i></li> </ul>

## prefix (Address-Assignment Pools)

---

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

## 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 number;
                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;
                sessions-limit-group limit-group-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 | pcrf);
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;
        interface-text-description;
        mac-address;
        nas-identifier;
        stacked-vlan;
        vlan;
    }
    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;
    pw-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;
service-activation {
    dynamic-profile (optional-at-login | required-at-login);
    extensible-service (optional-at-login | required-at-login);
}
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-idle-timeout-ingress-only;  
    client-session-timeout minutes;  
    strip-user-name {  
      delimiter [ delimiter ];  
      parse-direction (left-to-right | right-to-left);  
    }  
  }  
}
```

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><b><i>profile-name</i></b>—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"><li>• <i>Configuring the PPP Authentication Protocol</i></li><li>• <i>Configuring Access Profiles for L2TP or PPP Parameters</i></li><li>• <i>Configuring L2TP Properties for a Client-Specific Profile</i></li><li>• <i>Configuring an L2TP LNS with Inline Service Interfaces</i></li><li>• <i>Configuring PPP Properties for a Client-Specific Profile</i></li><li>• <i>Configuring Service Accounting with JSRC</i></li><li>• <i>AAA Service Framework Overview</i></li><li>• <a href="#">show network-access aaa statistics on page 388</a></li><li>• <i>clear network-access aaa statistics</i></li></ul>

## 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;  
  }  
}
```

<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ], [edit routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	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. <b>mpls</b> and <b>rsvp</b> options added in Junos OS Release 15.1. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	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.

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

## proxy-arp

---

<b>Syntax</b>	proxy-arp (restricted   unrestricted);
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.6 for EX Series switches. <b>restricted</b> added in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 12.2 for the QFX Series.
<b>Description</b>	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> <b>NOTE:</b> You must configure the IP address and the inet family for the interface when you enable proxy ARP.</div>	
<b>Default</b>	Proxy ARP is not enabled. The router or switch responds to an ARP request only if the destination IP address is its own.
<b>Options</b>	<ul style="list-style-type: none"><li>• <b>none</b>—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.</li><li>• <b>restricted</b>—(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.</li><li>• <b>unrestricted</b>—(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.</li></ul>
<b>Default:</b> unrestricted	
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Restricted and Unrestricted Proxy ARP</i></li><li>• <i>Configuring Proxy ARP (CLI Procedure)</i></li><li>• <i>Example: Configuring Proxy ARP on an EX Series Switch</i></li><li>• <i>Configuring Gratuitous ARP</i></li></ul>

## proxy-arp (Dynamic Profiles)


<b>Syntax</b>	<code>proxy-arp;</code>
<b>Hierarchy Level</b>	[edit <code>dynamic-profiles profile-name interfaces interface-name unit logical-unit-number</code> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.5.
<b>Description</b>	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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Restricted and Unrestricted Proxy ARP</i></li> <li>• <i>Configuring Gratuitous ARP</i></li> </ul>

## push (Dynamic VLANs)

<b>Syntax</b>	<code>push;</code>
<b>Hierarchy Level</b>	[edit <code>dynamic-profiles profile-name interfaces interface-name unit logical-unit-number input-vlan-map</code> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	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 <code>push</code> statement in the configuration, you must also include the <code>pop</code> statement at the [edit <code>dynamic-profiles profile-name interfaces interface-name unit logical-unit-number output-vlan-map</code> ] hierarchy level.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution on page 76</a></li> </ul>

## push-backup-to-master (Accounting Options)

---

<b>Syntax</b>	push-backup-to-master;
<b>Hierarchy Level</b>	[edit accounting-options <i>file filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	<p>Configure the router to save the accounting files from the new backup Routing Engine to the new master Routing Engine when a change in mastership occurs. The files are saved to the <code>/var/log/pfedBackup</code> directory on the router. The master Routing Engine includes these accounting files with its own current accounting files when it transfers the files from the backup directory to the archive site at the next transfer interval. Use this statement when the new backup Routing Engine is not able to connect to the archive site; for example, when site is not connected by means of an out-of-band interface or the path to the site is routed through a line card.</p>
	<div> <b>NOTE:</b> The backup Routing Engine's files on the master Routing Engine are sent at each interval even though the files remain the same. If this is more activity than you want, consider using the <code>backup-on-failure master-and-slave</code> statement instead.</div>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Accounting-Data Log Files</a></li><li>• <a href="#">Configuring Flat-File Accounting for Layer 2 Wholesale on page 166</a></li><li>• <a href="#">Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170</a></li><li>• <a href="#">Flat-File Accounting Overview on page 163</a></li></ul>

## 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;
service-activation {
    dynamic-profile (optional-at-login | required-at-login);
    extensible-service (optional-at-login | required-at-login);
}
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

<b>Syntax</b>	<pre>radius-server server-address {     accounting-port <i>port-number</i>;     accounting-retry <i>number</i>;     accounting-timeout <i>seconds</i>;     dynamic-request-port     max-outstanding-requests <i>value</i>;     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>;     <b>secret</b> <i>password</i>;     source-address <i>source-address</i>;     timeout <i>seconds</i>; }</pre>
<b>Hierarchy Level</b>	<p>[edit access],</p> <p>[edit access <b>profile</b> <i>profile-name</i>]</p>
<b>Release Information</b>	<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><b>dynamic-request-port</b> option added in Junos OS Release 14.2 for MX Series routers.</p> <p><b>preauthentication-port</b> and <b>preauthentication-secret</b> options added in Junos OS Release 15.1 for MX Series routers.</p> <p>Support for IPv6 <b>server-address</b> introduced in Junos OS Release 16.1.</p>
<b>Description</b>	<p>Configure RADIUS for subscriber access management, L2TP, or PPP.</p> <p>To configure multiple RADIUS servers, include multiple <b>radius-server</b> 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>
<b>Options</b>	<p><b>server-address</b>—IPv4 or IPv6 address of the RADIUS server.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring RADIUS Authentication for L2TP</i></li> <li>• <i>Configuring the PPP Authentication Protocol</i></li> <li>• <i>Configuring Router or Switch Interaction with RADIUS Servers</i></li> <li>• <i>Configuring Authentication and Accounting Parameters for Subscriber Access</i></li> <li>• <a href="#">show network-access aaa statistics on page 388</a></li> <li>• <i>clear network-access aaa statistics</i></li> </ul>

## range (Address-Assignment Pools)

---

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

## ranges (Dynamic VLAN)

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

## remote-id (VLAN Authentication Username)

<b>Syntax</b>	<code>remote-id;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> auto-configure vlan-ranges authentication <b>username-include</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Include the agent remote identifier (ARI) in the username sent to RADIUS for authentication of the dynamic VLAN. The ARI is conveyed by the Access-Loop-Remote-ID TLV in an out-of-band ANCP Port Up message.



**NOTE:** This statement is not supported for stacked VLANs.

<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs on page 155</a></li> <li>• <a href="#">Configuring VLAN Interface Username Information for AAA Authentication</a></li> <li>• <a href="#">Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93</a></li> </ul>

## route-distinguisher

<b>Syntax</b>	<code>route-distinguisher (as-number:id   ip-address:id);</code>
<b>Hierarchy Level</b>	<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>
<b>Release Information</b>	<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>
<b>Description</b>	<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 <b>vrf</b>, the <b>route-distinguisher</b> statement is required.</p> <p>For Layer 2 VPNs and VPLS, if you configure the <b>l2vpn-use-bgp-rules</b> 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)

---

<b>Syntax</b>	<code>routing-instances <i>routing-instance-name</i> { ... }</code>
<b>Hierarchy Level</b>	[edit], [edit logical-systems <i>logical-system-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<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"><li>• A set of routing tables</li><li>• A set of interfaces that belong to these routing tables</li><li>• A set of routing option configurations</li></ul> <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 <b>my-instance</b>, its corresponding IP unicast table is my-instance.inet.0. All routes for <b>my-instance</b> 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>
<b>Default</b>	Routing instances are disabled for the router.
<b>Options</b>	<i>routing-instance-name</i> —Name of the routing instance. This must be a non-reserved string of not more than 128 characters.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Configuring Interprovider Layer 3 VPN Option A</i></li><li>• <i>Example: Configuring Interprovider Layer 3 VPN Option B</i></li></ul>

- *Example: Configuring Interprovider Layer 3 VPN Option C*

## schema-version (Flat-File Accounting Options)

---

<b>Syntax</b>	<code>schema-version <i>schema-name</i>;</code>
<b>Hierarchy Level</b>	[edit accounting-options <a href="#">flat-file-profile</a> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Specify the name of the XML schema that defines the contents and format of the accounting file, and appears in the accounting record header.
<b>Options</b>	<i>schema-name</i> —Name of the schema.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Flat-File Accounting for Layer 2 Wholesale on page 166</a></li><li>• <a href="#">Configuring Flat File Accounting for Extensible Subscriber Services Manager on page 170</a></li><li>• <a href="#">Flat-File Accounting Overview on page 163</a></li></ul>

## secret

---

<b>Syntax</b>	<code>secret password;</code>
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> <b>radius-server</b> <i>server-address</i> ], [edit access radius-disconnect <i>client-address</i> ], [edit access <b>radius-server</b> <i>server-address</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	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.
<b>Options</b>	<b>password</b> —Password to use; it can include spaces if the character string is enclosed in quotation marks.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Authentication and Accounting Parameters for Subscriber Access</i></li><li>• <i>Configuring Router or Switch Interaction with RADIUS Servers</i></li><li>• <i>Example: Configuring CHAP Authentication with RADIUS</i></li><li>• <i>Configuring RADIUS Authentication for L2TP</i></li><li>• <i>Configuring the RADIUS Disconnect Server for L2TP</i></li></ul>

## server (Dynamic PPPoE)

---

<b>Syntax</b>	<code>server;</code>
<b>Hierarchy Level</b>	[edit <b>dynamic-profiles</b> <i>profile-name</i> <b>interfaces</b> <i>pp0</i> unit "\$junos-interface-unit" <b>pppoe-options</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	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.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring a PPPoE Dynamic Profile</i></li><li>• <i>Subscriber Interfaces and PPPoE Overview</i></li></ul>

## server-group

<b>Syntax</b>	<pre>server-group {     server-group-name {         server-ip-address;     } }</pre>
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay],          [edit forwarding-options <a href="#">dhcp-relay</a> dhcpv6],          [edit logical-systems <i>logical-system-name</i> forwarding-options <a href="#">dhcp-relay</a>],          [edit logical-systems <i>logical-system-name</i> forwarding-options <a href="#">dhcp-relay</a> dhcpv6],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay</a>],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay</a> dhcpv6],          [edit routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay</a>],          [edit routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay</a> dhcpv6]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.          Support at the <b>[edit ... dhcpv6]</b> hierarchy levels introduced in Junos OS Release 11.4.          Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
<b>Description</b>	<p>Specify the name of a group of DHCP server addresses for use by the extended DHCP relay agent. Use the statement at the <b>[edit ... dhcpv6]</b> hierarchy levels to configure DHCPv6 support.</p>
<b>Options</b>	<p><b>server-group-name</b>—Name of the group of DHCP or DHCPv6 server addresses.</p> <p><b>server-ip-address</b>—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.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.          interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">dhcp-relay on page 219</a></li> <li>• <i>Extended DHCP Relay Agent Overview</i></li> <li>• <i>Configuring Server Groups</i></li> </ul>

## services (System Services)

---

```
Syntax  services {
        dhcp { # DHCP is not supported on a DCF
            dhcp_services;
        }
        extension-service {
            request-response {
                thrift {
                    (clear-text | ssl) {
                        address ip-address;
                        local-certificate local-certificate; # Not specified for clear-text option
                        port port;
                    }
                    max-connections max-connections;
                    scripts {
                        file script-name {
                            arguments arguments;
                            checksum (md5 | sha-256 | sha1) hash;
                            daemonize;
                            username username;
                        }
                    }
                }
            }
        }
        notification {
            port port;
            max-connections max-connections;
            allow-clients {
                address ip-address;
            }
        }
        traceoptions {
            file <filename> <files number> <match regex> <size size> <world-readable |
                no-world-readable>;
            flag flag;
            no-remote-trace;
        }
        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];
  ciphers [ cipher-1 cipher-2 cipher-3 ...];
  client-alive-count-max seconds;
  client-alive-interval seconds;
  connection-limit limit;
  fingerprint-hash (md5 | sha2-256);
  hostkey-algorithm (algorithm | no-algorithm);
  key-exchange [algorithm];
  macs [algorithm];
  max-sessions-per-connection <number>;
  no-passwords;
  no-public-keys;
  no-tcp-forwarding;
  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]

<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Option <b>extension-service</b> introduced in Junos OS Release 16.1 for MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series.
<b>Description</b>	<p>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. Also, enable configuration of third-party applications developed using the Juniper Extension Toolkit (JET) to run on Junos OS.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring clear-text or SSL Service for Junos XML Protocol Client Applications</i></li><li>• <i>Configuring the Junos OS to Work with SRC Software</i></li><li>• <i>JETUnderstanding JET Interaction with Junos OS</i></li></ul>

## site (VPLS Multihoming for FEC 128)

<b>Syntax</b>	<pre> site <i>site-name</i> {   mac-pinning;   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>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">vpls</a> ], [edit routing-instances <i>routing-instance-name</i> protocols <a href="#">vpls</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the site name and site identifier for a site. Allows you to configure a remote site ID for remote sites.
<b>Options</b>	<p><b><i>site-name</i></b>—Name of the site.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## site-identifier (VPLS)

<b>Syntax</b>	<pre> site-identifier <i>identifier</i>; </pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">vpls</a> site <i>site-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols <a href="#">vpls</a> site <i>site-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the numerical identifier for the local VPLS site.
<b>Options</b>	<p><b><i>identifier</i></b>—Specify the numerical identifier for the local VPLS site. The identifier must be an unsigned 16-bit number greater than zero.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Configuring VPLS Routing Instances</li> </ul>

## site-range

---

<b>Syntax</b>	<code>site-range <i>number</i>;</code>
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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 <b>show vpls connections</b> command, such sites are displayed as OR (out of range).
<b>Options</b>	<b><i>number</i></b> —Maximum number of site identifiers. We recommend using the default value. <b>Range:</b> 1 through 65,534 <b>Default:</b> 65,534
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring VPLS Routing Instances</i></li></ul>

## 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);
            access-profile vlan-dynamic-profile-name;
            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

---

<b>Syntax</b>	stacked-vlan-tagging;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
<b>Description</b>	<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>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview</i></li><li>• <a href="#">vlan-tags (Stacked VLAN Tags) on page 379</a></li></ul>

## system

---

<b>Syntax</b>	system { ... }
<b>Hierarchy Level</b>	[edit]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure system management properties.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>System Management Configuration Statements</i></li></ul>

## traceoptions (DHCP)

<b>Syntax</b>	<pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;match <i>regular-expression</i> &gt; &lt;size <i>maximum-file-size</i>&gt;     &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i>;     level (all   error   info   notice   verbose   warning);     no-remote-trace; } </pre>
<b>Hierarchy Level</b>	[edit system processes dhcp-service]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
<b>Description</b>	<p>Define global tracing operations for extended DHCP local server and extended DHCP relay agent processes.</p> <p>This statement replaces the deprecated <b>traceoptions</b> statements at the <b>[edit forwarding-options dhcp-relay]</b> and <b>[edit system services dhcp-local-server]</b> hierarchy levels.</p>
<b>Options</b>	<p><b>file <i>filename</i></b>—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 <b>/var/log</b>.</p> <p><b>files <i>number</i></b>—(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 <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements:</p> <ul style="list-style-type: none"> <li>• <b>all</b>—Trace all events.</li> <li>• <b>auth</b>—Trace authentication events.</li> <li>• <b>database</b>—Trace database events.</li> <li>• <b>fwd</b>—Trace firewall process events.</li> <li>• <b>general</b>—Trace miscellaneous events.</li> <li>• <b>ha</b>—Trace high availability-related events.</li> <li>• <b>interface</b>—Trace interface operations.</li> <li>• <b>io</b>—Trace I/O operations.</li> <li>• <b>liveness-detection</b>—Trace liveness detection operations.</li> <li>• <b>packet</b>—Trace packet and option decoding operations.</li> <li>• <b>performance</b>—Trace performance measurement operations.</li> </ul>

- **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.

<b>Required Privilege Level</b>	<b>trace</b> —To view this statement in the configuration. <b>trace-control</b> —To add this statement to the configuration.
---------------------------------	---

**Related Documentation**

- *Tracing Extended DHCP Operations*

## underlying-interface (demux0)

<b>Syntax</b>	<code>underlying-interface <i>underlying-interface-name</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">dynamic-profiles</a> <i>profile-name</i> <a href="#">interfaces</a> <b>demux0</b> <i>interface-name</i> <a href="#">unit</a> <i>unit logical-unit-number</i> <a href="#">demux-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3. Support for aggregated Ethernet introduced in Junos OS Release 9.4.
<b>Description</b>	Configure the underlying interface on which the demultiplexing (demux) interface is running.
<b>Options</b>	<p><b><i>underlying-interface-name</i></b>—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"> <li>• <b>\$junos-underlying-interface</b> when configuring dynamic IP demux interfaces.</li> <li>• <b>\$junos-interface-ifd-name</b> when configuring dynamic VLAN demux interfaces.</li> </ul> <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><b>NOTE:</b> Logical demux interfaces are currently supported on Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet interfaces.</p> </div>	
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles</i></li> <li>• <i>Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles</i></li> <li>• <i>Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview</i></li> <li>• For information about static underlying interfaces, see the <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul>

## underlying-interface (Dynamic PPPoE)

---

<b>Syntax</b>	<code>underlying-interface <i>interface-name</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">dynamic-profiles</a> <i>profile-name</i> <a href="#">interfaces</a> pp0 unit "\$junos-interface-unit" <a href="#">ppoe-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	In a dynamic profile, configure the underlying interface on which the router creates the dynamic PPPoE logical interface.
<b>Options</b>	<b><i>interface-name</i></b> —Variable used to specify the name of the underlying interface on which the PPPoE logical interface is dynamically created. In the <b>underlying-interface <i>interface-name</i></b> statement for dynamic PPPoE logical interfaces, you must use the predefined variable <b>\$junos-underlying-interface</b> in place of <b><i>interface-name</i></b> . When the router creates the dynamic PPPoE interface, the <b>\$junos-underlying-interface</b> predefined variable is dynamically replaced with the name of the underlying interface supplied by the network when the subscriber logs in.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring a PPPoE Dynamic Profile</i></li><li>• <i>Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview</i></li></ul>

## 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 {
    mru size;
    mtu (size | use-lower-layer);
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
    dynamic-profile profile-name;
    ipcp-suggest-dns-option;
    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;
    }
}
}
}
}

```

<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.
<b>Options</b>	<p><b>logical-unit-number</b>—Number of the logical unit.</p> <p><b>Range:</b> 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** interface—To view this statement in the configuration.

**Level** interface-control—To add this statement to the configuration.

- Related Documentation**
- *Configuring Logical Interface Properties*
  - *Example: Configuring E-LINE and E-LAN Services for a PBB Network on MX Series Routers*
  - *Junos OS Services Interfaces Library for Routing Devices*

## 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** • *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*

## 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;
                    shared-name filter-shared-name;
                }
                output filter-name {
                    precedence precedence;
                    shared-name filter-shared-name;
                }
            }
            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>;
keepalives {
    interval seconds;
}
ppp-options {
    chap;
    pap;
}
vlan-id number;
vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
}
filter {
    input filter-name {
        shared-name filter-shared-name;
    }
    output filter-name {
        shared-name filter-shared-name;
    }
}
}

```

**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)

<b>Syntax</b>	<code>unnumbered-address interface-name &lt;preferred-source-address address&gt;;</code>
<b>Hierarchy Level</b>	[edit <code>dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family family</code> ], [edit <code>dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family</code> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2. Support for the <code>\$junos-preferred-source-address</code> and <code>\$junos-preferred-source-ipv6-address</code> predefined variables introduced in Junos OS Release 9.6. Support for the <code>\$junos-loopback-interface</code> predefined variable introduced in Junos OS Release 9.6.
<b>Description</b>	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 <code>\$junos-loopback-interface-address</code> predefined variable.  You can configure unnumbered address support on Ethernet interfaces for IPv4 and IPv6 address families.
<b>Options</b>	<b>interface-name</b> —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 <code>\$junos-loopback-interface</code> predefined variable.  When defining the <code>unnumbered-address</code> statement using a static interface, keep the following in mind: <ul style="list-style-type: none"> <li>If you choose to include the <code>routing-instance</code> statement at the [edit <code>dynamic-profiles</code>] hierarchy level, that statement must be configured with a dynamic value by using the <code>\$junos-routing-instance</code> predefined variable. In addition, whatever static unnumbered interface you specify must belong to that routing instance; otherwise, the profile instantiation fails.</li> <li>If you choose to not include the <code>routing-instance</code> statement at the [edit <code>dynamic-profiles</code>] hierarchy level, the <code>unnumbered-address</code> 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.</li> </ul>



**NOTE:** When you specify a static logical interface for the unnumbered interface in a dynamic profile that includes the `$junos-routing-instance` predefined variable, you must not configure a preferred source address, whether with the `$junos-preferred-source-address` predefined variable, the `$junos-preferred-source-ipv6-address` predefined variable, or the `preferred-source-address` statement. Configuring the preferred source address in this circumstance causes a commit failure.

When defining the **unnumbered-address** statement using the **\$junos-loopback-interface** predefined variable, keep the following in mind:

- To use the **\$junos-loopback-interface** predefined variable, the dynamic profile must also contain the **routing-instance** statement configured with the **\$junos-routing-instance** predefined 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** predefined variable for the inet family, or the **\$junos-preferred-source-ipv6-address** predefined 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 the **\$junos-preferred-source-ipv6-address** predefined variables, keep the following in mind:

- You must configure the **unnumbered-address** statement using the **\$junos-loopback-interface** predefined variable.
- You must configure the **routing-instance** statement using the **\$junos-routing-instance** predefined 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.

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

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Dynamic Profiles Overview</i></li></ul>
------------------------------	--

## unnumbered-address (Ethernet)

<b>Syntax</b>	<code>unnumbered-address interface-name &lt;preferred-source-address address&gt;;</code>
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">family</a> <i>family</i> ], [edit logical-systems <i>logical-system-name</i> <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">family</a> <i>family</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2. <b>preferred-source-address</b> option introduced in Junos OS Release 9.0.
<b>Description</b>	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.
<b>Options</b>	<b>interface-name</b> —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 <b>preferred-source-address</b> statement is explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring an Unnumbered Interface</i></li> <li>• <a href="#">address on page 188</a></li> <li>• <i>Junos System Basics Configuration Guide</i></li> </ul>

## username-include

---

Syntax	<pre>username-include {     circuit-id;     circuit-type;     delimiter <i>delimiter-character</i>;     domain-name <i>domain-name-string</i>;     interface-name;     mac-address;     option-18;     option-37;     option-82 &lt;circuit-id&gt; &lt;remote-id&gt;;     radius-realm <i>radius-realm-string</i>;     remote-id;     user-prefix <i>user-prefix-string</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> auto-configure <a href="#">vlan-ranges authentication</a> ], [edit interfaces <i>interface-name</i> auto-configure <a href="#">stacked-vlan-ranges authentication</a> ]
Release Information	Statement introduced in Junos OS Release 10.0.
Description	<p>Configure the username that the router passes to the external AAA server. You must include at least one of the optional statements for the username to be valid. If you do not configure a username, the router accesses the local authentication service only and does not use external authentication services, such as RADIUS.</p> <p>The username takes the format <b><i>user-prefix mac-address circuit-type circuit-id remote-id option-82 interface-name domain-name radius-realm</i></b>. By default, each component is separated by a period (.), but you can specify a different delimiter with the <b>delimiter</b> statement.</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"><li>• <a href="#">Configuring VLAN Interface Username Information for AAA Authentication</a></li><li>• <a href="#">Using DHCP Option 82 Suboptions in Authentication Usernames for Autosense VLANs</a></li><li>• <a href="#">Using DHCP Option 18 and Option 37 in Authentication Usernames for DHCPv6 Autosense VLANs</a></li><li>• <a href="#">Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs on page 155</a></li></ul>

## user-prefix (DHCP Local Server)

<b>Syntax</b>	<code>user-prefix <i>user-prefix-string</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication</b> username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dhcpv6 authentication</b> 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 <b>group group-name authentication</b> username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication</b> username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services <b>dhcp-local-server authentication</b> username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <b>dhcpv6 authentication</b> username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication</b> username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <b>group group-name authentication</b> username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication</b> username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dhcpv6 authentication</b> 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 <b>group group-name authentication</b> username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication</b> username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication</b> username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dhcpv6 authentication</b> username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication</b> username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication</b> username-include],</p> <p>[edit system services <b>dhcp-local-server authentication</b> username-include],</p> <p>[edit system services dhcp-local-server <b>dhcpv6 authentication</b> username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 <b>group group-name authentication</b> username-include],</p> <p>[edit system services dhcp-local-server <b>group group-name authentication</b> username-include]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Specify the user prefix that is concatenated with the username during the subscriber authentication or DHCP client authentication process.
<b>Options</b>	<i>user-prefix-string</i> —User prefix string.
<b>Required Privilege Level</b>	<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)

---

<b>Syntax</b>	<code>vlan-id number;</code>
<b>Hierarchy Level</b>	[edit <a href="#">dynamic-profiles profile-name interfaces interface-name unit logical-unit-number input-vlan-map</a> ], [edit <a href="#">dynamic-profiles profile-name interfaces interface-name unit logical-unit-number output-vlan-map</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	<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 <b>vlan-id</b> statement with the <b>swap</b> statement, <b>swap-push</b> statement, <b>push-push</b> statement, or <b>push-swap</b> statement at the [edit <a href="#">dynamic-profiles profile-name interfaces interface-name unit logical-unit-number output-vlan-map</a>] 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 <b>vlan-id</b> statement that you include at the [edit <a href="#">dynamic-profiles profile-name interfaces interface-name unit logical-unit-number</a>] hierarchy level.</p>
<b>Options</b>	<b>number</b> —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.
<b>Required Privilege Level</b>	<b>interface</b> —To view this statement in the configuration. <b>interface-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Rewriting the VLAN Tag on Tagged Frames</i></li><li>• <i>Binding VLAN IDs to Logical Interfaces</i></li></ul>

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

<b>Syntax</b>	<code>vlan-id <i>number</i>;</code>
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	For Fast Ethernet, Gigabit Ethernet, and Aggregated Ethernet interfaces only, bind a 802.1Q VLAN tag ID to a logical interface.
<b>Options</b>	<p><i>number</i>—A valid VLAN identifier.</p> <p><b>Range:</b> 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>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Mixed Tagging</i></li> </ul>

## vlan-model

---

<b>Syntax</b>	vlan-model one-to-one;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ], [edit routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2.
<b>Description</b>	Define the network VLAN model.
<b>Options</b>	<b>one-to-one</b> —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 <b>instance-role</b> statement is defined as <b>access</b> , or the outer VLAN tag only, when the <b>instance-role</b> statement is defined as <b>nni</b> .
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers on page 81</a></li><li>• <a href="#">Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers on page 84</a></li></ul>

## vlan-ranges

```
Syntax  vlan-ranges {
        access-profile profile-name;
        authentication {
            packet-types [packet-types];
            password password-string;
            username-include {
                circuit-type;
                circuit-id;
                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;
                remote-id;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name {
            accept (any | dhcp-v4 | inet);
            accept-out-of-band protocol;
            access-profile vlan-dynamic-profile-name;
            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

<b>Syntax</b>	<code>vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];</code>
<b>Hierarchy Level</b>	<code>[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.5. VLAN demux interface support introduced in Junos OS Release 10.2.
<b>Description</b>	For Gigabit Ethernet IQ and IQE interfaces only, binds TPIDs and 802.1Q VLAN tag IDs to a logical interface. You must include the <b>stacked-vlan-tagging</b> statement at the <code>[edit interfaces interface-name]</code> hierarchy level.



**NOTE:** The inner-range *vid1–vid2* option is supported on IQE PICs only.

<b>Options</b>	<b>inner [tpid].vlan-id</b> —A TPID (optional) and a valid VLAN identifier in the format <i>tpid.vlan-id</i> . When used in the <b>dynamic-profiles</b> hierarchy, specify the <code>\$junos-vlan-id</code> predefined variable to dynamically obtain the VLAN ID.
----------------	--



**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 (optional) and a valid VLAN identifier in the format *tpid.vlan-id*. When used in the **dynamic-profiles** hierarchy, specify the `$junos-stacked-vlan-id` predefined variable.

**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.

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

<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Dual VLAN Tags</a></li> <li>• <a href="#">stacked-vlan-tagging on page 352</a></li> </ul>
------------------------------	--

## vlan-tags (Stacked VLAN Tags)

<b>Syntax</b>	<code>vlan-tags inner <i>tpid.vlan-id</i> inner-list <i>value</i> inner-range <i>vid1—vid2</i> outer <i>tpid.vlan-id</i>;</code>
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.
<b>Description</b>	For Gigabit Ethernet IQ and IQE interfaces only, bind TPIDs and 802.1Q VLAN tag IDs to a logical interface.
<b>Options</b>	<b>inner <i>tpid.vlan-id</i></b> —A TPID and a valid VLAN identifier. <b>Range:</b> (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. <b>inner-list <i>value</i></b> — List or a set of VLAN identifiers.



