



---

Junos<sup>®</sup> OS

## Hierarchical Class of Service Feature Guide



---

Modified: 2018-03-14

Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://www.juniper.net)

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be property of their respective owners.

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

*Junos® OS Hierarchical Class of Service Feature Guide*  
Copyright © 2018 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 <https://www.juniper.net/support/eula/>. 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>Hierarchical Class of Service</b>	
<b>Chapter 1</b>	<b>Configuring Hierarchical Class of Service on MX Series 3D Universal Edge Routers . . . . .</b>	<b>3</b>
	Hierarchical Class of Service Overview . . . . .	3
	Hierarchical Class of Service Network Scenarios . . . . .	6
	Services to Subscribers . . . . .	6
	Services to Businesses . . . . .	7
	Wireless Backhaul . . . . .	7
	Understanding Hierarchical Scheduling . . . . .	7
	Hierarchical Scheduling Terminology . . . . .	8
	Scheduler Node-Level Designations in Hierarchical Scheduling . . . . .	8
	Hierarchical Scheduling at Non-Leaf Nodes . . . . .	9
	Priority Propagation . . . . .	10
	Configuring Hierarchical Schedulers for CoS . . . . .	12
	Hierarchical Schedulers and Traffic Control Profiles . . . . .	13
	Example: Building a Four-Level Hierarchy of Schedulers . . . . .	15
	Configuring the Interface Sets . . . . .	16
	Configuring the Interfaces . . . . .	16
	Configuring the Traffic Control Profiles . . . . .	16
	Configuring the Schedulers . . . . .	17
	Configuring the Drop Profiles . . . . .	18
	Configuring the Scheduler Maps . . . . .	18
	Applying the Traffic Control Profiles . . . . .	19
	Scheduling and Shaping in Hierarchical CoS Queues for Traffic Routed to GRE Tunnels . . . . .	20
	Understanding Scheduling and Shaping of Traffic Routed to GRE Tunnels . . . . .	20
	Configuration Overview . . . . .	20

	Configuration Caveats . . . . .	20
	Example: Performing Output Scheduling and Shaping in Hierarchical CoS Queues for Traffic Routed to GRE Tunnels . . . . .	21
	Configuring Ingress Hierarchical CoS on Enhanced Queuing DPCs . . . . .	35
<b>Chapter 2</b>	<b>Configuring Hierarchical Class of Service on MICs, MPCs, MLCs, and Aggregated Ethernet Interfaces . . . . .</b>	<b>39</b>
	Understanding Hierarchical Scheduling for MIC and MPC Interfaces . . . . .	39
	Scheduler Node Scaling for MIC and MPC Interfaces . . . . .	39
	Hierarchical Scheduling Priority Levels for MIC and MPC Interfaces . . . . .	40
	Guaranteed Bandwidth and Weight of an Interface Node on MIC and MPC Interfaces . . . . .	40
	Hierarchical Scheduling for MIC and MPC Interfaces in Oversubscribed PIR Mode . . . . .	41
	Configuring Ingress Hierarchical CoS on MIC and MPC Interfaces . . . . .	41
	Per-Unit Queuing and Hierarchical Queuing for MIC and MPC Interfaces . . . . .	43
	Queuing Models Supported for MIC and MPC Interfaces . . . . .	43
	Limited Scale Per-Unit Queuing MPCs . . . . .	43
	Hierarchical Queuing MICs and MPCs . . . . .	44
	Scheduler Node Levels for MIC and MPC Interfaces . . . . .	44
	Scheduler Node Levels for Per-Unit Queuing MPCs . . . . .	44
	Scheduler Node Levels for Hierarchical Queuing MICs and MPCs . . . . .	45
	Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview . . . . .	46
	Queue Scaling for MPCs . . . . .	46
	Managing Remaining Queues . . . . .	47
	Jitter Reduction in Hierarchical CoS Queues . . . . .	48
	Queue Jitter as a Function of the Maximum Number of Queues . . . . .	48
	Default Maximum Queues for Hierarchical Queuing MICs and MPCs . . . . .	48
	Shaping Rate Granularity as a Function of the Rate Wheel Update Period . . . . .	49
	Example: Reducing Jitter in Hierarchical CoS Queues . . . . .	50
	Hierarchical Schedulers on Aggregated Ethernet Interfaces Overview . . . . .	56
	Configuring Hierarchical Schedulers on Aggregated Ethernet Interfaces . . . . .	57
	Example: Configuring Scheduling Modes on Aggregated Interfaces . . . . .	58
	Increasing Available Bandwidth on Rich-Queuing MPCs by Bypassing the Queuing Chip . . . . .	63
<b>Part 2</b>	<b>Hierarchical CoS for Subscriber Management</b>	
<b>Chapter 3</b>	<b>Hierarchical Class of Service for Subscriber Management . . . . .</b>	<b>69</b>
	Hierarchical Class of Service for Subscriber Management Overview . . . . .	69
	Understanding Hierarchical CoS for Subscriber Interfaces . . . . .	70
	Two-Level Hierarchical Scheduling . . . . .	71
	Three-Level Hierarchical Scheduling . . . . .	72
	Interface Hierarchy Versus CoS Hierarchy . . . . .	74

	Four-Level Hierarchical Scheduling . . . . .	75
	Hardware Requirements for Dynamic Hierarchical CoS . . . . .	77
	Configuring Static Hierarchical Scheduling in a Dynamic Profile . . . . .	78
	Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links . . . . .	79
	Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface . . . . .	80
	Example: Maintaining a Constant Traffic Flow by Configuring a Static VLAN Interface with a Dynamic Profile for Subscriber Access . . . . .	81
<b>Chapter 4</b>	<b>Applying CoS to Groups of Subscriber Interfaces . . . . .</b>	<b>95</b>
	CoS for Interface Sets of Subscribers Overview . . . . .	95
	Guidelines for Configuring Dynamic Interface Sets in a Subscriber Access Network . . . . .	96
	Configuring an Interface Set of Subscribers in a Dynamic Profile . . . . .	98
	Example: Configuring a Dynamic Interface Set of VLAN Subscribers . . . . .	99
	Example: Configuring a Dynamic Service VLAN Interface Set of Subscribers in a Dynamic Profile . . . . .	112
<b>Chapter 5</b>	<b>Configuring Hierarchical Scheduling for MPLS Pseudowire Interfaces . . .</b>	<b>117</b>
	Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview . . . . .	117
	CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces . . . . .	118
	CoS Two-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces . . . . .	120
	Configuring CoS Two-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces . . . . .	122
	CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces . . . . .	123
	Three-Level Scheduling Hierarchy: Pseudowire Logical Interfaces over a Transport Logical Interface . . . . .	124
	Three-Level Scheduling Hierarchy : Pseudowire Service Logical Interfaces over a Pseudowire Service Interface Set . . . . .	125
	Three-Level Scheduling Hierarchy Combined Deployment Scenario . . . . .	125
	Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Transport Logical Interface) . . . . .	127
	Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Pseudowire Interface Set) . . . . .	128
<b>Chapter 6</b>	<b>Configuring Hierarchical Scheduling for L2TP . . . . .</b>	<b>131</b>
	CoS for L2TP LAC Subscriber Interfaces Overview . . . . .	132
	Traffic from LAC to LNS . . . . .	132
	LAC Tunnels: Traffic from LNS to LAC . . . . .	133
	Configuring Dynamic CoS for an L2TP LAC Tunnel . . . . .	134
	CoS for L2TP LNS Inline Services Overview . . . . .	135
	Guidelines for Applying CoS to the LNS . . . . .	135
	Hardware Requirements for Inline Services on the LNS . . . . .	136
	Configuring an Inline Service Interface for L2TP LNS . . . . .	136
	Configuring Dynamic CoS for an L2TP LNS Inline Service . . . . .	137

<b>Chapter 7</b>	<b>Preventing Bandwidth Contention on Subscriber Interfaces . . . . .</b>	<b>141</b>
	Hierarchical CoS Shaping-Rate Adjustments Overview . . . . .	141
	Types of Shaping-Rate Adjustments . . . . .	142
	Levels of Shaping-Rate Adjustments . . . . .	142
	Shaping Rate Adjustments for Subscriber Local Loops Overview . . . . .	143
	Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops . . . . .	144
	Configuring the Minimum Adjusted Shaping Rate on Scheduler Nodes for Subscribers . . . . .	145
	Overview . . . . .	145
	Configuring a Static Minimum Adjusted Shaping Rate on Scheduler Nodes . . . . .	146
	Configuring a Dynamic Minimum Adjusted Shaping Rate on Scheduler Nodes . . . . .	146
	Configuring Shaping-Rate Adjustments on Queues . . . . .	146
	Overview . . . . .	147
	Configuring a Static Shaping-Rate Adjustment for Queues . . . . .	147
	Configuring a Dynamic Shaping-Rate Adjustment for Queues . . . . .	148
	Enabling Shaping-Rate Adjustments for Subscriber Local Loops . . . . .	148
	Configuring Static Logical Interface Sets to Serve as CoS Hierarchical Scheduler Nodes for Subscriber Loops . . . . .	149
	Configuring the Logical Interfaces That Compose the Static Logical Interface Sets . . . . .	150
	Configuring Hierarchical CoS on the Static Logical Interface Sets That Serve as Hierarchical Scheduler Nodes for Subscriber Local Loops . . . . .	150
	Configuring ANCP Functionality That Supports and Drives Shaping-Rate Adjustments for Subscriber Local Loops . . . . .	152
	Disabling Shaping-Rate Adjustments for Subscriber Local Loops . . . . .	154
	Disabling Hierarchical Bandwidth Adjustment for Subscriber Interfaces with Reverse-OIF Mapping . . . . .	154
	Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops . . . . .	155
	Verifying the Configuration of Shaping-Rate Adjustments for Subscriber Local Loops . . . . .	158
	Verifying the Configuration of ANCP for Shaping-Rate Adjustments . . . . .	158
	Using Hierarchical CoS to Adjust Shaping Rates Based on Multicast Traffic . . . . .	159
<b>Chapter 8</b>	<b>Configuring Targeted Distribution of Subscribers on Aggregated Ethernet     Interfaces . . . . .</b>	<b>163</b>
	Distribution of Demux Subscribers in an Aggregated Ethernet Interface . . . . .	163
	Distribution Models . . . . .	163
	Sample Targeted Distribution Topology . . . . .	164
	Redundancy and Redistribution Mechanisms . . . . .	164
	Considerations and Best Practices . . . . .	165
	Providing Accurate Scheduling for a Demux Subscriber Interface of Aggregated Ethernet Links . . . . .	166
	Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces . . . . .	167

	Configuring Link and Module Redundancy for Demux Subscribers in an Aggregated Ethernet Interface . . . . .	168
	Configuring Rebalancing of Demux Subscribers in an Aggregated Ethernet Interface . . . . .	168
	Configuring Periodic Rebalancing of Subscribers in an Aggregated Ethernet Interface . . . . .	169
	Configuring Manual Rebalancing of Subscribers on an Aggregated Ethernet Interface . . . . .	169
	Example: Separating Targeted Multicast Traffic for Demux Subscribers on Aggregated Ethernet Interfaces . . . . .	170
	Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface . . . . .	181
	Configuring the Distribution Type for PPPoE Subscribers on Aggregated Ethernet Interfaces . . . . .	181
	Verifying the Distribution of PPPoE Subscribers in an Aggregated Ethernet Interface . . . . .	182
<b>Chapter 9</b>	<b>Applying CoS Using Parameters Received from RADIUS . . . . .</b>	<b>183</b>
	Subscriber Interfaces That Provide Initial CoS Parameters Dynamically Obtained from RADIUS . . . . .	183
	Dynamic Configuration of Initial CoS in Access Profiles . . . . .	184
	Predefined Variables for Dynamic Configuration of Initial Traffic Shaping . . . . .	184
	Predefined Variables for Dynamic Configuration of Initial Scheduling and Queuing . . . . .	185
	Changing CoS Services Overview . . . . .	187
	Types of CoS Variables Used in a Service Profile . . . . .	187
	Static and Dynamic CoS Configurations . . . . .	188
	Scenarios for Static and Dynamic Configuration of CoS Parameters . . . . .	188
	CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview . . . . .	190
	Supported Network Configurations . . . . .	190
	Traffic-Control Profiles in Subscriber Interface Dynamic Profiles . . . . .	190
	CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets and Member Subscriber Sessions . . . . .	191
	Guidelines for Configuring CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions . . . . .	192
	Configuring Initial CoS Parameters Dynamically Obtained from RADIUS . . . . .	193
	Configuring Static Default Values for Traffic Scheduling and Shaping . . . . .	194
	Applying CoS Traffic-Shaping Attributes to Dynamic Interface Sets and Member Subscriber Sessions . . . . .	195
	CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets . . . . .	198
	Example: Configuring Dynamic Hierarchical Scheduling for Subscribers . . . . .	203
<b>Part 3</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 10</b>	<b>Configuration Statements . . . . .</b>	<b>213</b>
	accept . . . . .	218
	address-assignment (Address-Assignment Pools) . . . . .	220
	authentication (DHCP Local Server) . . . . .	222

authentication (DHCP Relay Agent) .....	223
authentication-order .....	224
authentication-server .....	225
auto-configure .....	226
adjust-minimum .....	227
adjust-percent .....	228
atm-service .....	229
buffer-size (Schedulers) .....	230
cbr .....	232
classifiers (Definition) .....	233
classifiers (Logical Interface) .....	234
classifiers (Physical Interface) .....	235
code-points (CoS) .....	236
delay-buffer-rate .....	237
demux0 (Dynamic Interface) .....	238
demux-options (Dynamic Interface) .....	239
demux-source (Dynamic IP Demux Interface) .....	240
demux-source (Dynamic Underlying Interface) .....	241
demux-source (Underlying Interface) .....	242
dhcp-attributes (Address-Assignment Pools) .....	243
dhcp-local-server .....	245
dhcp-relay .....	252
drop-probability (Interpolated Value) .....	264
drop-probability (Percentage) .....	265
drop-profile (Schedulers) .....	266
drop-profile-map (Schedulers) .....	267
drop-profiles .....	268
dscp (CoS Classifiers) .....	269
dscp (Rewrite Rules) .....	270
dscp-ipv6 (CoS Rewrite Rules) .....	272
dynamic-profile (DHCP Local Server) .....	273
dynamic-profile (VLAN) .....	274
dynamic-profiles .....	275
egress-shaping-overhead .....	284
excess-bandwidth-share .....	286
excess-priority .....	287
excess-rate (Dynamic Traffic Shaping) .....	288
excess-rate-high (Dynamic Traffic Shaping) .....	289
excess-rate-low (Dynamic Traffic Shaping) .....	290
family .....	291
family (Address-Assignment Pools) .....	296
family (Dynamic Demux Interface) .....	297
family (Dynamic PPPoE) .....	298
family (Dynamic Standard Interface) .....	299
fill-level (Drop Profiles) .....	301
fill-level (Interpolated Value) .....	302
filter (Configuring) .....	303
flexible-vlan-tagging .....	304
forwarding-classes (Class-of-Service) .....	305



group (DHCP Local Server) . . . . .	306
guaranteed-rate . . . . .	309
hierarchical-scheduler . . . . .	310
hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) . . . . .	311
ieee-802.1ad . . . . .	312
inet-precedence (CoS Classifiers) . . . . .	313
input-excess-bandwidth-share . . . . .	314
input-scheduler-map . . . . .	315
input-shaping-rate (Logical Interface) . . . . .	316
input-shaping-rate (Physical Interface) . . . . .	317
input-traffic-control-profile . . . . .	318
input-traffic-control-profile-remaining . . . . .	319
interface (DHCP Local Server) . . . . .	320
interface-set (Ethernet Interfaces) . . . . .	321
interface-set (Hierarchical Schedulers) . . . . .	322
interfaces . . . . .	323
interfaces (CoS) . . . . .	324
interfaces (Static and Dynamic Subscribers) . . . . .	326
internal-node . . . . .	330
interpolate . . . . .	331
loss-priority (BA Classifiers) . . . . .	332
loss-priority (Scheduler Drop Profiles) . . . . .	333
max-queues . . . . .	334
max-queues-per-interface . . . . .	336
maximum-lease-time . . . . .	337
member-link-scheduler . . . . .	338
mode (Layer 2 Tunneling Protocol Shaping) . . . . .	339
output-traffic-control-profile . . . . .	340
output-traffic-control-profile-remaining . . . . .	342
overhead-accounting . . . . .	343
pap (Dynamic PPP) . . . . .	344
peak-rate . . . . .	345
per-unit-scheduler . . . . .	346
pool (Address-Assignment Pools) . . . . .	348
pppoe-options (Dynamic PPPoE) . . . . .	349
ppp-options (Dynamic PPP) . . . . .	350
priority (Schedulers) . . . . .	352
profile (Access) . . . . .	353
protocols . . . . .	358
proxy-arp . . . . .	361
qos-adjust-hierarchical . . . . .	362
queue (Global Queues) . . . . .	363
radius (Access Profile) . . . . .	364
radius-server . . . . .	367
range (Address-Assignment Pools) . . . . .	368
ranges (Dynamic VLAN) . . . . .	369
rewrite-rules (Definition) . . . . .	370
rewrite-rules (Interfaces) . . . . .	371
routing-instances (Multiple Routing Entities) . . . . .	373

scheduler (Scheduler Map) .....	374
scheduler-map (Interfaces and Traffic-Control Profiles) .....	375
scheduler-maps (For Most Interface Types) .....	376
schedulers (CoS) .....	377
secret .....	378
server (Dynamic PPPoE) .....	379
server-group .....	380
shaping-rate (Applying to an Interface) .....	381
shaping-rate (Schedulers) .....	383
shaping-rate (Oversubscribing an Interface) .....	385
shaping-rate-excess-high .....	387
shaping-rate-excess-low .....	388
shaping-rate-priority-high .....	389
shaping-rate-priority-low .....	390
shaping-rate-priority-medium .....	391
shared-bandwidth-policer (Configuring) .....	392
shared-instance .....	393
shared-scheduler .....	394
simple-filter (Applying to an Interface) .....	395
simple-filter .....	396
site-identifier (VPLS) .....	397
stacked-vlan-tagging .....	398
system .....	398
term (Simple Filter) .....	399
three-color-policer (Applying) .....	400
three-color-policer (Configuring) .....	401
traffic-control-profiles .....	403
traffic-manager .....	404
transmit-rate (Schedulers) .....	408
transparent .....	410
underlying-interface (demux0) .....	411
underlying-interface (Dynamic PPPoE) .....	412
unit .....	413
unit (Dynamic Demux Interface) .....	414
unit (Dynamic PPPoE) .....	416
unnumbered-address (Dynamic PPPoE) .....	418
unnumbered-address (Dynamic Profiles) .....	419
unnumbered-address (Ethernet) .....	421
user-prefix (DHCP Relay Agent) .....	422
vlan-id (Dynamic Profiles) .....	423
vlan-id (VLAN ID to Be Bound to a Logical Interface) .....	424
vlan-model .....	425
vlan-ranges .....	426
vlan-tag .....	427
vlan-tags .....	428
vlan-tags (Stacked VLAN Tags) .....	429