**NOTE:** This is supported on MX Series with Trio based FPCs.

**inner-range *tpid. vid1—vid2***—Specify a TPID and a range of VLAN IDs where *vid1* is the start of the range and *vid2* 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;  
}
```

---

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Dual VLAN Tags</i></li><li>• <i>Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers</i></li><li>• <a href="#">stacked-vlan-tagging on page 352</a></li></ul>

## vpls (Routing Instance)

```
Syntax  vpls {
    mac-pinning;
    active-interface {
        any;
        primary interface-name;
    }
    community COMM;
    connectivity-type (ce | irb);
    control-word;
    encapsulation-type ethernet;
    ignore-encapsulation-mismatch;
    ignore-mtu-mismatch;
    import-labeled-routes [ routing-instance-name ];
    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 {
        interface interface-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 hot-standby-vc-on;
        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. <b>mac-flush</b> option introduced in Junos OS Release 10.0. <b>hot-standby-vc-on</b> , <b>import-labeled-routes</b> [ <i>routing-instance-name</i> ], and <b>interface</b> <i>interface</i> options introduced in Junos OS Release 15.1R2.
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"><li>• <i>Configuring VPLS Routing Instances</i></li></ul>

## vrf-export

<b>Syntax</b>	<code>vrf-export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<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>
<b>Release Information</b>	<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>
<b>Description</b>	<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 <b>vrf</b> is specified for the <b>instance-type</b> 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>
<b>Default</b>	If the instance-type is <b>vrf</b> , <b>vrf-export</b> is a required statement. The default action is to reject.
<b>Options</b>	<i>policy-names</i> —Names for the export policies.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Implementing EVPN with VXLAN for Data Centers</i></li> <li>• <a href="#">instance-type on page 304</a></li> <li>• <i>Configuring Policies for the VRF Table on PE Routers in VPNs</i></li> </ul>

## vrf-import

---

<b>Syntax</b>	<code>vrf-import [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<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>
<b>Release Information</b>	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.
<b>Description</b>	<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 <b>vrf</b> is specified for the <b>instance-type</b> 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>
<b>Default</b>	If the instance type is <b>vrf</b> , <b>vrf-import</b> is a required statement. The default action is to accept.
<b>Options</b>	<i>policy-names</i> —Names for the import policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Implementing EVPN with VXLAN for Data Centers</i></li><li>• <a href="#">instance-type on page 304</a></li><li>• <i>Configuring Policies for the VRF Table on PE Routers in VPNs</i></li></ul>

## vrf-target

<b>Syntax</b>	<pre>vrf-target {     community;     auto     import community-name;     export community-name; }</pre>
<b>Hierarchy Level</b>	<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> protocols evpn vni-options],          [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>
<b>Release Information</b>	<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. <b>auto option</b> was also added at this time.</p>
<b>Description</b>	<p>Specify a virtual routing and forwarding (VRF) target community. If you configure the <b>community</b> option only, default VRF import and export policies are generated that accept and tag routes with the specified target community. The purpose of the <b>vrf-target</b> 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 <b>import</b> and <b>export</b> options.</p>
<b>Options</b>	<p><b>community</b>—Community name.</p> <p><b>auto</b>—Automatically derives the route target (RT) for QFX5100 switches.</p> <p><b>import community-name</b>—Allowed communities accepted from neighbors.</p> <p><b>export community-name</b>—Allowed communities sent to neighbors.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Implementing EVPN with VXLAN for Data Centers</i></li> <li>• <i>Configuring Policies for the VRF Table on PE Routers in VPNs</i></li> </ul>

- *Example: Configuring FEC 129 BGP Autodiscovery for VPWS*

## CHAPTER 12

# Operational Commands

- [AAA and Address Assignment Pool Commands on page 387](#)
- [ANCP and ANCP-Triggered Wholesale Commands on page 404](#)
- [DHCPv4 Local Server Commands on page 421](#)
- [DHCPv4 Relay Commands on page 435](#)
- [DHCPv6 Local Server Commands on page 449](#)
- [Interface Commands on page 461](#)
- [Subscriber Commands on page 561](#)
- [VPLS Commands on page 587](#)

### 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	<b>show network-access aaa statistics</b> <b>&lt;accounting (detail)&gt;</b> <b>&lt;address-assignment (client   pool <i>pool-name</i>)&gt;</b> <b>&lt;dynamic-requests&gt;</b> <b>&lt;radius&gt;</b>
Release Information	Command introduced in Junos OS Release 9.1. Option <b>address-assignment</b> introduced in Junos OS Release 10.0. Option <b>radius</b> introduced in Junos OS Release 11.4. Option <b>detail</b> introduced in Junos OS Release 13.3.
Description	Display AAA accounting, address-assignment, dynamic request statistics, and RADIUS settings and statistics.
Options	<b>accounting (detail)</b> —(Optional) Display AAA accounting statistics. The <b>detail</b> keyword displays additional accounting information  <b>address-assignment (client   pool <i>pool-name</i>)</b> —(Optional) Display AAA address-assignment client and pool statistics.  <b>dynamic-requests</b> —(Optional) Display AAA dynamic requests.  <b>radius</b> —(Optional) Display RADIUS settings and statistics.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"><li>• <i>Verifying and Managing Subscriber AAA Information</i></li></ul>
List of Sample Output	<a href="#">show network-access aaa statistics accounting on page 394</a> <a href="#">show network-access aaa statistics accounting detail on page 395</a> <a href="#">show network-access aaa statistics address-assignment client on page 395</a> <a href="#">show network-access aaa statistics address-assignment pool on page 395</a> <a href="#">show network-access aaa statistics dynamic-requests on page 396</a> <a href="#">show network-access aaa statistics radius on page 396</a>
Output Fields	<a href="#">Table 14 on page 389</a> lists the output fields for the <b>show network-access aaa statistics</b> command. Output fields are listed in the approximate order in which they appear.

Table 14: show network-access aaa statistics Output Fields

Field Name	Field Description	Level of Output
<b>Requests received</b>	<ul style="list-style-type: none"> <li>Number of accounting requests generated by the AAA framework.</li> <li>Number of dynamic requests received from the external server.</li> </ul> <p>Does not include requests sent from backup accounting.</p>	All levels
<b>Accounting request failures</b>	<p>Number of accounting requests that failed to be sent or queued from a client to a RADIUS accounting server.</p> <p>Does not include requests sent from backup accounting.</p>	detail
<b>Accounting request success</b>	<p>Number of accounting requests successfully sent or queued from a client to a RADIUS accounting server.</p> <p>Does not include requests sent from backup accounting.</p>	detail
<b>Account on requests</b>	<p>Number of accounting on requests sent from a client to a RADIUS accounting server.</p>	detail
<b>Accounting start requests</b>	<p>Number of accounting start requests sent from a client to a RADIUS accounting server.</p>	detail
<b>Accounting interim requests</b>	<p>Number of accounting interim requests sent from a client to a RADIUS accounting server.</p>	detail
<b>Accounting stop requests</b>	<p>Number of accounting stop requests sent from a client to a RADIUS accounting server.</p> <p>Does not include requests sent from backup accounting.</p>	detail
<b>Accounting request timeouts</b>	<p>Number of accounting requests to the accounting server that timed out. This field was named <b>Timed out requests</b> in releases before Junos OS Release 16.1.</p> <p>Does not include requests sent from backup accounting.</p>	All levels
<b>Accounting Response failures</b>	<p>Number of accounting requests not acknowledged (NAK) by the accounting server.</p> <p>Does not include requests sent from backup accounting.</p>	All levels

Table 14: show network-access aaa statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Accounting response success</b>	Number of accounting requests acknowledged by the accounting server.  Does not include requests sent from backup accounting.	All levels
<b>Account on responses</b>	Number of accounting on requests acknowledged by the RADIUS accounting server.	<b>detail</b>
<b>Accounting start responses</b>	Number of accounting start requests acknowledged by the RADIUS accounting server.	<b>detail</b>
<b>Accounting interim responses</b>	Number of accounting interim requests acknowledged by the RADIUS accounting server.	<b>detail</b>
<b>Accounting stop responses</b>	Number of accounting stop requests acknowledged by the RADIUS accounting server.  Does not include requests sent from backup accounting.	<b>detail</b>
<b>Accounting rollover requests</b>	Number of accounting requests coming to a RADIUS accounting server after a previous server timing out.	<b>detail</b>
<b>Accounting unknown requests</b>	Number of unknown accounting requests sent from a client to a RADIUS accounting server (for example, when the header has invalid or unsupported information).	<b>detail</b>
<b>Accounting radius pending requests</b>	Number of accounting requests sent from a client to a RADIUS accounting server that are waiting for a response from the server.	<b>detail</b>
<b>Accounting malformed responses</b>	Number of accounting responses from a RADIUS accounting server that have invalid or unexpected attributes.	<b>detail</b>
<b>Accounting retransmissions</b>	Number of accounting requests made by a client to the RADIUS sever that were retransmitted.  Does not include requests sent from backup accounting.	<b>detail</b>
<b>Accounting bad authenticators</b>	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).	<b>detail</b>

Table 14: show network-access aaa statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Accounting packets dropped</b>	Number of accounting responses from a RADIUS accounting server that are dropped by a client.	<b>detail</b>
<b>Accounting backup record creation requests</b>	Number of accounting stop requests from a client to a RADIUS accounting server that were forwarded to be backed up.	<b>detail</b>
<b>Accounting backup replay request success</b>	Number of backup accounting stop requests successfully created by clients after each timeout for replay to a RADIUS accounting server.	<b>detail</b>
<b>Accounting backup request failures</b>	Number of backup accounting requests that failed to be sent or queued from a client to a RADIUS accounting server.	<b>detail</b>
<b>Accounting backup request success</b>	Number of backup accounting requests successfully sent or queued from a client to a RADIUS accounting server.	<b>detail</b>
<b>Accounting backup timeouts</b>	Number of backup accounting requests that timed out after being sent to a RADIUS accounting server.	<b>detail</b>
<b>Accounting backup in-flight requests</b>	<p>Number of backup accounting requests that were successfully sent or queued to a RADIUS accounting server for which no response or error has been received yet.</p> <p>Backup requests are replayed only in the following circumstances:</p> <ul style="list-style-type: none"> <li>• When the request being replayed receives a positive response, the next request can be replayed.</li> <li>• When the request being replayed receives a timeout response, it can be replayed again.</li> </ul> <p>Consequently this intermediate timer displays 1 or 0. The value eventually drops to 0 as requests are responded to positively or fail due to error.</p>	<b>detail</b>
<b>Accounting backup responses success</b>	Number of backup records that were successfully acknowledged with a positive response from a RADIUS accounting server.	<b>detail</b>

Table 14: show network-access aaa statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Accounting backup radius requests</b>	<p>Number of backup requests sent to UDP level.</p> <p>This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. An observation that the value is increasing is more significant than the exact value of the counter.</p>	<b>detail</b>
<b>Accounting backup radius responses</b>	<p>Number of responses received at the UDP level for backup requests.</p> <p>This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. Observation that the value is increasing is more significant than the exact value of the counter.</p>	<b>detail</b>
<b>Accounting backup radius timeouts</b>	<p>Number of backup requests that timed out after being sent to UDP.</p> <p>This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. Observation that the value is increasing is more significant than the exact value of the counter.</p>	<b>detail</b>
<b>Accounting backup radius pending requests</b>	<p>Number of backup requests sent to a RADIUS accounting server that are waiting for a response from the server.</p> <p>This is an intermediate state counter that eventually drops to zero as requests are responded to or failed due to error.</p>	<b>detail</b>
<b>Accounting backup radius retransmissions</b>	<p>Sum of backup request retransmissions for each RADIUS accounting server.</p> <p>This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. Observation that the value is increasing is more significant than the exact value of the counter.</p>	<b>detail</b>
<b>Accounting backup malformed responses</b>	Sum of malformed responses received for backup requests sent to each RADIUS accounting server at the UDP level.	<b>detail</b>
<b>Accounting backup bad authenticators</b>	Sum of responses received for backup accounting requests for each RADIUS accounting server where authenticators were mismatched.	<b>detail</b>

Table 14: show network-access aaa statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Accounting backup responses dropped</b>	Sum of responses for backup accounting requests for each RADIUS accounting server that were dropped due to various sanity checks.	<b>detail</b>
<b>Accounting backup rollover requests</b>	Sum of backup accounting requests rolled over for each RADIUS accounting server.	<b>detail</b>
<b>Accounting backup unknown responses</b>	Sum of unknown responses for backup accounting requests for each RADIUS accounting server.	<b>detail</b>
<b>Client</b>	Client type; for example, DHCP, Mobile IP, PPP.	none specified
<b>Out of Memory</b>	Number of times an address was not given to the client due to memory issues.	none specified
<b>No Matches</b>	Number of times there were no network matches for the pool.	none specified
<b>Pool Name</b>	Name of the address-assignment pool for this client.	none specified
<b>Out of Addresses</b>	Number of times there were no available addresses in the pool.	none specified
<b>Address total</b>	Number of addresses in the pool.	none specified
<b>Addresses in use</b>	Number of addresses in use.	none specified
<b>Address Usage (percent)</b>	Percentage of total addresses in use.	none specified
<b>Pool Usage</b>	Percentage of allocated addresses in the specified address pool.	none specified
<b>processed successfully</b>	Number of dynamic requests processed successfully by the AAA framework.	All levels
<b>errors during processing</b>	Number of dynamic requests that resulted in processing errors by the AAA framework.	All levels
<b>Link Name</b>	Name of the secondary address-assignment pool to which the primary pool is linked.	

Table 14: show network-access aaa statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>silently dropped</b>	Number of dynamic requests dropped by the AAA framework due to multiple back-to-back or duplicate requests.	All levels
<b>RADIUS Server</b>	IPv4 or IPv6 address of the RADIUS server to which the router is sending requests.	All levels
<b>Profile</b>	Name of the RADIUS profile associated with the RADIUS server. A RADIUS server can be associated with more than one RADIUS profile.	All levels
<b>Configured</b>	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.	All levels
<b>Current</b>	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.	All levels
<b>Peak</b>	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.  <b>NOTE:</b> If the value of this field is equal to the value of the <b>Configured</b> field, you may want to increase the value of the <b>Configured</b> field.	All levels
<b>Exceeded</b>	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.  <b>NOTE:</b> If the value of this field is nonzero, you may want to increase the value of the <b>Configured</b> field.	All levels

## Sample Output

### show network-access aaa statistics accounting

```

user@host> show network-access aaa statistics accounting
Accounting module statistics
Accounting module statistics
Requests received: 5000
Accounting request timeouts: 2000
Accounting response failures: 0
Accounting response success: 3000

```

**show network-access aaa statistics accounting detail**

```

user@host> show network-access aaa statistics accounting detail
Accounting module statistics
Accounting module statistics
  Requests received: 5000
  Accounting request failures: 0
  Accounting request success: 5000
    Account on requests: 0
    Accounting start requests: 3000
    Accounting interim requests: 0
    Accounting stop requests: 2000
  Accounting request timeouts: 2000
  Accounting response failures: 0
  Accounting response success: 3000
    Account on responses: 0
    Accounting start responses: 3000
    Accounting interim responses: 0
    Accounting stop responses: 0
  Accounting rollover requests: 0
  Accounting unknown responses: 0
  Accounting radius pending requests: 0
  Accounting malformed responses: 0
  Accounting retransmissions: 6000
  Accounting bad authenticators: 0
  Accounting packets dropped: 0

  Accounting backup record creation requests: 3000
  Accounting backup request replay success: 9808
  Accounting backup request failures: 0
  Accounting backup request success: 3006
  Accounting backup timeouts: 6
  Accounting backup in-flight requests: 0
  Accounting backup responses success: 3000
  Accounting backup radius requests: 3006
  Accounting backup radius responses: 3000
  Accounting backup radius timeouts: 99
  Accounting backup radius pending requests: 0
  Accounting backup radius retransmissions: 99
  Accounting backup malformed responses: 0
  Accounting backup bad authenticators: 0
  Accounting backup responses dropped: 0
  Accounting backup rollover requests: 0
  Accounting backup unknown responses: 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
198.51.100.239     prof1        1000         0         1000   14
                   prof2        500          17        432    0
198.51.100.211     myprof       200          0         200    27
203.0.113.254      pppoe-auth   111          0         1       0
2001:db8:0:f101::2 xyz-profile11 1000        10        135    0
```

## show network-access aaa statistics authentication

<b>Syntax</b>	<b>show network-access aaa statistics authentication</b> <b>&lt;detail&gt;</b>
<b>Release Information</b>	Command introduced in Junos OS Release 9.1. Option <b>detail</b> introduced in Junos OS Release 12.1.
<b>Description</b>	Display AAA authentication statistics.
<b>Options</b>	<b>detail</b> —(Optional) Displays detailed information about authentication.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Verifying and Managing Subscriber AAA Information</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show network-access aaa statistics authentication on page 399</a> <a href="#">show network-access aaa statistics authentication detail on page 399</a>
<b>Output Fields</b>	Table 15 on page 397 lists the output fields for the <b>show network-access aaa statistics authentication</b> command. Output fields are listed in the approximate order in which they appear.

Table 15: 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 15: show network-access aaa statistics authentication Output Fields** (*continued*)

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

<b>Syntax</b>	<b>show network-access aaa subscribers</b> <b>&lt;logical-system <i>logical-system-name</i>&gt;</b> <b>&lt;routing-instance <i>routing-instance-name</i>&gt;</b> <b>&lt;statistics&gt;</b> <b>&lt;username&gt;</b>
<b>Release Information</b>	Command introduced in Junos OS Release 9.1.
<b>Description</b>	Display subscriber-specific AAA statistics.
<b>Options</b>	<p><b>logical-system <i>logical-system-name</i></b>—(Optional) List subscribers in the specific logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(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><b>statistics</b>—(Optional) Display statistics for the subscriber events.</p> <p><b>username</b>—(Optional) Display information for the specified subscriber.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Verifying and Managing Subscriber AAA Information</i></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show network-access aaa subscribers logical-system on page 401</a></p> <p><a href="#">show network-access aaa subscribers logical-system routing-instance on page 401</a></p> <p><a href="#">show network-access aaa subscribers statistics username on page 402</a></p> <p><a href="#">show network-access aaa subscribers username on page 402</a></p>
<b>Output Fields</b>	<a href="#">Table 16 on page 400</a> lists the output fields for the <b>show network-access aaa subscribers</b> command. Output fields are listed in the approximate order in which they appear.

**Table 16: 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 16: show network-access aaa subscribers Output Fields (*continued*)

Field Name	Field Description
<b>INTERIM sent successfully</b>	Number of accounting interim requests generated by the AAA framework for this subscriber.
<b>INTERIM send failures</b>	Number of accounting interim requests that failed to make it to the accounting server for this subscriber.
<b>INTERIM ack received</b>	Number of accounting interim requests acknowledged by the accounting server for this subscriber.
<b>Requests received</b>	Number of reauthentication requests received by the authentication server.
<b>Successful responses</b>	Number of successful reauthentication requests granted by the authentication server.
<b>Aborts handled</b>	Number of reauthentication requests aborted by the authentication server.
<b>Service name</b>	Name of the subscriber service.
<b>Creation requests</b>	Number of requests to create the service.
<b>Deletion requests</b>	Number of requests to delete the service.
<b>Request timeouts</b>	Number of times the service request was timed out.
<b>Client type</b>	Type of client; for example, DHCP, Mobile IP, PPP.
<b>Session-ID</b>	ID of the subscriber session.
<b>Session uptime</b>	How long the session has been up, in <i>HH:MM:SS</i> .
<b>Accounting</b>	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
user61@example.net  ppp             default
00010e020304.1231  dhcp            isp-bos-metro-12:isp-cmbrg-12
user54@example.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
user54@example.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 user80@example.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

<b>Syntax</b>	<code>show network-access address-assignment pool <i>pool-name</i></code> <code>&lt;logical-system <i>logical-system-name</i>&gt;</code> <code>&lt;routing-instance <i>routing-instance-name</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0.
<b>Description</b>	Display state information for each address-assignment pool.
<b>Options</b>	<p><b>none</b>—Display information about clients that have obtained addresses from the address-assignment pool.</p> <p><b>pool <i>pool-name</i></b>—Display information about the specified address-assignment pool.</p> <p><b>logical-system <i>logical-system-name</i></b>—(Optional) Perform this operation on the specified logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Perform this operation on the specified routing instance.</p>
<b>Required Privilege Level</b>	view and system
<b>List of Sample Output</b>	<a href="#">show network-access address-assignment pool on page 403</a>
<b>Output Fields</b>	<a href="#">Table 17 on page 403</a> lists the output fields for the <b>show network-access address-assignment pool</b> command. Output fields are listed in the approximate order in which they appear.

**Table 17: 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     00:00:5e:00:53:01 DHCP
192.168.2.2     00:00:5e:00:53:02 DHCP
192.168.2.3     00:00:5e:00:53:03 DHCP
192.168.2.4     00:00:5e:00:53:04 DHCP

```

## ANCP and ANCP-Triggered Wholesale Commands

---

- `clear ancp access-loop`
- `clear ancp neighbor`
- `clear network-access aaa subscriber`
- `request ancp oam port-down`
- `request ancp oam port-up`
- `request auto-configuration reconnect-pending`
- `show ancp neighbor`
- `show auto-configuration out-of-band pending`

## clear ancp access-loop

<b>Syntax</b>	clear ancp access-loop (neighbor <i>ip-address</i>   subscriber-interface <i>physical-interface-name</i> ) circuit-id <i>aci</i> remote-id <i>ari</i> outer-vlan-id <i>vlan-id</i>
<b>Release Information</b>	Command introduced in Junos OS Release 16.1.
<b>Description</b>	Clear the connection for the subscriber on the specified access loop for an ANCP-triggered, autosensed dynamic VLAN. The autoconfiguration daemon (autoconfd) deletes any existing cached information about the subscriber. This command simulates a CPE connection reset as seen when the access node sends a Port Down message followed by a Port Up message.
<b>Options</b>	<p><i>aci</i>—ANCP Access-Loop-Circuit-ID TLV that identifies the subscriber-side access loop logical port and partially identifies an access loop to clear.</p> <p><i>ari</i>—ANCP Access-Loop-Remote-ID TLV that uniquely identifies the subscriber on the access loop and partially identifies an access loop to clear.</p> <p><i>ip-address</i>—ANCP neighbor's IP address that specifies an access loop to clear.</p> <p><i>physical-interface-name</i>—Subscriber interface that specifies an access loop to clear.</p> <p><i>vlan-id</i>—ANCP Access-Aggregation-Circuit-ID-Binary TLV that identifies the logical circuit identifier on the NAS side and partially identifies an access loop to clear.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Clearing ANCP Access Loops on page 160</a></li> <li>• <a href="#">Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages on page 157</a></li> <li>• <a href="#">Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear ancp access-loop on page 405</a>
<b>Output Fields</b>	When you enter this command, you are provided no feedback on the status of your request. You can enter the <b>show ancp neighbor</b> command before and after clearing the access loop to verify the clear operation.

## Sample Output

### clear ancp access-loop

```
user@host> clear ancp-access-loop neighbor 192.168.25.31 circuit-id line-aci-1 remote-id line-ari-1
outer-vlan-id 126
```

## clear ancp neighbor

---

<b>Syntax</b>	<b>clear ancp neighbor</b> <b>&lt;ip-address <i>ip-address</i>&gt;</b> <b>&lt;system-name <i>mac-address</i>&gt;</b>
<b>Release Information</b>	Command introduced in Junos OS Release 9.4.
<b>Description</b>	<p>Clear the ANCP agent connection with all ANCP neighbors or with the specified ANCP neighbor. This command deletes information for subscribers associated with the neighbor, causing the adjusted traffic rates to revert to the configured rate for the subscriber interfaces. The neighbor remains configured (its administrative state is <i>enabled</i>) and can reestablish adjacencies.</p> <p>This command initiates logout of ANCP-triggered dynamic VLAN sessions on the physical interface associated with the specified neighbor; conventionally autosensed dynamic VLAN sessions and their associated logical interfaces are not affected.</p>
<b>Options</b>	<p><b>none</b>—Clear all ANCP neighbors.</p> <p><b>ip-address <i>ip-address</i></b>—(Optional) Clear the ANCP neighbor specified by the IP address.</p> <p><b>system-name <i>mac-address</i></b>—(Optional) Clear the ANCP neighbor specified by the MAC address.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show ancp neighbor on page 413</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear ancp neighbor on page 406</a> <a href="#">show ancp neighbor on page 406</a>
<b>Output Fields</b>	When you enter this command, you are provided no feedback on the status of your request. You can enter the <b>show ancp neighbor</b> command before and after clearing the ANCP neighbors to verify the clear operation.

## Sample Output

### clear ancp neighbor

```
user@host> clear ancp neighbor
```

### show ancp neighbor

The following sample output displays the connections with ANCP neighbors before and after the **clear ancp neighbor** command was issued.

```
user@host> show ancp neighbor
```

IP Address	MAC Address	State	Subscriber	Capabilities
------------	-------------	-------	------------	--------------

			Count	
203.0.113.102	00:00:5e:00:53:10	Established	5	Topo
203.0.113.122	00:00:5e:00:53:12	Established	5	Topo
203.0.113.132	00:00:5e:00:53:13	Established	5	Topo
203.0.113.142	00:00:5e:00:53:14	Established	5	Topo

```
user@host> clear ancp neighbor ip-address 203.0.113.102
```

```
user@host> show ancp neighbor
```

IP Address	MAC Address	State	Subscriber Count	Capabilities
203.0.113.122	00:00:5e:00:53:12	Established	5	Topo
203.0.113.132	00:00:5e:00:53:13	Established	5	Topo
203.0.113.142	00:00:5e:00:53:14	Established	5	Topo

## clear network-access aaa subscriber

---

<b>Syntax</b>	<code>clear network-access aaa subscriber</code> <code>&lt;session-id <i>identifier</i> &lt;reconnect&gt;&gt;</code> <code>&lt;statistics username <i>username</i>&gt;</code> <code>&lt;username <i>username</i> &lt;reconnect&gt;&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.1. <b>reconnect</b> and <b>session-id</b> options added in Junos OS Release 16.1.
<b>Description</b>	Clear AAA subscriber statistics and log out subscribers. You can log out subscribers based on the username or on the subscriber session identifier. Use the session identifier when more than one session has the same username string.
<b>Options</b>	<b>reconnect</b> —(Optional) Reconnect as a Layer 2 wholesale session when the subscriber session has been fully logged out. This option is equivalent to issuing a RADIUS-initiated disconnect with reconnect semantics; that is, when the message includes Acct-Terminate-Cause (RADIUS attribute 49) with a value of callback (16).  <b>session-id <i>identifier</i></b> —(Optional) Log out the subscriber based on the subscriber session identifier.  <b>statistics username <i>username</i></b> —(Optional) Clear AAA subscriber statistics and log out the subscriber.  <b>username <i>username</i></b> —(Optional) Log out the AAA subscriber.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Verifying and Managing Subscriber AAA Information</i></li></ul>
<b>List of Sample Output</b>	<a href="#">clear network-access aaa subscriber statistics username on page 408</a> <a href="#">clear network-access aaa subscriber username on page 408</a> <a href="#">clear network-access aaa subscriber username on page 408</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

clear network-access aaa subscriber statistics username

```
user@host> clear network-access aaa subscriber statistics username user22@example.com
```

clear network-access aaa subscriber username

```
user@host> clear network-access aaa subscriber username user22@example.com
```

clear network-access aaa subscriber session-id

```
user@host> clear network-access aaa subscriber session-id 18367425
```

## request ancp oam port-down

<b>Syntax</b>	<b>request ancp oam port-down</b> (neighbor <i>ip-address</i>   subscriber-interface <i>physical-interface-name</i> ) circuit-id <i>aci</i> remote-id <i>ari</i> outer-vlan-id <i>vlan-id</i>
<b>Release Information</b>	Command introduced in Junos OS Release 16.1.
<b>Description</b>	Simulate an ANCP Port Down message on the specified access loop for troubleshooting or to mitigate an abnormal condition. Triggers removal of the corresponding out-of-band triggered, autosensed dynamic VLAN session for which no ANCP-sourced information exists. You must specify an ACI, an ARI, and an outer VLAN tag. This command is overridden by a genuine ANCP Port-Up message, meaning that you cannot use this command to initiate a Port Down condition when the access node has already reported a Port Up condition.
<b>Options</b>	<p><b>aci</b>—ANCP Access-Loop-Circuit-ID TLV that corresponds to a subscriber interface on the access node; used to identify the access node from which the message is simulated.</p> <p><b>ari</b>—ANCP Access-Loop-Remote-ID TLV that identifies the subscriber associated with an interface on the access node; used to identify the access node from which the message is simulated.</p> <p><b>ip-address</b>—IP address that specifies the access node from which the message is simulated.</p> <p><b>physical-interface-name</b>—Name of the access-facing subscriber interface that specifies the access node on whose local loop the loopback test is run.</p> <p><b>vlan-id</b>—ANCP Access-Aggregation-Circuit-ID-Binary TLV, the outer VLAN tag inserted by the access node on upstream traffic; used to identify the access node from which the message is simulated.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages on page 157</a></li> <li>• <a href="#">Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">request ancp oam port-down neighbor circuit-id remote-id outer-vlan-id on page 410</a>
<b>Output Fields</b>	When you enter this command, you are provided no feedback on the status of your request. You can enter the <b>show ancp neighbor detail</b> , <b>show subscribers client-type vlan-oob detail</b> , and the <b>show subscribers summary</b> commands before and after initiating the Port Down message to verify the operation.

## Sample Output

request ancp oam port-down neighbor circuit-id remote-id outer-vlan-id