<b>Chapter 11</b>	<b>Operational Commands . . . . .</b>	<b>431</b>
	show access-cac interface-set . . . . .	432
	show class-of-service adjustment-control-profile . . . . .	434
	show class-of-service classifier . . . . .	436
	show class-of-service interface . . . . .	439
	show class-of-service scheduler-map . . . . .	468
	show class-of-service scheduler-hierarchy interface . . . . .	471
	show class-of-service scheduler-hierarchy interface-set . . . . .	473
	show interfaces . . . . .	475
	show interfaces (PPPoE) . . . . .	553
	show interfaces demux0 (Demux Interfaces) . . . . .	563
	show interfaces queue . . . . .	573
	show subscribers . . . . .	615



# List of Figures

<b>Part 1</b>	<b>Hierarchical Class of Service</b>	
<b>Chapter 1</b>	<b>Configuring Hierarchical Class of Service on MX Series 3D Universal Edge Routers</b>	<b>3</b>
	Figure 1: Hierarchical Scheduling Architecture	4
	Figure 2: Hierarchical Schedulers and Priorities	12
	Figure 3: Building a Scheduler Hierarchy	15
<b>Chapter 2</b>	<b>Configuring Hierarchical Class of Service on MICs, MPCs, MLCs, and Aggregated Ethernet Interfaces</b>	<b>39</b>
	Figure 4: Scheduler Node Levels for Per-Unit Queuing MPCs	45
	Figure 5: Scheduler Node Levels for Interfaces on Hierarchical Queuing and Scheduling MICs and MPCs	45
	Figure 6: Scaled Mode for Aggregated Ethernet Interfaces	61
	Figure 7: Replicated Mode for Aggregated Ethernet Interfaces	63
<b>Part 2</b>	<b>Hierarchical CoS for Subscriber Management</b>	
<b>Chapter 3</b>	<b>Hierarchical Class of Service for Subscriber Management</b>	<b>69</b>
	Figure 8: Two-Level Hierarchical Scheduling	71
	Figure 9: Three-Level Hierarchical Scheduling—Logical Interfaces at Level 3 with Underlying Logical Interfaces at Level 2	73
	Figure 10: Three-Level Hierarchical Scheduling—Logical Interfaces at Level 2 with Interface Sets at Level 3	73
	Figure 11: Logical Interfaces at Level 2 and Interface Sets at Level 3	74
	Figure 12: Logical Interfaces at Level 3 and Underlying Logical Interfaces at Level 2	75
	Figure 13: VLAN Subscriber Interfaces	82
<b>Chapter 5</b>	<b>Configuring Hierarchical Scheduling for MPLS Pseudowire Interfaces</b>	<b>117</b>
	Figure 14: MPLS Pseudowire Subscriber Interface Two-Level Scheduler Configuration	121
	Figure 15: Three-Level Scheduling Hierarchy Case 1: Pseudowire Service Logical Interfaces over a Transport Logical Interface	124
	Figure 16: Three-Level Scheduling Hierarchy Case 2: Pseudowire Service Logical Interfaces over a Pseudowire Service Interface Set	125
	Figure 17: Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces—Deployment Scenario	126
<b>Chapter 6</b>	<b>Configuring Hierarchical Scheduling for L2TP</b>	<b>131</b>
	Figure 18: CoS Configuration for L2TP LAC Topology	132
	Figure 19: Processing of CoS Parameters in an L2TP LNS Inline Service	135

<b>Chapter 7</b>	<b>Preventing Bandwidth Contention on Subscriber Interfaces . . . . .</b>	<b>141</b>
	Figure 20: Scheduler Node and Queues with Adjusted Shaping Rates . . . . .	142
	Figure 21: Queue with Adjusted Shaping Rate . . . . .	143
	Figure 22: Typical Broadcast Television Network Topology. . . . .	160
	Figure 23: Typical Subscriber Scheduler Hierarchy with M-VLANs . . . . .	161
<b>Chapter 8</b>	<b>Configuring Targeted Distribution of Subscribers on Aggregated Ethernet Interfaces . . . . .</b>	<b>163</b>
	Figure 24: Targeted Subscriber Links . . . . .	164
	Figure 25: Multicast Traffic Separation Using OIF Mapping . . . . .	171

# List of Tables

	<b>About the Documentation</b> . . . . .	xvii
	Table 1: Notice Icons . . . . .	xix
	Table 2: Text and Syntax Conventions . . . . .	xx
<b>Part 1</b>	<b>Hierarchical Class of Service</b>	
<b>Chapter 1</b>	<b>Configuring Hierarchical Class of Service on MX Series 3D Universal Edge Routers</b> . . . . .	<b>3</b>
	Table 3: Resources for Learning More About HCoS . . . . .	5
	Table 4: Node Levels Designations in Hierarchical Scheduling . . . . .	9
	Table 5: Queue Priority . . . . .	10
	Table 6: Internal Node Queue Priority for CIR Mode . . . . .	11
	Table 7: Internal Node Queue Priority for PIR-Only Mode . . . . .	11
<b>Chapter 2</b>	<b>Configuring Hierarchical Class of Service on MICs, MPCs, MLCs, and Aggregated Ethernet Interfaces</b> . . . . .	<b>39</b>
	Table 8: Dedicated Queues for MPCs . . . . .	46
	Table 9: Default Maximum Queues and Corresponding Rate Wheel Update Periods . . . . .	49
	Table 10: Default Shaping Granularities on Non-Enhanced Queuing MICs and MPCs . . . . .	50
<b>Part 2</b>	<b>Hierarchical CoS for Subscriber Management</b>	
<b>Chapter 3</b>	<b>Hierarchical Class of Service for Subscriber Management</b> . . . . .	<b>69</b>
	Table 11: Two-Level Hierarchical Scheduling—Interface Hierarchy Versus Scheduling Nodes . . . . .	72
	Table 12: Three-Level Hierarchical Scheduling—Interface Hierarchy Versus CoS Scheduling Node Levels . . . . .	72
	Table 13: Four-Level Hierarchical Scheduling—Interface Hierarchy Versus CoS Scheduling Node Levels . . . . .	76
	Table 14: Hardware Required for Dynamic Hierarchical CoS Configurations . . . . .	77
<b>Chapter 5</b>	<b>Configuring Hierarchical Scheduling for MPLS Pseudowire Interfaces</b> . . . . .	<b>117</b>
	Table 15: Two-Level Hierarchical Scheduling—Interface Hierarchy Versus Scheduling Nodes . . . . .	120
	Table 16: Three-Level Hierarchical Scheduling—Interface Hierarchy Versus CoS Scheduling Node Levels . . . . .	124
<b>Chapter 6</b>	<b>Configuring Hierarchical Scheduling for L2TP</b> . . . . .	<b>131</b>
	Table 17: Ingress LAC Tunnel Classifier Options . . . . .	132

	Table 18: Sample Result for the Classifier and Rewrite Rules for a VLAN Interface . . . . .	133
	Table 19: Hardware Requirements for L2TP LNS Inline Services . . . . .	136
<b>Chapter 9</b>	<b>Applying CoS Using Parameters Received from RADIUS . . . . .</b>	<b>183</b>
	Table 20: CoS Predefined Variables for Scheduler Map and Traffic Shaping . . . . .	184
	Table 21: CoS Predefined Variables for Scheduling and Queuing . . . . .	185
	Table 22: CoS Services and Variables . . . . .	188
	Table 23: Junos OS CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets . . . . .	199
	Table 24: Initial Scheduler Map and Shaping Values at Subscriber Login . . . . .	205
	Table 25: Initial CoS Values for the Voice Scheduler at Subscriber Login . . . . .	205
	Table 26: Initial CoS Values for the Data Scheduler at Subscriber Login . . . . .	206
	Table 27: Upgraded CoS Values for the Video Service . . . . .	208
	Table 28: Upgraded CoS Values for the Video Scheduler . . . . .	209
	Table 29: Initial CoS Values for the Expedited Forwarding Scheduler at Subscriber Login . . . . .	209
	Table 30: Initial CoS Values for the Best Effort Scheduler at Subscriber Login . . . . .	210
<b>Part 3</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 11</b>	<b>Operational Commands . . . . .</b>	<b>431</b>
	Table 31: show class-of-service interface-set Output Fields . . . . .	432
	Table 32: show class-of-service adjustment-control-profile Output Fields . . . . .	434
	Table 33: show class-of-service classifier Output Fields . . . . .	436
	Table 34: show class-of-service interface Output Fields . . . . .	440
	Table 35: show class-of-service scheduler-map Output Fields . . . . .	468
	Table 36: show class-of-service scheduler-hierarchy interface Output Fields . . . . .	471
	Table 37: show class-of-service scheduler-hierarchy interface-set Output Fields . . . . .	473
	Table 38: show interfaces (Gigabit Ethernet) Output Fields . . . . .	479
	Table 39: Gigabit and 10 Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type . . . . .	507
	Table 40: show interfaces Output Fields . . . . .	508
	Table 41: show interfaces (PPPoE) Output Fields . . . . .	553
	Table 42: show interfaces demux0 (Demux Interfaces) Output Fields . . . . .	563
	Table 43: Layer 2 Overhead and Transmitted Packets or Byte Counts . . . . .	574
	Table 44: show interfaces queue Output Fields . . . . .	577
	Table 45: Byte Count by Interface Hardware . . . . .	581
	Table 46: show subscribers Output Fields . . . . .	619



# 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 <https://www.juniper.net/documentation/>.

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 <https://www.juniper.net/books>.

## Supported Platforms

---

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

- MX Series
- T Series
- M 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 xx](#) 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>
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 <b>ospf area area-id</b>] hierarchy level.</li> <li>The console port is labeled <b>CONSOLE</b>.</li> </ul>
< > (angle brackets)	Encloses optional keywords or variables.	<b>stub</b> <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	<b>broadcast</b>   <b>multicast</b>  ( <i>string1</i>   <i>string2</i>   <i>string3</i> )
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<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> [ <b>community-ids</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.	

#### GUI Conventions

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

Convention	Description	Examples
<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>
> (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 <https://www.juniper.net/documentation/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 <https://www.juniper.net/documentation/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 <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://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: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <https://www.juniper.net/cm/>

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

## 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 <https://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 <https://www.juniper.net/support/requesting-support.html>.

## PART 1

# Hierarchical Class of Service

- [Configuring Hierarchical Class of Service on MX Series 3D Universal Edge Routers on page 3](#)
- [Configuring Hierarchical Class of Service on MICs, MPCs, MLCs, and Aggregated Ethernet Interfaces on page 39](#)





## CHAPTER 1

# Configuring Hierarchical Class of Service on MX Series 3D Universal Edge Routers

- [Hierarchical Class of Service Overview on page 3](#)
- [Hierarchical Class of Service Network Scenarios on page 6](#)
- [Understanding Hierarchical Scheduling on page 7](#)
- [Priority Propagation on page 10](#)
- [Configuring Hierarchical Schedulers for CoS on page 12](#)
- [Hierarchical Schedulers and Traffic Control Profiles on page 13](#)
- [Example: Building a Four-Level Hierarchy of Schedulers on page 15](#)
- [Scheduling and Shaping in Hierarchical CoS Queues for Traffic Routed to GRE Tunnels on page 20](#)
- [Example: Performing Output Scheduling and Shaping in Hierarchical CoS Queues for Traffic Routed to GRE Tunnels on page 21](#)
- [Configuring Ingress Hierarchical CoS on Enhanced Queuing DPCs on page 35](#)

## Hierarchical Class of Service Overview

---

Hierarchical class of service (HCoS) is the ability to apply traffic schedulers and shapers to a hierarchy of *scheduler nodes*. Each level of the scheduler hierarchy can be used to shape traffic based on different criteria such as application, user, VLAN, and physical port.

This allows you to support the requirements of different services, applications, and users on the same physical device and physical infrastructure.


HCoS is implemented primarily using traffic classifiers at the ingress and hierarchical schedulers and shapers at the egress.

A classifier is a filter that labels traffic at the device ingress based on configurable parameters such as application or destination. Traffic is classified into what is called a forwarding equivalence class (FEC). The FEC defines a class of traffic that receives common treatment.

Schedulers, and their associated shapers, is the function that controls the traffic bandwidth, jitter (delay variation), and packet loss priority at the egress of the device.

Hierarchical schedulers are used to apply multiple levels of scheduling and shaping with each level applied to different classifications such as forwarding equivalence class, VLAN, and physical interface (port) as shown in [Figure 1 on page 4](#).

**Figure 1: Hierarchical Scheduling Architecture**



**NOTE:** Hierarchical class of service is also referred to as Hierarchical Quality of Service (HQoS) in other vendor's documentation.

A typical application of HCoS is to configure multiple levels of egress schedulers and shapers, at the subscriber edge, using dynamic profiles to provide traffic shaping and prioritization at the subscriber VLAN level and for multiple classes of traffic.

Dynamic profiles are a mechanism that allows you to dynamically apply schedulers and shapers to individual subscribers or groups of subscribers.

To learn more about HCoS, the following topics are very helpful:

- *Junos CoS on MX Series 3D Universal Edge Routers Overview*
- *CoS Features and Limitations on MX Series Routers*
- *CoS Features of the Router Hardware, PIC, MIC, and MPC Interface Families*
- *How Schedulers Define Output Queue Properties*
- *Subscriber Access Network Overview*

- *CoS for Subscriber Access Overview*
- [Hierarchical Class of Service for Subscriber Management Overview on page 69](#)

The Junos OS hierarchical schedulers support up to five levels of scheduler hierarchies on MX Series devices when using enhanced queuing Dense Port Concentrators (DPCs) or fine-grained queuing Modular Port Concentrators (MPCs), and Modular Interface Cards (MICs). It is important to know the capabilities of your hardware with respect to HCoS. The following are a few tips to help you:

- Only certain hardware supports the five-level scheduler hierarchy of HCoS.
- The number of queues and logical interfaces supported is dependent upon exactly what hardware you are using.
- The MX Series Packet Forwarding Engine handles guaranteed bandwidth and scheduler node weight differently than other Packet Forwarding Engines.
- The fine-grained queuing MPCs and MICs have a certain granularity with respect to the shaping and delay buffer values. The values used are not necessarily exactly the values configured.

To learn more about platform support for HCoS, use the Juniper Networks Feature Explorer (<https://pathfinder.juniper.net/feature-explorer/>). In the Feature Explorer, search on *hierarchical schedulers*.

In addition, it is important to note the following:

- HCoS is most frequently used to enforce service level agreements at the subscriber edged using dynamic traffic control profiles.
- Hierarchical schedulers can also be applied to Ethernet pseudowire interfaces, aggregated Ethernet interfaces, Layer 2 Tunnel Protocol (L2TP) network server (LNS) inline services, and GRE tunnels.
- Hierarchical ingress policing is a feature that is complimentary to and often used in conjunction with HCoS.
- There are other features in Junos OS that have similar sounding names.



**NOTE:** The *hierarchical scheduler and shaper* feature supported on the SRX Series devices is not the HCoS feature described here.

Before planning HCoS for you network, you should learn about HCoS, define you needs, plan how you want to implement HCoS, and test the operation in a simulated environment.

**Table 3: Resources for Learning More About HCoS**

Document	Description
<a href="#">Day One: Deploying Basic QoS</a> Juniper Networks Books	This book is a good resource for learning the basics of CoS on Juniper Networks devices.

Table 3: Resources for Learning More About HCoS (*continued*)

Document	Description
Juniper MX-Series O'Reilly Media	Learn about the advanced features of HCoS. This book provides an in-depth description of how HCoS works and how it can be deployed. It also provides a lab tested topology and configuration example.
<a href="#">Day One: Dynamic Subscriber Management</a> Juniper Networks Books	Learn how to use HCoS in conjunction with dynamic traffic control profiles for subscriber management. This book also includes troubleshooting.
QoS Enabled Networks John Wiley & Sons	This book is an additional source for studying QoS.

Documentation related to HCoS is consolidated in the *Hierarchical Class of Service Feature Guide*.

#### Related Documentation

- [Hierarchical Class of Service for Subscriber Management Overview on page 69](#)
- [Hierarchical Class of Service Network Scenarios on page 6](#)
- [Understanding Hierarchical Scheduling on page 7](#)

## Hierarchical Class of Service Network Scenarios

Hierarchical class of service (HCoS) can be used to provide granular control of traffic for a variety of different applications.



**NOTE:** Hierarchical class of service is also referred to as Hierarchical Quality of Service (HQoS) in other vendor's documentation.

Hierarchical class of service is most frequently used in the following scenarios:

### Services to Subscribers

Multiservice network operators face a challenge to provide different types of services on the same infrastructure to residential and business subscribers. The network operator needs to make sure each subscriber gets the network resources they paid for and each service gets the network resources it needs to operate properly.

If no CoS is applied, one service could consume most of the bandwidth of the transmission infrastructure and starve the other services.

Using hierarchical class of service, the network edge device can have up to five levels of scheduling and prioritization. So the traffic can be shaped and prioritized per customer and per service type. Controlling traffic in this way provides the ability to deliver the required service level for each subscriber for each service type.

By allowing network operators to consolidate different services and multiple customers on the same physical infrastructure, hierarchical class of service helps maximize the ability to offer revenue generating services while simultaneously minimizing capital cost.

## Services to Businesses

Hierarchical class of service is a valuable tool for service providers that support business customers who are running applications with different prioritization and scheduling requirements over the same infrastructure. In this scenario hierarchical class of service allows lower priority traffic to fully utilize the available bandwidth on a port, while simultaneously ensuring low latency and guaranteed bandwidth to higher priority traffic on the same port.

This allows a provider to consolidate different services on the same physical device and physical infrastructure thus optimizing network resources while maintaining the required level of service.

All of this maximizes revenue and minimizes cost

## Wireless Backhaul

In a cellular network the operator might want to offer business services along with its cell tower traffic. One of the main challenges is to make sure that the time-sensitive cell traffic is not affected by the business services running on the same infrastructure. Each type of traffic has its own priority flows and bandwidth constraints. For example, wireless backhaul is very sensitive to fluctuations in the packet stream (Jitter) because it relies on synchronization.

In this scenario, hierarchical class of service allows each type of traffic to receive the required resources and quality of service while being delivered over the same infrastructure.

By consolidate different services on the same physical infrastructure, HCoS helps maximize revenue and minimize cost.

### Related Documentation

- [Hierarchical Class of Service Overview on page 3](#)
- [Hierarchical Class of Service for Subscriber Management Overview on page 69](#)

---

## Understanding Hierarchical Scheduling

Hierarchical class of service (HCoS) is a set of capabilities that enable you to apply unique CoS treatment for network traffic based on criteria such as user, application, VLAN, and physical port.

This allows you to support the requirements of different services, applications, and users on the same physical device and physical infrastructure.

This topic covers the following information:

- [Hierarchical Scheduling Terminology on page 8](#)
- [Scheduler Node-Level Designations in Hierarchical Scheduling on page 8](#)
- [Hierarchical Scheduling at Non-Leaf Nodes on page 9](#)

## Hierarchical Scheduling Terminology

Hierarchical scheduling introduces some new CoS terms and also uses some familiar terms in different contexts:

- **Customer VLAN (C-VLAN)**—A C-VLAN, defined by IEEE 802.1ad. A stacked VLAN contains an outer tag corresponding to the S-VLAN, and an inner tag corresponding to the C-VLAN. A C-VLAN often corresponds to CPE. Scheduling and shaping is often used on a C-VLAN to establish minimum and maximum bandwidth limits for a customer. See also *S-VLAN*.
- **Interface set**—A logical group of interfaces that describe the characteristics of set of service VLANs, logical interfaces, customer VLANs, or aggregated Ethernet interfaces. Interface sets establish the set and name the traffic control profiles. See also *Service VLAN*.
- **Scheduler**—A scheduler defines the scheduling and queuing characteristics of a queue. Transmit rate, scheduler priority, and buffer size can be specified. In addition, a drop profile may be referenced to describe WRED congestion control aspects of the queue. See also *Scheduler map*.
- **Scheduler map**—A scheduler map is referenced by traffic control profiles to define queues. The scheduler map establishes the queues that comprise a scheduler node and associates a forwarding class with a scheduler. See also *Scheduler*.
- **Stacked VLAN**—An encapsulation on an S-VLAN with an outer tag corresponding to the S-VLAN, and an inner tag corresponding to the C-VLAN. See also *Service VLAN* and *Customer VLAN*.
- **Service VLAN (S-VLAN)**—An S-VLAN, defined by IEEE 802.1ad, often corresponds to a network aggregation device such as a DSLAM. Scheduling and shaping is often established for an S-VLAN to provide CoS for downstream devices with little buffering and simple schedulers. See also *Customer VLAN*.
- **Traffic control profile**—Defines the characteristics of a scheduler node. Traffic control profiles are used at several levels of the CLI, including the physical interface, interface set, and logical interface levels. Scheduling and queuing characteristics can be defined for the scheduler node using the **shaping-rate**, **guaranteed-rate**, and **delay-buffer-rate** statements. Queues over these scheduler nodes are defined by referencing a scheduler map. See also *Scheduler* and *Scheduler map*.
- **VLAN**—Virtual LAN, defined on an Ethernet logical interface.

## Scheduler Node-Level Designations in Hierarchical Scheduling

Scheduler hierarchies are composed of nodes and queues. Queues terminate the hierarchy. Nodes can be either root nodes, leaf nodes, or internal (non-leaf) nodes. Internal nodes are nodes that have other nodes as “children” in the hierarchy.

Scheduler hierarchies consist of levels, starting with Level 1 at the physical port. This chapter establishes a four-level scheduler hierarchy which, when fully configured, consists of the physical interface (Level 1), the interface set (Level 2), one or more logical interfaces (Level 3), and one or more queues (Level 4).



**NOTE:** Beginning with Junos OS Release 16.1, certain MPCs on MX Series devices support up to five levels of scheduler hierarchies. The concepts presented in this topic apply similarly to five scheduler hierarchy levels.

Table 4 on page 9 describes the possible combinations of scheduler nodes and their corresponding node level designations for a hierarchical queuing MIC or MPC.

**Table 4: Node Levels Designations in Hierarchical Scheduling**

Scheduler Configuration for Hierarchical CoS	Hierarchical CoS Scheduler Nodes			
	Root Node	Internal (Non-Leaf) Nodes		Leaf Node
	Level 1	Level 2	Level 3	Level 4
One or more traffic control profiles configured on logical interfaces, but no interface-sets configured	Physical interface	—	One or more logical interfaces	One or more queues
Interface-sets (collections of logical interfaces) configured, but no traffic-control profiles configured on logical interfaces	Physical interface	—	Interface-set	One or more queues
Fully configured scheduler nodes	Physical interface	Interface-set	One or more logical interfaces	One or more queues

The table illustrates how the configuration of an interface set or logical interface affects the terminology of hierarchical scheduler nodes. For example, suppose you configure an **interface-set** statement with logical interfaces (such as **unit 0** and **unit 2**) and a queue. In this case, the interface-set is an internal node at Level 2 of the scheduler node hierarchy. However, if there are no traffic control profiles attached to logical interfaces, then the interface set is at Level 3 of the hierarchy.

## Hierarchical Scheduling at Non-Leaf Nodes

Whereas standard CoS scheduling is based on the scheduling and queuing characteristics of a router's egress ports and their queues, hierarchical CoS scheduling is based on the scheduling and queuing characteristics that span a hierarchy of *scheduler nodes* over a port. The hierarchy begins at Level 1, a *root node* at the physical interface (port) level of the CLI hierarchy and terminates at Level 4, a *leaf node* at the queue level. Between the root and leaf nodes of any scheduler hierarchy are one or more *internal nodes*, which are non-root nodes that have other nodes as “children” in the hierarchy.

Whereas you configure standard CoS scheduling by applying a scheduler map to each egress port to specify a forwarding class and a queue priority level, you configure hierarchical CoS scheduling with additional parameters. To configure hierarchical CoS scheduling, you apply a scheduler map to the queue level (Level 4) of a scheduler hierarchy, and you can apply a different traffic control profile at each of the other levels. A traffic control profile specifies not only a scheduler map (forwarding class and queue

priority level) but also optional shaping rate (PIR), guaranteed transmit rate (CIR), burst rate, delay buffer rate, and drop profile.

#### Release History Table

Release	Description
16.1	Beginning with Junos OS Release 16.1, certain MPCs on MX Series devices support up to five levels of scheduler hierarchies.

## Priority Propagation

Priority propagation is performed for MX Series router output Interfaces on Enhanced Queuing DPCs, MICs, and MPCs, and for M Series and T Series router output interfaces on IQ2E PICs. Priority propagation is useful for mixed traffic environments when, for example, you want to make sure that the voice traffic of one customer does not suffer due to the data traffic of another customer. Nodes and queues are serviced in the order of their priority. The default priority of a queue is low, and you can explicitly configure a queue priority by including the **priority** statement at the **[edit class-of-service schedulers scheduler-name]** hierarchy level.

You cannot directly configure the priorities of all hierarchical scheduling elements. The priorities of internal nodes, for example, are determined as follows:

- The highest priority of an active child, that is, a child currently containing traffic. (Interface sets only take the highest priority of their active children.)
- Whether the node is above its configured guaranteed rate (CIR) or not (this is only relevant if the physical interface is in CIR mode).

Each queue has a configured priority and a hardware priority. The usual mapping between the configured priority and the hardware priority is shown in [Table 5 on page 10](#).

**Table 5: Queue Priority**

Configured Priority	Hardware Priority
Strict-high	0
High	0
Medium-high	1
Medium-low	1
Low	2
MPCs also have configurable CLI priorities of <b>excess-priority high</b> , <b>excess-priority medium-high</b> , <b>excess-priority medium-low</b> , and <b>excess-priority low</b> . These priorities only take effect above the guaranteed rate.	



In CIR mode, the priority for each internal node depends on whether the highest active child node is above or below the guaranteed rate. The mapping between the highest active child's priority and the hardware priority below and above the guaranteed rate is shown in [Table 6 on page 11](#).

**Table 6: Internal Node Queue Priority for CIR Mode**

Configured Priority of Highest Active Child Node	Hardware Priority Below Guaranteed Rate	Hardware Priority Above Guaranteed Rate
Strict-high	0	0
High	0	3
Medium-high	1	3
Medium-low	1	3
Low	2	3
Excess-priority high*	N/A	3
Excess-priority medium-high*	N/A	3
Excess-priority medium-low*	N/A	4
Excess-priority low*	N/A	4
* MPCs only		

In PIR-only mode, nodes cannot send if they are above the configured shaping rate. The mapping between the configured priority and the hardware priority is for PIR-only mode is shown in [Table 7 on page 11](#).


**Table 7: Internal Node Queue Priority for PIR-Only Mode**

Configured Priority	Hardware Priority
Strict-high	0
High	0
Medium-high	1
Medium-low	1
Low	2

A physical interface with hierarchical schedulers configured is shown in [Figure 2 on page 12](#). The configured priorities are shown for each queue at the top of the figure. The hardware priorities for each node are shown in parentheses. Each node also shows any configured

shaping rate (PIR) or guaranteed rate (CIR) and whether or not the queues is above or below the CIR. The nodes are shown in one of three states: above the CIR (clear), below the CIR (dark), or in a condition where the CIR does not matter (gray).

**Figure 2: Hierarchical Schedulers and Priorities**



In the figure, the strict-high queue for customer VLAN 0 (cvlan 0) receives service first, even though the customer VLAN is above the configured CIR (see [Table 6 on page 11](#) for the reason: strict-high always has hardware priority 0 regardless of CIR state). Once that queue has been drained, and the priority of the node has become 3 instead of 0 (due to the lack of strict-high traffic), the system moves on to the medium queues next (cvlan 1 and cvlan 3), draining them in a round robin fashion (empty queue lose their hardware priority). The low queue on cvlan 4 (priority 2) is sent next, because that node is below the CIR. Then the high queues on cvlan 0 and cvlan 2 (both now with priority 3) are drained in a round robin fashion, and finally the low queue on cvlan 0 is drained (thanks to svlan 0 having a priority of 3).

#### Related Documentation

- [CoS on Enhanced IQ2 PICs Overview](#)
- [Enhanced Queuing DPC CoS Properties](#)
- [CoS Features and Limitations on MIC and MPC Interfaces](#)
- [Understanding Hierarchical Scheduling for MIC and MPC Interfaces on page 39](#)

## Configuring Hierarchical Schedulers for CoS

In metro Ethernet environments, a virtual LAN (VLAN) typically corresponds to a customer premises equipment (CPE) device and the VLANs are identified by an inner VLAN tag on Ethernet frames (called the customer VLAN, or C-VLAN, tag). A set of VLANs can be grouped at the DSL access multiplexer (DSLAM) and identified by using the same outer VLAN tag (called the service VLAN, or S-VLAN, tag). The service VLANs are typically gathered at the Broadband Remote Access Server (B-RAS) level. Hierarchical schedulers let you provide shaping and scheduling at the service VLAN level as well as other levels, such as the physical interface. In other words, you can group a set of logical interfaces and then apply scheduling and shaping parameters to the logical interface set as well as to other levels.

On Juniper Networks MX Series 3D Universal Edge Routers and systems with Enhanced IQ2 (IQ2E) PICs, you can apply CoS shaping and scheduling at one of four different levels, including the VLAN set level. You can only use this configuration on MX Series routers or IQ2E PICs. Beginning with Junos OS Release 16.1, certain MPCs support up to five levels of scheduler hierarchies.

The supported scheduler hierarchy is as follows:

- The physical interface (level 1)
- The service VLAN (level 2 is unique to MX Series routers)
- The logical interface or customer VLAN (level 3)
- The queue (level 4)

Users can specify a traffic control profile (**output-traffic-control-profile**) that can specify a shaping rate, a guaranteed rate, and a scheduler map with transmit rate and buffer delay. The scheduler map contains the mapping of queues (forwarding classes) to their respective schedulers (schedulers define the properties for the queue). Queue properties can specify a transmit rate and buffer management parameters such as buffer size and drop profile.

To configure CoS hierarchical scheduling, you must enable hierarchical scheduling by including the **hierarchical-scheduler** statement at the physical interface.

**Related Documentation**

- [Understanding Hierarchical Scheduling on page 7](#)
- [Understanding Hierarchical Scheduling for MIC and MPC Interfaces on page 39](#)
- [CoS on Enhanced IQ2 PICs Overview](#)
- [Understanding Hierarchical CoS for Subscriber Interfaces on page 70](#)

---

## Hierarchical Schedulers and Traffic Control Profiles

When used, the interface set level of the hierarchy falls between the physical interface level (Level 1) and the logical interface (Level 3). Queues are always Level 4 of the hierarchy.



.....

**NOTE:** Beginning with Junos OS Release 16.1, certain MPCs on MX Series devices support up to five levels of scheduler hierarchies. The concepts presented in this topic apply similarly to five scheduler hierarchy levels.

.....

Hierarchical schedulers add CoS parameters to the interface-set level of the configuration. They use traffic control profiles to set values for parameters such as shaping rate (the peak information rate [PIR]), guaranteed rate (the committed information rate [CIR] on these interfaces), scheduler maps (assigning queues and resources to traffic), and so on.