```
user@host> request ancp oam port-down neighbor 192.168.25.31 circuit-id line-aci-1 remote-id  
line-ari-1 outer-vlan-id 126
```

## request ancp oam port-up

<b>Syntax</b>	<b>request ancp oam port-up</b> (neighbor <i>ip-address</i>   subscriber-interface <i>physical-interface-name</i> ) circuit-id <i>aci</i> remote-id <i>ari</i> outer-vlan-id <i>vlan-id</i>
<b>Release Information</b>	Command introduced in Junos OS Release 16.1.
<b>Description</b>	Simulate an ANCP Port Up message on the specified access loop for troubleshooting or to mitigate an abnormal condition. You must specify an ACI, an ARI, and an outer VLAN tag. This command is overridden by a genuine ANCP Port Down message, meaning that you cannot use this command to initiate a Port Up condition when the access node has already reported a Port Down condition.
<b>Options</b>	<p><b>aci</b>—ANCP Access-Loop-Circuit-ID TLV that corresponds to a subscriber interface on the access node; used to identify the access node from which the message is simulated.</p> <p><b>ip-address</b>—IP address that specifies the access node from which the message is simulated.</p> <p><b>ari</b>—ANCP Access-Loop-Remote-ID TLV that identifies the subscriber associated with an interface on the access node; used to identify the access node from which the message is simulated.</p> <p><b>physical-interface-name</b>—Name of the access-facing subscriber interface that specifies the access node on whose local loop the loopback test is run.</p> <p><b>vlan-id</b>—ANCP Access-Aggregation-Circuit-ID-Binary TLV, the outer VLAN tag inserted by the access node on upstream traffic; used to identify the access node from which the message is simulated.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages on page 157</a></li> <li>• <a href="#">Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">request ancp oam port-up neighbor circuit-id remote-id outer-vlan-id on page 411</a>
<b>Output Fields</b>	When you enter this command, you are provided no feedback on the status of your request. You can enter the <b>show ancp neighbor detail</b> , <b>show subscribers client-type vlan-oob detail</b> , and the <b>show subscribers summary</b> commands before and after initiating the Port Up message to verify the operation.

## Sample Output

### request ancp oam port-up neighbor circuit-id remote-id outer-vlan-id

```
user@host> request ancp oam port-up neighbor 192.168.25.31 circuit-id line-aci-1 remote-id
line-ari-1 outer-vlan-id 126
```

## request auto-configuration reconnect-pending

---

<b>Syntax</b>	request auto-configuration reconnect-pending
<b>Release Information</b>	Command introduced in Junos OS Release 16.1.
<b>Description</b>	Initiate reestablishment of Layer 2 wholesale sessions that correspond to access lines that are in the pending state. Ordinarily, the most likely situations with pending sessions are handled automatically. This statement is intended to be used only when an uncommon condition might prevent automatic reestablishment. This command has no effect when no pending sessions are present.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Reestablishing Pending Access Line Sessions for Layer 2 Wholesale on page 160</a></li><li>• <a href="#">Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93</a></li></ul>
<b>List of Sample Output</b>	<a href="#">request auto-configuration reconnect-pending on page 412</a>
<b>Output Fields</b>	When you enter this command, you are provided no feedback on the status of your request. You can enter the <b>show ancp neighbor detail</b> , <b>show subscribers client-type vlan-oob detail</b> , and the <b>show subscribers summary</b> commands before and after initiating the Port Up message to verify the operation.

### Sample Output

#### request auto-configuration reconnect-pending

```
user@host> request auto-configuration reconnect-pending
```

## show ancp neighbor

<b>Syntax</b>	show ancp neighbor <brief   detail> <ip-address <i>ip-address</i> <system-name <i>mac-address</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 9.4.
<b>Description</b>	Display information about all ANCP neighbors or the specified ANCP neighbor, regardless of operational state.
<b>Options</b>	<p><b>brief   detail</b>—(Optional) Display the specified level of detail.</p> <p><b>ip-address <i>ip-address</i></b>—(Optional) Display information about the neighbor (access node) specified by the IP address.</p> <p><b>system-name <i>mac-address</i></b>—(Optional) Display information about the neighbor (access node) specified by the MAC address.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>show ancp cos</i></li> <li>• <i>show ancp subscriber</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ancp neighbor on page 416</a> <a href="#">show ancp neighbor detail on page 416</a> <a href="#">show ancp neighbor ip-address on page 418</a> <a href="#">show ancp neighbor system-name on page 419</a>
<b>Output Fields</b>	Table 18 on page 413 lists the output fields for the <b>show ancp neighbor</b> command. Output fields are listed in the approximate order in which they appear.

Table 18: show ancp neighbor Output Fields

Field Name	Field Description	Level of Output
Version	Version of the ANCP implementation: <ul style="list-style-type: none"> <li>• <b>0x31</b>—General Switch Management Protocol (GSMP) version 3, sub-version 1; ANCP version before <i>RFC 6320, Protocol for Access Node Control Mechanism in Broadband Networks</i>.</li> <li>• <b>0x32</b>—ANCP version 1, defined in <i>RFC 6320, Protocol for Access Node Control Mechanism in Broadband Networks</i>.</li> </ul>	<b>brief detail</b> none
IP Address	IP address of the ANCP neighbor.	<b>brief detail</b> none
Partid	Number that associates the ANCP message with a specific partition.	<b>brief</b> none

Table 18: show ancp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>State</b>	Operational state of the ANCP adjacency: <ul style="list-style-type: none"> <li>Configured—The neighbor has been configured, but has never been in the <b>Established</b> state. An asterisk (*) is prefixed to the neighbor entry for this state.</li> <li>Establishing—Adjacency negotiations are in progress for the neighbor. An asterisk (*) is prefixed to the neighbor entry for this state. This state is rarely seen because the adjacency is established so quickly.</li> <li><b>Established</b>—Adjacency negotiations have succeeded for the neighbor and an ANCP session has been established.</li> <li>Not Estblshed—Not Established; adjacency negotiations are ready to begin. Indicates that this neighbor previously had been in the <b>Established</b> state; that is, it has lost a previously established adjacency. An asterisk (*) is prefixed to the neighbor entry for this state.</li> </ul>	All levels
<b>Time</b>	How long the adjacency has been up in one of the following formats: <ul style="list-style-type: none"> <li><i>nwndnh</i>—number of weeks, days, and hours</li> <li><i>nd hh:mm:ss</i>—number of days, hours, minutes, and seconds</li> </ul>	<b>brief detail</b> none
<b>Subscriber Count</b>	Number of subscribers associated with the ANCP neighbor (access local loop).	<b>brief</b> none
<b>Capabilities</b>	Negotiated ANCP capability: <ul style="list-style-type: none"> <li><b>Topo</b>—Topology discovery.</li> <li><b>OAM</b>—Performance of local Operations Administration Maintenance (OAM) procedures on an access loop controlled by the router.</li> </ul>	All levels
<b>System Name</b>	MAC address of the ANCP neighbor.	<b>detail</b>
<b>TCP Port</b>	TCP port on which ANCP messages are exchanged.	<b>detail</b>
<b>System Instance</b>	Number identifying the ANCP link instance from the edge device's perspective.	<b>detail</b>
<b>Peer Instance</b>	Number identifying the ANCP instance from the access node's perspective. This number is unique and changes when the node or link comes back up after going down.	<b>detail</b>
<b>Timer</b>	Adjacency timer value advertised by the ANCP peer in 100 ms increments; the interval between ANCP ACK messages. This value remains constant for the duration of an ANCP session.	<b>detail</b>
<b>Partition Type</b>	Number that identifies whether partitions are used and how the ID is negotiated: <ul style="list-style-type: none"> <li><b>0</b>—No partition.</li> <li><b>1</b>—Fixed partition requested.</li> <li><b>2</b>—Fixed partition assigned.</li> </ul>	<b>detail</b>
<b>Partition Flag</b>	Number that specifies the type of partition requested: 1 (new adjacency) or 2 (recovered adjacency).	<b>detail</b>

Table 18: show ancp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Partition Identifier</b>	Number that associates the ANCP message with a specific partition.	<b>detail</b>
<b>Partition Adjacencies</b>	Number of adjacencies that share the partition.	<b>detail</b>
<b>Dead Timer</b>	Remaining period that the edge device waits for adjacency packets from a neighbor before declaring the neighbor to be down. The maximum dead time value is three times the configured adjacency timer value. This field displays the current value based on the time that the last adjacency packet was received.	<b>detail</b>
<b>Received Syn Count</b>	Number of synchronization messages received from neighbors to maintain adjacencies.	<b>detail</b>
<b>Received Synack Count</b>	Number of synchronization acknowledgment messages received from neighbors in response to the node's synchronization messages.	<b>detail</b>
<b>Received Rstack Count</b>	Number of messages received from neighbors indicating that the link to the neighbor needs to be reset.	<b>detail</b>
<b>Received Ack Count</b>	Number of acknowledgment messages periodically received from neighbors after an adjacency has been established.	<b>detail</b>
<b>Received Port Up Count</b>	Number of status messages received from neighbors indicating that a port has transitioned to the up state.	<b>detail</b>
<b>Received Port Down Count</b>	Number of status messages received from neighbors indicating that a port has transitioned to the down state.	<b>detail</b>
<b>Received Generic Resp Count</b>	Number of generic response messages received from neighbors.	<b>detail</b>
<b>Received Adjacency Update Count</b>	Number of adjacency update messages received from neighbors.	<b>detail</b>
<b>Received OAM Count</b>	Number of OAM responses received from neighbors in reply to <b>request</b> commands.	<b>detail</b>
<b>Received Other Count</b>	Number of all other ANCP message packets received from neighbors that do not fit into one of the other categories.	<b>detail</b>
<b>Sent Syn Count</b>	Number of synchronization messages sent to neighbors to maintain adjacencies.	<b>detail</b>
<b>Sent Synack Count</b>	Number of synchronization acknowledgment messages sent to neighbors in response to the their synchronization messages.	<b>detail</b>
<b>Sent Rstack Count</b>	Number of messages sent to neighbors indicating that the link to the neighbor needs to be reset.	<b>detail</b>
<b>Sent Ack Count</b>	Number of acknowledgment messages periodically sent to neighbors after an adjacency has been established.	<b>detail</b>

Table 18: show ancp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Sent Generic Resp Count</b>	Number of generic response messages sent to neighbors.	<b>detail</b>
<b>Sent OAM Count</b>	Number of OAM <b>request</b> commands sent to neighbors.	<b>detail</b>
<b>Max Discovery Limit Exceed Count</b>	Number of times that the maximum number of discovery table entries accepted from the neighbor has been exceeded.	<b>detail</b>
<b>Result Codes</b>	<p>Number of generic response messages sent to neighbors that include each of the following result codes:</p> <ul style="list-style-type: none"> <li>• <b>Invalid Request Message Count</b>—A properly formed request message violated the protocol because of timing (such as a race condition) or direction of transmission.</li> <li>• <b>Specified Port(s) Down Count</b>—One or more of the specified ports are down because of a state mismatch between the router and an ANCP control application.</li> <li>• <b>Out of Resources Count</b>—ANCP is out of resources, probably not related to the access lines. This result code is sent only by an access node.</li> <li>• <b>Request Msg Not Implemented Count</b>—</li> <li>• <b>Malformed Msg Count</b>—Message is malformed because it was corrupted in transit or there was an implementation error at either end of the connection.</li> <li>• <b>TLV Missing Count</b>—One or more mandatory TLVs was missing from a request.</li> <li>• <b>Invalid TLV Contents Count</b>—The contents of one or more TLVs in the request do not match its required specification.</li> <li>• <b>Non-Existent Port(s) Count</b>—One or more of the ports specified in a request do not exist, possibly because of a configuration mismatch between the access node and the router or AAA.</li> </ul>	<b>detail</b>

## Sample Output

### show ancp neighbor

```

user@host> show ancp neighbor
  Version IP Address      PartID  State      Time      Subscriber
Capabilities
  0x31    203.0.113.13          0       Established  11:24      Count
Topo      2
  0x31    203.0.113.15          0       Not Estblshd  2:45      2
Topo
* 0x0     198.51.100.102        0       Establishing  0          0
* 0x0     192.0.2.0             0       Configured   0          0
* 0x0     192.0.2.1             0       Configured   0          0

```

### show ancp neighbor detail

```

user@host> show ancp neighbor detail
  Neighbor Information
    Version      : 0x31
    IP Address    : 192.0.2.85
    System Name   : 00:00:5e:00:53:01

```

```

Up Time : 26
TCP Port : 32666
State : Established
Subscriber Count : 4
Capabilities : Topo
System Instance : 2
Peer Instance : 20
Adjacency Timer (in 100ms) : 100
Peer Adjacency Timer (in 100ms) : 100
Partition Type : 0
Partition Flag : 1
Partition Identifier : 0
Partition Adjacencies : 0
Dead Timer : 23
Received Syn Count : 1
Received Synack Count : 1
Received Rstack Count : 0
Received Ack Count : 4
Received Port Up Count : 10
Received Port Down Count : 0
Received Generic Resp Count : 0
Received Adjacency Update Count : 0
Received OAM Count : 0
Received Other Count : 0
Sent Syn Count : 1
Sent Synack Count : 2
Sent Rstack Count : 0
Sent Ack Count : 3
Sent Generic Resp Count : 0
Sent OAM Count : 0
Max Discovery Limit Exceed Count : 0
Result Codes:
Invalid Request Message Count : 0
Specified Port(s) Down Count : 0
Out of Resources Count : 0
Request Msg Not Implemented Count: 0
Malformed Msg Count : 0
TLV Missing Count : 0
Invalid TLV Contents Count : 0
Non-Existent Port(s) Count : 0

Version : 0x32
IP Address : 192.168.9.1
System Name : 00:00:5e:00:53:02
Up Time : 36
TCP Port : 61408
State : Not Established
Subscriber Count : 1
Capabilities : Topology Discovery
System Instance : 12
Peer Instance : 1
Adjacency Timer (in 100ms) : 50
Peer Adjacency Timer (in 100ms) : 100
Partition Type : 0
Partition Flag : 1
Partition Identifier : 0
Partition Adjacencies : 0
Dead Timer : 23
Received Syn Count : 24
Received Synack Count : 20

```

```

Received Rstack Count      : 2
Received Ack Count         : 9
Received Port Up Count     : 5
Received Port Down Count   : 0
Received Generic Resp Count : 0
Received Adjacency Update Count : 0
Received OAM Responses Count : 2
Received Other Count       : 0
Sent Syn Count             : 20
Sent Synack Count          : 24
Sent Rstack Count          : 1
Sent Generic Resp Count    : 0
Sent Ack Count             : 9
Sent OAM Requests Count    : 4
Max Discovery Limit Exceed Count : 0
Result Codes:
Invalid Request Message Count : 0
Specified Port(s) Down Count : 0
Out of Resources Count        : 0
Request Msg Not Implemented Count: 0
Malformed Msg Count          : 0
TLV Missing Count            : 0
Invalid TLV Contents Count    : 0
Non-Existent Port(s) Count    : 0
Received                     Sent

```

### show ancp neighbor ip-address

```
user@host> show ancp neighbor ip-address 192.0.2.85
```

```

Neighbor Information
Version           : 0x32
IP Address        : 192.0.2.85
System Name       : 00:00:5e:00:53:ba
Up Time           : 26
TCP Port          : 32666
State             : Established
Subscriber Count   : 4
Capabilities       : Topo
System Instance    : 2
Peer Instance      : 20
Adjacency Timer (in 100ms) : 100
Peer Adjacency Timer (in 100ms) : 100
Partition Type     : 0
Partition Flag     : 1
Partition Identifier : 0
Partition Adjacencies : 0
Dead Timer         : 23
Received Syn Count : 1
Received Synack Count : 1
Received Rstack Count : 0
Received Ack Count   : 4
Received Port Up Count : 10
Received Port Down Count : 0
Received Generic Resp Count : 0
Received Adjacency Update Count : 0
Received OAM Count    : 0
Received Other Count  : 0
Sent Syn Count        : 1
Sent Synack Count     : 2
Sent Rstack Count     : 0
Sent Ack Count        : 3

```

Sent Generic Resp Count	: 0	
Sent OAM Count	: 0	
Max Discovery Limit Exceed Count	: 0	
Result Codes:	Received	Sent
Invalid Request Message Count	: 0	0
Specified Port(s) Down Count	: 0	0
Out of Resources Count	: 0	0
Request Msg Not Implemented Count	: 0	0
Malformed Msg Count	: 0	0
TLV Missing Count	: 0	0
Invalid TLV Contents Count	: 0	0
Non-Existent Port(s) Count	: 0	0

### show ancp neighbor system-name

user@host> show ancp neighbor 00:00:5e:00:53:ba detail

#### Neighbor Information

Version	: 0x31	
IP Address	: 203.0.113.101	
System Name	: 00:00:5e:00:53:ba	
Up Time	: 19	
TCP Port	: 1028	
State	: Established	
Subscriber Count	: 2	
Capabilities	: Topology Discovery, OAM	
System Instance	: 1	
Peer Instance	: 10	
Adjacency Timer (in 100ms)	: 100	
Peer Adjacency Timer (in 100ms)	: 250	
Partition Type	: 0	
Partition Flag	: 1	
Partition Identifier	: 0	
Partition Adjacencies	: 0	
Dead Timer	: 55	
Received Syn Count	: 1	
Received Synack Count	: 1	
Received Rstack Count	: 0	
Received Ack Count	: 1	
Received Port Up Count	: 34	
Received Port Down Count	: 0	
Received Generic Resp Count	: 0	
Received Adjacency Update Count	: 0	
Received OAM Responses Count	: 2	
Received Other Count	: 0	
Sent Syn Count	: 1	
Sent Synack Count	: 1	
Sent Rstack Count	: 0	
Sent Ack Count	: 3	
Sent Generic Resp Count	: 0	
Sent OAM Requests Count	: 4	
Max Discovery Limit Exceed Count	: 3	
Result Codes:	Received	Sent
Invalid Request Message Count	: 0	0
Specified Port(s) Down Count	: 0	0
Out of Resources Count	: 0	0
Request Msg Not Implemented Count	: 0	0
Malformed Msg Count	: 0	0
TLV Missing Count	: 0	0

Invalid TLV Contents Count	: 0	0
Non-Existent Port(s) Count	: 0	0

## show auto-configuration out-of-band pending

<b>Syntax</b>	show auto-configuration out-of-band pending
<b>Release Information</b>	Command introduced Junos OS Release 16.1.
<b>Description</b>	Display a list of access lines that are in the pending state for each routing instance. Access lines transition to the pending state when the Layer 2 wholesale session passes authorized and is assigned to an existing, nondefault routing instance, but profile instantiation to create the dynamic VLAN logical interface fails.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show auto-configuration out-of-band pending on page 421</a>
<b>Output Fields</b>	<a href="#">Table 19 on page 421</a> lists the output fields for the <b>show auto-configuration out-of-band pending</b> command. Output fields are listed in the approximate order in which they appear.

Table 19: show auto-configuration out-of-band pending Output Fields

Field Name	Field Description
<b>Routing-Instance</b>	Name of the routing instance for which the pending count is displayed.
<b>Pending count</b>	Number of access lines in the routing instance that are in the pending state.

## Sample Output

### show auto-configuration out-of-band pending

```

user@host> show auto-configuration out-of-band pending
Routing-Instance: NSP1
Pending count: 12

Routing-Instance: NSP2
Pending count: 0

```

## 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 431](#)

**List of Sample Output**

[show dhcp server binding on page 425](#)  
[show dhcp server binding detail on page 426](#)  
[show dhcp server binding detail \(ACI Interface Set Configured\) on page 426](#)  
[show dhcp server binding interface <vlan-id> on page 427](#)  
[show dhcp server binding interface <svlan-id> on page 427](#)  
[show dhcp server binding <ip-address> on page 427](#)  
[show dhcp server binding <session-id> on page 427](#)  
[show dhcp server binding summary on page 427](#)  
[show dhcp server binding <interfaces-vlan> on page 427](#)  
[show dhcp server binding <interfaces-wildcard> on page 427](#)

**Output Fields** Table 20 on page 423 lists the output fields for the **show dhcp server binding** command. Output fields are listed in the approximate order in which they appear.

**Table 20: 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.	<b>summary</b>
IP address	IP address of the DHCP client.	<b>brief detail</b>
Session Id	Session ID of the subscriber session.	<b>brief detail</b>
Hardware address	Hardware address of the DHCP client.	<b>brief detail</b>
Expires	Number of seconds in which lease expires.	<b>brief detail</b>

Table 20: show dhcp server binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>State</b>	State of the address binding table on the extended DHCP local server: <ul style="list-style-type: none"> <li>• <b>BOUND</b>—Client has active IP address lease.</li> <li>• <b>FORCERENEW</b>—Client has received forcerenew message from server.</li> <li>• <b>INIT</b>—Initial state.</li> <li>• <b>RELEASE</b>—Client is releasing IP address lease.</li> <li>• <b>RENEWING</b>—Client sending request to renew IP address lease.</li> <li>• <b>REQUESTING</b>—Client requesting a DHCP server.</li> <li>• <b>SELECTING</b>—Client receiving offers from DHCP servers.</li> </ul>	<b>brief</b> <b>detail</b>
<b>Interface</b>	Interface on which the request was received.	<b>brief</b>
<b>Lease Expires</b>	Date and time at which the client's IP address lease expires.	<b>detail</b>
<b>Lease Expires in</b>	Number of seconds in which lease expires.	<b>detail</b>
<b>Lease Start</b>	Date and time at which the client's IP address lease started.	<b>detail</b>
<b>Lease time violated</b>	Lease time violation has occurred.	<b>detail</b>
<b>Last Packet Received</b>	Date and time at which the router received the last packet.	<b>detail</b>
<b>Incoming Client Interface</b>	Client's incoming interface.	<b>detail</b>
<b>Client Interface Svlan Id</b>	S-VLAN ID of the client's incoming interface.	<b>detail</b>
<b>Client Interface Vlan Id</b>	VLAN ID of the client's incoming interface.	<b>detail</b>
<b>Demux Interface</b>	Name of the IP demultiplexing (demux) interface.	<b>detail</b>
<b>Server IP Address or Server Identifier</b>	IP address of DHCP server.	<b>detail</b>
<b>Server Interface</b>	Interface of DHCP server.	<b>detail</b>
<b>Client Pool Name</b>	Name of address pool used to assign client IP address lease.	<b>detail</b>

Table 20: 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><b>NOTE:</b> 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"> <li><b>DOWN</b>—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 <b>DOWN</b> state. A liveness detection session that was previously in an <b>UP</b> state has transitioned to a <b>DOWN</b> state, beginning with a liveness detection failure, and ending with the deletion of the client binding. The <b>DOWN</b> state is reported only during this transition period of time.</li> <li><b>UNKNOWN</b>—Liveness detection has been enabled for a subscriber but the actual liveness detection state has not yet been determined. The <b>UNKNOWN</b> 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 <b>UP</b> state.</li> <li><b>UP</b>—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 <b>UP</b> state.</li> <li><b>WENT_DOWN</b>—State is functionally equivalent to the <b>DOWN</b> state. A liveness detection session that was previously in an <b>UP</b> state has transitioned to a <b>DOWN</b> state implying a liveness detection failure. The <b>WENT_DOWN</b> 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.</li> </ul>	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
198.51.100.15	6	00:00:5e:00:53:01	86180	BOUND	ge-1/0/0.0
198.51.100.16	7	00:00:5e:00:53:02	86180	BOUND	ge-1/0/0.0
198.51.100.17	8	00:00:5e:00:53:03	86180	BOUND	ge-1/0/0.0
198.51.100.18	9	00:00:5e:00:53:04	86180	BOUND	ge-1/0/0.0
198.51.100.19	10	00:00:5e:00:53:05	86180	BOUND	ge-1/0/0.0

### show dhcp server binding detail

```

user@host> show dhcp server binding detail
Client IP Address: 198.51.100.15
    Hardware Address:      00:00:5e:00:53:01
    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:     198.51.100.9
    Server Interface:      none
    Session Id:            6
    Client Pool Name:      6
    Liveness Detection State: UP
Client IP Address: 198.51.100.16
    Hardware Address:      00:00:5e:00:53:02
    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:     198.51.100.9
    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: 198.51.100.14
    Hardware Address:      00:00:5e:00:53:02
    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:     198.51.100.202
    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
198.51.100.15   6          00:00:5e:00:53:01 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
198.51.100.16   7          00:00:5e:00:53:02 86124    BOUND  ge-1/1/0:10-100

```

#### show dhcp server binding <ip-address>

```

user@host> show dhcp server binding 198100.19
IP address      Session Id  Hardware address  Expires  State  Interface
198.51.100.19   10         00:00:5e:00:53:05 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
198.51.100.15   6          00:00:5e:00:53:01 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         00:00:5e:00:53:02 86346    BOUND  ge-1/0/0.1073741827
192.168.0.16    41         00:00:5e:00:53:01 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         00:00:5e:00:53:04 86361    BOUND  ge-1/3/0.110
192.168.0.8     23         00:00:5e:00:53:03 86361    BOUND  ge-1/3/0.110
192.168.0.7     22         00:00:5e:00:53:02 86361    BOUND  ge-1/3/0.110

```

## show dhcp server statistics

---

Syntax	<b>show dhcp server statistics</b> <b>&lt;bulk-leasequery-connections&gt;</b> <b>&lt;logical-system <i>logical-system-name</i>&gt;</b> <b>&lt;routing-instance <i>routing-instance-name</i>&gt;</b>
Release Information	Command introduced in Junos OS Release 9.0. <b>bulk-leasequery-connections</b> option introduced in Junos OS Release 16.1.
Description	Display extended Dynamic Host Configuration Protocol (DHCP) local server statistics.
Options	<b>bulk-leasequery-connections</b> —(Optional) Display bulk leasequery statistics.  <b>bulk-leasequery-connections</b> —(Optional) Display information about bulk leasequery statistics.  <b>logical-system <i>logical-system-name</i></b> —(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.  <b>routing-instance <i>routing-instance-name</i></b> —(Optional) Display information about extended DHCP local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">clear dhcp server statistics on page 434</a></li></ul>
List of Sample Output	<a href="#">show dhcp server statistics on page 430</a>
Output Fields	<a href="#">Table 21 on page 429</a> lists the output fields for the <b>show dhcp server statistics</b> command. Output fields are listed in the approximate order in which they appear.

Table 21: show dhcp server statistics Output Fields

Field Name	Field Description
<b>Packets dropped</b>	<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"> <li>• <b>Total</b>—Total number of packets discarded by the extended DHCP local server</li> <li>• <b>Authentication</b>—Number of packets discarded because they could not be authenticated</li> <li>• <b>Bad hardware address</b>—Number of packets discarded because an invalid hardware address was specified</li> <li>• <b>Bad opcode</b>—Number of packets discarded because an invalid operation code was specified</li> <li>• <b>Bad options</b>—Number of packets discarded because invalid options were specified</li> <li>• <b>Dynamic profile</b>—Number of packets discarded due to dynamic profile information</li> <li>• <b>Invalid server address</b>—Number of packets discarded because an invalid server address was specified</li> <li>• <b>Lease Time Violation</b>—Number of packets discarded because of a lease time violation</li> <li>• <b>No available addresses</b>—Number of packets discarded because there were no addresses available for assignment</li> <li>• <b>No interface match</b>—Number of packets discarded because they did not belong to a configured interface</li> <li>• <b>No routing instance match</b>—Number of packets discarded because they did not belong to a configured routing instance</li> <li>• <b>No valid local address</b>—Number of packets discarded because there was no valid local address</li> <li>• <b>Packet too short</b>—Number of packets discarded because they were too short</li> <li>• <b>Read error</b>—Number of packets discarded because of a system read error</li> <li>• <b>Send error</b>—Number of packets that the extended DHCP local server could not send</li> </ul>
<b>Messages received</b>	<p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREQUEST</b>—Number of BOOTP protocol data units (PDUs) received</li> <li>• <b>DHCPDECLINE</b>—Number of DHCP PDUs of type DECLINE received</li> <li>• <b>DHCPDISCOVER</b>—Number of DHCP PDUs of type DISCOVER received</li> <li>• <b>DHCPINFORM</b>—Number of DHCP PDUs of type INFORM received</li> <li>• <b>DHCPRELEASE</b>—Number of DHCP PDUs of type RELEASE received</li> <li>• <b>DHCPREQUEST</b>—Number of DHCP PDUs of type REQUEST received</li> <li>• <b>DHCPLEASEQUERY</b>—Number of DHCP leasequery messages received.</li> <li>• <b>DHCPBULKLEASEQUERY</b>—Number of DHCP bulk leasequery messages received.</li> </ul>
<b>Messages sent</b>	<p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREPLY</b>—Number of BOOTP PDUs transmitted</li> <li>• <b>DHCPOFFER</b>—Number of DHCP OFFER PDUs transmitted</li> <li>• <b>DHCPACK</b>—Number of DHCP ACK PDUs transmitted</li> <li>• <b>DHCPNACK</b>—Number of DHCP NACK PDUs transmitted</li> <li>• <b>DHCPFORCERENEW</b>—Number of DHCP FORCERENEW PDUs transmitted</li> <li>• <b>DHCPLEASEUNASSIGNED</b>—Number of DHCP leases that are managed by the server but have not yet been assigned</li> <li>• <b>DHCPLEASEUNKNOWN</b>—Number of unknown DHCP leases</li> <li>• <b>DHCPLEASEACTIVE</b>—Number of active DHCP leases</li> <li>• <b>DHCPLEASEQUERYDONE</b>—The leasequery is complete</li> </ul>

Table 21: show dhcp server statistics Output Fields (*continued*)

Field Name	Field Description
<b>Total Accepted Connections</b>	Total number of bulk leasequery connections accepted by the server.
<b>Total Not-Accepted Connections</b>	Total number of bulk leasequery connections not accepted by the server.
<b>Connections Closed due to Errors</b>	Number of bulk leasequery connections that the server closed due to an internal error.
<b>Connections Closed due to max-empty-replies</b>	Number of bulk leasequery connections that the server closed because the maximum number of empty replies was reached.
<b>In-flight Connections</b>	Number of bulk leasequery connections on the server.

## 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 422](#)

**List of Sample Output**

- [clear dhcp server binding <ip-address> on page 432](#)
- [clear dhcp server binding all on page 432](#)
- [clear dhcp server binding interface on page 433](#)
- [clear dhcp server binding <interfaces-vlan> on page 433](#)
- [clear dhcp server binding <interfaces-wildcard> on page 433](#)

**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
198.51.100.1	00:00:5e:00:53:01	active	2007-01-17 11:38:47 PST
198.51.100.3	00:00:5e:00:53:02	active	2007-01-17 11:38:41 PST