The following CoS configuration places the following parameters in traffic control profiles at various levels:

- Traffic control profile at the port level (**tcp-port-level1**):
  - A shaping rate (PIR) of 100 Mbps
  - A delay buffer rate of 100 Mbps
- Traffic control profile at the interface set level (**tcp-interface-level2**):
  - A shaping rate (PIR) of 60 Mbps
  - A guaranteed rate (CIR) of 40 Mbps
- Traffic control profile at the logical interface level (**tcp-unit-level3**):
  - A shaping rate (PIR) of 50 Mbps
  - A guaranteed rate (CIR) of 30 Mbps
  - A scheduler map called **smap1** to hold various queue properties (level 4)
  - A delay buffer rate of 40 Mbps

In this case, the traffic control profiles look like this:

```
[edit class-of-service traffic-control-profiles]
tcp-port-level1 { # This is the physical port level
  shaping-rate 100m;
  delay-buffer-rate 100m;
}
tcp-interface-level2 { # This is the interface set level
  shaping-rate 60m;
  guaranteed-rate 40m;
}
tcp-unit-level3 { # This is the logical interface level
  shaping-rate 50m;
  guaranteed-rate 30m;
  scheduler-map smap1;
  delay-buffer-rate 40m;
}
```

Once configured, the traffic control profiles must be applied to the proper places in the CoS interfaces hierarchy.

```
[edit class-of-service interfaces]
interface-set level-2 {
  output-traffic-control-profile tcp-interface-level-2;
}
ge-0/1/0 {
  output-traffic-control-profile tcp-port-level-1;
  unit 0 {
    output-traffic-control-profile tcp-unit-level-3;
  }
}
```

In all cases, the properties for level 4 of the hierarchical schedulers are determined by the scheduler map.

## Release History Table

Release	Description
16.1	Beginning with Junos OS Release 16.1, certain MPCs on MX Series devices support up to five levels of scheduler hierarchies.


## Related Documentation

- [Oversubscribing Interface Bandwidth](#)
- [Providing a Guaranteed Minimum Rate](#)
- [Configuring Scheduler Maps](#)
- [Configuring Traffic Control Profiles for Shared Scheduling and Shaping](#)

## Example: Building a Four-Level Hierarchy of Schedulers

This section provides a more complete example of building a 4-level hierarchy of schedulers. The configuration parameters are shown in [Figure 3 on page 15](#). The queues are shown at the top of the figure with the other three levels of the hierarchy below.

Figure 3: Building a Scheduler Hierarchy



The figure's PIR values are configured as the shaping rates and the CIRs are configured as the guaranteed rate on the Ethernet interface **ge-1/0/0**. The PIR can be oversubscribed (that is, the sum of the children PIRs can exceed the parent's, as in **svlan 1**, where 200 + 200 + 100 exceeds the parent rate of 400). However, the sum of the children node level's CIRs must never exceed the parent node's CIR, as shown in all the service VLANs (otherwise, the guaranteed rate could never be provided in all cases).

This configuration example presents all details of the CoS configuration for the interface in the figure (**ge-1/0/0**), including:

- [Configuring the Interface Sets on page 16](#)
- [Configuring the Interfaces on page 16](#)
- [Configuring the Traffic Control Profiles on page 16](#)
- [Configuring the Schedulers on page 17](#)

- [Configuring the Drop Profiles on page 18](#)
- [Configuring the Scheduler Maps on page 18](#)
- [Applying the Traffic Control Profiles on page 19](#)

## Configuring the Interface Sets

```
[edit interfaces]
interface-set svlan-0 {
  interface ge-1/0/0 {
    unit 0;
    unit 1;
  }
}
interface-set svlan-1 {
  interface ge-1/0/0 {
    unit 2;
    unit 3;
    unit 4;
  }
}
```

## Configuring the Interfaces

The keyword to configure hierarchical schedulers is at the physical interface level, as is VLAN tagging and the VLAN IDs. In this example, the interface sets are defined by logical interfaces (units) and not outer VLAN tags. All VLAN tags in this example are customer VLAN tags.

```
[edit interface ge-1/0/0]
hierarchical-scheduler;
vlan-tagging;
unit 0 {
  vlan-id 100;
}
unit 1 {
  vlan-id 101;
}
unit 2 {
  vlan-id 102;
}
unit 3 {
  vlan-id 103;
}
unit 4 {
  vlan-id 104;
}
```

## Configuring the Traffic Control Profiles

The traffic control profiles hold parameters for levels above the queue level of the scheduler hierarchy. This section defines traffic control profiles for both the service VLAN level (logical interfaces) and the customer VLAN (VLAN tag) level.

```
[edit class-of-service traffic-control-profiles]
tcp-500m-shaping-rate {
  shaping-rate 500m;
```

```

}
tcp-svlan0 {
    shaping-rate 200m;
    guaranteed-rate 100m;
    delay-buffer-rate 300m; # This parameter is not shown in the figure.
}
tcp-svlan1 {
    shaping-rate 400m;
    guaranteed-rate 300m;
    delay-buffer-rate 100m; # This parameter is not shown in the figure.
}
tcp-cvlan0 {
    shaping-rate 100m;
    guaranteed-rate 60m;
    scheduler-map tcp-map-cvlan0; # Applies scheduler maps to customer VLANs.
}
tcp-cvlan1 {
    shaping-rate 100m;
    guaranteed-rate 40m;
    scheduler-map tcp-map-cvlan1; # Applies scheduler maps to customer VLANs.
}
tcp-cvlan2 {
    shaping-rate 200m;
    guaranteed-rate 100m;
    scheduler-map tcp-map-cvlanx; # Applies scheduler maps to customer VLANs.
}
tcp-cvlan3 {
    shaping-rate 200m;
    guaranteed-rate 150m;
    scheduler-map tcp-map-cvlanx; # Applies scheduler maps to customer VLANs
}
tcp-cvlan4 {
    shaping-rate 100m;
    guaranteed-rate 50m;
    scheduler-map tcp-map-cvlanx; # Applies scheduler maps to customer VLANs
}
}

```

## Configuring the Schedulers

The schedulers hold the information about the queues, the last level of the hierarchy. Note the consistent naming schemes applied to repetitive elements in all parts of this example.

```

[edit class-of-service schedulers]
sched-cvlan0-qx {
    priority low;
    transmit-rate 20m;
    buffer-size temporal 100ms;
    drop-profile loss-priority low dp-low;
    drop-profile loss-priority high dp-high;
}
sched-cvlan1-q0 {
    priority high;
    transmit-rate 20m;
    buffer-size percent 40;
    drop-profile loss-priority low dp-low;
}

```

```
    drop-profile loss-priority high dp-high;
}
sched-cvlanx-qx {
    transmit-rate percent 30;
    buffer-size percent 30;
    drop-profile loss-priority low dp-low;
    drop-profile loss-priority high dp-high;
}
sched-cvlan1-qx {
    transmit-rate 10m;
    buffer-size temporal 100ms;
    drop-profile loss-priority low dp-low;
    drop-profile loss-priority high dp-high;
}
```

## Configuring the Drop Profiles

This section configures the drop profiles for the example. For more information about interpolated drop profiles, see *Managing Congestion Using RED Drop Profiles and Packet Loss Priorities*.

```
[edit class-of-service drop-profiles]
dp-low {
    interpolate fill-level 80 drop-probability 80;
    interpolate fill-level 100 drop-probability 100;
}
dp-high {
    interpolate fill-level 60 drop-probability 80;
    interpolate fill-level 80 drop-probability 100;
}
```

## Configuring the Scheduler Maps

This section configures the scheduler maps for the example. Each one references a scheduler configured in [“Configuring the Schedulers” on page 17](#).

```
[edit class-of-service scheduler-maps]
tcp-map-cvlan0 {
    forwarding-class voice scheduler sched-cvlan0-qx;
    forwarding-class video scheduler sched-cvlan0-qx;
    forwarding-class data scheduler sched-cvlan0-qx;
}
tcp-map-cvlan1 {
    forwarding-class voice scheduler sched-cvlan1-q0;
    forwarding-class video scheduler sched-cvlan1-qx;
    forwarding-class data scheduler sched-cvlan1-qx;
}
tcp-map-cvlanx {
    forwarding-class voice scheduler sched-cvlanx-qx;
    forwarding-class video scheduler sched-cvlanx-qx;
    forwarding-class data scheduler sched-cvlanx-qx;
}
```



## Applying the Traffic Control Profiles

This section applies the traffic control profiles to the proper levels of the hierarchy.



**NOTE:** Although a shaping rate can be applied directly to the physical interface, hierarchical schedulers must use a traffic control profile to hold this parameter.

```
[edit class-of-service interfaces]
ge-1/0/0 {
  output-traffic-control-profile tcp-500m-shaping-rate;
  unit 0 {
    output-traffic-control-profile tcp-cvlan0;
  }
  unit 1 {
    output-traffic-control-profile tcp-cvlan1;
  }
  unit 2 {
    output-traffic-control-profile tcp-cvlan2;
  }
  unit 3 {
    output-traffic-control-profile tcp-cvlan3;
  }
  unit 4 {
    output-traffic-control-profile tcp-cvlan4;
  }
}
interface-set svlan0 {
  output-traffic-control-profile tcp-svlan0;
}
interface-set svlan1 {
  output-traffic-control-profile tcp-svlan1;
}
```



**NOTE:** You should be careful when using a `show interfaces queue` command that references nonexistent class-of-service logical interfaces. When multiple logical interfaces (units) are not configured under the same interface set or physical interface, but are referenced by a command such as `show interfaces queue ge-10/0/1.12 forwarding-class be` or `show interfaces queue ge-10/0/1.13 forwarding-class be` (where logical units 12 and 13 are not configured as a class-of-service interfaces), these interfaces display the same traffic statistics for each logical interface. In other words, even if there is no traffic passing through a particular unconfigured logical interface, as long as one or more of the other unconfigured logical interfaces under the same interface set or physical interface is passing traffic, this particular logical interface displays statistics counters showing the total amount of traffic passed through all other unconfigured logical interfaces together.

## Scheduling and Shaping in Hierarchical CoS Queues for Traffic Routed to GRE Tunnels

This topic covers the following information:

- [Understanding Scheduling and Shaping of Traffic Routed to GRE Tunnels on page 20](#)
- [Configuration Overview on page 20](#)
- [Configuration Caveats on page 20](#)

### Understanding Scheduling and Shaping of Traffic Routed to GRE Tunnels

On MX Series routers running Junos OS Release 12.3R4 or later revisions, 13.2R2 or later revision, or 13.3R1 or later, you can manage CoS scheduling and shaping of traffic routed to generic route encapsulation (GRE) tunnel interfaces configured on *MPC1 Q*, *MPC2 Q*, or *MPC2 EQ* modules.

A single egress logical interface can be converted to multiple GRE tunnel interfaces. A GRE tunnel physical interface can support many logical interfaces, but one or more of those logical interfaces might not have an output traffic control profiles attached. If a GRE tunnel logical interface is not attached to an output traffic control profile, the router does not assign the interface a dedicated scheduler. Instead, the interface uses a reserved scheduler intended for all *unshaped tunnel traffic* (traffic entering a GRE tunnel logical interface that does not have an explicit traffic control profile configuration).

### Configuration Overview

At GRE tunnel interfaces, the **output-traffic-control-profile** configuration statement can apply an output traffic scheduling and shaping profile at the physical or logical interface level, while the **output-traffic-control-profile-remaining** configuration statement can apply an output traffic scheduling and shaping profile for remaining traffic at the physical interface level only. Interface sets (sets of interfaces used to configure hierarchical CoS schedulers on supported Ethernet interfaces) are not supported on GRE tunnel interfaces.

By default—if you do not attach an output traffic control profile to the GRE tunnel physical interface—traffic entering the interface is scheduled and shaped using the default 95/5 scheduler with parameters as specified in the **tunnel-services** configuration.

If you use an output traffic control profile to configure the shaping rate at the GRE tunnel physical interface, the **shaping-rate** specified by the attached traffic control profile overrides the **bandwidth** specified as the tunnel services default value.

### Configuration Caveats

When configuring hierarchical CoS scheduling and shaping of traffic routed to GRE tunnels, keep the following guidelines in mind:

- You must first configure and commit a hierarchical scheduler on the GRE tunnel physical interface, specifying a maximum of two hierarchical scheduling levels for node scaling. After you commit the **hierarchical-scheduler** configuration, you can configure scheduling and queuing parameters at the GRE tunnel physical or logical interfaces.

- GRE tunnel interfaces support eight egress queues only. For interfaces on MPC1 Q, MPC2 Q, and MPC2 EQ modules, you can include the [max-queues-per-interface 4](#) statement at the `[edit fpc slot-number pic pic-number]` hierarchy level to configure four-queue mode for the interface. However, any GRE tunnel interfaces configured on those ports have eight queues.
- Queuing and scheduling calculations include Layer 3 fields. For GRE interfaces, Layer 3 fields include the delivery header (the outer IP header), the 4-byte GRE header, and the payload protocol header and data.

**Related Documentation**

- [Example: Performing Output Scheduling and Shaping in Hierarchical CoS Queues for Traffic Routed to GRE Tunnels on page 21](#)
- [Per-Unit Queuing and Hierarchical Queuing for MIC and MPC Interfaces on page 43](#)
- [Understanding Hierarchical Scheduling for MIC and MPC Interfaces on page 39](#)

## Example: Performing Output Scheduling and Shaping in Hierarchical CoS Queues for Traffic Routed to GRE Tunnels

---

This example shows how to configure a generic routing encapsulation (GRE) tunnel device to perform CoS output scheduling and shaping of IPv4 traffic routed to GRE tunnels. This feature is supported on MX Series routers running Junos OS Release 12.3R4 or later revisions, 13.2R2 or later revision, or 13.3R1 or later, with GRE tunnel interfaces configured on *MPC1 Q*, *MPC2 Q*, or *MPC2 EQ* modules.

- [Requirements on page 21](#)
- [Overview on page 22](#)
- [Configuration on page 22](#)
- [Verification on page 33](#)

### Requirements

This example uses the following Juniper Networks hardware and Junos OS software:

- Transport network—An IPv4 network running Junos OS Release 13.3.
- GRE tunnel device—One MX80 router installed as an ingress provider edge (PE) router.
- Input and output logical interfaces configurable on two ports of the built-in 10-Gigabit Ethernet Modular Interface Card (MIC):
  - Input logical interface **ge-1/1/0.0** for receiving traffic that is to be transported across the network.
  - Output logical interfaces **ge-1/1/1.0**, **ge-1/1/1.1**, and **ge-1/1/1.2** to convert to GRE tunnel source interfaces **gr-1/1/10.1**, **gr-1/1/10.2**, and **gr-1/1/10.3**.

For information about interfaces hosted on modules in MX80 routers, see the following topics:

- *MX5, MX10, MX40, and MX80 Modular Interface Card Description*

- *MX5, MX10, MX40, and MX80 Port and Interface Numbering*

## Overview

In this example, you configure the router with input and output logical interfaces for IPv4 traffic, and then you convert the output logical interface to four GRE tunnel source interfaces. You also install static routes in the routing table so that input traffic is routed to the four GRE tunnels.



**NOTE:** Before you apply a traffic control profile with a scheduler-map and shaping rate to a GRE tunnel interface, you must configure and commit a hierarchical scheduler on the GRE tunnel physical interface, specifying a maximum of two hierarchical scheduling levels for node scaling.