```
user@host> clear dhcp server binding 198.51.100.1
```

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

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

IP address	Hardware address	Type	Lease expires at
198.51.100.3	00:00:5e:00:53:02	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>&lt;bulk-leasequery-connections&gt;</code> <code>&lt;logical-system <i>logical-system-name</i>&gt;</code> <code>&lt;routing-instance <i>routing-instance-name</i>&gt;</code>
Release Information	Command introduced in Junos OS Release 9.0. <code>bulk-leasequery-connections</code> option introduced in Junos OS Release 16.1.
Description	Clear all extended Dynamic Host Configuration Protocol (DHCP) local server statistics.
Options	<code>bulk-leasequery-connections</code> —(Optional) Clear bulk leasequery statistics.  <code>logical-system <i>logical-system-name</i></code> —(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.  <code>routing-instance <i>routing-instance-name</i></code> —(Optional) Clear the statistics for DHCP clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.
Required Privilege Level	view
List of Sample Output	<a href="#">clear dhcp server statistics on page 434</a>
Output Fields	See <a href="#">show dhcp server statistics</a> 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**    **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>

**Release Information**    Command introduced in Junos OS Release 8.3.  
                              Options **interface** and **mac-address** added in Junos OS Release 8.4.  
                              Options **interfaces-vlan** and **interfaces-wildcard** added in Junos OS Release 12.1.  
                              Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.

**Description**    Display the address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.

**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**—(Optional) Display brief information about the active client bindings. This is the default, and produces the same output as **show dhcp relay binding**.

**detail**—(Optional) Display detailed client binding information.

**interface *interface-name***—(Optional) Perform this operation 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 binding state information. This option supports the use of the wildcard character (\*).

**logical-system *logical-system-name***—(Optional) Perform this operation on the specified logical system.

**routing-instance *routing-instance-name***—(Optional) Perform this operation on the specified routing instance.

**summary**—(Optional) Display a summary of DHCP client information.

**Required Privilege Level** view

**Related Documentation**

- [Clearing DHCP Bindings for Subscriber Access](#)
- [clear dhcp relay binding on page 445](#)

**List of Sample Output**

- [show dhcp relay binding on page 438](#)
- [show dhcp relay binding detail on page 439](#)
- [show dhcp relay binding interface on page 439](#)
- [show dhcp relay binding interface vlan-id on page 439](#)
- [show dhcp relay binding interface svlan-id on page 439](#)
- [show dhcp relay binding ip-address on page 440](#)
- [show dhcp relay binding mac-address on page 440](#)
- [show dhcp relay binding session-id on page 440](#)
- [show dhcp relay binding <interfaces-vlan> on page 440](#)
- [show dhcp relay binding <interfaces-wildcard> on page 440](#)
- [show dhcp relay binding summary on page 440](#)

**Output Fields** Table 22 on page 437 lists the output fields for the **show dhcp relay binding** command. Output fields are listed in the approximate order in which they appear.

**Table 22: 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.	<b>summary</b>
IP address	IP address of the DHCP client.	<b>briefdetail</b>
Session Id	Session ID of the subscriber session.	<b>briefdetail</b>
Generated Remote ID	Remote ID generated by the Option 82 Agent Remote ID (suboption 1)	<b>detail</b>
Hardware address	Hardware address of the DHCP client.	<b>briefdetail</b>
Expires	Number of seconds in which the lease expires.	<b>briefdetail</b>

Table 22: show dhcp relay binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>State</b>	State of the DHCP relay address binding table on the DHCP client: <ul style="list-style-type: none"> <li>• <b>BOUND</b>—Client has an active IP address lease.</li> <li>• <b>INIT</b>—Initial state.</li> <li>• <b>REBINDING</b>—Client is broadcasting a request to renew the IP address lease.</li> <li>• <b>RELEASE</b>—Client is releasing the IP address lease.</li> <li>• <b>RENEWING</b>—Client is sending a request to renew the IP address lease.</li> <li>• <b>REQUESTING</b>—Client is requesting a DHCP server.</li> <li>• <b>SELECTING</b>—Client is receiving offers from DHCP servers.</li> </ul>	<b>briefdetail</b>
<b>Interface</b>	Incoming client interface.	brief
<b>Lease Expires</b>	Date and time at which the client's IP address lease expires.	detail
<b>Lease Expires in</b>	Number of seconds in which the lease expires.	detail
<b>Lease Start</b>	Date and time at which the client's IP address lease started.	detail
<b>Lease time violated</b>	Lease time violation has occurred.	detail
<b>Incoming Client Interface</b>	Client's incoming interface.	detail
<b>Server IP Address</b>	IP address of the DHCP server.	detail
<b>Server Interface</b>	Interface of the DHCP server.	detail
<b>Bootp Relay Address</b>	IP address of BOOTP relay.	detail
<b>Type</b>	Type of DHCP packet processing performed on the router: <ul style="list-style-type: none"> <li>• <b>active</b>—Router actively processes and relays DHCP packets.</li> <li>• <b>passive</b>—Router passively snoops DHCP packets passing through the router.</li> </ul>	All levels
<b>Lease expires at</b>	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
198.51.100.11   41         00:00:5e:00:53:01 86371     BOUND   ge-1/0/0.0
198.51.100.12   42         00:00:5e:00:53:02 86371     BOUND   ge-1/0/0.0

```

198.51.100.13	43	00:00:5e:00:53:03	86371	BOUND	ge-1/0/0.0
198.51.100.14	44	00:00:5e:00:53:04	86371	BOUND	ge-1/0/0.0
198.51.100.15	45	00:00:5e:00:53:05	86371	BOUND	ge-1/0/0.0

### show dhcp relay binding detail

```
user@host> show dhcp relay binding detail
```

```
Client IP Address: 198.51.100.11
  Hardware Address: 00:00:5e:00:53:01
  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: 198.51.100.22
  Server Interface: none
  Bootp Relay Address: 198.51.100.32
  Session Id: 41
```

```
Client IP Address: 198.51.100.12
  Hardware Address: 00:00:5e:00:53:02
  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: 198.51.100.22
  Server Interface: none
  Bootp Relay Address: 198.51.100.32
  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
198.51.100.1	00:00:5e:00:53:01	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
198.51.100.15	6	00:00:5e:00:53:94	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
------------	------------	------------------	---------	-------	-----------

198.51.100.16	7	00:00:5e:00:53:92	86124	BOUND
ge-1/1/0:10-100				

#### show dhcp relay binding ip-address

```
user@host> show dhcp relay binding 198.51.100.13
```

IP address	Session Id	Hardware address	Expires	State	Interface
198.51.100.13	43	00:00:5e:00:53:03	86293	BOUND	ge-1/0/0.0

#### show dhcp relay binding mac-address

```
user@host> show dhcp relay binding 00:00:5e:00:53:05
```

IP address	Session Id	Hardware address	Expires	State	Interface
198.51.100.15	45	00:00:5e:00:53:05	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
198.51.100.11	41	00:00:5e:00:53:53	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	00:00:5e:00:53:02	86346	BOUND	ge-1/0/0.1073741827
192.168.0.16	41	00:00:5e:00:53:01	86346	BOUND	

#### 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	00:00:5e:00:53:04	86361	BOUND	ge-1/3/0.110
192.168.0.8	23	00:00:5e:00:53:03	86361	BOUND	
192.168.0.7	22	00:00:5e:00:53:02	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

<b>Syntax</b>	<pre>show dhcp relay statistics &lt;bulk-leasequery-connections&gt; &lt;logical-system <i>logical-system-name</i>&gt; &lt;routing-instance <i>routing-instance-name</i>&gt;</pre>
<b>Syntax</b>	<p>Syntax for EX Series switches:</p> <pre>show dhcp relay statistics &lt;routing-instance <i>routing-instance-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.3.</p> <p>Command 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> <p><b>bulk-leasequery-connections</b> option introduced in Junos OS Release 16.1.</p>
<b>Description</b>	Display Dynamic Host Configuration Protocol (DHCP) relay statistics.
<b>Options</b>	<p><b>bulk-leasequery-connections</b>—(Optional) Display information about bulk leasequery operations.</p> <p><b>logical-system <i>logical-system-name</i></b>—(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear dhcp relay statistics on page 447</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show dhcp relay statistics on page 443</a></p> <p><a href="#">show dhcp relay statistics bulk-leasequery-connections on page 444</a></p>
<b>Output Fields</b>	<p><a href="#">Table 23 on page 442</a> lists the output fields for the <b>show dhcp relay statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>

Table 23: show dhcp relay statistics Output Fields

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

Table 23: show dhcp relay statistics Output Fields (*continued*)

Field Name	Field Description
External Server Response	State of the external DHCP server responsiveness.
Packets forwarded	Number of packets forwarded. <ul style="list-style-type: none"> <li><b>BOOTREQUEST</b>—Number of BOOTREQUEST protocol data units (PDUs) forwarded</li> <li><b>BOOTREPLY</b>—Number of BOOTREPLY protocol data units (PDUs) forwarded</li> </ul>
External Server Response	State of the external DHCP server responsiveness.
Total Requested Servers	Total number of servers with which the DHCP relay agent has requested a bulk leasequery connection.
Total Attempted Servers	Total number of servers with which the DHCP relay agent has attempted to create a bulk leasequery connection.
Total Connected	Total number of servers that have formed a bulk leasequery connection with the DHCP relay agent.
Total Terminated by Server	Total number of servers that have terminated a bulk leasequery connection with the DHCP relay agent.
Total Max Attempted	Total number of servers where the DHCP relay agent reached the maximum retry limit when it attempted to create a bulk leasequery connection.
Total Closed due to Errors	Total number of bulk leasequery connections that closed due to an internal error on the DHCP relay agent.
In-Flight Connected	Number of current bulk leasequery connections on the DHCP relay agent.
Bulk Leasequery Reply Packet Retries	Number of bulk leasequery reply packets that the DHCP relay agent has retried.

## 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
DHCPCDECLINE       0
DHCPDISCOVER       11
DHCPINFORM         0
DHCPRELEASE        0
DHCPREQUEST        105
DHCPLEASEACTIVE    0
DHCPLEASEUNASSIGNED 0
DHCPLEASEUNKNOWN   0
DHCPLEASEQUERYDONE 0

Messages sent:
BOOTREPLY          0
DHCPOFFER          2
DHCPACK            1
DHCPNAK            0
DHCPFORCERENEW     0
DHCPLEASEQUERY     0
DHCPBULKLEASEQUERY 0

Packets forwarded:
Total              4
BOOTREQUEST        2
BOOTREPLY          2

External Server Response:
State              Responding
```

#### **show dhcp relay statistics bulk-leasequery-connections**

```
user@host> show dhcp relay statistics bulk-leasequery-connections

Total Requested Servers: 0
Total Attempted Servers: 0
Total Connected: 0
Total Terminated by Server: 0
Total Max Attempted: 0
Total Closed due to Errors: 0
In-Flight Connected: 0
Bulk Leasequery Reply Packet Retries: 0
```

## clear dhcp relay binding

<b>Syntax</b>	<pre>clear dhcp relay binding &lt;address&gt; &lt;all&gt; &lt;interface interface-name&gt; &lt;interfaces-vlan&gt; &lt;interfaces-wildcard&gt; &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.3.</p> <p>Options <b>all</b> and <b>interface</b> 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>
<b>Description</b>	Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table.
<b>Options</b>	<p><b>address</b>—(Optional) Clear the binding state for the DHCP client, using one of the following entries:</p> <ul style="list-style-type: none"> <li><i>ip-address</i>—The specified IP address.</li> <li><i>mac-address</i>—The specified MAC address.</li> <li><i>session-id</i>—The specified session ID.</li> </ul> <p><b>all</b>—(Optional) Clear the binding state for all DHCP clients.</p> <p><b>interface interface-name</b>—(Optional) Clear the binding state for DHCP clients on the specified interface.</p> <p><b>interfaces-vlan</b>—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.</p> <p><b>interfaces-wildcard</b>—(Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).</p> <p><b>logical-system logical-system-name</b>—(Optional) Clear the binding state for DHCP clients on the specified logical system.</p> <p><b>routing-instance routing-instance-name</b>—(Optional) Clear the binding state for DHCP clients on the specified routing instance.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Clearing DHCP Bindings for Subscriber Access</li> <li><a href="#">show dhcp relay binding on page 436</a></li> </ul>

**List of Sample Output** [clear dhcp relay binding on page 446](#)  
[clear dhcp relay binding all on page 446](#)  
[clear dhcp relay binding interface on page 446](#)  
[clear dhcp relay binding <interfaces-vlan> on page 446](#)  
[clear dhcp relay binding <interfaces-wildcard> on page 446](#)

**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
198.51.100.32    00:00:5e:00:53:01 active    2007-02-08 16:41:17 EST
192.168.14.8     00:00:5e:00:53:02 active    2007-02-10 10:01:06 EST
```

```
user@host> clear dhcp relay binding 198.51.100.32
```

```
user@host> show dhcp relay binding
IP address      Hardware address  Type    Lease expires at
192.168.14.8     00:00:5e:00:53:02 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

<b>Syntax</b>	<pre>clear dhcp relay statistics &lt;bulk-leasequery-connections&gt; &lt;logical-system <i>logical-system-name</i>&gt; &lt;routing-instance <i>routing-instance-name</i>&gt;</pre>
<b>Syntax</b>	<p>Syntax for EX Series switches:</p> <pre>show dhcp relay statistics &lt;routing-instance <i>routing-instance-name</i>&gt;</pre>
<b>Release Information</b>	<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> <p><b>bulk-leasequery-connections</b> option introduced in Junos OS Release 16.1.</p>
<b>Description</b>	Clear all Dynamic Host Configuration Protocol (DHCP) relay statistics.
<b>Options</b>	<p><b>bulk-leasequery-connections</b>—(Optional) Clear bulk leasequery statistics.</p> <p><b>logical-system <i>logical-system-name</i></b>—(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><b>routing-instance <i>routing-instance-name</i></b>—(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>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show dhcp relay statistics on page 441</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear dhcp relay statistics on page 448</a>
<b>Output Fields</b>	<a href="#">Table 24 on page 448</a> lists the output fields for the <b>clear dhcp relay statistics</b> command.

Table 24: clear dhcp relay statistics Output Fields

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

## 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 &lt;address&gt; &lt;brief   detail   summary&gt; &lt;interface interface-name&gt; &lt;interfaces-vlan&gt; &lt;interfaces-wildcard&gt; &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt;</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><b>address</b>—(Optional) One of the following identifiers for the DHCPv6 client whose binding state you want to show:</p> <ul style="list-style-type: none"><li>• <i>CID</i>—The specified Client ID (CID).</li><li>• <i>ipv6-prefix</i>—The specified IPv6 prefix.</li><li>• <i>session-id</i>—The specified session ID.</li></ul> <p><b>brief   detail   summary</b>—(Optional) Display the specified level of output about active client bindings. The default is <b>brief</b>, which produces the same output as <b>show dhcpv6 server binding</b>.</p> <p><b>interface interface-name</b>—(Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.</p> <p><b>interfaces-vlan</b>—(Optional) Interface VLAN ID or S-VLAN ID interface on which to show binding state information.</p> <p><b>interfaces-wildcard</b>—(Optional) Set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).</p> <p><b>logical-system logical-system-name</b>—(Optional) Display information about active client bindings for DHCPv6 clients on the specified logical system.</p> <p><b>routing-instance routing-instance-name</b>—(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"><li>• <i>Clearing DHCP Bindings for Subscriber Access</i></li><li>• <a href="#">clear dhcpv6 server binding on page 459</a></li></ul>

**List of Sample Output**

- [show dhcpv6 server binding on page 452](#)
- [show dhcpv6 server binding detail on page 453](#)
- [show dhcpv6 server binding interface on page 453](#)
- [show dhcpv6 server binding interface detail on page 453](#)
- [show dhcpv6 server binding \(IPv6 Prefix\) on page 454](#)
- [show dhcpv6 server binding \(Session ID\) on page 454](#)
- [show dhcpv6 server binding \(Interfaces VLAN\) on page 454](#)
- [show dhcpv6 server binding \(Interfaces Wildcard\) on page 454](#)
- [show dhcpv6 server binding \(Interfaces Wildcard\) on page 454](#)
- [show dhcpv6 server binding summary on page 455](#)

**Output Fields** Table 25 on page 451 lists the output fields for the **show dhcpv6 server binding** command. Output fields are listed in the approximate order in which they appear.

**Table 25: 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.	<b>summary</b>
<b>Prefix</b>	Client's DHCPv6 prefix, or prefix used to support multiple address assignment.	<b>brief detail</b>
<b>Session Id</b>	Session ID of the subscriber session.	<b>brief detail</b>
<b>Expires</b>	Number of seconds in which lease expires.	<b>brief detail</b>
<b>State</b>	State of the address binding table on the extended DHCPv6 local server: <ul style="list-style-type: none"> <li>• <b>BOUND</b>—Client has active IP address lease.</li> <li>• <b>INIT</b>—Initial state.</li> <li>• <b>RECONFIGURE</b>—Server has sent reconfigure message to client.</li> <li>• <b>RELEASE</b>—Client is releasing IP address lease.</li> <li>• <b>RENEWING</b>—Client sending request to renew IP address lease.</li> <li>• <b>REQUESTING</b>—Client requesting a DHCPv6 server.</li> <li>• <b>SELECTING</b>—Client receiving offers from DHCPv6 servers.</li> </ul>	<b>brief detail</b>
<b>Interface</b>	Interface on which the DHCPv6 request was received.	<b>brief</b>
<b>Client IPv6 Address</b>	Client's IPv6 address.	<b>detail</b>
<b>Client IPv6 Prefix</b>	Client's IPv6 prefix.	<b>detail</b>
<b>Client DUID</b>	Client's DHCP Unique Identifier (DUID).	<b>brief detail</b>
<b>Lease expires</b>	Date and time at which the client's IP address lease expires.	<b>detail</b>

Table 25: 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
Server Id	DHCP unique identifier (DUID) for the DHCPv6 server.	detail

## Sample Output

### show dhcpv6 server binding

```

user@host> show dhcpv6 server binding
Prefix          Session Id Expires State Interface Client DUID
2001:db8:1111:2222::/64 6      86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:01
2001:db8:1111:2222::/64 7      86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
2001:db8:1111:2222::/64 8      86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03
2001:db8:1111:2222::/64 9      86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:04
2001:db8:1111:2222::/64 10     86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:05
2001:db8:2002::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:db8:1111:2222::/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:         2001:db8:2002::1/128
  Client IPv6 Prefix:         2001:db8:1111:2222::/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:db8:1111:2222::/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:db8:1111:2222::/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:db8:1111:2222::/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:db8:19::/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:db8:19::/32  31      79681  BOUND demux0.1073741825
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:db8:C9::/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:db8:19::/32      33      79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:db8:C9::/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	<b>show dhcpv6 server statistics</b> <b>&lt;bulk-leasequery-connections&gt;</b> <b>&lt;logical-system <i>logical-system-name</i>&gt;</b> <b>&lt;routing-instance <i>routing-instance-name</i>&gt;</b>
Release Information	Command introduced in Junos OS Release 9.6. <b>bulk-leasequery-connections</b> option introduced in Junos OS Release 16.1.
Description	Display extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.
Options	<b>bulk-leasequery-connections</b> —(Optional) Display bulk leasequery statistics.  <b>logical-system <i>logical-system-name</i></b> —(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.  <b>routing-instance <i>routing-instance-name</i></b> —(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"><li>• <a href="#">clear dhcpv6 server statistics on page 461</a></li></ul>
List of Sample Output	<a href="#">show dhcpv6 server statistics on page 458</a> <a href="#">show dhcpv6 server statistics bulk-leasequery-connections on page 458</a>
Output Fields	<a href="#">Table 26 on page 457</a> lists the output fields for the <b>show dhcpv6 server statistics</b> command. Output fields are listed in the approximate order in which they appear.

Table 26: show dhcpv6 server statistics Output Fields

Field Name	Field Description
<b>Packets dropped</b>	<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"> <li>• <b>Total</b>—Total number of packets discarded by the extended DHCPv6 local server</li> <li>• <b>Strict Reconfigure</b>—Number of solicit messages discarded because the client does not support reconfiguration</li> <li>• <b>Bad hardware address</b>—Number of packets discarded because an invalid hardware address was specified</li> <li>• <b>Bad opcode</b>—Number of packets discarded because an invalid operation code was specified</li> <li>• <b>Bad options</b>—Number of packets discarded because invalid options were specified</li> <li>• <b>Invalid server address</b>—Number of packets discarded because an invalid server address was specified</li> <li>• <b>Lease Time Violation</b>—Number of packets discarded because of a lease time violation</li> <li>• <b>No available addresses</b>—Number of packets discarded because there were no addresses available for assignment</li> <li>• <b>No interface match</b>—Number of packets discarded because they did not belong to a configured interface</li> <li>• <b>No routing instance match</b>—Number of packets discarded because they did not belong to a configured routing instance</li> <li>• <b>No valid local address</b>—Number of packets discarded because there was no valid local address</li> <li>• <b>Packet too short</b>—Number of packets discarded because they were too short</li> <li>• <b>Read error</b>—Number of packets discarded because of a system read error</li> <li>• <b>Send error</b>—Number of packets that the extended DHCPv6 local server could not send</li> </ul>
<b>Messages received</b>	<p>Number of DHCPv6 messages received.</p> <ul style="list-style-type: none"> <li>• <b>DHCPV6_CONFIRM</b>—Number of DHCPv6 CONFIRM PDUs received.</li> <li>• <b>DHCPV6_DECLINE</b>—Number of DHCPv6 DECLINE PDUs received.</li> <li>• <b>DHCPV6_INFORMATION_REQUEST</b>—Number of DHCPv6 INFORMATION-REQUEST PDUs received.</li> <li>• <b>DHCPV6_REBIND</b>—Number of DHCPv6 REBIND PDUs received.</li> <li>• <b>DHCPV6_RELAY_FORW</b>—Number of DHCPv6 RELAY-FORW PDUs received.</li> <li>• <b>DHCPV6_RELAY_REPL</b>—Number of DHCPv6 RELAY-REPL PDUs received.</li> <li>• <b>DHCPV6_RELEASE</b>—Number of DHCPv6 RELEASE PDUs received.</li> <li>• <b>DHCPV6_RENEW</b>—Number of DHCPv6 RENEW PDUs received.</li> <li>• <b>DHCPV6_REQUEST</b>—Number of DHCPv6 REQUEST PDUs received.</li> <li>• <b>DHCPV6_SOLICIT</b>—Number of DHCPv6 SOLICIT PDUs received.</li> <li>• <b>DHCPV6_LEASEQUERY</b>—Number of DHCPv6 leasequery messages received.</li> </ul>
<b>Messages sent</b>	<p>Number of DHCPv6 messages sent.</p> <ul style="list-style-type: none"> <li>• <b>DHCPV6_ADVERTISE</b>—Number of DHCPv6 ADVERTISE PDUs transmitted.</li> <li>• <b>DHCPV6_REPLY</b>—Number of DHCPv6 ADVERTISE PDUs transmitted.</li> <li>• <b>DHC6_RECONFIGURE</b>—Number of DHCPv6 RECONFIGURE PDUs transmitted.</li> <li>• <b>DHCPV6_LEASEQUERY_REPLY</b>—Number of DHCPv6 leasequery replies transmitted to the DHCPv6 relay agent</li> <li>• <b>DHCPV6_LEASEQUERY_DATA</b>—xxxx</li> <li>• <b>DHCPV6_LEASEQUERY_DONE</b>—The leasequery is complete</li> </ul>

## 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
DHCPV6_LEASEQUERY	0

```
Messages sent:
```

DHCPV6_ADVERTISE	9
DHCPV6_REPLY	5
DHCPV6_RECONFIGURE	0
DHCPV6_LEASEQUERY_REPLY	0
DHCPV6_LEASEQUERY_DATA	0
DHCPV6_LEASEQUERY_DONE	0

### show dhcpv6 server statistics bulk-leasequery-connections

```
user@host> show dhcpv6 server statistics bulk-leasequery-connections
```

Total Accepted Connections:	0
Total Not-Accepted Connections:	0
Connections Closed due to Errors:	0
Connections Closed due to max-empty-replies:	0
In-flight Connections:	0

## clear dhcpv6 server binding

<b>Syntax</b>	<pre>clear dhcpv6 server binding &lt;address&gt; &lt;all&gt; &lt;interface interface-name&gt; &lt;interfaces-vlan&gt; &lt;interfaces-wildcard&gt; &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt;</pre>
<b>Release Information</b>	<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>
<b>Description</b>	Clear the binding state of a Dynamic Host Configuration Protocol for IPv6 (DHCPv6) client from the client table on the extended DHCPv6 local server.
<b>Options</b>	<p><b>address</b>—(Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:</p> <ul style="list-style-type: none"> <li>• <i>CID</i>—The specified Client ID (CID).</li> <li>• <i>ipv6-prefix</i>—The specified IPv6 prefix.</li> <li>• <i>session-id</i>—The specified session ID.</li> </ul> <p><b>all</b>—(Optional) Clear the binding state for all DHCPv6 clients.</p> <p><b>interface interface-name</b>—(Optional) Clear the binding state for DHCPv6 clients on the specified interface.</p> <p><b>interfaces-vlan</b>—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.</p> <p><b>interfaces-wildcard</b>—(Optional) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (*).</p> <p><b>logical-system logical-system-name</b>—(Optional) Clear the binding state for DHCPv6 clients on the specified logical system.</p> <p><b>routing-instance routing-instance-name</b>—(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Clearing DHCP Bindings for Subscriber Access</a></li> <li>• <a href="#">show dhcpv6 server binding on page 450</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">clear dhcpv6 server binding all on page 460</a></p> <p><a href="#">clear dhcpv6 server binding &lt;ipv6-prefix&gt; on page 460</a></p>

[clear dhcpv6 server binding interface on page 460](#)  
[clear dhcpv6 server binding <interfaces-vlan> on page 460](#)  
[clear dhcpv6 server binding <interfaces-wildcard> on page 460](#)

**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

<b>Syntax</b>	clear dhcpv6 server statistics <bulk-leasequery-connections> <interface <i>interface-name</i> > <logical-system <i>logical-system-name</i> > <routing-instance <i>routing-instance-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 9.6. <b>bulk-leasequery-connections</b> option introduced in Junos OS Release 16.1.
<b>Description</b>	Clear all extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.
<b>Options</b>	<b>bulk-leasequery-connections</b> —(Optional) Clear bulk leasequery statistics.  <b>logical-system <i>logical-system-name</i></b> —(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.  <b>routing-instance <i>routing-instance-name</i></b> —(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.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show dhcpv6 server statistics on page 456</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear dhcpv6 server statistics on page 461</a>
<b>Output Fields</b>	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 l2-routing-instance`
- `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 466](#)  
[show interfaces brief \(Loopback\) on page 467](#)  
[show interfaces detail \(Loopback\) on page 467](#)  
[show interfaces extensive \(Loopback\) on page 468](#)

**Output Fields** [Table 27 on page 463](#) lists the output fields for the **show interfaces** (loopback) command. Output fields are listed in the approximate order in which they appear.

**Table 27: Loopback show interfaces Output Fields**

Field Name	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Physical Interface</b>	Name of the physical interface.	All levels

Table 27: Loopback show interfaces Output Fields (*continued*)

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

Table 27: Loopback show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Traffic statistics</b>	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface.</li> <li>• <b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Input errors</b>	<ul style="list-style-type: none"> <li>• <b>Errors</b>—Input errors on the interface.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Frames received smaller than the runt threshold.</li> <li>• <b>Giants</b>—Frames received larger than the giant threshold.</li> <li>• <b>Policed Discards</b>—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.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>
<b>Output errors</b>	<ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. 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.</li> <li>• <b>Errors</b>—Sum of outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—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.</li> <li>• <b>MTU errors</b>—Number of packets larger than the MTU threshold.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>
<b>Logical Interface</b>		
<b>Logical interface</b>	Name of the logical interface.	All levels
<b>Index</b>	Logical interface index number, which reflects its initialization sequence.	<b>detail extensive</b>
<b>SNMP ifIndex</b>	Logical interface SNMP interface index number.	<b>detail extensive</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Flags</b>	Information about the logical interface; values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	<b>brief detail extensive</b>
<b>Encapsulation</b>	Encapsulation on the logical interface.	<b>brief detail extensive</b>
<b>Input packets</b>	Number of packets received on the logical interface.	None specified

Table 27: Loopback show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Output packets</b>	Number of packets transmitted on the logical interface.	None specified
<b>Traffic statistics</b>	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.	<b>detail extensive</b>
<b>Local statistics</b>	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.	<b>detail extensive</b>
<b>Protocol</b>	Protocol family configured on the logical interface (such as <b>iso</b> or <b>inet6</b> ).	<b>detail extensive none</b>
<b>MTU</b>	MTU size on the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route Table</b>	Route table in which this address exists; for example, <b>Route table:0</b> refers to inet.0.	<b>detail extensive</b>
<b>Flags</b>	Information about the protocol family flags. Possible values are described in the "Family Flags" section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>Addresses, Flags</b>	Information about the address flags. Possible values are described in the "Addresses Flags" section under <i>Common Output Fields Description</i> .	<b>detail extensive</b>
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive none</b>
<b>Local</b>	IP address of the logical interface.	<b>detail extensive none</b>
<b>Broadcast</b>	Broadcast address on the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>

## 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: 203.0.113.1
    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  203.0.113.1      --> 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: 203.0.113.1, 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: 203.0.113.1, 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)

<b>Syntax</b>	<pre>show interfaces <i>aenumber</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	(M Series, T Series, and MX Series routers only) Display status information about the specified aggregated Fast Ethernet or Gigabit Ethernet interface.
<b>Options</b>	<p><b><i>aenumber</i></b>—Display standard information about the specified aggregated Fast Ethernet or Gigabit Ethernet interface.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Ethernet Interfaces Feature Guide for Routing Devices</i></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show interfaces (Aggregated Ethernet) on page 475</a></p> <p><a href="#">show interfaces brief (Aggregated Ethernet) on page 476</a></p> <p><a href="#">show interfaces detail (Aggregated Ethernet) on page 476</a></p> <p><a href="#">show interfaces extensive (Aggregated Ethernet) on page 477</a></p> <p><a href="#">show interfaces extensive (Aggregated Ethernet with VLAN Stacking) on page 478</a></p>
<b>Output Fields</b>	Table 28 on page 470 lists the output fields for the <b>show interfaces</b> (Aggregated Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 28: Aggregated Ethernet show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
<b>Physical interface</b>	Name of the physical interface and state of the interface.	All levels
<b>Enabled</b>	State of the physical interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels

Table 28: Aggregated Ethernet show interfaces Output Fields (*continued*)

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

Table 28: Aggregated Ethernet show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Traffic statistics</b>	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes and rate, in bps, at which bytes are received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes and rate, in bps, at which bytes are transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets and rate, in pps, at which packets are received on the interface.</li> <li>• <b>Output packets</b>—Number of packets and rate, in pps, at which packets are transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Input errors</b>	<p>Input errors on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of incoming frame aborts and frame check sequence (FCS) errors.</li> <li>• <b>Drops</b>—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.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant threshold.</li> <li>• <b>Policed discards</b>—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.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>detail extensive</b>
<b>Output errors</b>	<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"> <li>• <b>Carrier transitions</b> —Number of times the interface has gone from <b>down</b> to <b>up</b>. 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.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—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.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>detail extensive</b>

Table 28: 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"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	detail extensive
Queue counters	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul> <p><b>NOTE:</b> 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 <b>show interfaces</b> 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
<b>Logical Interface</b>		
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"> <li>• <b>Source Family Inet</b></li> <li>• <b>Destination Family Inet</b></li> </ul>	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels

Table 28: 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"> <li>• <b>Bundle</b>—Information about input and output bundle rates.</li> <li>• <b>Link</b>—(<b>detail</b> and <b>extensive</b> only) Information about specific links in the aggregate, including link state and input and output rates.</li> <li>• <b>Adaptive Statistics</b>—(<b>extensive</b> only) Information about adaptive load balancing counter statistics. <ul style="list-style-type: none"> <li>• <b>Adaptive Adjusts</b>—Number of times traffic flow imbalance was corrected by implementation of adaptive load balancing.</li> <li>• <b>Adaptive Scans</b>—Number of times the link utilization on each member link of the AE bundle was scanned by for adaptive load balancing</li> <li>• <b>Adaptive Tolerance</b>—Tolerance level, in percentage, for load imbalance on link utilization on each member link of the AE bundle.</li> <li>• <b>Adaptive Updates</b>—Number of times traffic flow loads have been updated on an AE bundle.</li> </ul> </li> <li>• <b>Marker Statistics</b>—(<b>detail</b> and <b>extensive</b> only) Information about 802.3ad marker protocol statistics on the specified links. <ul style="list-style-type: none"> <li>• <b>Marker Rx</b>—Number of valid marker protocol data units (PDUs) received on this aggregation port.</li> <li>• <b>Resp Tx</b>—Number of marker response PDUs transmitted on this aggregation port.</li> <li>• <b>Unknown Rx</b>—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.</li> <li>• <b>Illegal Rx</b>—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).</li> </ul> </li> </ul>	<b>detail extensive</b> none
LACP info	<p>Link Aggregation Control Protocol (LACP) information for each aggregated interface.</p> <ul style="list-style-type: none"> <li>• <b>Role</b> can be one of the following: <ul style="list-style-type: none"> <li>• <b>Actor</b>—Local device participating in LACP negotiation.</li> <li>• <b>Partner</b>—Remote device participating in LACP negotiation.</li> </ul> </li> <li>• <b>System priority</b>—Priority assigned to the system (by management or administrative policy), encoded as an unsigned integer.</li> <li>• <b>System identifier</b>—Actor or partner system ID, encoded as a MAC address.</li> <li>• <b>Port priority</b>—Priority assigned to the port by the actor or partner (by management or administrative policy), encoded as an unsigned integer.</li> <li>• <b>Unknown Rx</b>—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.</li> <li>• <b>Port key</b>—Operational key value assigned to the port by the actor or partner, encoded as an unsigned integer.</li> </ul>	

Table 28: 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"> <li>• <b>LACP Rx</b>—LACP received counter that increments for each normal hello.</li> <li>• <b>LACP Tx</b>—Number of LACP transmit packet errors logged.</li> <li>• <b>Unknown Rx</b>—Number of unrecognized packet errors logged.</li> <li>• <b>Illegal Rx</b>—Number of invalid packets received.</li> </ul> <p><b>NOTE:</b> For <b>LACP Rx</b> and <b>LACP Tx</b>, 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: 00:00:5e:00:53:f0, Hardware address: 00:00:5e:00:53: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 Packets pps Bytes bps
Bundle:
  Input : 0 0 0 0
  Output: 0 0 0 0
Protocol inet, MTU: 1500
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
Destination: 203.0.113/24, Local: 203.0.113.2, 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 203.0.113.2/24

```

#### 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: 00:00:5e:00:53:f0, Hardware address: 00:00:5e:00:53:f0
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Queue counters: Queued packets Transmitted packets Dropped packets

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

```

```

Logical interface ae0.0 (Index 72) (SNMP ifIndex 60) (Generation 18)
Flags: SNMP-Traps 16384 Encapsulation: ENET2
Statistics          Packets          pps          Bytes          bps
Bundle:
  Input :           0           0           0           0
  Output:           0           0           0           0
Link:
  fe-0/1/0.0
    Input :           0           0           0           0
    Output:           0           0           0           0
  fe-0/1/2.0
    Input :           0           0           0           0
    Output:           0           0           0           0
  fe-0/1/3.0
    Input :           0           0           0           0
    Output:           0           0           0           0
Marker Statistics:  Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
fe-0/1/0.0          0           0           0           0
fe-0/1/2.0          0           0           0           0
fe-0/1/3.0          0           0           0           0
Protocol inet, MTU: 1500, Generation: 37, Route table: 0
Flags: Is-Primary, Mac-Validate-Strict
Mac-Validate Failures: Packets: 0, Bytes: 0
  Destination: 203.0.113/24, Local: 203.0.113.2, Broadcast: 203.0.113.255,

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: 00:00:5e:00:53:f0, Hardware address: 00:00:5e:00:53: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, 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
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  00:00:5e:00:53:85      127      2      1
  fe-1/0/3.0      Partner    127  00:00:5e:00:53:c3      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: 203.0.113/24, Local: 203.0.113.2, Broadcast: 203.0.113.255,
  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: 00:00:5e:00:53:3f, Hardware address: 00:00:5e:00:53:3f
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes :           2406875           40152 bps
Output bytes :          1124470           22056 bps
Input packets:           5307             5 pps
Output packets:          13295            21 pps

```

```

IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Giants: 0, Policed discards:
0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
Ingress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort      0              859777              0

  1 expedited-fo     0              0              0

  2 assured-forw     0              0              0

  3 network-cont     0              0              0

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

  0 best-effort      0              1897615             0

  1 expedited-fo     0              0              0

  2 assured-forw     0              0              0

  3 network-cont     0              662505             0

Logical interface ae0.451 (Index 69) (SNMP ifIndex 167) (Generation 601)
Flags: SNMP-Traps VLAN-Tag [ 0x8100.451 ] Encapsulation: VLAN-VPLS
Statistics      Packets      pps      Bytes      bps
Bundle:
  Input :      289      0      25685      376
  Output:      1698      4      130375      3096
Link:
  ge-1/2/0.451
  Input :      289      0      25685      376
  Output:      0      0      0      0
  ge-1/2/1.451
  Input :      0      0      0      0
  Output:      1698      4      130375      3096
Marker Statistics:  Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
  ge-1/2/0.451      0      0      0      0
  ge-1/2/1.451      0      0      0      0
Protocol vpls, MTU: 1518, Generation: 849, Route table: 3
Flags: Is-Primary

Logical interface ae0.452 (Index 70) (SNMP ifIndex 170) (Generation 602)
Flags: SNMP-Traps VLAN-Tag [ 0x8100.452 ] Encapsulation: VLAN-VPLS
Statistics      Packets      pps      Bytes      bps
Bundle:
  Input :      293      1      26003      1072
  Output:      1694      3      130057      2400
Link:
  ge-1/2/0.452

```

```
      Input :          293          1          26003          1072
      Output:         1694          3          130057          2400
ge-1/2/1.452
      Input :           0           0           0           0
      Output:           0           0           0           0
Marker Statistics:  Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
ge-1/2/0.452              0              0              0              0
ge-1/2/1.452              0              0              0              0
Protocol vpls, MTU: 1518, Generation: 850, Route table: 3
Flags: None
...
```

## show interfaces (Fast Ethernet)

<b>Syntax</b>	<pre>show interfaces <i>interface-type</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Display status information about the specified Fast Ethernet interface.
<b>Options</b>	<p><b><i>interface-type</i></b>—On M Series and T Series routers, the interface type is <b><i>fe-fpc/pic/port</i></b>.</p> <p><b><i>brief   detail   extensive   terse</i></b>—(Optional) Display the specified level of output.</p> <p><b><i>descriptions</i></b>—(Optional) Display interface description strings.</p> <p><b><i>media</i></b>—(Optional) Display media-specific information about network interfaces.</p> <p><b><i>snmp-index snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b><i>statistics</i></b>—(Optional) Display static interface statistics.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show interfaces (Fast Ethernet) on page 494</a></p> <p><a href="#">show interfaces brief (Fast Ethernet) on page 495</a></p> <p><a href="#">show interfaces detail (Fast Ethernet) on page 495</a></p> <p><a href="#">show interfaces extensive (Fast Ethernet) on page 495</a></p>
<b>Output Fields</b>	Table 29 on page 481 lists the output fields for the <b>show interfaces</b> (Fast Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 29: show interfaces Fast Ethernet Output Fields

Field Name	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Physical interface</b>	Name of the physical interface.	All levels
<b>Enabled</b>	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
<b>Interface index</b>	Index number of the physical interface, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>

Table 29: 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: <b>Full-duplex</b> or <b>Half-duplex</b>	<b>extensive</b>
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .	All levels
Source filtering	Source filtering status: <b>Enabled</b> or <b>Disabled</b> .	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: <b>Enabled</b> or <b>Disabled</b> for parent interface; <b>Rx-only</b> or <b>Tx-only</b> for child interfaces.	All levels
Flow control	Flow control status: <b>Enabled</b> or <b>Disabled</b> .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> <li>• <b>Online</b>—Autonegotiation is manually configured as online.</li> <li>• <b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul>	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 29: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>CoS queues</b>	Number of CoS queues configured.	<b>detail extensive none</b>
<b>Schedulers</b>	(GigabitEthernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	<b>extensive</b>
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds.	<b>detail extensive</b>
<b>Current address</b>	Configured MAC address.	<b>detail extensive none</b>
<b>Hardware address</b>	Hardware MAC address.	<b>detail extensive none</b>
<b>Last flapped</b>	Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .	<b>detail extensive none</b>
<b>Input Rate</b>	Input rate in bits per second (bps) and packets per second (pps).	None specified
<b>Output Rate</b>	Output rate in bps and pps.	None specified
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>
<b>Traffic statistics</b>	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <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>	<b>detail extensive</b>

Table 29: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Input errors</b>	<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"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—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.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—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.</li> <li>• <b>L3 incompletes</b>—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 <b>ignore-l3-incompletes</b> statement.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>FIFO errors</b>—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.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>

Table 29: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Output errors</b>	<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"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. 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.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—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.</li> <li>• <b>Collisions</b>—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.</li> <li>• <b>Aged packets</b>—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.</li> <li>• <b>FIFO errors</b>—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.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>
<b>Egress queues</b>	<p>Total number of egress queues supported on the specified interface.</p> <p><b>NOTE:</b> In DPCs that are not of the enhanced type, such as DPC 40x1GE R, DPCE 20x1GE + 2x10GE R, or DPCE 40x1GE R, you might notice a discrepancy in the output of the <b>show interfaces</b> 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>	<b>detail extensive</b>
<b>Queue counters (Egress)</b>	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>	<b>detail extensive</b>

Table 29: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Ingress queues</b>	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	<b>extensive</b>
<b>Queue counters (Ingress)</b>	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>	<b>extensive</b>
<b>Active alarms and Active defects</b>	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 <b>None</b> or <b>Link</b> . <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—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.</li> </ul>	<b>detail extensive none</b>
<b>OTN FEC statistics</b>	The forward error correction (FEC) counters provide the following statistics: <ul style="list-style-type: none"> <li>• <b>Corrected Errors</b>—The count of corrected errors in the last second.</li> <li>• <b>Corrected Error Ratio</b>—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.</li> </ul>	
<b>PCS statistics</b>	(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"> <li>• <b>Bit errors</b>—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode.</li> <li>• <b>Errored blocks</b>—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.</li> </ul>	<b>detail extensive</b>

Table 29: 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"> <li>• <b>Total octets</b> and <b>total packets</b>—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.</li> <li>• <b>Unicast packets</b>, <b>Broadcast packets</b>, and <b>Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—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).</li> <li>• <b>FIFO error</b>—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.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—Number of frames that exceed 1518 octets.</li> <li>• <b>Jabber frames</b>—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.</li> <li>• <b>Fragment frames</b>—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.</li> <li>• <b>VLAN tagged frames</b>—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.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."</li> </ul>	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 29: show interfaces Fast Ethernet Output Fields (*continued*)

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

Table 29: show interfaces Fast Ethernet Output Fields (*continued*)

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

Table 29: show interfaces Fast Ethernet Output Fields (*continued*)

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

Table 29: 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"> <li>• <b>Negotiation status:</b> <ul style="list-style-type: none"> <li>• <b>Incomplete</b>—Ethernet interface has the speed or link mode configured.</li> <li>• <b>No autonegotiation</b>—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.</li> <li>• <b>Complete</b>—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> </ul> </li> <li>• <b>Link partner status</b>—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> <li>• <b>Link partner:</b> <ul style="list-style-type: none"> <li>• <b>Link mode</b>—Depending on the capability of the attached Ethernet device, either <b>Full-duplex</b> or <b>Half-duplex</b>.</li> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is <b>None</b>. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports both <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> receive).</li> <li>• <b>Remote fault</b>—Remote fault information from the link partner—<b>Failure</b> indicates a receive link error. <b>OK</b> indicates that the link partner is receiving. <b>Negotiation error</b> indicates a negotiation error. <b>Offline</b> indicates that the link partner is going offline.</li> </ul> </li> <li>• <b>Local resolution</b>—Information from the link partner: <ul style="list-style-type: none"> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports both <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> receive).</li> <li>• <b>Remote fault</b>—Remote fault information. <b>Link OK</b> (no error detected on receive), <b>Offline</b> (local interface is offline), and <b>Link Failure</b> (link error detected on receive).</li> </ul> </li> </ul>	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"> <li>• <b>Destination slot</b>—FPC slot number.</li> </ul>	extensive

Table 29: show interfaces Fast Ethernet Output Fields (*continued*)

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

Table 29: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Demux:</b>	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"> <li>Source Family Inet</li> <li>Destination Family Inet</li> </ul>	<b>detail extensive none</b>
<b>Encapsulation</b>	Encapsulation on the logical interface.	All levels
<b>Protocol</b>	Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>MTU</b>	Maximum transmission unit size on the logical interface.	<b>detail extensive none</b>
<b>Maximum labels</b>	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	<b>detail extensive none</b>
<b>Traffic statistics</b>	Number and rate of bytes and packets received and transmitted on the specified interface set. <ul style="list-style-type: none"> <li><b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface set</li> <li><b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface set.</li> </ul>	<b>detail extensive</b>
<b>IPv6 transit statistics</b>	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	<b>extensive</b>
<b>Local statistics</b>	Number and rate of bytes and packets destined to the routing device.	<b>extensive</b>
<b>Transit statistics</b>	Number and rate of bytes and packets transiting the switch. <p><b>NOTE:</b> 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 <b>Output bytes</b> and <b>Output packets</b> 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>	<b>extensive</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route Table</b>	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	<b>detail extensive none</b>
<b>Flags</b>	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive</b>
<b>Donor interface</b>	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	<b>detail extensive none</b>

Table 29: show interfaces Fast Ethernet Output Fields (*continued*)

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

## 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: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:38
  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: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255

```

### 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 203.0.113.1/24

```

### 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: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:38
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: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255,
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: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:38
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          0          64
Total packets         0          1
Unicast packets       0          0
Broadcast packets     0          1
Multicast packets     0          0
CRC/Align errors      0          0
FIFO errors           0          0
MAC control frames    0          0
MAC pause frames      0          0
Oversized frames      0
Jabber frames         0
Fragment frames       0
VLAN tagged frames    0
Code violations        0
Filter statistics:
Input packet count    0
Input packet rejects  0
Input DA rejects      0
Input SA rejects      0
Output packet count   1
Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link partner:
  Link partner: Full-duplex, Flow control: None, Remote fault: Ok
Local resolution:
Packet Forwarding Engine configuration:
Destination slot: 0
CoS information:
      Bandwidth      %      Buffer Priority  Limit
      bps           %           ussec
0 best-effort      95    950000000  95         0    low  none
3 network-control   5     50000000   5         0    low  none
Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
Flags: SNMP-Traps Encapsulation: ENET2
Protocol inet, MTU: 1500, Generation: 105, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255,

Generation: 136

```

## show interfaces (Gigabit Ethernet)

<b>Syntax</b>	<pre>show interfaces <i>ge-fpc/pic/port</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Display status information about the specified Gigabit Ethernet interface.
<b>Options</b>	<p><b><i>ge-fpc/pic/port</i></b>—Display standard information about the specified Gigabit Ethernet interface.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p>
<b>Additional Information</b>	In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show interfaces (Gigabit Ethernet) on page 513</a></p> <p><a href="#">show interfaces (Gigabit Ethernet on MX Series Routers) on page 514</a></p> <p><a href="#">show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration) on page 514</a></p> <p><a href="#">show interfaces brief (Gigabit Ethernet) on page 514</a></p> <p><a href="#">show interfaces detail (Gigabit Ethernet) on page 515</a></p> <p><a href="#">show interfaces extensive (Gigabit Ethernet IQ2) on page 516</a></p> <p><a href="#">show interfaces (Gigabit Ethernet Unnumbered Interface) on page 519</a></p> <p><a href="#">show interfaces (ACI Interface Set Configured) on page 520</a></p>
<b>Output Fields</b>	<p><a href="#">Table 30 on page 498</a> describes the output fields for the <b>show interfaces</b> (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 <a href="#">Table 31 on page 512</a>.</p>

Table 30: show interfaces (Gigabit Ethernet) Output Fields

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

Table 30: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

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

Table 30: 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"> <li><b>Input bytes</b>—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.</li> <li><b>Output bytes</b>—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.</li> <li><b>Input packets</b>—Number of packets received on the interface.</li> <li><b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <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"> <li><b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li><b>Drops</b>—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.</li> <li><b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li><b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li><b>Policed discards</b>—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.</li> <li><b>L3 incompletes</b>—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 <b>ignore-l3-incompletes</b> statement.</li> <li><b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li><b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li><b>FIFO errors</b>—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.</li> <li><b>Resource errors</b>—Sum of transmit drops.</li> </ul>	extensive

Table 30: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Output errors</b>	<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"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. 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.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—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.</li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the <b>Drops</b> 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"> <li>• <b>Collisions</b>—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.</li> <li>• <b>Aged packets</b>—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.</li> <li>• <b>FIFO errors</b>—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.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>
<b>Egress queues</b>	<p>Total number of egress queues supported on the specified interface.</p> <p><b>NOTE:</b> 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 <b>show interfaces</b> 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>	<b>detail extensive</b>

Table 30: 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"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the <b>Dropped packets</b> 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"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>	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 <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—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.</li> </ul>	detail extensive none
Interface transmit statistics	<p>(On MX Series devices) Status of the <b>interface-transmit-statistics</b> configuration: Enabled or Disabled.</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—When the <b>interface-transmit-statistics</b> statement is included in the configuration. If this is configured, the interface statistics show the actual transmitted load on the interface.</li> <li>• <b>Disabled</b>—When the <b>interface-transmit-statistics</b> statement is not included in the configuration. If this is not configured, the interface statistics show the offered load on the interface.</li> </ul>	detail extensive
OTN FEC statistics	<p>The forward error correction (FEC) counters provide the following statistics:</p> <ul style="list-style-type: none"> <li>• <b>Corrected Errors</b>—Count of corrected errors in the last second.</li> <li>• <b>Corrected Error Ratio</b>—Corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.</li> </ul>	detail extensive

Table 30: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>PCS statistics</b>	(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"><li>• <b>Bit errors</b>—Number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode.</li><li>• <b>Errored blocks</b>—Number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.</li></ul>	<b>detail extensive</b>

Table 30: 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"> <li>• <b>Total octets</b> and <b>total packets</b>—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.</li> <li>• <b>Unicast packets</b>, <b>Broadcast packets</b>, and <b>Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—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).</li> <li>• <b>FIFO error</b>—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.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—There are two possible conditions regarding the number of oversized frames: <ul style="list-style-type: none"> <li>• Packet length exceeds 1518 octets, or</li> <li>• Packet length exceeds MRU</li> </ul> </li> <li>• <b>Jabber frames</b>—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.</li> <li>• <b>Fragment frames</b>—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.</li> <li>• <b>VLAN tagged frames</b>—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><b>NOTE:</b> The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the <b>VLAN tagged frames</b> field displays 0 when the <b>show interfaces</b> 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> </li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."</li> </ul>	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 30: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

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

Table 30: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

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

Table 30: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

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

Table 30: 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"> <li>• <b>Negotiation status:</b> <ul style="list-style-type: none"> <li>• <b>Incomplete</b>—Ethernet interface has the speed or link mode configured.</li> <li>• <b>No autonegotiation</b>—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.</li> <li>• <b>Complete</b>—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> </ul> </li> <li>• <b>Link partner status</b>—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> <li>• <b>Link partner</b>—Information from the remote Ethernet device: <ul style="list-style-type: none"> <li>• <b>Link mode</b>—Depending on the capability of the link partner, either <b>Full-duplex</b> or <b>Half-duplex</b>.</li> <li>• <b>Flow control</b>—Types of flow control supported by the link partner. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), <b>Symmetric/Asymmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> on transmit), and <b>None</b> (link partner does not support flow control).</li> <li>• <b>Remote fault</b>—Remote fault information from the link partner—<b>Failure</b> indicates a receive link error. <b>OK</b> indicates that the link partner is receiving. <b>Negotiation error</b> indicates a negotiation error. <b>Offline</b> indicates that the link partner is going offline.</li> </ul> </li> <li>• <b>Local resolution</b>—Information from the local Ethernet device: <ul style="list-style-type: none"> <li>• <b>Flow control</b>—Types of flow control supported by the local device. For Gigabit Ethernet interfaces, advertised capabilities are <b>Symmetric/Asymmetric</b> (local device supports <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> on receive) and <b>None</b> (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 <b>Symmetric</b> (local device supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (local device supports <b>PAUSE</b> on receive), and <b>None</b> (local device does not support flow control).</li> <li>• <b>Remote fault</b>—Remote fault information. <b>Link OK</b> (no error detected on receive), <b>Offline</b> (local interface is offline), and <b>Link Failure</b> (link error detected on receive).</li> </ul> </li> </ul>	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"> <li>• <b>Destination slot</b>—FPC slot number.</li> </ul>	extensive

Table 30: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>CoS information</b>	<p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li>• <b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li>• <b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li>• <b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li>• <b>Buffer usec</b>—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.</li> <li>• <b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li>• <b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul>	<b>extensive</b>
<b>Logical Interface</b>		
<b>Logical interface</b>	Name of the logical interface.	All levels
<b>Index</b>	Index number of the logical interface, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP interface index number for the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Flags</b>	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 30: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>VLAN-Tag</b>	<p>Rewrite profile applied to incoming or outgoing frames on the outer (<b>Out</b>) VLAN tag or for both the outer and inner (<b>In</b>) VLAN tags.</p> <ul style="list-style-type: none"> <li><b>push</b>—An outer VLAN tag is pushed in front of the existing VLAN tag.</li> <li><b>pop</b>—The outer VLAN tag of the incoming frame is removed.</li> <li><b>swap</b>—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information.</li> <li><b>push</b>—An outer VLAN tag is pushed in front of the existing VLAN tag.</li> <li><b>push-push</b>—Two VLAN tags are pushed in from the incoming frame.</li> <li><b>swap-push</b>—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.</li> <li><b>swap-swap</b>—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value.</li> <li><b>pop-swap</b>—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.</li> <li><b>pop-pop</b>—Both the outer and inner VLAN tags of the incoming frame are removed.</li> </ul>	<b>brief detail extensive none</b>
<b>Demux</b>	<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"> <li>Source Family Inet</li> <li>Destination Family Inet</li> </ul>	<b>detail extensive none</b>
<b>Encapsulation</b>	Encapsulation on the logical interface.	All levels
<b>ACI VLAN: Dynamic Profile</b>	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.	<b>brief detail extensive none</b>
<b>Protocol</b>	Protocol family. Possible values are described in the "Protocol Field" section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>MTU</b>	Maximum transmission unit size on the logical interface.	<b>detail extensive none</b>
<b>Neighbor Discovery Protocol (NDP) Queue Statistics</b>	<p>NDP statistics for protocol <b>inet6</b> under logical interface statistics.</p> <ul style="list-style-type: none"> <li><b>Max nh cache</b>—Maximum interface neighbor discovery nexthop cache size.</li> <li><b>New hold nh limit</b>—Maximum number of new unresolved nexthops.</li> <li><b>Curr nh cnt</b>—Current number of resolved nexthops in the NDP queue.</li> <li><b>Curr new hold cnt</b>—Current number of unresolved nexthops in the NDP queue.</li> <li><b>NH drop cnt</b>—Number of NDP requests not serviced.</li> </ul>	All levels
<b>Dynamic Profile</b>	Name of the dynamic profile that was used to create this interface configured with a Point-to-Point Protocol over Ethernet (PPPoE) family.	<b>detail extensive none</b>
<b>Service Name Table</b>	Name of the service name table for the interface configured with a PPPoE family.	<b>detail extensive none</b>

Table 30: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Max Sessions</b>	Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	<b>detail extensive none</b>
<b>Duplicate Protection</b>	State of PPPoE duplicate protection: <b>On</b> or <b>Off</b> . 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.	<b>detail extensive none</b>
<b>Direct Connect</b>	State of the configuration to ignore DSL Forum VSAs: <b>On</b> or <b>Off</b> . When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	<b>detail extensive none</b>
<b>AC Name</b>	Name of the access concentrator.	<b>detail extensive none</b>
<b>Maximum labels</b>	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	<b>detail extensive none</b>
<b>Traffic statistics</b>	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes, Output bytes</b>—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.</li> <li>• <b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface set.</li> </ul>	<b>detail extensive</b>
<b>IPv6 transit statistics</b>	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	<b>extensive</b>
<b>Local statistics</b>	Number and rate of bytes and packets destined to the router.	<b>extensive</b>
<b>Transit statistics</b>	<p>Number and rate of bytes and packets transiting the switch.</p> <p><b>NOTE:</b> 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 <b>Output bytes</b> and <b>Output packets</b> 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>	<b>extensive</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route Table</b>	Route table in which the logical interface address is located. For example, <b>0</b> refers to the routing table inet.0.	<b>detail extensive none</b>
<b>Flags</b>	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive</b>

Table 30: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Donor interface</b>	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	<b>detail extensive none</b>
<b>Preferred source address</b>	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	<b>detail extensive none</b>
<b>Input Filters</b>	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.	<b>detail extensive</b>
<b>Output Filters</b>	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.	<b>detail extensive</b>
<b>Mac-Validate Failures</b>	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	<b>detail extensive none</b>
<b>Addresses, Flags</b>	Information about the address flags. Possible values are described in the "Addresses Flags" section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b><i>protocol-family</i></b>	Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.	<b>brief</b>
<b>Flags</b>	Information about the address flag. Possible values are described in the "Addresses Flags" section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive none</b>
<b>Local</b>	IP address of the logical interface.	<b>detail extensive none</b>
<b>Broadcast</b>	Broadcast address of the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>

Table 31: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound physical interface	<b>show interfaces ge-0/3/0 extensive</b>	<p>Traffic statistics:</p> <p>Input bytes: 496 bytes per packet, representing the Layer 2 packet</p> <p>MAC statistics:</p> <p>Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes</p>	The additional 4 bytes are for the CRC.

Table 31: 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	<b>show interfaces ge-0/3/0.50 extensive</b>	Traffic statistics:  Input bytes: 478 bytes per packet, representing the Layer 3 packet	
Outbound physical interface	<b>show interfaces ge-0/0/0 extensive</b>	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	<b>show interfaces ge-0/0/0.50 extensive</b>	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: 00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
  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: 00:00:5e:00:53:c0, Hardware address: 00:00:5e:00:53:76
  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: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255
    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:db8:1::/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: 00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
  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: 00:00:5e:00:53:74, Hardware address: 00:00:5e:00:53:74
Last flapped    : 2007-11-07 21:31:41 PST (02:03:33 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes :      38910844056      7952 bps
Output bytes :      7174605      8464 bps
Input packets:      418398473      11 pps
Output packets:      78903      12 pps
IPv6 transit statistics:
Input bytes :      0
Output bytes :      0
Input packets:      0
Output packets:      0
Ingress traffic statistics at Packet Forwarding Engine:
Input bytes :      38910799145      7952 bps
Input packets:      418397956      11 pps
Drop bytes :      0      0 bps
Drop packets:      0      0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

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

0 best-effort      418390823      418390823      0

1 expedited-fo      0      0      0

2 assured-forw      0      0      0

3 network-cont      7133      7133      0

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

0 best-effort      1031      1031      0

1 expedited-fo      0      0      0

2 assured-forw      0      0      0

3 network-cont      77872      77872      0

Active alarms : None
Active defects : None
MAC statistics:
Total octets      Receive      Transmit
Total packets      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       50000000      5          0
low none
Direction : Input
CoS transmit queue          Bandwidth      Buffer      Priority      Limit
                             %          bps          %          usec
0 best-effort                95      950000000    95          0
low none
3 network-control            5       50000000      5          0
low none

Logical interface 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: 203.0.113/24, Local: 203.0.113.2, Broadcast: 203.0.113.255,
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 30 on page 498](#).