## Configuration

To configure scheduling and shaping in hierarchical CoS queues for traffic routed to GRE tunnel interfaces configured on MPC1Q, MPC2Q, or MPC2 EQ modules on an MX Series router, perform these tasks:

- [Configuring Interfaces, Hierarchical Scheduling on the GRE Tunnel Physical Interface, and Static Routes on page 24](#)
- [Measuring GRE Tunnel Transmission Rates Without Shaping Applied on page 27](#)
- [Configuring Output Scheduling and Shaping at GRE Tunnel Physical and Logical Interfaces on page 28](#)

### CLI Quick Configuration

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

### Configuring Interfaces, Hierarchical Scheduling on the GRE Tunnel Physical Interface, and Static Routes

```
set chassis fpc 1 pic 1 tunnel-services bandwidth 1g
set interfaces ge-1/1/0 unit 0 family inet address 10.6.6.1/24
set interfaces ge-1/1/1 unit 0 family inet address 10.70.1.1/24 arp 10.70.1.3 mac 00:00:03:00:04:00
set interfaces ge-1/1/1 unit 0 family inet address 10.80.1.1/24 arp 10.80.1.3 mac 00:00:03:00:04:01
set interfaces ge-1/1/1 unit 0 family inet address 10.90.1.1/24 arp 10.90.1.3 mac 00:00:03:00:04:02
set interfaces ge-1/1/1 unit 0 family inet address 10.100.1.1/24 arp 10.100.1.3 mac 00:00:03:00:04:04
set interfaces gr-1/1/10 unit 1 family inet address 10.100.1.1/24
set interfaces gr-1/1/10 unit 1 tunnel source 10.70.1.1 destination 10.70.1.3
set interfaces gr-1/1/10 unit 2 family inet address 10.200.1.1/24
set interfaces gr-1/1/10 unit 2 tunnel source 10.80.1.1 destination 10.80.1.3
set interfaces gr-1/1/10 unit 3 family inet address 10.201.1.1/24
set interfaces gr-1/1/10 unit 3 tunnel source 10.90.1.1 destination 10.90.1.3
set interfaces gr-1/1/10 unit 4 family inet address 10.202.1.1/24
set interfaces gr-1/1/10 unit 4 tunnel source 10.100.1.1 destination 10.100.1.3
```

```

set interfaces gr-1/1/10 hierarchical-scheduler
set routing-options static route 10.2.2.0/24 next-hop gr-1/1/10.1
set routing-options static route 10.3.3.0/24 next-hop gr-1/1/10.2
set routing-options static route 10.4.4.0/24 next-hop gr-1/1/10.3
set routing-options static route 10.5.5.0/24 next-hop gr-1/1/10.4

```

### Configuring Output Scheduling and Shaping at GRE Tunnel Physical and Logical Interfaces

```

set class-of-service forwarding-classes queue 0 be
set class-of-service forwarding-classes queue 1 ef
set class-of-service forwarding-classes queue 2 af
set class-of-service forwarding-classes queue 3 nc
set class-of-service forwarding-classes queue 4 be1
set class-of-service forwarding-classes queue 5 ef1
set class-of-service forwarding-classes queue 6 af1
set class-of-service forwarding-classes queue 7 nc1
set class-of-service classifiers inet-precedence gr-inet forwarding-class be
loss-priority low code-points 000
set class-of-service classifiers inet-precedence gr-inet forwarding-class ef
loss-priority low code-points 001
set class-of-service classifiers inet-precedence gr-inet forwarding-class af
loss-priority low code-points 010
set class-of-service classifiers inet-precedence gr-inet forwarding-class nc
loss-priority low code-points 011
set class-of-service classifiers inet-precedence gr-inet forwarding-class be1
loss-priority low code-points 100
set class-of-service classifiers inet-precedence gr-inet forwarding-class ef1
loss-priority low code-points 101
set class-of-service classifiers inet-precedence gr-inet forwarding-class af1
loss-priority low code-points 110
set class-of-service classifiers inet-precedence gr-inet forwarding-class nc1
loss-priority low code-points 111
set class-of-service interfaces ge-1/1/0 unit 0 classifiers inet-precedence gr-inet
set class-of-service schedulers be_sch transmit-rate percent 30
set class-of-service schedulers ef_sch transmit-rate percent 40
set class-of-service schedulers af_sch transmit-rate percent 25
set class-of-service schedulers nc_sch transmit-rate percent 5
set class-of-service schedulers be1_sch transmit-rate percent 60
set class-of-service schedulers be1_sch priority low
set class-of-service schedulers ef1_sch transmit-rate percent 40
set class-of-service schedulers ef1_sch priority medium-low
set class-of-service schedulers af1_sch transmit-rate percent 10
set class-of-service schedulers af1_sch priority strict-high
set class-of-service schedulers nc1_sch shaping-rate percent 10
set class-of-service schedulers nc1_sch priority high
set class-of-service scheduler-maps sch_map_1 forwarding-class be scheduler be_sch
set class-of-service scheduler-maps sch_map_1 forwarding-class ef scheduler ef_sch
set class-of-service scheduler-maps sch_map_1 forwarding-class af scheduler af_sch
set class-of-service scheduler-maps sch_map_1 forwarding-class nc scheduler nc_sch
set class-of-service scheduler-maps sch_map_2 forwarding-class be scheduler be1_sch
set class-of-service scheduler-maps sch_map_2 forwarding-class ef scheduler ef1_sch
set class-of-service scheduler-maps sch_map_3 forwarding-class af scheduler af_sch
set class-of-service scheduler-maps sch_map_3 forwarding-class nc scheduler nc_sch
set class-of-service traffic-control-profiles gr-ift-tcp3 guaranteed-rate 5m
set class-of-service traffic-control-profiles gr-ift-tcp shaping-rate 10m
set class-of-service traffic-control-profiles gr-ift-tcp-remain shaping-rate 7m
set class-of-service traffic-control-profiles gr-ift-tcp-remain guaranteed-rate 4m
set class-of-service traffic-control-profiles gr-ift-tcp1 scheduler-map sch_map_1
set class-of-service traffic-control-profiles gr-ift-tcp1 shaping-rate 8m
set class-of-service traffic-control-profiles gr-ift-tcp1 guaranteed-rate 3m

```

```
set class-of-service traffic-control-profiles gr-ifl-tcp2 scheduler-map sch_map_2
set class-of-service traffic-control-profiles gr-ifl-tcp2 guaranteed-rate 2m
set class-of-service traffic-control-profiles gr-ifl-tcp3 scheduler-map sch_map_3
set class-of-service interfaces gr-1/1/10 output-traffic-control-profile gr-ifd-tcp
set class-of-service interfaces gr-1/1/10 output-traffic-control-profile-remaining
gr-ifd-remain
set class-of-service interfaces gr-1/1/10 unit 1 output-traffic-control-profile
gr-ifl-tcp1
set class-of-service interfaces gr-1/1/10 unit 2 output-traffic-control-profile
gr-ifl-tcp2
set class-of-service interfaces gr-1/1/10 unit 3 output-traffic-control-profile
gr-ifl-tcp3
```

### Configuring Interfaces, Hierarchical Scheduling on the GRE Tunnel Physical Interface, and Static Routes

---

#### Step-by-Step Procedure

To configure GRE tunnel interfaces (including enabling hierarchical scheduling) and static routes:

1. Configure the amount of bandwidth for tunnel services on the physical interface.

[edit]

```
user@host# set chassis fpc 1 pic 1 tunnel-services bandwidth 1g
```

2. Configure the GRE tunnel device output logical interface.

[edit]

```
user@host# set interfaces ge-1/1/0 unit 0 family inet address 10.6.6.1/24
```

3. Configure the GRE tunnel device output logical interface.

[edit]

```
user@host# set interfaces ge-1/1/1 unit 0 family inet address 10.70.1.1/24 arp 10.70.1.3
mac 00:00:03:00:04:00
```

```
user@host# set interfaces ge-1/1/1 unit 0 family inet address 10.80.1.1/24 arp 10.80.1.3
mac 00:00:03:00:04:01
```

```
user@host# set interfaces ge-1/1/1 unit 0 family inet address 10.90.1.1/24 arp 10.90.1.3
mac 00:00:03:00:04:02
```

```
user@host# set interfaces ge-1/1/1 unit 0 family inet address 10.100.1.1/24 arp
10.100.1.3 mac 00:00:03:00:04:04
```

4. Convert the output logical interface to four GRE tunnel interfaces.

[edit]

```
user@host# set interfaces gr-1/1/10 unit 1 family inet address 10.100.1.1/24
```

```
user@host# set interfaces gr-1/1/10 unit 1 tunnel source 10.70.1.1 destination 10.70.1.3
```

```
user@host# set interfaces gr-1/1/10 unit 2 family inet address 10.200.1.1/24
```

```
user@host# set interfaces gr-1/1/10 unit 2 tunnel source 10.80.1.1 destination 10.80.1.3
```

```
user@host# set interfaces gr-1/1/10 unit 3 family inet address 10.201.1.1/24
```

```
user@host# set interfaces gr-1/1/10 unit 3 tunnel source 10.90.1.1 destination 10.90.1.3
```

```
user@host# set interfaces gr-1/1/10 unit 4 family inet address 10.202.1.1/24
```

```
user@host# set interfaces gr-1/1/10 unit 4 tunnel source 10.100.1.1 destination
10.100.1.3
```

5. Enable the GRE tunnel interfaces to use hierarchical scheduling.

```
[edit]
user@host# set interfaces gr-1/1/10 hierarchical-scheduler
```

6. Install static routes in the routing table so that the device routes IPv4 traffic to the GRE tunnel source interfaces.

Traffic destined to the subnets 10.2.2.0/24, 10.3.3.0/24, 10.4.4.0/24, and 10.5.5.0/24 is routed to the tunnel interfaces at IP addresses 10.70.1.1, 10.80.1.1, 10.90.1.1, and 10.100.1.1, respectively.

```
[edit]
user@host# set routing-options static route 10.2.2.0/24 next-hop gr-1/1/10.1
user@host# set routing-options static route 10.3.3.0/24 next-hop gr-1/1/10.2
user@host# set routing-options static route 10.4.4.0/24 next-hop gr-1/1/10.3
user@host# set routing-options static route 10.5.5.0/24 next-hop gr-1/1/10.4
```

7. If you are done configuring the device, commit the configuration.

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

**Results** From configuration mode, confirm your configuration by entering the **show chassis fpc 1 pic 1**, **show interfaces ge-1/1/0**, **show interfaces ge-1/1/1**, **show interfaces gr-1/1/10**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Confirm the configuration of interfaces, hierarchical scheduling on the GRE tunnel physical interface, and static routes.

```
user@host# show chassis fpc 1 pic 1
tunnel-services {
  bandwidth 1g;
}

user@host# show interfaces ge-1/1/0
unit 0 {
  family inet {
    address 10.6.6.1/24;
  }
}

user@host# show interfaces ge-1/1/1
unit 0 {
  family inet {
    address 10.70.1.1/24 {
      arp 10.70.1.3 mac 00:00:03:00:04:00;
    }
    address 10.80.1.1/24 {
      arp 10.80.1.3 mac 00:00:03:00:04:01;
    }
    address 10.90.1.1/24 {
```

```
        arp 10.90.1.3 mac 00:00:03:00:04:02;
    }
    address 10.100.1.1/24 {
        arp 10.100.1.3 mac 00:00:03:00:04:04;
    }
}
}
```

user@host# show interfaces gr-1/1/10

```
hierarchical-scheduler;
unit 1 {
    tunnel {
        destination 10.70.1.3;
        source 10.70.1.1;
    }
    family inet {
        address 10.100.1.1/24;
    }
}
unit 2 {
    tunnel {
        destination 10.80.1.3;
        source 10.80.1.1;
    }
    family inet {
        address 10.200.1.1/24;
    }
}
unit 3 {
    tunnel {
        destination 10.90.1.3;
        source 10.90.1.1;
    }
    family inet {
        address 10.201.1.1/24;
    }
}
unit 4 {
    tunnel {
        destination 10.100.1.3;
        source 10.100.1.1;
    }
    family inet {
        address 10.202.1.1/24;
    }
}
```

user@host# show routing-options

```
static {
    route 10.2.2.0/24 next-hop gr-1/1/10.1;
    route 10.3.3.0/24 next-hop gr-1/1/10.2;
    route 10.4.4.0/24 next-hop gr-1/1/10.3;
    route 10.5.5.0/24 next-hop gr-1/1/10.4;
}
```



## Measuring GRE Tunnel Transmission Rates Without Shaping Applied

### Step-by-Step Procedure

To establish a baseline measurement, note the transmission rates at each GRE tunnel source.

1. Pass traffic through the GRE tunnel at logical interfaces **gr-1/1/10.1**, **gr-1/1/10.2**, and **gr-1/1/10.3**.
2. To display the traffic rates at each GRE tunnel source, use the **show interfaces queue** operational mode command.

The following example command output shows detailed CoS queue statistics for logical interface gr-1/1/10.1 (the GRE tunnel from source IP address 10.70.1.1 to destination IP address 10.70.1.3).

```
user@host> show interfaces queue gr-1/1/10.1
Logical interface gr-1/1/10.1 (Index 331) (SNMP ifIndex 4045)
Forwarding classes: 16 supported, 8 in use
Egress queues: 8 supported, 8 in use
Burst size: 0
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :          31818312          102494 pps
    Bytes        :          6522753960        168091936 bps
  Transmitted:
    Packets      :          1515307           4879 pps
    Bytes        :          310637935         8001632 bps
    Tail-dropped packets :          21013826         68228 pps
    RED-dropped packets :          9289179         29387 pps
    Low          :          9289179         29387 pps
    Medium-low   :              0              0 pps
    Medium-high  :              0              0 pps
    High         :              0              0 pps
    RED-dropped bytes :          1904281695        48194816 bps
    Low          :          1904281695        48194816 bps
    Medium-low   :              0              0 bps
    Medium-high  :              0              0 bps
    High         :              0              0 bps
  ...
```



**NOTE:** This step shows command output for queue 0 (forwarding class be) only.

The command output shows that the GRE tunnel device transmits traffic from queue 0 at a rate of 4879 pps. Allowing for 182 bytes per Layer 3 packet, preceded by 24 bytes of GRE overhead (a 20-byte delivery header consisting of the IPv4 packet header followed by 4 bytes for GRE flags plus encapsulated protocol type), the traffic rate received at the tunnel destination device is 8,040,592 bps:

$4879 \text{ packets/second} \times 206 \text{ bytes/packet} \times 8 \text{ bits/byte} = 8,040,592 \text{ bits/second}$

## Configuring Output Scheduling and Shaping at GRE Tunnel Physical and Logical Interfaces

**Step-by-Step Procedure** To configure the GRE tunnel device with scheduling and shaping at GRE tunnel physical and logical interfaces:

1. Define eight transmission queues.

```
[edit]
user@host# set class-of-service forwarding-classes queue 0 be
user@host# set class-of-service forwarding-classes queue 1 ef
user@host# set class-of-service forwarding-classes queue 2 af
user@host# set class-of-service forwarding-classes queue 3 nc
user@host# set class-of-service forwarding-classes queue 4 be1
user@host# set class-of-service forwarding-classes queue 5 ef1
user@host# set class-of-service forwarding-classes queue 6 af1
user@host# set class-of-service forwarding-classes queue 7 nc1
```



**NOTE:** To configure up to eight forwarding classes with one-to-one mapping to output queues for interfaces on M120, M320, MX Series, and T Series routers and EX Series switches, use the `queue` statement at the `[edit class-of-service forwarding-classes]` hierarchy level.

If you need to configure up to 16 forwarding classes with multiple forwarding classes mapped to single queues for those interface types, use the `class` statement instead.

2. Configure BA classifier `gr-inet` that, based on IPv4 precedence bits set in an incoming packet, sets the forwarding class, loss-priority value, and DSCP bits of the packet.

```
[edit]
user@host# set class-of-service classifiers inet-precedence gr-inet forwarding-class
be loss-priority low code-points 000
user@host# set class-of-service classifiers inet-precedence gr-inet forwarding-class
ef loss-priority low code-points 001
user@host# set class-of-service classifiers inet-precedence gr-inet forwarding-class
af loss-priority low code-points 010
user@host# set class-of-service classifiers inet-precedence gr-inet forwarding-class
nc loss-priority low code-points 011
user@host# set class-of-service classifiers inet-precedence gr-inet forwarding-class
be1 loss-priority low code-points 100
user@host# set class-of-service classifiers inet-precedence gr-inet forwarding-class
ef1 loss-priority low code-points 101
user@host# set class-of-service classifiers inet-precedence gr-inet forwarding-class
af1 loss-priority low code-points 110
user@host# set class-of-service classifiers inet-precedence gr-inet forwarding-class
nc1 loss-priority low code-points 111
```

3. Apply BA classifier **gr-inet** to the GRE tunnel device input at logical interface ge-1/1/0.0.

```
[edit]
user@host# set class-of-service interfaces ge-1/1/0 unit 0 classifiers inet-precedence
gr-inet
```

4. Define a scheduler for each forwarding class.

```
[edit]
user@host# set class-of-service schedulers be_sch transmit-rate percent 30
user@host# set class-of-service schedulers ef_sch transmit-rate percent 40
user@host# set class-of-service schedulers af_sch transmit-rate percent 25
user@host# set class-of-service schedulers nc_sch transmit-rate percent 5
user@host# set class-of-service schedulers be1_sch transmit-rate percent 60
user@host# set class-of-service schedulers be1_sch priority low
user@host# set class-of-service schedulers ef1_sch transmit-rate percent 40
user@host# set class-of-service schedulers ef1_sch priority medium-low
user@host# set class-of-service schedulers af1_sch transmit-rate percent 10
user@host# set class-of-service schedulers af1_sch priority strict-high
user@host# set class-of-service schedulers nc1_sch shaping-rate percent 10
user@host# set class-of-service schedulers nc1_sch priority high
```

5. Define a scheduler map for each of three GRE tunnels.

```
[edit]
user@host# set class-of-service scheduler-maps sch_map_1 forwarding-class be
scheduler be_sch
user@host# set class-of-service scheduler-maps sch_map_1 forwarding-class ef
scheduler ef_sch
user@host# set class-of-service scheduler-maps sch_map_1 forwarding-class af
scheduler af_sch
user@host# set class-of-service scheduler-maps sch_map_1 forwarding-class nc
scheduler nc_sch
user@host# set class-of-service scheduler-maps sch_map_2 forwarding-class be
scheduler be1_sch
user@host# set class-of-service scheduler-maps sch_map_2 forwarding-class ef
scheduler ef1_sch
user@host# set class-of-service scheduler-maps sch_map_3 forwarding-class af
scheduler af_sch
user@host# set class-of-service scheduler-maps sch_map_3 forwarding-class nc
scheduler nc_sch
```