### 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: 00:00:5e:00:53:f8, Hardware address: 00:00:5e:00:53: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: 203.0.113.22

```

**show interfaces (ACI Interface Set Configured)**

```
user@host> show interfaces ge-1/0/0.4001
Logical interface ge-1/0/0.4001 (Index 340) (SNMP ifIndex 548)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4001 ] Encapsulation: PPP-over-
  Ethernet
  ACI VLAN:
    Dynamic Profile: aci-vlan-set-profile
  PPPoE:
    Dynamic Profile: aci-vlan-pppoe-profile,
    Service Name Table: None,
    Max Sessions: 32000, Max Sessions VSA Ignore: Off,
    Duplicate Protection: On, Short Cycle Protection: Off,
    Direct Connect: Off,
    AC Name: nbc
  Input packets : 9
  Output packets: 8
  Protocol multiservice, MTU: Unlimited
```

## show interfaces (PPPoE)

<b>Syntax</b>	<pre>show interfaces pp0.logical &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index snmp-index&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	(M120 routers, M320 routers, and MX Series routers only). Display status information about the PPPoE interface.
<b>Options</b>	<p><b>pp0.logical</b>—Display standard status information about the PPPoE interface.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about PPPoE interfaces.</p> <p><b>snmp-index snmp-index</b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display PPPoE interface statistics.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show interfaces (PPPoE) on page 527</a></p> <p><a href="#">show interfaces (PPPoE over Aggregated Ethernet) on page 527</a></p> <p><a href="#">show interfaces brief (PPPoE) on page 527</a></p> <p><a href="#">show interfaces detail (PPPoE) on page 528</a></p> <p><a href="#">show interfaces extensive (PPPoE on M120 and M320 Routers) on page 529</a></p>
<b>Output Fields</b>	Table 32 on page 521 lists the output fields for the <b>show interfaces (PPPoE)</b> command. Output fields are listed in the approximate order in which they appear.

Table 32: show interfaces (PPPoE) Output Fields

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

Table 32: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Type</b>	Physical interface type (PPPoE).	All levels
<b>Link-level type</b>	Encapsulation on the physical interface (PPPoE).	All levels
<b>MTU</b>	MTU size on the physical interface.	All levels
<b>Clocking</b>	Reference clock source. It can be <b>Internal</b> or <b>External</b> .	All levels
<b>Speed</b>	Speed at which the interface is running.	All levels
<b>Device flags</b>	Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Interface flags</b>	Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Link type</b>	Physical interface link type: <b>full duplex</b> or <b>half duplex</b> .	All levels
<b>Link flags</b>	Information about the interface. Possible values are described in the "Link Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Input rate</b>	Input rate in bits per second (bps) and packets per second (pps).	None specified
<b>Output rate</b>	Output rate in bps and pps.	None specified
<b>Physical Info</b>	Physical interface information.	All levels
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds.	<b>detail extensive</b>
<b>Current address</b>	Configured MAC address.	<b>detail extensive</b>
<b>Hardware address</b>	MAC address of the hardware.	<b>detail extensive</b>
<b>Alternate link address</b>	Backup address of the link.	<b>detail extensive</b>
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>
<b>Traffic statistics</b>	Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>

Table 32: 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><b>NOTE:</b> These fields include dropped traffic and exception traffic, as those fields are not separately defined.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	detail extensive
Input errors	<p>Input errors on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—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.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant threshold.</li> <li>• <b>Policed discards</b>—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.</li> <li>• <b>Resource errors</b>—Sum of B chip Tx drops and IXP Tx net transmit drops.</li> </ul>	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"> <li>• <b>Carrier transitions</b> —Number of times the interface has gone from <b>down</b> to <b>up</b>. 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.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—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.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of B chip Tx drops and IXP Tx net transmit drops.</li> </ul>	extensive
<b>Logical Interface</b>		
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 32: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Flags</b>	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
<b>Encapsulation</b>	Type of encapsulation configured on the logical interface.	All levels
<b>PPP parameters</b>	PPP status: <ul style="list-style-type: none"> <li>• LCP restart timer—Length of time (in milliseconds) between successive Link Control Protocol (LCP) configuration requests.</li> <li>• NCP restart timer—Length of time (in milliseconds) between successive Network Control Protocol (NCP) configuration requests.</li> </ul>	<b>detail</b>
<b>PPPoE</b>	PPPoE status: <ul style="list-style-type: none"> <li>• <b>State</b>—State of the logical interface (<b>up</b> or <b>down</b>).</li> <li>• <b>Session ID</b>—PPPoE session ID.</li> <li>• <b>Service name</b>—Type of service required. Can be used to indicate an Internet service provider (ISP) name or a class or quality of service.</li> <li>• <b>Configured AC name</b>—Configured access concentrator name.</li> <li>• <b>Auto-reconnect timeout</b>—Time after which to try to reconnect after a PPPoE session is terminated, in seconds.</li> <li>• <b>Idle Timeout</b>—Length of time (in seconds) that a connection can be idle before disconnecting.</li> <li>• <b>Underlying interface</b>—Interface on which PPPoE is running.</li> </ul>	All levels
<b>Link</b>	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
<b>Traffic statistics</b>	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.	<b>detail extensive</b>
<b>IPv6 transit statistics</b>	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. <p><b>NOTE:</b> 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"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Local statistics</b>	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.	<b>detail extensive</b>

Table 32: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Transit statistics</b>	<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><b>NOTE:</b> The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p>	<b>detail extensive</b>
<b>Keepalive settings</b>	<p>(PPP and HDLC) Configured settings for keepalives.</p> <ul style="list-style-type: none"> <li><b>interval seconds</b>—The time in seconds between successive keepalive requests. The range is <b>10</b> seconds through <b>32,767</b> seconds, with a default of <b>10</b> seconds.</li> <li><b>down-count number</b>—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is <b>1</b> through <b>255</b>, with a default of <b>3</b>.</li> <li><b>up-count number</b>—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is <b>1</b> through <b>255</b>, with a default of <b>1</b>.</li> </ul>	<b>detail extensive</b>
<b>Keepalive statistics</b>	<p>(PPP and HDLC) Information about keepalive packets.</p> <ul style="list-style-type: none"> <li><b>Input</b>—Number of keepalive packets received by PPP. <ul style="list-style-type: none"> <li><b>(last seen 00:00:00 ago)</b>—Time the last keepalive packet was received, in the format <i>hh:mm:ss</i>.</li> </ul> </li> <li><b>Output</b>—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"> <li><b>(last seen 00:00:00 ago)</b>—Time the last keepalive packet was sent, in the format <i>hh:mm:ss</i>.</li> </ul> </li> </ul> <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>	<b>detail extensive</b>
<b>Input packets</b>	Number of packets received on the logical interface.	None specified
<b>Output packets</b>	Number of packets transmitted on the logical interface.	None specified
<b>LCP state</b>	<p>(PPP) Link Control Protocol state.</p> <ul style="list-style-type: none"> <li><b>Conf-ack-received</b>—Acknowledgement was received.</li> <li><b>Conf-ack-sent</b>—Acknowledgement was sent.</li> <li><b>Conf-req-sent</b>—Request was sent.</li> <li><b>Down</b>—LCP negotiation is incomplete (not yet completed or has failed).</li> <li><b>Not-configured</b>—LCP is not configured on the interface.</li> <li><b>Opened</b>—LCP negotiation is successful.</li> </ul>	none <b>detail extensive</b>

Table 32: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>NCP state</b>	(PPP) Network Control Protocol state.  <ul style="list-style-type: none"> <li>• <b>Conf-ack-received</b>—Acknowledgement was received.</li> <li>• <b>Conf-ack-sent</b>—Acknowledgement was sent.</li> <li>• <b>Conf-req-sent</b>—Request was sent.</li> <li>• <b>Down</b>—NCP negotiation is incomplete (not yet completed or has failed).</li> <li>• <b>Not-configured</b>—NCP is not configured on the interface.</li> <li>• <b>Opened</b>—NCP negotiation is successful.</li> </ul>	<b>detail extensive none</b>
<b>CHAP state</b>	(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction.  <ul style="list-style-type: none"> <li>• <b>Chap-Chal-received</b>—Challenge was received but response not yet sent.</li> <li>• <b>Chap-Chal-sent</b>—Challenge was sent.</li> <li>• <b>Chap-Resp-received</b>—Response was received for the challenge sent, but CHAP has not yet moved into the Success state. (Most likely with RADIUS authentication.)</li> <li>• <b>Chap-Resp-sent</b>—Response was sent for the challenge received.</li> <li>• <b>Closed</b>—CHAP authentication is incomplete.</li> <li>• <b>Failure</b>—CHAP authentication failed.</li> <li>• <b>Not-configured</b>—CHAP is not configured on the interface.</li> <li>• <b>Success</b>—CHAP authentication was successful.</li> </ul>	<b>none detail extensive</b>
<b>Protocol</b>	Protocol family configured on the logical interface.	<b>detail extensive none</b>
<b><i>protocol-family</i></b>	Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.	<b>brief</b>
<b>MTU</b>	MTU size on the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route table</b>	Routing table in which the logical interface address is located. For example, <b>0</b> refers to the routing table <b>inet.0</b> .	<b>detail extensive none</b>
<b>Flags</b>	Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>Addresses, Flags</b>	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> .	<b>detail extensive none</b>
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive none</b>
<b>Local</b>	IP address of the logical interface.	<b>detail extensive none</b>
<b>Broadcast</b>	Broadcast address.	<b>detail extensive none</b>

## 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: 00:00:5e:00:53:01,
    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: 203.0.113.1

```

### 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, 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 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: 00:00:5e:00:53:12,
    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: 203.0.113.2, Local: 203.0.113.1, Broadcast: Unspecified,
Generation: 206
```

## show interfaces demux0 (Demux Interfaces)

<b>Syntax</b>	show interfaces demux0 <i>logical-interface-number</i> <brief   detail   extensive   terse> <descriptions> <media> <snmp-index <i>snmp-index</i> > <statistics>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0.
<b>Description</b>	(MX Series and M Series routers only) Display status information about the specified demux interface.
<b>Options</b>	<p><b>none</b>—Display standard information about the specified demux interface.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show interfaces demux0 (Demux) on page 537</a> <a href="#">show interfaces demux0 (PPPoE over Aggregated Ethernet) on page 538</a> <a href="#">show interfaces demux0 extensive (Targeted Distribution for Aggregated Ethernet Links) on page 539</a> <a href="#">show interfaces demux0 (ACI Interface Set Configured) on page 539</a>
<b>Output Fields</b>	Table 33 on page 531 lists the output fields for the <b>show interfaces demux0</b> (Demux Interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 33: 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 33: 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 ( <b>Up</b> or <b>Down</b> ).	detail extensive none
Admin	Administrative state of the interface ( <b>Up</b> or <b>Down</b> ).	terse
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
Link	Status of the physical link ( <b>Up</b> or <b>Down</b> ).	terse
Targeting summary	Status of aggregated Ethernet links that are configured with targeted distribution ( <b>primary</b> or <b>backup</b> )	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. <b>Software-Pseudo</b> 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: <b>Internal</b> (1) or <b>External</b> (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 33: 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 <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .	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"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> <li>• <b>IPv6 transit statistics</b>—Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled.</li> </ul> <p><b>NOTE:</b> These fields include dropped traffic and exception traffic, as those fields are not separately defined.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	detail extensive

Table 33: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Input errors</b>	Input errors on the interface whose definitions are as follows: <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—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.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant packet threshold.</li> <li>• <b>Policed discards</b>—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.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>
<b>Input Rate</b>	Input rate in bits per second (bps) and packets per second (pps).	<b>none</b>
<b>Output errors</b>	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. 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.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—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.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>
<b>Output Rate</b>	Output rate in bps and pps.	<b>none</b>
<b>Logical Interface</b>		
<b>Logical interface</b>	Name of the logical interface.	<b>brief detail extensive none</b>
<b>Index</b>	Index number of the logical interface, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP interface index number for the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail</b>
<b>Flags</b>	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> .	<b>brief detail extensive none</b>
<b>Encapsulation</b>	Encapsulation on the logical interface.	<b>brief extensive none</b>

Table 33: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>ACI VLAN: Dynamic Profile</b>	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.	<b>brief detail extensive</b> none
<b>Demux</b>	Specific IP demultiplexing (demux) values: <ul style="list-style-type: none"> <li>• <b>Underlying interface</b>—The underlying interface that the demux interface uses.</li> <li>• <b>Index</b>—Index number of the logical interface.</li> <li>• <b>Family</b>—Protocol family configured on the logical interface.</li> <li>• <b>Source prefixes, total</b>—Total number of source prefixes for the underlying interface.</li> <li>• <b>Destination prefixes, total</b>—Total number of destination prefixes for the underlying interface.</li> <li>• <b>Prefix</b>—inet family prefix.</li> </ul>	<b>detail extensive</b> none
<i>protocol-family</i>	Protocol family configured on the logical interface.	<b>brief</b>
<b>Traffic statistics</b>	Number and rate of bytes and packets received and transmitted on the specified interface set. <ul style="list-style-type: none"> <li>• <b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface set.</li> <li>• <b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface set.</li> <li>• <b>IPv6 transit statistics</b>—Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.</li> </ul> <p><b>NOTE:</b> 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"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Local statistics</b>	Number of transit bytes and packets received and transmitted on the local interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>

Table 33: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Transit statistics</b>	<p>Number and rate of bytes and packets transiting the switch.</p> <p><b>NOTE:</b> 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"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>IPv6 Transit statistics</b>	<p>Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.</p> <p><b>NOTE:</b> 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"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Input packets</b>	Number of packets received on the interface.	none
<b>Output packets</b>	Number of packets transmitted on the interface.	none
<b>Protocol</b>	Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .	<b>detail extensive</b> none
<b>MTU</b>	Maximum transmission unit size on the logical interface.	<b>detail extensive</b> none
<b>Maximum labels</b>	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	<b>detail extensive</b> none
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route table</b>	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	<b>detail extensive</b>
<b>Flags</b>	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive</b> none
<b>Mac-Validate Failures</b>	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	<b>detail extensive</b> none
<b>Addresses, Flags</b>	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive</b> none
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive</b> <b>statistics</b> none

Table 33: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Local</b>	IP address of the logical interface.	<b>detail extensive terse</b> none
<b>Remote</b>	IP address of the remote interface.	<b>terse</b>
<b>Broadcast</b>	Broadcast address of the logical interlace.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Link</b>	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.	<b>detail extensive none</b>
<b>Dynamic-profile</b>	Name of the PPPoE dynamic profile assigned to the underlying interface.	<b>detail extensive none</b>
<b>Service Name Table</b>	Name of the PPPoE service name table assigned to the PPPoE underlying interface.	<b>detail extensive none</b>
<b>Max Sessions</b>	Maximum number of dynamic PPPoE logical interfaces that the router can activate on the underlying interface.	<b>detail extensive none</b>
<b>Duplicate Protection</b>	State of duplicate protection: <b>On</b> or <b>Off</b> . 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.	<b>detail extensive none</b>
<b>Direct Connect</b>	State of the configuration to ignore DSL Forum VSAs: <b>On</b> or <b>Off</b> . When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	<b>detail extensive none</b>
<b>AC Name</b>	Name of the access concentrator.	<b>detail extensive none</b>

## 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: 203.0.113/24
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: 203.0.113/24, Local: 203.0.113.13, Broadcast: 203.0.113.255,

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: Sendbroadcast-pkt-to-re, Unnumbered
  Donor interface: lo0.0 (Index 322)
  Preferred source address: 203.0.113.202
  Addresses, Flags: Primary Is-Default Is-Primary
    Local: 203.0.113.119
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

<b>Syntax</b>	<code>show interfaces filters</code> <code>&lt;interface-name&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.
<b>Description</b>	Display all firewall filters that are installed on each interface in a system.
<b>Options</b>	<b>none</b> —Display filter information about all interfaces.  <b>interface-name</b> —(Optional) Display filter information about a particular interface.
<b>Additional Information</b>	For information about how to configure firewall filters, see the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide</i> . For related operational mode commands, see the <a href="#">CLI Explorer</a> .
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show interfaces filters on page 541</a> <a href="#">show interfaces filters (Interface-Name) on page 541</a> <a href="#">show interfaces filters (PTX Series Packet Transport Routers) on page 541</a>
<b>Output Fields</b>	<a href="#">Table 34 on page 540</a> lists the output fields for the <b>show interfaces filters</b> command. Output fields are listed in the approximate order in which they appear.

**Table 34: show interfaces filters Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of the interface.
<b>Admin</b>	Interface state: <b>up</b> or <b>down</b> .
<b>Link</b>	Link state: <b>up</b> or <b>down</b> .
<b>Proto</b>	Protocol configured on the interface.
<b>Input Filter</b>	Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.
<b>Output Filter</b>	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 l2-routing-instance

<b>Syntax</b>	<code>show interfaces l2-routing-instance <i>routing-instance-name</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 16.1.
<b>Description</b>	Display information about core-facing physical interfaces in the specified routing instance. The routing instance must be configured for Layer 2 wholesale operations or the command displays no output.
<b>Options</b>	<i>routing-instance-name</i> —Name of the routing instance.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Layer 2 Wholesale with ANCP-Triggered VLANs Overview on page 93</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show interfaces l2-routing-instance on page 542</a>
<b>Output Fields</b>	<a href="#">Table 35 on page 542</a> lists the output fields for the <b>show interfaces l2-routing-instance</b> command. Output fields are listed in the approximate order in which they appear.

**Table 35: show interfaces l2-routing-instance Output Fields**

Field Name	Field Description
Interface	Name of the core-facing physical interface.
VLAN Id	Trunk VLAN ID assigned to the core-facing physical interface within its input VLAN map.
Inner VLAN Id total	Aggregate number of VLAN IDs assigned to this physical interface based on the inner VLAN ID-swap-ranges configured for the interface.
Inner VLAN range	Inner VLAN ID swap range; numbers that can be used for swapping inner VLAN IDs.
Inner VLAN range total	Size of the inner VLAN ID swap range.
Inner VLAN id use count	Number of inner VLAN ID tags in use.
Inner VLAN id free count	Number of inner VLAN ID tags not in use.

## Sample Output

### show interfaces l2-routing-instance

```

user@host> show interfaces l2-routing-instance NSP1
Interface: ge-1/1/1.0
VLAN Id: 100
  Inner VLAN Id total: 6
  Inner VLAN range: 15-20

```

```
Inner VLAN range total   : 6
Inner VLAN id use count  : 1
Inner VLAN id free count : 5
```

## show interfaces routing

<b>Syntax</b>	show interfaces routing <brief   detail> <interface-name> <logical-system (all   <i>logical-system-name</i> )>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Options</b>	<p><b>none</b>—Display standard information about the state of all router interfaces on all logical systems.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface-name</b>—(Optional) Name of a specific interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Additional Information</b>	For information about how to configure routing protocols, see the <i>Junos OS Routing Protocols Library</i> . For information about related operational mode commands for routing instances and protocols, see the <a href="#">CLI Explorer</a> .
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show interfaces routing brief on page 545</a> <a href="#">show interfaces routing brief (TX Matrix Plus Router) on page 546</a> <a href="#">show interfaces routing detail on page 546</a> <a href="#">show interfaces routing detail (TX Matrix Plus Router) on page 547</a>
<b>Output Fields</b>	<a href="#">Table 36 on page 544</a> lists the output fields for the <b>show interfaces routing</b> command. Output fields are listed in the approximate order in which they appear.

**Table 36: show interfaces routing Output Fields**

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of the physical interface.	none <b>brief</b>
<b>State</b>	State of the physical interface: <b>Up</b> or <b>Down</b> .	none <b>brief</b>
<b>Addresses</b>	Protocols and addresses configured on the interface.	none <b>brief</b>
<b>Index</b>	Interface index number, which reflects its initialization sequence.	<b>detail</b>
<b>RefCount</b>	Number of references to the interface in the routing software.	<b>detail</b>

Table 36: show interfaces routing Output Fields (*continued*)

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

## 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
so-5/0/1.0      Up    MPLS  enabled
                  ISO   enabled
                  INET   192.168.2.130
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
1o0.0           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  203.0.113.91
                  ISO   enabled
                  MPLS  enabled
ge-23/0/3.0    Up    INET  203.0.113.81
                  ISO   enabled
                  MPLS  enabled
ge-23/0/2.0    Up    INET  203.0.113.71
                  ISO   enabled
                  MPLS  enabled
ge-23/0/1.0    Up    INET  203.0.113.61
                  ISO   enabled
                  MPLS  enabled
ge-23/0/0.0    Up    INET  203.0.113.51
                  ISO   enabled
                  MPLS  enabled
ge-31/0/7.599  Up    INET  192.0.2.93
ge-31/0/7.598  Up    INET  192.0.2.89
ge-31/0/7.597  Up    INET  192.0.2.85
ge-31/0/7.596  Up    INET  192.0.2.81
ge-31/0/7.595  Up    INET  192.0.2.77
ge-31/0/7.594  Up    INET  192.0.2.73
...
ixgbe1.0       Up    INET  203.0.113.34
                  INET  198.51.100.4
                  INET6  fe80::200:1ff:fe22:4
                  INET6  fec0::a:22:0:4
ixgbe0.0       Up    INET  203.0.113.34
                  INET  198.51.100
                  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 203.0.113.91
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 203.0.113.93
    Destination: 203.0.113.0/30
    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 203.0.113.81
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 2.8.1.3
    Destination: 203.0.113.80/30
    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 203.0.113.71
State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
Preference 0, metric 0, MTU 1500 bytes
Broadcast address 203.0.113.73
Destination: 203.0.113.70/30
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 203.0.113.61
State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
Preference 0, metric 0, MTU 1500 bytes
Broadcast address 203.0.113.63
...
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 203.0.113.34
State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
Preference 0, metric 0, MTU 1500 bytes
Broadcast address 203.0.113.255
Destination: 203.0.113.0/8
System flags: <Is-Preferred>
INET address 198.51.100.4
State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
Preference 0, metric 0, MTU 1500 bytes
Broadcast address 191.255.255.255
Destination: 192.0.2.0/2
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 203.0.113.34
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 203.0.113.255
    Destination: 203.0.113.0/8
    System flags: <Is-Preferred>
  INET address 198.51.100.4
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 191.255.255.255
    Destination: 192.0.2.0/2
    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

<b>Syntax</b>	show interfaces routing-instance ( <i>instance-name</i>   all) <brief   detail   extensive   terse>
<b>Release Information</b>	Command introduced in Junos OS Release 9.1.
<b>Description</b>	Display information about the interfaces configured for either a specific routing instance or for all of the routing instances.
<b>Options</b>	<p><b>all</b>—Display information about all of the interfaces configured for all of the routing instances on the router.</p> <p><b><i>instance-name</i></b>—Display information about the interfaces configured for the specified routing instance.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show interfaces routing-instance terse on page 550</a> <a href="#">show interfaces routing-instance all on page 550</a> <a href="#">show interfaces routing-instance extensive on page 550</a>
<b>Output Fields</b>	The output fields from the <b>show interfaces routing-instance</b> command are identical to those produced by the <b>show interfaces <i>interface-name</i></b> 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   203.0.113.1/24
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   203.0.113.1/32
```

### 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: 192.0.2/24, Local: 192.0.2.51, Broadcast: 192.0.2.255,
Generation: 263
```

## show ppp interface

<b>Syntax</b>	<code>show ppp interface <i>interface-name</i></code> <code>&lt;extensive  terse&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 7.5.
<b>Description</b>	Display information about PPP interfaces.
<b>Options</b>	<i>interface-name</i> —Name of a logical interface.  <b>extensive   terse</b> —(Optional) Display the specified level of output.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ppp interface on page 560</a> <a href="#">show ppp interface extensive on page 560</a> <a href="#">show ppp interface terse on page 560</a>
<b>Output Fields</b>	<a href="#">Table 37 on page 552</a> lists the output fields for the <b>show ppp interface</b> command. Output fields are listed in the approximate order in which they appear.

Table 37: show ppp interface Output Fields

Field Name	Field Description	Level of Output
<b>Session</b>	Name of the logical interface on which the session is running.	All levels
<b>Type</b>	Session type: PPP.	All levels
<b>Phase</b>	PPP process phase: <b>Authenticate</b> , <b>Pending</b> , <b>Establish</b> , <b>LCP</b> , <b>Network</b> , <b>Disabled</b> , and <b>Tunneled</b> .	All levels
<b>Session flags</b>	Special conditions present in the session: <b>Bundled</b> , <b>TCC</b> , <b>No-keepalives</b> , <b>Looped</b> , <b>Monitored</b> , and <b>NCP-only</b> .	All levels
<b><i>protocol</i> State</b>	Protocol state information. See specific protocol state fields for information.	None specified
<b>AUTHENTICATION</b>	Challenge-Handshake Authentication Protocol (CHAP) authentication state information or Password Authentication Protocol (PAP) state information. See the <b>Authentication</b> field description for further information.	None specified

Table 37: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Keepalive settings</b>	<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"> <li>• <b>Interval</b>—Time in seconds between successive keepalive requests. Keepalive aging timeout is calculated as a product of the <b>interval</b> and <b>Down-count</b> 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.</li> <li>• <b>Up-count</b>—The number of keepalive packets a destination must receive to change a link's status from down to up.</li> <li>• <b>Down-count</b>—The number of keepalive packets a destination must fail to receive before the network takes down a link.</li> </ul>	<b>extensive</b>
<b>RE Keepalive statistics</b>	<p>Keepalive statistics for the packets handled by the Routing Engine.</p> <ul style="list-style-type: none"> <li>• <b>LCP echo req Tx</b>—LCP echo requests sent from the Routing Engine.</li> <li>• <b>LCP echo req Rx</b>—LCP echo requests received at the Routing Engine.</li> <li>• <b>LCP echo rep Tx</b>—LCP echo responses sent from the Routing Engine.</li> <li>• <b>LCP echo rep Rx</b>—LCP echo responses received at the Routing Engine.</li> <li>• <b>LCP echo req timeout</b>—Number of keepalive packets where the keepalive aging timer has expired.</li> <li>• <b>LCP Rx echo req Magic Num Failures</b>—LCP echo requests where the magic numbers shared between the PPP peers during LCP negotiation did not match.</li> <li>• <b>LCP Rx echo rep Magic Num Failures</b>—LCP echo responses where the magic numbers shared between the PPP peers during LCP negotiation did not match.</li> </ul>	<b>extensive</b>

Table 37: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
LCP	<p><b>LCP information:</b></p> <ul style="list-style-type: none"> <li>• <b>State</b>—LCP protocol state (all platforms except M120 and M320 routers): <ul style="list-style-type: none"> <li>• <b>Ack-rcvd</b>—A Configure-Request has been sent and a Configure-Ack has been received.</li> <li>• <b>Ack-sent</b>—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.</li> <li>• <b>Closed</b>—Link is not available for traffic.</li> <li>• <b>Opened</b>—Link is administratively available for traffic.</li> <li>• <b>Req-sent</b>—An attempt has been made to configure the connection.</li> </ul> </li> <li>• <b>State</b>—LCP protocol state (M120 and M320 routers): <ul style="list-style-type: none"> <li>• <b>Ack-rcvd</b>—A Configure-Request has been sent and a Configure-Ack has been received.</li> <li>• <b>Ack-sent</b>—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.</li> <li>• <b>Closed</b>—Link is available (up), but no Open has occurred.</li> <li>• <b>Closing</b>—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.</li> <li>• <b>Opened</b>—Link is administratively available for traffic. A Configure-Ack has been both sent and received.</li> <li>• <b>Req-sent</b>—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received.</li> <li>• <b>Starting</b>—An administrative Open has been initiated, but the lower layer is still unavailable (Down).</li> <li>• <b>Stopped</b>—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack.</li> <li>• <b>Stopping</b>—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.</li> </ul> </li> <li>• <b>Last started</b>—LCP state start time.</li> <li>• <b>Last completed</b>—LCP state completion time.</li> </ul>	extensive

Table 37: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<ul style="list-style-type: none"> <li>• <b>Negotiated options:</b> <ul style="list-style-type: none"> <li>• <b>ACFC</b>—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.</li> <li>• <b>Asynchronous map</b>—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.</li> <li>• <b>Authentication protocol</b>—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.</li> <li>• <b>Authentication algorithm</b>—Type of authentication algorithm. The Message Digest algorithm (MD5) is the only algorithm supported.</li> <li>• <b>Endpoint discriminator class</b>—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.</li> <li>• <b>Magic number</b>—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.</li> <li>• <b>MRU</b>—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.</li> <li>• <b>MRRU</b>—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.</li> <li>• <b>Multilink header suspendable classes</b>—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.</li> <li>• <b>Multilink header format classes</b>—For MLPPP, an LCP option that advises the peer that the implementation wishes to receive fragments with a format given by the code number.</li> <li>• <b>PFC</b>—Protocol-Field-Compression. A configuration option that provides a method to negotiate the compression of the PPP Protocol field.</li> <li>• <b>short sequence</b>—For MLPPP, an option that advises the peer that the implementation wishes to receive fragments with short, 12-bit sequence numbers.</li> </ul> </li> </ul>	

Table 37: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Authentication</b>	<p>CHAP or PAP authentication state information. For CHAP authentication:</p> <ul style="list-style-type: none"> <li>• <b>Chap-ans-rcvd</b>—Packet was sent from the peer, indicating that the peer received the <b>Chap-resp-sent</b> packet.</li> <li>• <b>Chap-ans-sent</b>—Packet was sent from the authenticator, indicating that the authenticator received the peer's <b>Chap-resp-rcvd</b> packet.</li> <li>• <b>Chap-chal-rcvd</b>—Challenge packet has been received by the peer.</li> <li>• <b>Chap-chal-sent</b>—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.</li> <li>• <b>Chap-resp-rcvd</b>—CHAP response packet has been received by the authenticator.</li> <li>• <b>Chap-resp-sent</b>—CHAP response packet has been sent to the authenticator.</li> <li>• <b>Closed</b>—Link is not available for authentication.</li> <li>• <b>Failure</b>—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.</li> <li>• <b>Success</b>—Authenticator compares the response value in the response packet from the peer with its own response value, and the value matches. Authentication is successful.</li> </ul> <p>For PAP authentication:</p> <ul style="list-style-type: none"> <li>• <b>Pap-resp-sent</b>—PAP response sent to peer (ACK/NACK).</li> <li>• <b>Pap-req-rcvd</b>—PAP request packet received from peer.</li> <li>• <b>Pap-resp-rcvd</b>—PAP response received from the peer (ACK/NACK).</li> <li>• <b>Pap-req-sent</b>—PAP request packet sent to the peer.</li> <li>• <b>Closed</b>—Link is not available for authentication.</li> <li>• <b>Failure</b>—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.</li> <li>• <b>Success</b>—Authenticator compares the response value in the response packet from the peer with its own response value, and the value matches. Authentication is successful.</li> </ul>	None specified

Table 37: 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"> <li>• <b>State</b>—(All platforms except M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> <li>• <b>Ack-rcvcd</b>—A Configure-Request has been sent and a Configure-Ack has been received.</li> <li>• <b>Ack-sent</b>—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.</li> <li>• <b>Closed</b>—Link is not available for traffic.</li> <li>• <b>Opened</b>—Link is administratively available for traffic.</li> <li>• <b>Req-sent</b>—An attempt has been made to configure the connection.</li> </ul> </li> <li>• <b>State</b>—(M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> <li>• <b>Ack-rcvcd</b>—A Configure-Request has been sent and a Configure-Ack has been received.</li> <li>• <b>Ack-sent</b>—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.</li> <li>• <b>Closed</b>—Link is available (up), but no Open has occurred.</li> <li>• <b>Closing</b>—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.</li> <li>• <b>Opened</b>—Link is administratively available for traffic. A Configure-Ack has been both sent and received.</li> <li>• <b>Req-sent</b>—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received.</li> <li>• <b>Starting</b>—An administrative Open has been initiated, but the lower layer is still unavailable (Down).</li> <li>• <b>Stopped</b>—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack.</li> <li>• <b>Stopping</b>—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.</li> </ul> </li> <li>• <b>Last started</b>—IPCP state start time.</li> <li>• <b>Last completed</b>—IPCP state authentication completion time.</li> <li>• <b>Negotiated options</b>: <ul style="list-style-type: none"> <li>• <b>compression protocol</b>—Negotiate the use of a specific compression protocol. By default, compression is not enabled.</li> <li>• <b>local address</b>—Desired local address of the sender of a Configure-Request. If all four octets are set to zero, the peer provides the IP address.</li> <li>• <b>primary DNS server</b>—Negotiate with the remote peer to select the address of the primary DNS server to be used on the local end of the link.</li> <li>• <b>primary WINS server</b>—Negotiate with the remote peer to select the address of the primary WINS server to be used on the local end of the link.</li> <li>• <b>remote address</b>—IP address of the remote end of the link in dotted quad notation.</li> <li>• <b>secondary DNS server</b>—Negotiate with the remote peer to select the address of the secondary DNS server to be used on the local end of the link.</li> <li>• <b>secondary WINS server</b>—Negotiate with the remote peer to select the address of the secondary WINS server to be used on the local end of the link.</li> </ul> </li> <li>• <b>Negotiation mode</b>—PPP Network Control Protocol (NCP) negotiation mode configured for IPCP: <b>Active</b> or <b>Passive</b></li> </ul>	extensive

Table 37: 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"> <li>• <b>State</b>—(All platforms except M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> <li>• <b>Ack-rcvcd</b>—A Configure-Request has been sent and a Configure-Ack has been received.</li> <li>• <b>Ack-sent</b>—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.</li> <li>• <b>Closed</b>—Link is not available for traffic.</li> <li>• <b>Opened</b>—Link is administratively available for traffic.</li> <li>• <b>Req-sent</b>—An attempt has been made to configure the connection.</li> </ul> </li> <li>• <b>State</b>—(M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> <li>• <b>Ack-rcvcd</b>—A Configure-Request has been sent and a Configure-Ack has been received.</li> <li>• <b>Ack-sent</b>—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.</li> <li>• <b>Closed</b>—Link is available (up), but no Open has occurred.</li> <li>• <b>Closing</b>—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.</li> <li>• <b>Opened</b>—Link is administratively available for traffic. A Configure-Ack has been both sent and received.</li> <li>• <b>Req-sent</b>—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received.</li> <li>• <b>Starting</b>—An administrative Open has been initiated, but the lower layer is still unavailable (Down).</li> <li>• <b>Stopped</b>—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack.</li> <li>• <b>Stopping</b>—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.</li> </ul> </li> <li>• <b>Last started</b>—IPv6CP state start time.</li> <li>• <b>Last completed</b>—IPv6CP state authentication completion time.</li> <li>• <b>Negotiated options</b>: <ul style="list-style-type: none"> <li>• <b>local interface identifier</b>—Desired local address of the sender of a Configure-Request. If all four octets are set to zero, the peer provides the IP address.</li> <li>• <b>remote interface identifier</b>—IP address of the remote end of the link in dotted quad notation.</li> </ul> </li> <li>• <b>Negotiation mode</b>—PPP Network Control Protocol (NCP) negotiation mode configured for IPv6CP: <b>Active</b> or <b>Passive</b></li> </ul>	<b>extensive</b>

Table 37: 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"> <li>• <b>State:</b> <ul style="list-style-type: none"> <li>• <b>Ack-rcvd</b>—Configure-Request has been sent and Configure-Ack has been received.</li> <li>• <b>Ack-sent</b>—Configure-Request and Configure-Ack have both been sent, but Configure-Ack has not yet been received.</li> <li>• <b>Closed</b>—Link is not available for traffic.</li> <li>• <b>Opened</b>—Link is administratively available for traffic.</li> <li>• <b>Req-sent</b>—Attempt has been made to configure the connection.</li> </ul> </li> <li>• <b>Last started</b>—OSINLCP state start time.</li> <li>• <b>Last completed</b>—OSINLCP state completion time.</li> </ul>	extensive
TAGCP	<p>TAGCP information.</p> <ul style="list-style-type: none"> <li>• <b>State</b>—(All platforms except M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> <li>• <b>Ack-rcvd</b>—A Configure-Request has been sent and a Configure-Ack has been received.</li> <li>• <b>Ack-sent</b>—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.</li> <li>• <b>Closed</b>—Link is not available for traffic.</li> <li>• <b>Opened</b>—Link is administratively available for traffic.</li> <li>• <b>Req-sent</b>—An attempt has been made to configure the connection.</li> </ul> </li> <li>• <b>State</b>—(M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> <li>• <b>Ack-rcvd</b>—A Configure-Request has been sent and a Configure-Ack has been received.</li> <li>• <b>Ack-sent</b>—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.</li> <li>• <b>Closed</b>—Link is available (up), but no Open has occurred.</li> <li>• <b>Closing</b>—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.</li> <li>• <b>Opened</b>—Link is administratively available for traffic. A Configure-Ack has been both sent and received.</li> <li>• <b>Req-sent</b>—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received.</li> <li>• <b>Starting</b>—An administrative Open has been initiated, but the lower layer is still unavailable (Down).</li> <li>• <b>Stopped</b>—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack.</li> <li>• <b>Stopping</b>—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.</li> </ul> </li> <li>• <b>Last started</b>—TAGCP state start time.</li> <li>• <b>Last completed</b>—TAGCP state authentication completion time.</li> </ul>	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: 203.0.113.21, Remote address: 203.0.113.22
  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,

192.0.2.0). If you specify the IP address as a prefix with a netmask (for example, 192.0.2.0/32), 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 582](#)
- *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 569](#)  
[show subscribers \(IPv6\) on page 569](#)  
[show subscribers \(IPv4 and IPv6 Dual Stack\) on page 569](#)  
[show subscribers \(LNS on MX Series Routers\) on page 570](#)  
[show subscribers \(L2TP Switched Tunnels\) on page 570](#)  
[show subscribers client-type dhcp detail on page 570](#)  
[show subscribers client-type vlan-oob detail on page 570](#)  
[show subscribers count on page 571](#)  
[show subscribers address detail \(IPv6\) on page 571](#)

[show subscribers detail \(IPv4\) on page 571](#)  
[show subscribers detail \(IPv6\) on page 572](#)  
[show subscribers detail \(IPv6 Static Demux Interface\) on page 572](#)  
[show subscribers detail \(L2TP LNS Subscribers on MX Series Routers\) on page 572](#)  
[show subscribers detail \(L2TP Switched Tunnels\) on page 573](#)  
[show subscribers detail \(Tunneled Subscriber\) on page 573](#)  
[show subscribers detail \(IPv4 and IPv6 Dual Stack\) on page 573](#)  
[show subscribers detail \(ACI Interface Set Session\) on page 574](#)  
[show subscribers detail \(PPPoE Subscriber Session with ACI Interface Set\) on page 574](#)  
[show subscribers extensive on page 575](#)  
[show subscribers extensive \(RPF Check Fail Filter\) on page 575](#)  
[show subscribers extensive \(L2TP LNS Subscribers on MX Series Routers\) on page 575](#)  
[show subscribers extensive \(IPv4 and IPv6 Dual Stack\) on page 576](#)  
[show subscribers extensive \(ADF Rules \) on page 577](#)  
[show subscribers extensive \(Effective Shaping-Rate\) on page 577](#)  
[show subscribers aci-interface-set-name detail \(Subscriber Sessions Using Specified ACI Interface Set\) on page 577](#)  
[show subscribers agent-circuit-identifier detail \(Subscriber Sessions Using Specified ACI Substring\) on page 578](#)  
[show subscribers interface extensive on page 578](#)  
[show subscribers logical-system terse on page 579](#)  
[show subscribers physical-interface count on page 579](#)  
[show subscribers routing-instance inst1 count on page 579](#)  
[show subscribers stacked-vlan-id detail on page 579](#)  
[show subscribers stacked-vlan-id vlan-id detail \(Combined Output\) on page 580](#)  
[show subscribers stacked-vlan-id vlan-id interface detail \(Combined Output for a Specific Interface\) on page 580](#)  
[show subscribers user-name detail on page 580](#)  
[show subscribers vlan-id on page 580](#)  
[show subscribers vlan-id detail on page 580](#)  
[show subscribers vpi vci extensive \(PPPoE-over-ATM Subscriber Session\) on page 581](#)  
[show subscribers address detail \(Enhanced Subscriber Management\) on page 581](#)

**Output Fields** [Table 38 on page 565](#) lists the output fields for the **show subscribers** command. Output fields are listed in the approximate order in which they appear.

**Table 38: show subscribers Output Fields**

Field Name	Field Description
<b>Interface</b>	<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>
<b>IP Address/VLAN ID</b>	<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 <b>Tunnel-switched</b>.</p>
<b>User Name</b>	Name of subscriber.

Table 38: 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.  (MX Series) This field displays 255.255.255.255 by default. For tunneled or terminated PPP subscribers only, this field displays the actual value of Framed-IP-Netmask when the SDB_FRAMED_PROTOCOL attribute in the session database is equal to AUTHD_FRAMED_PROTOCOL_PPP. This occurs in the use case where the LNS generates access-internal routes when it receives Framed-IP-Netmask from RADIUS during authorization. When it receives Framed-Pool from RADIUS, the pool mask is ignored and the default /32 mask is used.
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 <b>demux0.nnnn</b> (for example, <b>demux0.3221225472</b> ), to which access-internal and framed subscriber routes are mapped.
Interface Type	Whether the subscriber interface is <b>Static</b> or <b>Dynamic</b> .
Interface Set	Internally generated name of the dynamic ACI interface set used by the subscriber session.

Table 38: show subscribers Output Fields (*continued*)

Field Name	Field Description
Interface Set Type	Interface type of the ACI interface set: <b>Dynamic</b> . 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.
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 ( <b>Init</b> , <b>Configured</b> , <b>Active</b> , <b>Terminating</b> , <b>Tunneled</b> ).
L2TP State	Current state of the L2TP session, <b>Tunneled</b> or <b>Tunnel-switched</b> . When the value is <b>Tunnel-switched</b> , 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 <b>dhcp</b> 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 <b>vlan-oob</b> 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 <b>dhcp</b> 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 <b>vlan-oob</b> 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.

Table 38: show subscribers Output Fields (*continued*)

Field Name	Field Description
<b>ATM VPI</b>	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual path identifier (VPI) on the subscriber's physical interface.
<b>ATM VCI</b>	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual circuit identifier (VCI) for each VPI configured on the subscriber interface.
<b>Login Time</b>	Date and time at which the subscriber logged in.
<b>Effective shaping-rate</b>	Actual downstream traffic shaping rate for the subscriber, in kilobits per second.
<b>IPv4 rpf-check Fail Filter Name</b>	Name of the filter applied by the dynamic profile to IPv4 packets that fail the RPF check.
<b>IPv6 rpf-check Fail Filter Name</b>	Name of the filter applied by the dynamic profile to IPv6 packets that fail the RPF check.
<b>DHCP Options</b>	<b>len</b> = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options, as defined in RFC 2132.
<b>Session ID</b>	ID number for a subscriber service session.
<b>Underlying Session ID</b>	For DHCPv6 subscribers on a PPPoE network, displays the session ID of the underlying PPPoE interface.
<b>Service Sessions</b>	Number of service sessions (that is, a service activated using RADIUS CoA) associated with the subscribers.
<b>Service Session Name</b>	Service session profile name.
<b>Session Timeout (seconds)</b>	Number of seconds of access provided to the subscriber before the session is automatically terminated.
<b>Idle Timeout (seconds)</b>	Number of seconds subscriber can be idle before the session is automatically terminated.
<b>IPv6 Delegated Address Pool</b>	Name of the pool used for DHCPv6 prefix delegation.
<b>IPv6 Delegated Network Prefix Length</b>	Length of the prefix configured for the IPv6 delegated address pool.
<b>IPv6 Interface Address</b>	Address assigned by the Framed-Ipv6-Prefix AAA attribute.
<b>IPv6 Framed Interface Id</b>	Interface ID assigned by the Framed-Interface-Id AAA attribute.
<b>ADF IPv4 Input Filter Name</b>	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.

Table 38: show subscribers Output Fields (*continued*)

Field Name	Field Description
<b>ADF IPv4 Output Filter Name</b>	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.
<b>ADF IPv6 Input Filter Name</b>	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.
<b>ADF IPv6 Output Filter Name</b>	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.
<b>IPv4 Input Filter Name</b>	Name assigned to the IPv4 input filter (client or service session).
<b>IPv4 Output Filter Name</b>	Name assigned to the IPv4 output filter (client or service session).
<b>IPv6 Input Filter Name</b>	Name assigned to the IPv6 input filter (client or service session).
<b>IPv6 Output Filter Name</b>	Name assigned to the IPv6 output filter (client or service session).
<b>IFL Input Filter Name</b>	Name assigned to the logical interface input filter (client or service session).
<b>IFL Output Filter Name</b>	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 10                   WHOLESALE-CLIENT  default:default
demux0.1073741824   203.0.113.10        RETAILER1-CLIENT  default:default
demux0.1073741825   203.0.113.3         RETAILER1-CLIENT  test1:retailer1
demux0.1073741826   203.0.113.3         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 WHOLESALE-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      203.0.113.13      dualstackuser1@example1.com

```

```
default:ASP-1
*                2001:db8:1::/48
*                2001:db8:1:1::/64
pp0.1073741837   203.0.113.33          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     192.0.2.0           user@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     user@example.com default:default

si-2/1/0.1073741843 Tunnel-switched     user@example.com default:default
```

#### show subscribers client-type dhcp detail

```
user@host> show subscribers client-type dhcp detail
Type: DHCP
IP Address: 203.0.113.29
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux
MAC Address: 00:00:5e:00:53:98
State: Active
Radius Accounting ID: user :2304
Login Time: 2009-08-25 14:43:52 PDT

Type: DHCP
IP Address: 203.0.113.27
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: 00:00:5e:00:53:f3
State: Active
Radius Accounting ID: 1234 :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: 1234
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 203.0.113.137 detail
Type: PPPoE
User Name: pppoeTerV6User1Svc
IP Address: 203.0.113.137
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: 00:00:5e:00:53:53
Session Timeout (seconds): 31622400
Idle Timeout (seconds): 86400
State: Active
Radius Accounting ID: example 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: 203.0.113.29
IP Netmask: 255.255.0.0
Primary DNS Address: 192.0.2.0
Secondary DNS Address: 192.0.2.1
Primary WINS Address: 192.0.2.3
Secondary WINS Address: 192.0.2.4
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:00:5e:00:53:98
State: Active

```

```
Radius Accounting ID: example :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:ffff:1::/32
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:00:5e:00:53:03
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: user@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: user@example.net
IP Address: 203.0.113.58
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: user@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: 203.0.113.51
Remote IP Address: 192.0.2.0
Radius Accounting ID: 21
Session ID: 21
Login Time: 2013-01-18 03:01:11 PST

Type: L2TP
User Name: user@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: 203.0.113.31
Remote IP Address: 192.0.2.1
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: 203.0.113.13
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: 00:00:5e:00:53:02
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: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: test :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: 203.0.113.15
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: 00:00:5e:00:53:02
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:ffff:1::/32
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:00:5e:00:53:03
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: user@example.net
IP Address: 203.0.113.58
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: 203.0.113.13
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: 00:00:5e:00:53:02
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:db8:2016:1:1::1/64
IPv6 Framed Interface Id: 1:1:2:2
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::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: test :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
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 203.0.113.232;
        destination-address 198.51.100.0/24;
        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: 203.0.113.17

```

```
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:
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: 203.0.113.17
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: 00:00:5e:00:53:52
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: user@test.example.com
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Dynamic
```

```
Dynamic Profile Name: profile-vdemux-relay-23qos
MAC Address: 00:00:5e:00:53:04
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: user@test.example.com
IP Address: 192.0.2.0
IP Netmask: 255.255.255.0
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Static
MAC Address: 00:00:5e:00:53:04
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  203.0.113.3         RETAILER1-CLIENT   test1:retailer1
demux0.1073741826  203.0.113.4         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: 203.0.113.37
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: 00:00:5e:00:53:01
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: 203.0.113.2
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: pp0.0
Interface type: Static
MAC Address: 00:00:5e:00:53:02
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 203.0.113.111 detail
Type: DHCP
User Name: simple_filters_service
IP Address: 203.0.113.111
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: 00:00:5e:00:53:0f
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

---

<b>Syntax</b>	<code>show subscribers summary</code> <code>&lt;all&gt;</code> <code>&lt;detail   extensive   terse&gt;</code> <code>&lt;count&gt;</code> <code>&lt;physical-interface <i>physical-interface-name</i>&gt;</code> <code>&lt;logical-system <i>logical-system</i> pic   port   routing-instance <i>routing-instance</i>   slot&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.2.
<b>Description</b>	Display summary information for subscribers.
<b>Options</b>	<p><b>none</b>—Display summary information by state and client type for all subscribers.</p> <p><b>all</b>—(Optional) Display summary information by state, client type, and LS:RI.</p> <p><b>detail   extensive   terse</b>—(Not supported on MX Series routers) (Optional) Display the specified level of output.</p> <p><b>count</b>—(Not supported on MX Series routers) (Optional) Display the count of total subscribers and active subscribers for any specified option.</p> <p><b>logical-system <i>logical-system</i></b>—(Optional) Display subscribers whose logical system matches the specified logical system.</p> <p><b>physical-interface <i>physical-interface-name</i></b>—(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><b>pic</b>—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by PIC number and the total number of subscribers.</p> <p><b>port</b>—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by port number and the total number of subscribers.</p> <p><b>routing-instance <i>routing-instance</i></b>—(Optional) Display subscribers whose routing instance matches the specified routing instance.</p> <p><b>slot</b>—(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> <b>NOTE:</b> Due to display limitations, logical system and routing instance output values are truncated when necessary.</div>
<b>Required Privilege Level</b>	view

**Related Documentation**

- [show subscribers on page 562](#)

**List of Sample Output**

- [show subscribers summary on page 584](#)
- [show subscribers summary all on page 584](#)
- [show subscribers summary physical-interface on page 585](#)
- [show subscribers summary physical-interface pic on page 585](#)
- [show subscribers summary physical-interface port on page 585](#)
- [show subscribers summary physical-interface slot on page 586](#)
- [show subscribers summary pic on page 586](#)
- [show subscribers summary pic \(Aggregated Ethernet Interfaces\) on page 586](#)
- [show subscribers summary port on page 586](#)
- [show subscribers summary port extensive on page 586](#)
- [show subscribers summary slot on page 587](#)
- [show subscribers summary terse on page 587](#)

**Output Fields**

[Table 39 on page 583](#) lists the output fields for the **show subscribers summary** command. Output fields are listed in the approximate order in which they appear.

**Table 39: show subscribers summary Output Fields**

Field Name	Field Description	Level of Output
<b>Subscribers by State</b>	Number of subscribers summarized by state. The summary information includes the following: <ul style="list-style-type: none"> <li>• Init—Number of subscriber currently in the initialization state.</li> <li>• Configured—Number of configured subscribers.</li> <li>• Active—Number of active subscribers.</li> <li>• Terminating—Number of subscribers currently terminating.</li> <li>• Terminated—Number of terminated subscribers.</li> <li>• Total—Total number of subscribers for all states.</li> </ul>	<b>detail none</b>
<b>Subscribers by Client Type</b>	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).	<b>detail extensive none</b>
<b>Subscribers by LS:RI</b>	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).	<b>detail none</b>
<b>Subscribers by Connection Type</b>	Number of subscribers summarized by connection type, <b>Cross-connected</b> or <b>Terminated</b> .	<b>extensive</b>
<b>Interface</b>	Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface.  The * character indicates a continuation of addresses for the same session.  For aggregated Ethernet interfaces, the output of the <b>summary (pic   port   slot)</b> options prefixes the interface name with ae0:.	All levels

Table 39: show subscribers summary Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Count</b>	Count of subscribers displayed for each PIC, port, or slot when those options are specified with the <b>summary</b> 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.	<b>detail extensive none</b>
<b>Total Subscribers</b>	Total number of subscribers for all physical interfaces, all PICS, all ports, or all LS:RI slots.	<b>detail extensive none</b>
<b>IP Address/VLAN ID</b>	Subscriber IP address or VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i>	<b>terse</b>
<b>User Name</b>	Name of subscriber.	<b>terse</b>
<b>LS:RI</b>	Logical system and routing instance associated with the subscriber.	<b>terse</b>

## 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-5/0/1       201
ge-5/0/2       301

Total Subscribers: 502
```

### 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-OOB: 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-OOB: 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                  default:default
demux0.1073741824   203.0.113.10        WHOLESALER-CLIENT default:default
demux0.1073741825   203.0.113.13        RETAILER1-CLIENT  test1:retailer1
demux0.1073741826   203.0.113.213       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

---

<b>Syntax</b>	<code>show vpls connections</code> <code>&lt;brief   extensive&gt;</code> <code>&lt;down   up   up-down&gt;</code> <code>&lt;history&gt;</code> <code>&lt;instance <i>instance-name</i> local-site <i>local-site-name</i> remote-site <i>remote-site-name</i>&gt;</code> <code>&lt;instance-history&gt;</code> <code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code> <code>&lt;status&gt;</code> <code>&lt;summary&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. <b>instance-history</b> option introduced in Junos OS Release 12.3R2.
<b>Description</b>	(T Series and M Series routers, except for the M160 router) Display virtual private LAN service (VPLS) connection information.
<b>Options</b>	<b>none</b> —Display information about all VPLS connections for all routing instances.  <b>brief   extensive</b> —(Optional) Display the specified level of output.  <b>down   up   up-down</b> —(Optional) Display nonoperational, operational, or both types of connections.  <b>history</b> —(Optional) Display information about connection history.  <b>instance <i>instance-name</i></b> —(Optional) Display the VPLS connections for the specified routing instance only.  <b>instance-history</b> —(Optional) Display information about connection history for a particular instance.  <b>local-site <i>local-site-name</i></b> —(Optional) Display the VPLS connections for the specified local site name or ID only.  <b>remote-site <i>remote-site-name</i></b> —(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.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.  <b>status</b> —(Optional) Display information about the connection and interface status.  <b>summary</b> —(Optional) Display summary of all VPLS connections information.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show vpls connections on page 593</a> <a href="#">show vpls connections (with multiple pseudowires) on page 595</a>

[show vpls connections extensive \(Static VPLS Neighbors\) on page 596](#)

**Output Fields** [Table 40 on page 589](#) lists the output fields for the **show vpls connections** command. Output fields are listed in the approximate order in which they appear.

**Table 40: show vpls connections Output Fields**

Field Name	Field Description
<b>Instance</b>	Name of the VPLS instance.
<b>Local site</b>	Name of the local site.
<b>VPLS-id</b>	Identifier for the VPLS site.
<b>Number of local interfaces</b>	Number of interfaces configured for the local site.
<b>Number of local interfaces up</b>	Number of interfaces configured for the local site that are currently up.
<b>IRB interface present</b>	Indicates whether or not an integrated routing and bridging (IRB) interface is present ( <b>yes</b> or <b>no</b> ).
<b>Intf</b>	<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 <b>vt-fpc/pic/port.nnnnn</b> format. Label-switched interfaces are displayed using the <b>lsi.nnnnn</b> format. In both cases, <b>nnnnn</b> 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"> <li>• Identification as a VPLS interface</li> <li>• Name of the associated VPLS routing instance</li> <li>• Local site number</li> <li>• Remote site number</li> <li>• VPLS neighbor address</li> <li>• VPLS identifier</li> </ul>
<b>Interface flags</b>	<p>Flag associated with the interface. Can include the following:</p> <ul style="list-style-type: none"> <li>• <b>VC-Down</b>—The virtual circuit associated with this interface is down.</li> </ul>
<b>Label-base</b>	First label in a block of labels. A remote PE router uses this first label when sending traffic toward the advertising PE router.
<b>Offset</b>	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 40: show vpls connections Output Fields (*continued*)

Field Name	Field Description
<b>Size</b>	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: <b>2</b> , <b>4</b> , <b>8</b> and <b>16</b> . The default value is <b>2</b> . A value of <b>0</b> will be displayed when using the <b>remote-site</b> option.
<b>Range</b>	Label block range. A value that keeps track of the numbers of remote sites discovered within each label block.
<b>Preference</b>	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.
<b>status-vector</b>	Bit vector advertising the state of local PE-CE circuits to remote PE routers. A bit value of <b>0</b> indicates that the local circuit and LSP tunnel to the remote PE router are up, whereas a value of <b>1</b> indicates either one or both are down.
<b>connection-site</b>	Name of the connection site.
<b>Neighbor</b>	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, 203.0.113.144 (vpls-id 200).
<b>Type</b>	Type of connection: <b>loc</b> (local) or <b>rmt</b> (remote).

Table 40: show vpls connections Output Fields (*continued*)

Field Name	Field Description
<b>St</b>	<p>Status of the VPLS connection (corresponds with Legend for Connection Status):</p> <ul style="list-style-type: none"> <li>• <b>EI</b>—The local VPLS interface is configured with an encapsulation that is not supported.</li> <li>• <b>EM</b>—The encapsulation type received on this VPLS connection from the neighbor does not match the local VPLS connection interface encapsulation type.</li> <li>• <b>VC-Dn</b>—The virtual circuit is currently down.</li> <li>• <b>CM</b>—The two routers do not agree on a control word, which causes a control word mismatch.</li> <li>• <b>CN</b>—The virtual circuit is not provisioned properly.</li> <li>• <b>OR</b>—The label associated with the virtual circuit is out of range.</li> <li>• <b>OL</b>—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit.</li> <li>• <b>LD</b>—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.</li> <li>• <b>RD</b>—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.</li> <li>• <b>LN</b>—The local site has lost path selection to the remote site and therefore no pseudowires can be established from this local site.</li> <li>• <b>RN</b>—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 <b>LN</b>, and the other multihomed PE site displays the state <b>RN</b> in the following circumstances: <ul style="list-style-type: none"> <li>• The multihomed links are both configured to be the backup site.</li> <li>• 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.</li> </ul> </li> <li>• <b>XX</b>—The VPLS connection is down for an unknown reason. This is a programming error.</li> <li>• <b>MM</b>—The MTU for the local site and the remote site do not match.</li> <li>• <b>BK</b>—The router is using a backup connection.</li> <li>• <b>PF</b>—Profile parse failure.</li> <li>• <b>RS</b>—The remote site is in a standby state.</li> <li>• <b>NC</b>—The interface encapsulation is not configured as an appropriate CCC, TCC, or VPLS encapsulation.</li> <li>• <b>WE</b>—The encapsulation configured for the interface does not match the encapsulation configured for the associated connection within the VPLS routing instance.</li> </ul>

Table 40: show vpls connections Output Fields (*continued*)

Field Name	Field Description
	<ul style="list-style-type: none"> <li>• <b>NP</b>—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.</li> <li>• <b>--&gt;</b>—Only the outbound connection is up.</li> <li>• <b>&lt;--</b>—Only the inbound connection is up.</li> <li>• <b>Up</b>—The VPLS connection is operational.</li> <li>• <b>Dn</b>—The VPLS connection is down.</li> <li>• <b>CF</b>—The router cannot find enough bandwidth to the remote router to satisfy the VPLS connection bandwidth requirement.</li> <li>• <b>SC</b>—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.</li> <li>• <b>LM</b>—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.</li> <li>• <b>RM</b>—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.</li> <li>• <b>IL</b>—The incoming packets for the VPLS connection have no MPLS label.</li> <li>• <b>MI</b>—The configured mesh group identifier is in use by another system in the network.</li> <li>• <b>ST</b>—The router has switched to a standby connection.</li> <li>• <b>PB</b>—Profile busy.</li> <li>• <b>SN</b>—The VPLS neighbor is static.</li> </ul>
<b>Time last up</b>	Time connection was last in the <b>Up</b> condition.
<b># Up trans</b>	Number of transitions from <b>Down</b> to <b>Up</b> condition.
<b>Status</b>	Status of the (local or remote circuit) local interface: <ul style="list-style-type: none"> <li>• <b>Up</b>—Operational</li> <li>• <b>Dn</b>—Down</li> <li>• <b>NP</b>—Not present</li> <li>• <b>DS</b>—Disabled</li> <li>• <b>WE</b>—Wrong encapsulation</li> <li>• <b>UN</b>—Uninitialized</li> </ul>
<b>Encapsulation</b>	Type of encapsulation: <b>VPLS</b> .
<b>Remote PE</b>	Address of the remote provider edge router.