6. Define traffic control profiles for three GRE tunnel interfaces.

```
[edit]
user@host# set class-of-service traffic-control-profiles gr-iftcp1 scheduler-map
sch_map_1
user@host# set class-of-service traffic-control-profiles gr-iftcp1 shaping-rate 8m
user@host# set class-of-service traffic-control-profiles gr-iftcp1 guaranteed-rate
3m
user@host# set class-of-service traffic-control-profiles gr-iftcp2 scheduler-map
sch_map_2
```

```
user@host# set class-of-service traffic-control-profiles gr-ift-tcp2 guaranteed-rate
2m
user@host# set class-of-service traffic-control-profiles gr-ift-tcp3 scheduler-map
sch_map_3
user@host# set class-of-service traffic-control-profiles gr-ift-tcp3 guaranteed-rate
5m
user@host# set class-of-service traffic-control-profiles gr-ift-tcp shaping-rate 10m
user@host# set class-of-service traffic-control-profiles gr-ift-tcp-remain
shaping-rate 7m
user@host# set class-of-service traffic-control-profiles gr-ift-tcp-remain
guaranteed-rate 4m
```

7. Apply CoS scheduling and shaping to the output traffic at the physical interface and logical interfaces.

```
[edit]
user@host# set class-of-service interfaces gr-1/1/10 output-traffic-control-profile
gr-ift-tcp
user@host# set class-of-service interfaces gr-1/1/10
output-traffic-control-profile-remaining gr-ift-remain
user@host# set class-of-service interfaces gr-1/1/10 unit 1
output-traffic-control-profile gr-ift-tcp1
user@host# set class-of-service interfaces gr-1/1/10 unit 2
output-traffic-control-profile gr-ift-tcp2
user@host# set class-of-service interfaces gr-1/1/10 unit 2
output-traffic-control-profile gr-ift-tcp3
```

8. If you are done configuring the device, commit the configuration.

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

**Results** From configuration mode, confirm your configuration by entering the **show class-of-service forwarding-classes**, **show class-of-service classifiers**, **show class-of-service interfaces ge-1/1/0**, **show class-of-service schedulers**, **show class-of-service scheduler-maps**, **show class-of-service traffic-control-profiles**, and **show class-of-service interfaces gr-1/1/10** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Confirm the configuration of output scheduling and shaping at the GRE tunnel physical and logical interfaces.

```
user@host# show class-of-service forwarding-classes
queue 0 be;
queue 1 ef;
queue 2 af;
queue 3 nc;
queue 4 be1;
queue 5 ef1;
queue 6 af1;
queue 7 nc1;

user@host# show class-of-service classifiers
```

```
inet-precedence gr-inet {
  forwarding-class be {
    loss-priority low code-points 000;
  }
  forwarding-class ef {
    loss-priority low code-points 001;
  }
  forwarding-class af {
    loss-priority low code-points 010;
  }
  forwarding-class nc {
    loss-priority low code-points 011;
  }
  forwarding-class be1 {
    loss-priority low code-points 100;
  }
  forwarding-class ef1 {
    loss-priority low code-points 101;
  }
  forwarding-class af1 {
    loss-priority low code-points 110;
  }
  forwarding-class nc1 {
    loss-priority low code-points 111;
  }
}

user@host# show class-of-service interfaces ge-1/1/0
unit 0 {
  classifiers {
    inet-precedence gr-inet;
  }
}

user@host# show class-of-service schedulers
be_sch {
  transmit-rate percent 30;
}
ef_sch {
  transmit-rate percent 40;
}
af_sch {
  transmit-rate percent 25;
}
nc_sch {
  transmit-rate percent 5;
}
be1_sch {
  transmit-rate percent 60;
  priority low;
}
ef1_sch {
  transmit-rate percent 40;
  priority medium-low;
}
af1_sch {
```

```
    transmit-rate percent 10;
    priority strict-high;
}
nc1_sch {
    shaping-rate percent 10;
    priority high;
}

user@host# show class-of-service scheduler-maps
sch_map_1 {
    forwarding-class be scheduler be_sch;
    forwarding-class ef scheduler ef_sch;
    forwarding-class af scheduler af_sch;
    forwarding-class nc scheduler nc_sch;
}
sch_map_2 {
    forwarding-class be scheduler be1_sch;
    forwarding-class ef scheduler ef1_sch;
}
sch_map_3 {
    forwarding-class af scheduler af_sch;
    forwarding-class nc scheduler nc_sch;
}

user@host# show class-of-service traffic-control-profiles
gr-ifl-tcp1 {
    scheduler-map sch_map_1;
    shaping-rate 8m;
    guaranteed-rate 3m;
}
gr-ifl-tcp2 {
    scheduler-map sch_map_2;
    guaranteed-rate 2m;
}
gr-ifl-tcp3 {
    scheduler-map sch_map_3;
    guaranteed-rate 5m;
}
gr-ifd-remain {
    shaping-rate 7m;
    guaranteed-rate 4m;
}
gr-ifd-tcp {
    shaping-rate 10m;
}

user@host# show class-of-service interfaces gr-1/1/10
gr-1/1/10 {
    output-traffic-control-profile gr-ifd-tcp;
    output-traffic-control-profile-remaining gr-ifd-remain;
    unit 1 {
        output-traffic-control-profile gr-ifl-tcp1;
    }
    unit 2 {
        output-traffic-control-profile gr-ifl-tcp2;
    }
}
```

```

unit 3 {
    output-traffic-control-profile gr-ift-tcp3;
}
}

```

## Verification

Confirm that the configurations are working properly.

- [Verifying That Scheduling and Shaping Are Attached to the GRE Tunnel Interfaces on page 33](#)
- [Verifying That Scheduling and Shaping Are Functioning at the GRE Tunnel Interfaces on page 34](#)

### Verifying That Scheduling and Shaping Are Attached to the GRE Tunnel Interfaces

**Purpose** Verify the association of traffic control profiles with GRE tunnel interfaces.

**Action** Verify the traffic control profile attached to the GRE tunnel physical interface by using the `show class-of-service interface gr-1/1/10 detail` operational mode command.

```

user@host> show class-of-service interface gr-1/1/10 detail
Physical interface: gr-1/1/10, Enabled, Physical link is Up
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

Physical interface: gr-1/1/10, Index: 220
Queues supported: 8, Queues in use: 8
  Output traffic control profile: gr-ift-tcp, Index: 17721
  Output traffic control profile remaining: gr-ift-remain, Index: 58414
  Congestion-notification: Disabled

Logical interface gr-1/1/10.1
  Flags: Point-To-Point SNMP-Traps 0x4000 IP-Header
10.70.1.3:10.70.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
  Gre keepalives configured: Off, Gre keepalives adjacency state: down
  inet 10.100.1.1/24
Logical interface: gr-1/1/10.1, Index: 331
Object      Name                      Type      Index
Traffic-control-profile gr-ift-tcp1              Output    17849
Classifier   ipprec-compatibility    ip        13

Logical interface gr-1/1/10.2
  Flags: Point-To-Point SNMP-Traps 0x4000 IP-Header
10.80.1.3:10.80.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
  Gre keepalives configured: Off, Gre keepalives adjacency state: down
  inet 10.200.1.1/24
Logical interface: gr-1/1/10.2, Index: 332
Object      Name                      Type      Index
Traffic-control-profile gr-ift-tcp2              Output    17856
Classifier   ipprec-compatibility    ip        13

Logical interface gr-1/1/10.3
  Flags: Point-To-Point SNMP-Traps 0x4000 IP-Header
10.90.1.3:10.90.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
  Gre keepalives configured: Off, Gre keepalives adjacency state: down

```

```

      inet 10.201.1.1/24
Logical interface: gr-1/1/10.3, Index: 333
Object      Name      Type      Index
Traffic-control-profile gr-ifl-tcp3      Output    17863
Classifier   ipprec-compatibility ip         13

```

**Meaning** Ingress IPv4 traffic routed to GRE tunnels on the device is subject to CoS output scheduling and shaping.

### Verifying That Scheduling and Shaping Are Functioning at the GRE Tunnel Interfaces

**Purpose** Verify the traffic rate shaping at the GRE tunnel interfaces.

- Action**
1. Pass traffic through the GRE tunnel at logical interfaces **gr-1/1/10.1**, **gr-1/1/10.2**, and **gr-1/1/10.3**.
  2. To verify the rate shaping at each GRE tunnel source, use the [show interfaces queue](#) operational mode command.

The following example command output shows detailed CoS queue statistics for logical interface gr-1/1/10.1 (the GRE tunnel from source IP address 10.70.1.1 to destination IP address 10.70.1.3):

```

user@host> show interfaces queue gr-1/1/10.1
Logical interface gr-1/1/10.1 (Index 331) (SNMP ifIndex 4045)
Forwarding classes: 16 supported, 8 in use
Egress queues: 8 supported, 8 in use
Burst size: 0
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :          59613061          51294 pps
    Bytes        :      12220677505      84125792 bps
  Transmitted:
    Packets      :          2230632           3039 pps
    Bytes        :      457279560      4985440 bps
    Tail-dropped packets :          4471146           2202 pps
    RED-dropped packets :          52911283          46053 pps
      Low        :          49602496          46053 pps
      Medium-low :              0              0 pps
      Medium-high:              0              0 pps
      High       :          3308787              0 pps
    RED-dropped bytes :      10846813015      75528000 bps
      Low        :      10168511680      75528000 bps
      Medium-low :              0              0 bps
      Medium-high :              0              0 bps
      High       :          678301335              0 bps
Queue: 1, Forwarding classes: ef
  Queued:
    Packets      :          15344874          51295 pps
    Bytes        :      3145699170      84125760 bps
  Transmitted:
    Packets      :           366115           1218 pps
    Bytes        :       75053575      1997792 bps

```



```

Tail-dropped packets :          364489          1132 pps
RED-dropped packets :          14614270         48945 pps
  Low                  :          14614270         48945 pps
  Medium-low           :              0           0 pps
  Medium-high          :              0           0 pps
  High                 :              0           0 pps
RED-dropped bytes    :          2995925350       80270528 bps
  Low                  :          2995925350       80270528 bps
  Medium-low           :              0           0 bps
  Medium-high          :              0           0 bps
  High                 :              0           0 bps
...

```



**NOTE:** This step shows command output for queue 0 (forwarding class be) and queue 1 (forwarding class ef) only.

**Meaning** Now that traffic shaping is attached to the GRE tunnel interfaces, the command output shows that traffic shaping specified for the tunnel at logical interface gr-1/1/10.1 (**shaping-rate 8m** and **guaranteed-rate 3m**) is honored.

- For queue 0, the GRE tunnel device transmits traffic at a rate of 3039 pps. The traffic rate received at the tunnel destination device is 5,008,272 bps:

$3039 \text{ packets/second} \times 206 \text{ bytes/packet} \times 8 \text{ bits/byte} = 5,008,272 \text{ bits/second}$

- For queue 0, the GRE tunnel device transmits traffic at a rate of 1218 pps. The traffic rate received at the tunnel destination device is 2,007,264 bps:

$1218 \text{ packets/second} \times 206 \text{ bytes/packet} \times 8 \text{ bits/byte} = 2,007,264 \text{ bits/second}$

Compare these statistics to the baseline measurements taken without traffic shaping, as described in [“Measuring GRE Tunnel Transmission Rates Without Shaping Applied” on page 27](#).

**Related Documentation**

- [Scheduling and Shaping in Hierarchical CoS Queues for Traffic Routed to GRE Tunnels on page 20](#)
- [Configuring Traffic Control Profiles for Shared Scheduling and Shaping](#)

## Configuring Ingress Hierarchical CoS on Enhanced Queuing DPCs

You can configure ingress CoS parameters, including hierarchical schedulers, on MX Series routers with Enhanced Queuing DPCs (that is, line cards that have a queuing chip). In general, the supported configuration statements apply to per-unit schedulers or to hierarchical schedulers.



**NOTE:** Ingress CoS is not supported on line cards that do not contain a queuing chip.

To configure ingress CoS for per-unit schedulers, include the following statements at the `[edit class-of-service interfaces interface-name]` hierarchy level:



**NOTE:** The `input-scheduler-map` and `input-traffic-control-profile` statements are mutually exclusive at the same hierarchy level.

```
[edit class-of-service interfaces interface-name]
input-excess-bandwidth-share (proportional value | equal);
input-scheduler-map map-name;
input-shaping-rate rate;
input-traffic-control-profile profile-name shared-instance instance-name;
unit logical-unit-number;
    input-scheduler-map map-name;
    input-shaping-rate (percent percentage | rate);
    input-traffic-control-profile profile-name shared-instance instance-name;
}
```

To configure ingress CoS for hierarchical schedulers, include the `interface-set` `interface-set-name` statement at the `[edit class-of-service interfaces]` hierarchy level:

```
[edit class-of-service interfaces]
interface-set interface-set-name {
    input-excess-bandwidth-share (proportional value | equal);
    input-traffic-control-profile profile-name shared-instance instance-name;
    input-traffic-control-profile-remaining profile-name;
    interface interface-name {
        input-excess-bandwidth-share (proportional value | equal);
        input-traffic-control-profile profile-name shared-instance instance-name;
        input-traffic-control-profile-remaining profile-name;
        unit logical-unit-number;
        input-traffic-control-profile profile-name shared-instance instance-name;
    }
}
}
```

By default, ingress CoS features are disabled on the Enhanced Queuing DPC.

For an Enhanced Queuing (EQ) DPC on an MX Series router, CoS queuing and scheduling are enabled on the egress side but disabled on the ingress side by default. To enable ingress CoS on the EQ DPC, you must configure the `traffic-manager` statement with `ingress-and-egress` mode:

```
[edit chassis fpc slot-number pic pic-number]
traffic-manager mode ingress-and-egress;
```



**NOTE:** If you enable ingress CoS settings and inline services on the same FPC, the FPC moves to the offline state. This behavior is expected because traffic black hole detection is triggered that causes the FPC to move to the offline state.

Configured CoS features on the ingress are independent of CoS features on the egress, with the following exceptions:

- If you configure a per-unit or hierarchical scheduler at the **[edit class-of-service interfaces]** hierarchy level, the schedulers apply in both the ingress and egress directions.
- You cannot configure the same logical interface on an ingress and an egress interface set. A logical interface can only belong to one interface set.
- The DPC's frame buffer of 512 MB is shared between ingress and egress configurations.

The following behavior aggregate (BA) classification tables are supported on the ingress side of the Enhanced Queuing DPC:

- inet-precedence
- DSCP
- exp (MPLS)
- DSCP for IPv6
- IEEE 802.1p

**Related  
Documentation**

- *Configuring Traffic Control Profiles for Shared Scheduling and Shaping*
- *Enhanced Queuing DPC CoS Properties*



## CHAPTER 2

# Configuring Hierarchical Class of Service on MICs, MPCs, MLCs, and Aggregated Ethernet Interfaces

- [Understanding Hierarchical Scheduling for MIC and MPC Interfaces on page 39](#)
- [Configuring Ingress Hierarchical CoS on MIC and MPC Interfaces on page 41](#)
- [Per-Unit Queuing and Hierarchical Queuing for MIC and MPC Interfaces on page 43](#)
- [Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview on page 46](#)
- [Jitter Reduction in Hierarchical CoS Queues on page 48](#)
- [Example: Reducing Jitter in Hierarchical CoS Queues on page 50](#)
- [Hierarchical Schedulers on Aggregated Ethernet Interfaces Overview on page 56](#)
- [Configuring Hierarchical Schedulers on Aggregated Ethernet Interfaces on page 57](#)
- [Example: Configuring Scheduling Modes on Aggregated Interfaces on page 58](#)
- [Increasing Available Bandwidth on Rich-Queuing MPCs by Bypassing the Queuing Chip on page 63](#)

## Understanding Hierarchical Scheduling for MIC and MPC Interfaces

This topic covers the following information:

- [Scheduler Node Scaling for MIC and MPC Interfaces on page 39](#)
- [Hierarchical Scheduling Priority Levels for MIC and MPC Interfaces on page 40](#)
- [Guaranteed Bandwidth and Weight of an Interface Node on MIC and MPC Interfaces on page 40](#)
- [Hierarchical Scheduling for MIC and MPC Interfaces in Oversubscribed PIR Mode on page 41](#)

### Scheduler Node Scaling for MIC and MPC Interfaces

In per-unit scheduling, the logical interfaces share a common level 2 node (one per port). In hierarchical-scheduling, each logical interface has its own level 2 node. Thus, scaling is limited by the number of level 2 nodes.

To better control system resources in hierarchical-scheduling mode, you can limit the number of scheduler node levels to two. In this case, all logical interfaces and interface sets with CoS scheduling policy share a single level 2 node. Consequently, the maximum number of logical interfaces with CoS scheduling policies is increased (the interface sets must be at level 3).

To configure scheduler node scaling, include the **hierarchical-scheduler** statement and set the **maximum-hierarchy-levels** option to 2 at the **[edit interfaces xe-fpc/pic/port]** hierarchy level.

```
[edit interfaces]
xe-2/0/0 {
  hierarchical-scheduler {
    maximum-hierarchy-levels 2;
  }
}
```



**NOTE:** The **maximum-hierarchy-levels** option supports level 3 interface sets but not level 2 interface sets. If you configure level 2 interface sets with the **maximum-hierarchy-levels** option, you generate Packet Forwarding Engine errors.

---

## Hierarchical Scheduling Priority Levels for MIC and MPC Interfaces

The queuing model used by MIC and MPC interfaces supports three priority levels for guaranteed scheduling priority and two lower priority levels for excess scheduling priority. You can configure a queue with one guaranteed priority and one excess priority. For example, you can configure a queue for guaranteed low (GL) as the guaranteed priority and configure excess high (EH) as the excess priority.

You can associate a guaranteed level with only one excess level. You can associate an excess level with any number of guaranteed priority levels, including none.

Interface nodes maintain their guaranteed priority level (for example, guaranteed high, GH) as long as they do not exceed their guaranteed bandwidth. If the queue bandwidth exceeds the guaranteed rate, then the priority drops to the excess priority (for example, excess high, EH). Because excess level priorities are lower than their guaranteed counterparts, the bandwidth guarantees for each of the other levels can be maintained.

## Guaranteed Bandwidth and Weight of an Interface Node on MIC and MPC Interfaces

The queuing model used by MIC and MPC interfaces separates the concepts of *guaranteed bandwidth* and *weight* of an interface node, although the two terms are often used interchangeably. The guaranteed bandwidth for an interface node is the bandwidth the node can use, independent of what is happening at the other nodes of the scheduling hierarchy. The weight of an interface node, on the other hand, is a value that determines how *excess bandwidth* is used. The weight of a node comes into play when other nodes at the same hierarchical scheduling level use less than the sum of their guaranteed bandwidths

For some application traffic types (such as constant bit rate voice, where there is little concern about excess bandwidth), the guaranteed bandwidth dominates the node. For other types of application traffic (such as bursty data, where a well-defined bandwidth is not always possible), the concept of weight dominates the node.

### Hierarchical Scheduling for MIC and MPC Interfaces in Oversubscribed PIR Mode

In contrast to the Intelligent Queuing Enhanced (IQE) and Intelligent Queuing 2 Enhanced (IQ2E) PICs, the interfaces on MICs and MPCs set the guaranteed rate to zero in oversubscribed peak information rate (PIR) mode for the per-unit scheduler. Also, the configured rate is scaled down to fit the oversubscribed value. For example, if there are two logical interface units with a shaping rate of 1 Gbps each on a 1-Gbps port (which is, therefore, oversubscribed 2 to 1), then the guaranteed rate on each unit is scaled down to 500 Mbps (scaled down by 2).

With hierarchical schedulers in oversubscribed PIR mode, the guaranteed rate for every logical interface unit is set to zero. This means that the queue transmit rates are always oversubscribed.

Because in oversubscribed PIR mode the queue transmit rates are always oversubscribed, the following are true:

- If the queue transmit rate is set as a percentage, then the guaranteed rate of the queue is set to zero; but the excess rate (weight) of the queue is set correctly.
- If the queue transmit rate is set as an absolute value and if the queue has guaranteed high or medium priority, then traffic up to the queue's transmit rate is sent at that priority level. However, for guaranteed low traffic, that traffic is demoted to the excess low region. This means that best-effort traffic well within the queue's transmit rate gets a lower priority than out-of-profile excess high traffic. This differs from the IQE and IQ2E PICs.

#### Related Documentation

- [Per-Unit Queuing and Hierarchical Queuing for MIC and MPC Interfaces on page 43](#)
- [CoS Features and Limitations on MIC and MPC Interfaces](#)
- [Jitter Reduction in Hierarchical CoS Queues on page 48](#)
- [Scheduling and Shaping in Hierarchical CoS Queues for Traffic Routed to GRE Tunnels on page 20](#)
- [CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 123](#)

---

### Configuring Ingress Hierarchical CoS on MIC and MPC Interfaces

You can configure ingress CoS parameters, including hierarchical schedulers, on MIC and MPC interfaces on MX Series routers. In general, the supported configuration statements apply to per-unit schedulers or to hierarchical schedulers.



**NOTE:** Ingress CoS is not supported on AE interfaces on MPCs.

To configure ingress CoS for per-unit schedulers, include the following statements at the `[edit class-of-service interfaces interface-name]` hierarchy level:

```
[edit class-of-service interfaces interface-name]
input-excess-bandwidth-share (proportional value | equal);
input-scheduler-map map-name;
input-shaping-rate rate;
input-traffic-control-profile profile-name;
unit logical-unit-number;
  input-scheduler-map map-name;
  input-shaping-rate (percent percentage | rate);
  input-traffic-control-profile profile-name;
```

To configure ingress CoS for hierarchical schedulers, include the `interface-set` `interface-set-name` statement at the `[edit class-of-service interfaces]` hierarchy level:

```
[edit class-of-service interfaces]
interface-set interface-set-name {
  input-traffic-control-profile profile-name;
  input-traffic-control-profile-remaining profile-name;
  interface interface-name {
    input-excess-bandwidth-share (proportional value | equal);
    input-traffic-control-profile profile-name;
    input-traffic-control-profile-remaining profile-name;
    unit logical-unit-number;
  }
}
```

By default, ingress CoS features are disabled on MIC and MPC interfaces. To enable ingress CoS on a MIC or MPC interface, configure the `traffic-manager` statement with `ingress-and-egress` mode as shown in the following example:

```
chassis {
  fpc 7 {
    pic 0 {
      traffic-manager {
        mode ingress-and-egress;
      }
    }
  }
}
```

Configured CoS features on the ingress are independent of CoS features on the egress.



**NOTE:** Prior to Junos OS 16.1R1, for MIC-based MX80 and MX104 routers, only ten queues on one MIC can be configured for ingress CoS. Starting with Junos OS 16.1R1, MX80 and MX104 routers support up to 12 ingress queues on any combination of both MIC and built-in ports.



The following behavior aggregate (BA) classification tables are supported on the ingress side of MIC and MPC interfaces:

- DSCP
- DSCP for IPv6
- exp (MPLS)
- IEEE 802.1p
- inet-precedence

---

## Per-Unit Queuing and Hierarchical Queuing for MIC and MPC Interfaces

This topic covers the following information:

- [Queuing Models Supported for MIC and MPC Interfaces on page 43](#)
- [Scheduler Node Levels for MIC and MPC Interfaces on page 44](#)

### Queuing Models Supported for MIC and MPC Interfaces

Interfaces hosted on Modular Interface Card (MIC) and Modular Port Concentrator (MPC) line cards in MX Series 3D Universal Edge Routers support the following models of class-of-service (CoS) queuing, depending on MIC or MPC type:

- [Limited Scale Per-Unit Queuing MPCs on page 43](#)
- [Hierarchical Queuing MICs and MPCs on page 44](#)

---

#### Limited Scale Per-Unit Queuing MPCs

Per-unit CoS queuing features on a limited scale are supported for interfaces hosted on some MPCs that do not have a dedicated queuing chip, specifically the MPC3E, MPC4E, and MPC6E line cards and on the fixed-configuration 16-port 10-Gigabit Ethernet MPC in MX240, MX480, MX960, MX2010, and MX2020 routers.



**NOTE:** The nonqueuing MPC1, MPC2, and MPC5E line cards *do not* support per-unit queuing.

On MPCs that support per-unit queuing, the following queuing capabilities are available:

- Four or eight egress queues per unit.
- Delay buffer capacities of 100 ms by default, and up to 200 ms maximum delay.
- Rate shaping of the ports and their queues.
- Guaranteed rate enforced at the queues.

The per-unit CoS queuing features also support pre-classification of incoming packets to protect high priority packets in the event of congestion. Such features include ingress DSCP rewrite and per-VLAN classification, ingress and egress policing, and rewrites.

### **Hierarchical Queuing MICs and MPCs**

---

Hierarchical CoS queuing features are supported on interfaces hosted on MICs in MPC1 Q, MPC2 Q, MPC2 EQ, MPC5EQ, MPC7E, MPC8E, and MPC9E line cards in MX240, MX480, MX960, MX2010, and MX2020 routers and for interfaces hosted on 1-Gigabit and 10-Gigabit Ethernet MICs in MX5, MX10, MX40, MX80, or MX104 modular chassis routers. These MICs and MPCs provide a dedicated queuing chip that supports hierarchical queuing.

Hierarchical queuing MICs and MPCs support all per-unit queuing functionality plus fine-grained queuing abilities over four or five levels of hierarchical scheduling:

- Hierarchical scheduling with ports, interface sets, and logical interfaces.
- Shaping—Committed Information Rate (CIR) and a peak information rate (PIR)—at all scheduling levels, including queues.
- Three normal- priority levels and two excess- priority levels configurable at all scheduling levels, including queues.
- Per-priority shaping of traffic at Level 1 or Level 2.
- Shaping for unconfigured customer VLANs (C-VLANs) and for service VLANs (S-VLANs).

### **Scheduler Node Levels for MIC and MPC Interfaces**

Interfaces hosted on MICs and MPCs support different scheduler node levels, depending on MIC or MPC type:

- [Scheduler Node Levels for Per-Unit Queuing MPCs on page 44](#)
- [Scheduler Node Levels for Hierarchical Queuing MICs and MPCs on page 45](#)


#### **Scheduler Node Levels for Per-Unit Queuing MPCs**

---

For an interface hosted on a per-unit queuing MPC, each logical interface has its own dedicated level 3 node, and all logical interfaces share a common level 2 node (one per port).

[Figure 4 on page 45](#) illustrates scheduler node levels for an interface hosted on a per-unit queuing MPC.

Figure 4: Scheduler Node Levels for Per-Unit Queuing MPCs




For interfaces hosted on per-unit queuing MPCs, the level 2 node is always a dummy node.

Scheduler Node Levels for Hierarchical Queuing MICs and MPCs

With the exception of the 10-Gigabit Ethernet MPC with SFP+, the queuing model used by interfaces hosted on hierarchical queuing MICs and MPCs supports up to five levels of scheduler nodes: the queue itself (level 5), session logical interface (ppp or dhcp) (level 4), customer VLAN (C-VLAN) (level 3), the interface set or service VLAN (S-VLAN) collection (level 2), and the physical interface or port (level 1).

Figure 5 on page 45 illustrates the scheduler node levels for an interface hosted on a hierarchical queuing MIC or MPC.

Figure 5: Scheduler Node Levels for Interfaces on Hierarchical Queuing and Scheduling MICs and MPCs



The figure depicts scheduler nodes for an interface that does not include interface sets and for which traffic control profiles are applied to the logical interfaces only.



**NOTE:** If an interface set has a CoS scheduling policy but none of its child logical interfaces has a CoS scheduling policy, then the interface set is considered to be a leaf node and has one level 2 and one level 3 node.

- Related Documentation**
- [Understanding Hierarchical Scheduling for MIC and MPC Interfaces on page 39](#)
  - [CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 123](#)
  - *MX Series MPC Overview*
  - *MPCs Supported by MX Series Routers*
  - *MX Series MIC Overview*
  - *MICs Supported by MX Series Routers*
  - *MX5, MX10, MX40, and MX80 Modular Interface Card Description*

## Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview

Queuing Ethernet Modular Port Concentrators (MPCs) provide a set of dedicated queues for subscriber interfaces configured with hierarchical scheduling or per-unit scheduling.

The dedicated queues offered on these MPCs enable service providers to reduce costs through different scaling configurations. These queuing MPCs enable service providers to reduce the cost per subscriber by allowing many subscriber interfaces to be created with four or eight queues.

This topic describes the overall queue, scheduler node, and logical interface scaling for subscriber interfaces created on these MIC and MPC combinations.

### Queue Scaling for MPCs

Beginning with Junos OS Release 15.1, MPC2E-3D-NG-Q, MPC3E-3D-NG-Q, MPC5EQ-40G10G, and MPC5EQ-100G10G MPCs support up to five levels of hierarchical queuing. [Table 8 on page 46](#) lists the number of dedicated queues and nodes supported per MPC.

**Table 8: Dedicated Queues for MPCs**

MPC	Dedicated Queues	Level 4 Nodes	Level 3 Nodes	Level 2 Nodes	Level 1 Nodes (Ports)
MPC2E-3D-NG-Q	512,000	64,000	16,000	4000	384
MPC3E-3D-NG-Q					
MPC5EQ-40G10G	1 million	128,000	32,000	4000	384
MPC5EQ-100G10G					



**CAUTION:** The maximum scaling targets provided in [Table 8 on page 46](#) are based on system level design specifications. Actual realized subscriber or session scale is highly dependent upon the configuration and can be influenced by configuration variables including: the number of routes, the

number of enabled services, the number of policy and firewall filters, policers, counters, statistics and access model type. Once you define a configuration, your Juniper account team can help characterize the expected system level scale or scale range for your live deployment.

MPCs vary in the number of Packet Forwarding Engines on board. MPC2E-3D-NG-Q and MPC3E-3D-NG-Q MPCs each have one Packet Forwarding Engine, allowing all 64,000 level 4 (subscriber) nodes to be allocated to a single MIC. MPC5EQ MPCs have two Packet Forwarding Engines, one for each possible MIC, each supporting 64,000 level 4 (subscriber) nodes.



**NOTE:** The nonqueuing MPCs MPC2E-3D-NG, MPC3E-3D-NG, MPC5E-40G10G, and MPC5E-100G10G provide up to eight queues per port in standard configuration. However, each of these MPCs can be configured to provide limited-scale hierarchical class of service (HCoS) and up to 32,000 queues.

## Managing Remaining Queues

When the number of available dedicated queues on the MPC drops below 10 percent, an SNMP trap is generated to notify you.

When the maximum number of dedicated queues on the MPCs is reached, a system log message, **COSD\_OUT\_OF\_DEDICATED\_QUEUES**, is generated. The system does not provide subsequent subscriber interfaces with a dedicated set of queues. For per-unit scheduling configurations, there are no configurable queues remaining on the MPC.

For hierarchical scheduling configurations, remaining queues are available when the maximum number of dedicated queues is reached on the MPC. Traffic from these logical interfaces is considered unclassified and attached to a common set of queues that are shared by all subsequent logical interfaces. These common queues are the default port queues that are created for every port. You can configure a traffic-control profile and attach that to the interface to provide CoS parameters for the remaining queues.

These subscriber interfaces remain with this traffic-control profile, even if dedicated queues become available.

Release History Table

Release	Description
15.1	Beginning with Junos OS Release 15.1, MPC2E-3D-NG-Q, MPC3E-3D-NG-Q, MPC5EQ-40G10G, and MPC5EQ-100G10G MPCs support up to five levels of hierarchical queuing.

### Related Documentation

- *Hierarchical Class of Service Feature Guide*
- [Understanding Hierarchical Scheduling on page 7](#)

- [Managing Dedicated and Remaining Queues for Static CoS Configurations on MIC and MPC Interfaces](#)
- [Managing Dedicated and Remaining Queues for Dynamic CoS Configurations on MIC and MPC Interfaces](#)
- [Understanding Hierarchical Scheduling for MIC and MPC Interfaces on page 39](#)

---

## Jitter Reduction in Hierarchical CoS Queues

---

- [Queue Jitter as a Function of the Maximum Number of Queues on page 48](#)
- [Default Maximum Queues for Hierarchical Queuing MICs and MPCs on page 48](#)
- [Shaping Rate Granularity as a Function of the Rate Wheel Update Period on page 49](#)

### Queue Jitter as a Function of the Maximum Number of Queues

Each queuing chip on a Modular Interface Card (MIC) or Modular Port Concentrator (MPC) internally hosts a *rate wheel thread* that updates the *shaper credits* into the *shapers* available at each level of scheduling hierarchy. At each hierarchy level, the length of this update period determines two key characteristics of scheduling:

- The minimum buffer needed for the queue to pass packets without dropping.
- The degree of jitter encountered in the queue.

At each hierarchy level, the length of the rate wheel update period is dependent upon the number of entities enabled for that node level. Because traffic is queued at Level 5 (queues) and scheduled upwards to Level 1 (the port), the number of entities (queues) enabled at Level 5 determines the number of entities (logical interfaces, interface-sets, or ports) enabled at the other levels of the scheduling hierarchy. By extension, the number of queues enabled for a given scheduler node hierarchy determines the length of the update period at all hierarchy levels. Consequently, limiting the maximum number of queues supported by a hierarchical queuing MIC or MPC can reduce jitter in the queues. To configure the maximum number of queues allowed per hierarchical queuing MIC or MPC, include the `max-queues` statement at the `[edit chassis fpc slot-number]` hierarchy level.