Table 40: show vpls connections Output Fields (*continued*)

Field Name	Field Description
<b>Negotiated control-word</b>	Whether a control word has been negotiated: <b>Yes</b> or <b>No</b> .
<b>Incoming label</b>	Name of the incoming label.
<b>Outgoing label</b>	Name of the outgoing label.
<b>Negotiated PW status TLV</b>	Indicates whether or not the pseudowire status TLV has been negotiated for the VPLS connection.
<b>Local interface</b>	Provides the following information about the local interface configured for the VPLS neighbor: <ul style="list-style-type: none"> <li>• Name of the local interface</li> <li>• <b>Status</b>—Interface status (<b>Up</b> or <b>Down</b>)</li> <li>• <b>Encapsulation</b>—Interface encapsulation (for example, <b>ETHERNET</b>)</li> <li>• <b>Description</b>—Includes the VPLS instance name, the VPLS neighbor address, and the VPLS identifier</li> </ul>
<b>Time</b>	Date and time of VPLS connection event.
<b>Event</b>	Type of event.
<b>Interface/Lbl/PE</b>	Interface, label, or PE router.
<b>Connection History</b>	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: 192.0.2.1, 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: 192.0.2.71, 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: 192.0.2.41, 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: 192.0.2.61, 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: 192.0.2.31, 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 -- 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

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: 192.0.2.2
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 192.0.2.2 remote-id
192.0.2.4 neighbor 192.0.2.4
Remote-id              Type  St      Time last up          # Up trans
192.0.2.4              rmt   Up      Jan 31 13:49:52 2012      1
  Remote PE: 192.0.2.4, Negotiated control-word: No
  Incoming label: 262146, Outgoing label: 262146
  Local interface: lsi.101711873, Status: Up, Encapsulation: ETHERNET
  Description: Intf - vpls green local-id 192.0.2.2 remote-id 192.0.2.4
neighbor 192.0.2.4
  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
    192.0.2.4           100    0x0

```

## show vpls flood event-queue

<b>Syntax</b>	show vpls flood event-queue
<b>Release Information</b>	Command introduced in Junos OS Release 8.0.
<b>Description</b>	Display the pending events in the VPLS flood queue.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show vpls flood event-queue on page 599</a>
<b>Output Fields</b>	<a href="#">Table 41 on page 599</a> lists the output fields for the <b>show vpls flood event-queue</b> command. Output fields are listed in the approximate order in which they appear.

**Table 41: show vpls flood event-queue Output Fields**

Field Name	Field Description
<b>Current Pending Event</b>	Provides information on the current event in the VPLS flood event queue.
<b>Name</b>	Name of the event.
<b>Owner Name</b>	Name of the interface associated with the flood event.
<b>Pending Op</b>	Pending operation for the event.
<b>Last Error</b>	Name of the last error encountered.
<b>Number of Retries</b>	Number of attempts made to update the event queue.
<b>Pending Event List</b>	List of the events awaiting processing.
<b>Event Name</b>	Name of the event.
<b>Pending Op</b>	Pending operation for the event.
<b>Event Identifier</b>	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

<b>Syntax</b>	show vpls flood instance <brief   detail   extensive> <instance-name> <logical-system <i>logical-system-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 8.0.
<b>Description</b>	Display VPLS information related to the flood process.
<b>Options</b>	<p><b>none</b>—Display VPLS information related to the flood process for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>instance-name</b>—(Optional) Display VPLS information related to the flood process for the specified routing instance.</p> <p><b>logical-system <i>logical-system-name</i></b>—(Optional) Display VPLS information related to the flood process for the specified logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show vpls flood instance on page 602</a> <a href="#">show vpls flood instance logical-system-name on page 602</a> <a href="#">show vpls flood instance detail on page 602</a>
<b>Output Fields</b>	Table 42 on page 601 lists the output fields for the <b>show vpls flood instance</b> command. Output fields are listed in the approximate order in which they appear.

**Table 42: 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, <b>flood</b> for a flood route.

Table 42: 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: __example_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: __example_ls1__> show vpls flood instance example_ls1

Logical system: __example_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: __example_ls1__> show vpls flood instance detail

Logical system: __example_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

<b>Syntax</b>	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> )
<b>Release Information</b>	Command introduced in Junos OS Release 8.0.
<b>Description</b>	Display VPLS route information related to the flood process for either the specified routing instance or the specified interface.
<b>Options</b>	<p><b>all-ce-flood</b>—Display the flood next-hop route for all customer edge routers for traffic coming from the core of the network.</p> <p><b>ce-flood interface <i>interface-name</i></b>—Display the flood next-hop route for traffic coming from the specified customer edge interface.</p> <p><b>instance-name <i>instance-name</i></b>—Display the flood routes for the specified instance.</p> <p><b>logical-system-name <i>logical-system-name</i></b>—(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 <code>_example_ls1_</code> (the name must be entered in the command with the underscore characters).</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show vpls flood route all-ce-flood on page 604</a> <a href="#">show vpls flood route ce-flood on page 604</a>
<b>Output Fields</b>	Table 43 on page 603 lists the output for the <b>show vpls flood route</b> command. Output fields are listed in the approximate order in which they appear.

**Table 43: 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 <b>CE_FLOOD</b> or <b>ALL_CE_FLOOD</b> ).
Flood route owner	VPLS routing instance or interface associated with the flood route.
Nexthop type	Next-hop type. For example, <b>flood</b> 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 43: 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: __example_lsl__> show vpls flood route all-ce-flood logical-system-name
__example_lsl__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: __example_lsl__> 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

<b>Syntax</b>	<pre>show vpls mac-table   &lt;age&gt;   &lt;brief   detail   extensive   summary&gt;   &lt;bridge-domain <i>bridge-domain-name</i>&gt;   &lt;instance <i>instance-name</i>&gt;   &lt;interface <i>interface-name</i>&gt;   &lt;logical-system (all   <i>logical-system-name</i>)&gt;   &lt;mac-address&gt;   &lt;vlan-id <i>vlan-id-number</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 15.1.</p>
<b>Description</b>	Display learned virtual private LAN service (VPLS) media access control (MAC) address information.
<b>Options</b>	<p><b>none</b>—Display all learned VPLS MAC address information.</p> <p><b>age</b>— (Optional) Display age of a single mac-address.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>bridge-domain <i>bridge-domain-name</i></b>—(Optional) Display learned VPLS MAC addresses for the specified bridge domain.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display learned VPLS MAC addresses for the specified instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display learned VPLS MAC addresses for the specified instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Display learned VPLS MAC addresses for all logical systems or for the specified logical system.</p> <p><b>mac-address</b>—(Optional) Display the specified learned VPLS MAC address information..</p> <p><b>vlan-id <i>vlan-id-number</i></b>—(Optional) Display learned VPLS MAC addresses for the specified VLAN.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show vpls mac-table on page 606</a></p> <p><a href="#">show vpls mac-table (with Layer 2 Services over GRE Interfaces) on page 607</a></p> <p><a href="#">show vpls mac-table (with VXLAN enabled) on page 607</a></p> <p><a href="#">show vpls mac-table age (for GE interface) on page 607</a></p> <p><a href="#">show vpls mac-table age (for AE interface) on page 607</a></p> <p><a href="#">show vpls mac-table count on page 608</a></p> <p><a href="#">show vpls mac-table detail on page 608</a></p> <p><a href="#">show vpls mac-table extensive on page 609</a></p>

**Output Fields** Table 44 on page 606 describes the output fields for the **show vpls mac-table** command. Output fields are listed in the approximate order in which they appear.

**Table 44: show vpls mac-table Output fields**

Field Name	Field Description
Age	Age of a single mac-address.
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"> <li>• <b>S</b>—Static MAC address configured.</li> <li>• <b>D</b>—Dynamic MAC address learned.</li> <li>• <b>SE</b>—MAC accounting is enabled.</li> <li>• <b>NM</b>—Nonconfigured MAC.</li> </ul>
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
  00:00:5e:00:53:5d  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          MAC          Logical
  address      flags      interface
  00:00:5e:00:53:12  D      lsi.1051138
  00:00:5e:00:53: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          Logical
  address      flags      interface
  00:00:5e:00:53:f4  D,SE      ge-4/2/0.1000
  00:00:5e:00:53:33  D,SE      lsi.1052004
  00:00:5e:00:53:32  D,SE      lsi.1048840
  00:00:5e:00:53:14  D,SE      lsi.1052005
  00:00:5e:00:53: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: 233.252.0.1
  MAC          MAC          Logical
  address      flags      interface
  00:00:5e:00:53:f4  D,SE      ge-4/2/0.1000
  00:00:5e:00:53:33  D,SE      lsi.1052004
  00:00:5e:00:53:32  D,SE      lsi.1048840
  00:00:5e:00:53:14  D,SE      lsi.1052005
  00:00:5e:00:53:f7  D,SE      vtep.1052010
  00:00:5e:00:53:3f  D,SE      vtep.1052011

```

#### show vpls mac-table age (for GE interface)

```

user@host> show vpls mac-table age 00:00:5e:00:53:1a instance vpls_instance_1
MAC Entry Age information
Current Age: 4 seconds

```

#### show vpls mac-table age (for AE interface)

```

user@host> show vpls mac-table age 000:00:5e:00:53:1a instance vpls_instance_1
MAC Entry Age information
Current Age on FPC1: 102 seconds
Current Age on FPC2: 94 seconds

```

### show vpls mac-table count

```
user@host> show vpls mac-table count
0 MAC address learned in routing instance __example_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
lc-2/3/0.32769	0
lc-9/2/0.32769	0
lc-11/2/0.32769	0
lc-11/3/0.32769	0
lc-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: 00:00:5e:00:53:5d
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: 00:00:5e:00:53:5d
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: 00:00:5e:00:53::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: 00:00:5e:00:53:01
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: 00:00:5e:00:53:02
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: 00:00:5e:00:53:03
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

<b>Syntax</b>	show vpls statistics <instance <i>instance-name</i> > <logical-system (all   <i>logical-system-name</i> )>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	(T Series and M Series routers, except for the M160 router) Display virtual private LAN service (VPLS) statistics.
<b>Options</b>	<p><b>none</b>—Display VPLS statistics for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display VPLS statistics for a specific VPLS routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show vpls statistics on page 611</a> <a href="#">show vpls statistics instance on page 611</a>
<b>Output Fields</b>	<a href="#">Table 45 on page 610</a> lists the output fields for the <b>show vpls statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 45: 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, <b>vt-fpc/pic/port.nnnnn</b> , where <b>nnnnn</b> 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 45: 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 6

# Index

- [Index on page 615](#)



# Index

## Symbols

#, comments in configuration statements.....	xx
( ), in syntax descriptions.....	xx
802.1Q VLANs	
dynamic.....	200
mixed VLAN tagging.....	277
VLAN IDs.....	375
< >, in syntax descriptions.....	xx
[ ], in configuration statements.....	xx
{ }, in configuration statements.....	xx
(pipe), in syntax descriptions.....	xx

## A

AAA	
.....	412
subscriber statistics	
clearing.....	408
displaying.....	388, 397
subscribers	
displaying.....	400
logging out.....	408
AAA access messages	
supported attributes.....	141
AAA accounting messages	
supported attributes.....	147
accept statement.....	182
accept-out-of-band statement	
ANCP.....	183
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.....	184
accounting statements	
compress.....	204
accounting statistics	
extensible subscriber services manager	
(ESSM).....	170
Layer 2 wholesale.....	166
accounting-server statement.....	186
active-server-group statement	
DHCP relay agent.....	187
address assignment pool	
128-bit address	
DHCPv6 example.....	44
74-bit prefix	
DHCPv6 example.....	45
configuring.....	40
address assignment pools	
displaying.....	403
address server	
configuring.....	40
address statement.....	188
address-assignment pools	
DHCPv6 attributes	
specifying the preferred-lifetime.....	212
specifying the t1-percentage.....	212
specifying the t2-percentage.....	212
specifying the valid-lifetime.....	212
address-assignment statement	
address-assignment pools.....	191
adjacency timer	
ANCP neighbor configuration.....	153
adjacency-loss-hold-time statement	
ANCP.....	192
aggregated Ethernet interfaces	
status information, displaying.....	470
ANCP	
access loop	
clearing.....	405
access loops	
clearing.....	160
backwards compatibility neighbor	
configuration.....	153
discovery table neighbor configuration.....	153
displaying pending access lines.....	421
neighbor	
clearing.....	406

neighbor configuration	
adjacency timer.....	153
discovery table entries.....	153
ietf mode.....	153
IP address.....	153
pre-ietf mode.....	153
neighbors	
displaying.....	413
pre-ietf mode neighbor configuration.....	153
reestablishing pending access line	
sessions.....	160
triggering OAM	
Port Down and Port Up message	
simulation.....	157
ancp statement	
ANCP.....	193
ANCP statements	
accept-out-of-band.....	183
adjacency-loss-hold-time.....	192
ancp.....	193
auto-configure-trigger interface.....	201
neighbor	
for all neighbors.....	311
authentication statement	
DHCP local server.....	196
DHCP relay agent.....	197
dynamic VLAN authentication.....	195
vlan ranges.....	195
authentication-order statement	
access.....	198
authentication-server statement.....	199
auto-configure statement.....	200
auto-configure-trigger interface statement	
ANCP.....	201

## B

backup-on-failure statement	
flat-file accounting.....	202
backwards compatibility	
ANCP neighbor configuration.....	153
binding state of DHCP client	
clearing.....	431, 445
displaying.....	422, 436
binding state of DHCPv6 client	
clearing.....	459
displaying.....	450
braces, in configuration statements.....	xx

brackets	
angle, in syntax descriptions.....	xx
square, in configuration statements.....	xx
broadband subscriber management	
Layer 2 wholesale topology.....	71
Layer 3 wholesale topology	
DHCPv4.....	9
DHCPv6.....	29
PPPoE.....	53

## C

circuit-id statement	
dynamic VLAN authentication.....	203
clear ancp access-loop command.....	405
clear ancp neighbor command.....	406
clear dhcp relay binding command.....	445
clear dhcp relay statistics command.....	447
clear dhcp server binding command.....	431
clear dhcp server statistics command.....	434
clear dhcpv6 server binding command.....	459
clear dhcpv6 server statistics command.....	461
clear network-access aaa subscriber	
command.....	408
comments, in configuration statements.....	xx
compress statement	
accounting.....	204
flat-file accounting.....	204
connectivity-type statement.....	205
conventions	
text and syntax.....	xix
core-facing statement.....	206
curly braces, in configuration statements.....	xx
customer support.....	xxi
contacting JTAC.....	xxi

## D

demux interfaces	
unit statement.....	364
demux-options statement	
dynamic IP demux interface.....	208
demux-source statement.....	211
dynamic IP demux interfaces.....	209
dynamic underlying interface.....	210
demux0 statement	
dynamic IP demux interface.....	207
demux0 statements	
underlying-interface.....	355

- 
- DHCP
    - extended server binding
      - clearing.....431
      - displaying.....422
    - extended server statistics
      - clearing.....434
      - displaying.....428
    - relay binding
      - clearing.....445
    - relay binding state
      - displaying.....436
    - relay statistics
      - clearing.....447
      - displaying.....441
  - DHCP client
    - binding state
      - clearing.....431, 445
      - displaying.....422, 436
    - statistics
      - clearing.....434, 447
      - displaying.....428
  - DHCP local server
    - lease timers.....313
  - DHCP local server statements
    - dhcp-local-server.....213
    - dhcpv6.....226
    - dynamic-profile.....230
    - group.....282
    - interface.....291
    - ip-address-first.....306
    - pool-match-order.....318
    - traceoptions.....353
  - DHCP relay agent statements
    - active-server-group.....187
    - dhcp-relay.....219
    - dynamic-profile.....231
    - group.....285
    - interface.....293
    - server-group.....345
    - traceoptions.....353
    - user-prefix.....373
    - username-include.....197
  - dhcp-attributes statement
    - address-assignment pools.....212
  - dhcp-local-server statement.....213
  - dhcp-relay statement.....219
  - DHCPv6
    - extended server binding
      - clearing.....459
      - displaying.....450
    - extended server statistics
      - clearing.....461
      - displaying.....456
  - DHCPv6 client
    - binding state
      - clearing.....459
      - displaying.....450
    - statistics
      - clearing.....461
      - displaying.....456
  - DHCPv6 local server
    - lease timers.....313
  - dhcpv6 statement.....226
  - discovery table
    - ANCP neighbor configuration.....153
  - documentation
    - comments on.....xxi
  - domain-name statement
    - address-assignment pools.....229
  - dynamic IP demux interface statements
    - demux-source.....209
    - family.....268
  - dynamic IP demux statements
    - mac-validate.....309
  - dynamic PPP statements
    - pap.....316
    - ppp-options.....321
  - dynamic PPPoE statements
    - dynamic-profile.....232
    - family.....269
    - pppoe-options.....319
    - pppoe-underlying-options.....320
    - server.....344
    - underlying-interface.....356
    - unnumbered-address.....368
  - dynamic profiles
    - DHCPv4 retailer example.....24
    - DHCPv4 wholesaler example.....24
    - DHCPv6 retailer example.....43
    - Layer 2 wholesale retailer example.....87
    - PPPoE wholesaler example.....62
    - predefined variables.....109

retailer		encapsulation statement	
DHCPv4.....	19	dynamic profiles.....	244
DHCPv6.....	38	logical interfaces.....	247
Layer 2 wholesale.....	75	physical interface.....	251
wholesale network		Ethernet interfaces	
DHCPv4.....	17	dynamic VLANs.....	200
DHCPv6.....	37	mixed VLAN tagging.....	277
PPPoE.....	59	status information, displaying	
wholesaler		aggregated.....	470
DHCP.....	37	Fast Ethernet.....	481
DHCPv4.....	18	Gigabit Ethernet.....	497
PPPoE.....	60	VLAN IDs.....	375
dynamic profiles statements		exclude statement.....	257
dynamic-profiles.....	235	extended DHCP	
inner-vlan-id-swap-ranges.....	290	configuring	
interface.....	294	local server.....	42
interfaces.....	298	extensible subscriber services manager (ESSM)	
keepalives.....	307	flat-file accounting.....	170
routing-instances.....	340		
unnumbered-address.....	369	<b>F</b>	
vlan-tags.....	378	family statement.....	262
dynamic subscribers		address-assignment pools.....	267
interfaces statement.....	298	dynamic IP demux interface.....	268
pppoe-underlying-options statement.....	320	dynamic PPPoE.....	269
dynamic underlying interface statements		dynamic profiles.....	270
demux-source.....	210	Fast Ethernet interfaces	
dynamic VLAN authentication statements		dynamic VLANs.....	200
authentication.....	195	status information, displaying.....	481
auto-configure.....	200	VLAN IDs.....	375
circuit-id.....	203	fields statement	
remote-id.....	337	flat-file accounting.....	272
username-include.....	372	file statement	
dynamic VLANs.....	200	flat-file accounting.....	273
authenticating.....	195	firewall filters	
dynamic-profile statement		displaying.....	540
DHCP local server.....	230	flat-file accounting	
DHCP relay agent.....	231	attributes and options.....	163
dynamic PPPoE.....	232	ESSM.....	170
stacked VLAN ranges.....	233	extensible subscriber services manager	
VLAN ranges.....	234	(ESSM).....	170
dynamic-profiles		Layer 2 wholesale.....	166
interfaces statement.....	298	flat-file accounting statements	
dynamic IP demux.....	298	backup-on-failure.....	202
dynamic-profiles statement.....	235	compress.....	204
		egress-stats.....	243
		fields.....	272
		file.....	273
		format.....	278
		general-param.....	280
<b>E</b>			
egress-stats statement			
flat-file accounting.....	243		

- ingress-stats.....288
  - interval.....302
  - overall-packet.....314
  - push-backup-to-master.....332
  - schema-version.....343
  - flexible-vlan-tagging statement.....277
  - font conventions.....xix
  - format statement
    - flat-file accounting.....278
  - forwarding options
    - configuring.....22
    - DHCP relay agent.....436, 441, 445, 447
    - DHCPv4 example.....25
  - forwarding-options statement.....279
- G**
- general-param statement
    - flat-file accounting.....280
  - Gigabit Ethernet interfaces
    - demultiplexing interface information,
      - displaying.....531
    - dynamic VLANs.....200
    - status information, displaying.....497
    - VLAN IDs.....375
  - grace-period statement.....281
  - group statement
    - DHCP local server.....282
    - DHCP relay agent.....285
- I**
- ingress-stats statement
    - flat-file accounting.....288
  - inner-vlan-id statement
    - dynamic VLAN interfaces.....289
  - inner-vlan-id-swap-ranges statement
    - dynamic profiles.....290
  - input-vlan-map statement
    - dynamic interfaces.....290
  - instance-role statement.....303
  - instance-type statement.....304
  - interface
    - direct ISP-facing.....89
    - Layer 2 wholesale example.....87
    - NNI ISP-facing.....89
  - interface statement
    - DHCP local server.....291
    - DHCP relay agent.....293
    - dynamic profiles.....294
    - IS-IS.....295
    - OSPF.....295
    - VPNs.....295
  - interface-mac-limit statement.....296
  - interfaces
    - loopback
      - DHCPv4 Layer 3 wholesale.....11
      - DHCPv6 Layer 3 wholesale.....32
      - PPPoE Layer 3 wholesale.....55
    - mixed VLAN tagging.....277
    - unit statement.....366
  - interfaces statement.....297
    - dynamic profiles.....298
  - interval statement
    - flat-file accounting.....302
  - IP demultiplexing interface statements
    - unit.....364
  - ip-address-first statement.....306
- J**
- Juniper Networks VSAs
    - supported.....129
- K**
- keepalives statement
    - dynamic profiles.....307
- L**
- I2-stats statement.....308
  - Layer 2 VPNs
    - VPLS connections, displaying.....588
    - VPLS statistics, displaying.....610
  - Layer 2 wholesale
    - access.....86
    - access loops
      - clearing.....160
    - access profile.....86
    - basic topology.....71
    - configuration elements.....71
    - configuration options
      - DHCPv4.....4, 68
      - DHCPv6.....4, 68
    - direct interface.....89
    - displaying pending access lines.....421
    - dynamic profiles.....75
      - retailer example.....87
    - flat-file accounting.....166
    - interface.....87
    - NNI interface.....89
    - overview.....3, 49, 67

reestablishing pending access line sessions.....	160
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
vlans	
DHCPv4.....	12
DHCPv6.....	32
local server	
configuring DHCP.....	42
DHCPv6 example.....	45
logical interface statements	
family.....	270
logical interfaces	
VLAN IDs.....	375
logical systems	
subscriber relationship with.....	6, 50
loopback interface	
Layer 3 wholesale	
DHCPv4.....	11
DHCPv6.....	32
PPPoE.....	55
status information, displaying.....	463
<b>M</b>	
mac-validate statement	
dynamic IP demux interface.....	309
manuals	
comments on.....	xxi
maximum-lease-time statement.....	313
mixed VLAN tagging.....	277
multicast-replication statement.....	310
<b>N</b>	
neighbor	
ANCP.....	406
parameters for ANCP.....	153

- 
- neighbor statement
    - ANCP
      - for all neighbors.....311
  - no-tunnel-services statement.....312
  - O**
  - OAM
    - triggering ANCP
      - Port Down and Port Up message
        - simulation.....157
  - OSPF
    - enabling.....295
  - output-vlan-map statement
    - dynamic interfaces.....315
  - overall-packet statement
    - flat-file accounting.....314
  - P**
  - pap statement
    - dynamic PPP.....316
  - parentheses, in syntax descriptions.....xx
  - pending access line sessions
    - reestablishing for Layer 2 wholesale.....160
  - physical interfaces
    - dynamic VLANs.....200
    - mixed VLAN tagging.....277
  - pool statement
    - address-assignment pools.....317
  - pool-match-order statement.....318
  - pop statement
    - dynamic VLAN interfaces.....318
  - PPP
    - interfaces, displaying.....552
  - ppp-options statement
    - dynamic PPP.....321
  - PPPoE
    - interfaces, displaying.....521
  - pppoe-options statement
    - dynamic PPPoE.....319
  - pppoe-underlying-options statement
    - static and dynamic PPPoE.....320
  - pre-ietf mode
    - ANCP neighbor configuration.....153
  - predefined variables.....109
  - prefix statement
    - address-assignment pools.....322
  - profile statement
    - subscriber access.....323
  - protocols statement.....327
  - proxy-arp statement.....330
    - subscriber interfaces.....331
  - pseudowires
    - mixed VLAN tagging.....277
  - push statement
    - dynamic VLAN interfaces.....331
  - push-backup-to-master statement
    - flat-file accounting.....332
  - R**
  - RADIUS
    - wholesale VSA support.....6, 51, 69
  - RADIUS server
    - access configuration
      - configuring.....15, 35, 58, 86
  - radius statement
    - subscriber access.....333
  - radius-server statement.....335
  - range statement
    - address-assignment pools.....336
  - ranges statement
    - stacked VLAN.....186
    - VLAN.....185, 337
  - remote-id statement
    - dynamic VLAN authentication.....337
  - request ancp oam port-down command.....409
  - request ancp oam port-up command.....411
  - request auto-configuration reconnect-pending
    - command.....412
  - route-distinguisher statement.....338
  - routing information
    - interfaces
      - state, displaying.....544
  - 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.....342
    - dynamic profiles.....340

**S**

schema-version statement	
flat-file accounting.....	343
secret statement	
access.....	344
server statement	
dynamic PPPoE.....	344
server-group statement	
DHCP relay agent.....	345
services statement	
remote router access.....	346
show ancp neighbor command.....	413
show auto-configuration out-of-band pending	
command.....	421
show dhcp relay binding command.....	436
show dhcp relay statistics command.....	441
show dhcp server binding command.....	422
show dhcp server statistics command.....	428
show dhcpv6 server binding command.....	450
show dhcpv6 server statistics command.....	456
show interfaces (Aggregated Ethernet)	
command.....	470
show interfaces (Fast Ethernet) command.....	481
show interfaces (Gigabit Ethernet) command.....	497
show interfaces (Loopback) command.....	463
show interfaces (PPPoE) command.....	521
show interfaces demux0 (Demux Interfaces)	
command.....	531
show interfaces filters command.....	540
show interfaces routing command.....	544
show interfaces routing-instance command.....	550
show network-access aaa statistics authentication	
command.....	397
show network-access aaa statistics	
command.....	388
show network-access aaa subscribers	
command.....	400
show network-access address-assignment pool	
command.....	403
show ppp interface command.....	552
show subscribers command.....	562
show subscribers summary command.....	582
show vpls connections command.....	588
show vpls flood event-queue command.....	599
show vpls flood instance command.....	601
show vpls flood route command.....	603
show vpls mac-table command.....	605
show vpls statistics command.....	610
site statement.....	349

site-identifier statement.....	349
site-range statement.....	350
stacked VLAN-tag framing	
VLAN ID list.....	379
stacked-vlan-ranges statement.....	351
stacked-vlan-tagging statement.....	352
static PPPoE statements	
pppoe-underlying-options.....	320
static subscribers	
interfaces statement.....	298
pppoe-underlying-options statement.....	320
subscriber access	
subscriber information, displaying.....	562
subscriber summary information,	
displaying.....	582
subscriber interface statements	
demux-options.....	208
demux-source.....	209, 210
demux0.....	207
dynamic-profile.....	232
family.....	268, 269, 270
interfaces.....	298
mac-validate.....	309
pap.....	316
ppp-options.....	321
pppoe-options.....	319
pppoe-underlying-options.....	320
proxy-arp.....	331
server.....	344
underlying-interface.....	355, 356
unit.....	364, 366
unnumbered-address.....	368, 369
subscribers	
displaying.....	562
displaying summary.....	582
support, technical See technical support	
syntax conventions.....	xix
system statement.....	352

**T**

technical support	
contacting JTAC.....	xxi
topology	
Layer 2 wholesale network.....	71
Layer 3 wholesale network	
DHCPv4.....	9
DHCPv6.....	29
PPPoE.....	53

traceoptions statement		
DHCP local server.....	353	
DHCP relay agent.....	353	
<b>U</b>		
underlying-interface statement		
dynamic PPPoE.....	356	
dynamic profiles.....	355	
unit statement.....	357	
demux interfaces.....	364	
interfaces.....	366	
unnumbered-address statement		
dynamic PPPoE.....	368	
dynamic profiles.....	369	
Ethernet.....	371	
user-prefix statement		
DHCP local server.....	373	
username-include statement		
dynamic VLAN authentication.....	372	
<b>V</b>		
variables		
predefined.....	109	
vendor-specific attributes		
supported.....	129	
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.....	375	
dynamic VLAN map		
rewriting at ingress or egress.....	374	
Ethernet interfaces.....	375	
vlan-model statement.....	376	
vlan-ranges statement.....	377	
vlan-tags statement		
dynamic profiles.....	378	
stacked VLAN tags.....	379	
VLANs		
configuring VLAN ranges.....	351, 377	
VPLS		
connections, displaying.....	588	
statistics, displaying.....	610	
vpls statement		
routing-instances.....	381	
vrf-export statement.....	383	
vrf-import statement.....	384	
vrf-target statement.....	385	
VSAs		
supported.....	129	
<b>W</b>		
wholesale See Layer 2 wholesale See Layer 3		
wholesale		