### Default Maximum Queues for Hierarchical Queuing MICs and MPCs

The QX chip on a MIC or MPC consists of two symmetrical halves, and each half supports a maximum of 64 K queues (128 K queues per QX chip). The 2-port and 4-port 10-Gigabit Ethernet MICs with XFP and the MPC1\_Q line cards have one chipset and can support a maximum of 128 K queues, distributed across the two partitions of the single QX chip. The MPC2\_Q and MPC2\_EQ line cards have two chipsets and can support a maximum of 256 K queues, distributed across the four partitions of the two QX chips.

[Table 9 on page 49](#) lists the maximum number of queues supported by default and the corresponding rate wheel update period for each hierarchical queuing MIC or MPC.

**Table 9: Default Maximum Queues and Corresponding Rate Wheel Update Periods**

Router Model	Hierarchical Queuing MIC or MPC	Maximum Queues	Rate Wheel Update Period
MX5, MX10, MX40, and MX80 modular	<b>2-port 10-Gigabit Ethernet MIC with XFP</b> The chassis base board hosts one chipset-based Packet Forwarding Engine process that operates in standalone mode. The single QX chip is composed of two partitions that each support 64 K queues for egress ports.	128 K	1.6 ms
MX240, MX480, MX960, MX2010, and MX2020	<b>MPC1 Q</b> The MPC1 Q line card hosts one chipset-based Packet Forwarding Engine process that operates in fabric mode. The single QX chip is composed of two partitions that each support 64 K queues for egress ports.	128 K	1.6 ms
	<b>MPC2 Q</b> The MPC2 Q line card hosts two chipset-based Packet Forwarding Engine processes that operate in fabric mode. The two QX chips are composed of four partitions that each support 64 K queues for egress ports.	256 K	1.6 ms
	<b>MPC2 EQ</b> The MPC2 EQ line card hosts two chipset-based Packet Forwarding Engine processes that operate in fabric mode. The two QX chips are composed of four partitions that each support 64 K queues for egress ports.	256 K	2.6 ms

You can configure hierarchical queuing MICs and MPCs to support a reduced maximum number of queues. Doing so reduces the rate wheel update period used by the QX chip, which in turn reduces jitter in the queues for the egress interfaces hosted on the line card.

### Shaping Rate Granularity as a Function of the Rate Wheel Update Period

Reducing the length of the QX chip rate wheel update period, in addition to reducing jitter in the hierarchical scheduling queues, also indirectly increases the shaping granularity.

For a given port line rate and scheduling hierarchy level, the shaping granularity is a function of the minimum shaper credit size and the rate wheel update period in effect as a result of the number of queues supported by the line card.

$$\text{shaping granularity} = \text{minimum shaper credit size} / \text{rate wheel update period}$$

[Table 10 on page 50](#) shows how shaping granularity is calculated for non-enhanced hierarchical queuing MIC and MPC line cards with default values for minimum shaper credit size and for rate wheel update period.

Table 10: Default Shaping Granularities on Non-Enhanced Queuing MICs and MPCs

Port Type	Hierarchy Level	Non-Enhanced Queuing MIC or MPC Defaults		Calculation of Shaping Granularity
		Minimum Credit	Update Period	
1 Gbps Queuing	Level 1 (port), Level 4 (queues)	4 bytes = 32 bits	13.33 ms = 0.01333 sec	32 bits / 0.01333 sec = 2.4 Kbps
	Level 2, Level 3	16 bytes = 128 bits	1.66 ms = 0.00166 sec	128 bits / 0.01333 sec = 9.6 Kbps
10 Gbps Queuing	Level 1 (port), Level 4 (queues)	16 bytes = 128 bits	13.33 ms = 0.01333 sec	128 bits / 0.01333 sec = 9.6 Kbps
	Level 2, Level 3	64 bytes = 512 bits	1.66 ms = 0.00166 sec	512 bits / 0.01333 sec = 38.4 Kbps

- Related Documentation**
- [Example: Reducing Jitter in Hierarchical CoS Queues on page 50](#)
  - [Per-Unit Queuing and Hierarchical Queuing for MIC and MPC Interfaces on page 43](#)
  - [Understanding Hierarchical Scheduling for MIC and MPC Interfaces on page 39](#)

## Example: Reducing Jitter in Hierarchical CoS Queues

This example shows how to reduce jitter in the output queues for VLAN ports hosted on a hierarchical queuing MPC.

- [Requirements on page 50](#)
- [Overview on page 50](#)
- [Configuration on page 51](#)

### Requirements

This example uses the following Juniper Networks hardware and Junos OS software:

- MX960 router in an IPv4 network and running Junos OS Release 13.2 or later.
- Available Gigabit Ethernet port hosted on FPC slot 2, PIC slot 0, port 0.
- Available Gigabit Ethernet port hosted on port 0 of a Gigabit Ethernet Modular Interface Card (MIC) in PIC slot 0 of an MPC2 Q Modular Port Concentrator (MPC) in FPC slot 5.

Before you begin configuring this example, make sure that the maximum number of queues allowed for the hierarchical queuing MPC in slot 5 has not yet been configured. When you enter the **show chassis fpc 5** command from configuration mode, the **max-queues** statement should not display.

### Overview

In this example you configure hierarchical scheduling on a VLAN port hosted on a hierarchical queuing MPC. To reduce jitter in the queues for all egress ports hosted on the MPC, reduce the maximum number of queues allowed for MPC.



## Configuration

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

```
set interfaces xe-2/0/0 per-unit-scheduler
set interfaces xe-2/0/0 flexible-vlan-tagging
set interfaces xe-2/0/0 unit 0 vlan-id 1
set interfaces xe-2/0/0 unit 0 family inet address 10.1.1.1/24
set interfaces xe-2/0/0 unit * classifiers ieee-802.1 ieee_jitter
set interfaces xe-5/0/0 per-unit-scheduler
set interfaces xe-5/0/0 flexible-vlan-tagging
set interfaces xe-5/0/0 unit 0 vlan-id 1
set interfaces xe-5/0/0 unit 0 family inet address 10.2.1.1/24
set class-of-service interfaces xe-5/0/0 unit * output-traffic-control-profile tcp
set class-of-service forwarding-classes queue 0 be
set class-of-service forwarding-classes queue 1 ef
set class-of-service forwarding-classes queue 2 af
set class-of-service forwarding-classes queue 3 nc
set class-of-service schedulers be_sch priority low
set class-of-service schedulers ef_sch priority low
set class-of-service schedulers af_sch priority strict-high
set class-of-service schedulers nc_sch priority low
set class-of-service classifiers ieee_jitter forwarding-class be loss-priority low code-points
000
set class-of-service classifiers ieee_jitter forwarding-class ef loss-priority low code-points
001
set class-of-service classifiers ieee_jitter forwarding-class af loss-priority low code-points
010
set class-of-service classifiers ieee_jitter forwarding-class nc loss-priority low code-points
011
set class-of-service scheduler-maps smap_jitter forwarding-class be scheduler be_sch
set class-of-service scheduler-maps smap_jitter forwarding-class ef scheduler ef_sch
set class-of-service scheduler-maps smap_jitter forwarding-class af scheduler af_sch
set class-of-service scheduler-maps smap_jitter forwarding-class nc scheduler nc_sch
set class-of-service traffic-control-profiles tcp scheduler-map smap_jitter
set class-of-service traffic-control-profiles tcp shaping-rate 6g
```

---

### Baseline Configuration

**Step-by-Step Procedure** Configure hierarchical scheduling at xe-5.0.0.

1. To configure the VLAN 1 input and output at xe-2/0/0.0 and xe-5/0/0.0:

```
[edit]
user@host# set interfaces xe-2/0/0 per-unit-scheduler
user@host# set interfaces xe-2/0/0 flexible-vlan-tagging
user@host# set interfaces xe-2/0/0 unit 0 vlan-id 1
user@host# set interfaces xe-2/0/0 unit 0 family inet address 10.1.1.1/24

user@host# set interfaces xe-5/0/0 per-unit-scheduler
user@host# set interfaces xe-5/0/0 flexible-vlan-tagging
```

```
user@host# set interfaces xe-5/0/0 unit 0 vlan-id 1
user@host# set interfaces xe-5/0/0 unit 0 family inet address 10.2.1.1/24
```

2. Map each of four queues to a forwarding class.

```
[edit]
user@host# set class-of-service forwarding-classes queue 0 be
user@host# set class-of-service forwarding-classes queue 1 ef
user@host# set class-of-service forwarding-classes queue 2 af
user@host# set class-of-service forwarding-classes queue 3 nc
```

3. Assign a packet-scheduling priority value to each forwarding class.

```
[edit]
user@host# set class-of-service schedulers be_sch priority low
user@host# set class-of-service schedulers ef_sch priority low
user@host# set class-of-service schedulers af_sch priority strict-high
user@host# set class-of-service schedulers nc_sch priority low
```

4. Customize the default IEEE 802.1p classifier (BA classifier based on Layer 2 header) by defining different values for IEEE 802.1p code points.

```
[edit]
user@host# set class-of-service classifiers ieee_jitter forwarding-class be
  loss-priority low code-points 000
user@host# set class-of-service classifiers ieee_jitter forwarding-class ef loss-priority
  low code-points 001
user@host# set class-of-service classifiers ieee_jitter forwarding-class af loss-priority
  low code-points 010
user@host# set class-of-service classifiers ieee_jitter forwarding-class nc
  loss-priority low code-points 011
```

5. Apply the BA classifier to the input of the logical units on xe-2/0/0.

```
[edit]
user@host# set interfaces xe-2/0/0 unit * classifiers ieee-802.1 ieee_jitter
```

6. Configure the scheduler map **smap\_jitter** to map the forwarding classes to the schedulers.

```
[edit]
user@host# set class-of-service scheduler-maps smap_jitter forwarding-class be
  scheduler be_sch
user@host# set class-of-service scheduler-maps smap_jitter forwarding-class ef
  scheduler ef_sch
user@host# set class-of-service scheduler-maps smap_jitter forwarding-class af
  scheduler af_sch
user@host# set class-of-service scheduler-maps smap_jitter forwarding-class nc
  scheduler nc_sch
```

7. Configure the traffic control profile **tcp** to combine the scheduler map **smap\_jitter** (that maps the forwarding classes to the schedulers for port-based scheduling) with a shaping rate (for hierarchical scheduling).

```
[edit]
user@host# set class-of-service traffic-control-profiles tcp scheduler-map
smap_jitter
user@host# set class-of-service traffic-control-profiles tcp shaping-rate 6g
```

8. Apply the traffic control profile to the router output at **xe-5/0/0**.

```
[edit]
user@host# set class-of-service-interfaces xe-5/0/0 unit *
output-traffic-control-profile tcp
```

9. If you are done configuring the device, commit the configuration.

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

**Results** Confirm your configuration by entering **show interfaces** and **show class-of-service** commands from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
xe-2/0/0 {
  per-unit-scheduler;
  flexible-vlan-tagging;
  unit 0 {
    vlan-id 1;
    family inet {
      address 10.1.1.1/24;
    }
  }
}
xe-5/0/0 {
  per-unit-scheduler;
  flexible-vlan-tagging;
  unit 0 {
    vlan-id 1;
    family inet {
      address 10.2.1.1/24;
    }
  }
}

[edit]
user@host# show class-of-service
classifiers {
  ieee-802.1 ieee_jitter {
    forwarding-class be {
      loss-priority low code-points 000;
    }
  }
}
```

```
    forwarding-class ef {
        loss-priority low code-points 001;
    }
    forwarding-class af {
        loss-priority low code-points 010;
    }
    forwarding-class nc {
        loss-priority low code-points 011;
    }
}
forwarding-classes {
    queue 0 be;
    queue 1 ef;
    queue 2 af;
    queue 3 nc;
}
traffic-control-profiles {
    tcp {
        scheduler-map smap_jitter;
        shaping-rate 6g;
    }
}
interfaces {
    xe-2/0/0 {
        unit * {
            classifiers {
                ieee-802.1 ieee_jitter;
            }
        }
    }
    xe-5/0/0 {
        unit * {
            output-traffic-control-profile tcp;
        }
    }
}
scheduler-maps {
    smap_jitter {
        forwarding-class be scheduler be_sch;
        forwarding-class ef scheduler ef_sch;
        forwarding-class af scheduler af_sch;
        forwarding-class nc scheduler nc_sch;
    }
}
schedulers {
    be_sch {
        priority low;
    }
    ef_sch {
        priority low;
    }
    af_sch {
        priority strict-high;
    }
    nc_sch {
```

```
        priority low;  
    }  
}
```

### Verification

---

Confirm that the configuration is working properly

- [Measuring End-to-End Jitter to Establish the Baseline on page 55](#)
- [Configuring Jitter Reduction on page 55](#)
- [Measuring End-to-End Jitter to Verify Jitter Reduction on page 55](#)

#### *Measuring End-to-End Jitter to Establish the Baseline*

**Purpose** Establish a baseline measurement by noting the amount of jitter that occurs when the hierarchical queuing line card hosting the egress port is configured with the default maximum number of queues.

**Action** To measure jitter:

1. Pass traffic through the VLAN.
2. Measure the variation in packet delay for selected packets in the data flow.

#### *Configuring Jitter Reduction*

**Purpose** Reduce jitter in the VLAN port output queues.

**Action** 1. Configure a reduced maximum number of queues for egress ports on the hierarchical queuing MPC in slot 5, thereby reducing the jitter in the port queues.

[edit]

user@host# set chassis fpc 5 max-queue 64k

2. If you are done configuring the device, commit the configuration.

[edit]

user@host# commit

#### *Measuring End-to-End Jitter to Verify Jitter Reduction*

**Purpose** Measure the amount of jitter that occurs when the hierarchical queuing line card hosting the egress port is configured with a reduced maximum number of queues.

**Action** To measure jitter:

1. Pass traffic through the VLAN.

2. Measure the variation in packet delay for selected packets in the data flow.

**Related  
Documentation**

- [Jitter Reduction in Hierarchical CoS Queues on page 48](#)
- [max-queues on page 334](#)

---

## Hierarchical Schedulers on Aggregated Ethernet Interfaces Overview

On MX Series routers, you can apply hierarchical schedulers on aggregated ethernet bundles using interface sets. This feature enables you to configure a group of virtual LANs (VLANs) and control their bandwidth. This feature is supported at egress only.

You can configure interface sets for aggregated Ethernet (AE) interfaces created under static configurations. You can configure class-of-service parameters on AE interfaces, in either link-protect or non-link-protect mode. You can configure these parameters at the AE physical interface level. The CoS configuration is fully replicated for all AE member links in link-protect mode. You can control the way these parameters are applied to member links in non-link-protect mode by configuring the AE interface to operate in scaled mode or replicate mode.

The link membership list and scheduler mode of the interface set is inherited from the underlying aggregated Ethernet interface over which the interface set is configured. When an aggregated Ethernet interface operates in link protection mode, or if scheduler mode is configured to replicate member links, the scheduling parameters of the interface set are copied to each of the member links.

If the scheduler mode of the aggregated Ethernet interface is set to scale member links, the scheduling parameters are scaled based on the number of active member links (scaling factor is  $1/A$  where  $A$  is the number of active links in the bundle) and applied to each of the AE interface member links.

To configure an interface set, include the **interface-set** statement at the **[edit class-of-service interfaces]** hierarchy level.

To apply scheduling and queuing parameters to the interface set, include the **output-traffic-control-profile** *profile-name* statement at the **[edit class-of-service interfaces interface-name interface-set interface-set-name]** hierarchy level.

To apply an output traffic scheduling and shaping profile for the remaining traffic to the logical interface or interface set, include the **output-traffic-control-profile-remaining** *profile-name* statement at the **[edit class-of-service interfaces interface-name]** hierarchy level or the **[edit class-of-service interfaces interface-name interface-set interface-set-name]** hierarchy level.

**Related  
Documentation**

- [Configuring Hierarchical Schedulers on Aggregated Ethernet Interfaces on page 57](#)
- [output-traffic-control-profile-remaining on page 342](#)
- *Controlling Remaining Traffic*

## Configuring Hierarchical Schedulers on Aggregated Ethernet Interfaces

The following example shows the creation of an interface set for aggregated Ethernet interfaces in a static Ethernet configuration.

To configure interface sets for aggregated Ethernet (AE) interfaces created under static configurations:

1. Create the AE interfaces.

```
[edit]
user@host# show chassis | display set
set chassis aggregated-devices ethernet device-count 10
```

2. Configure the AE physical interfaces and member links.

```
user@host# show interfaces | display set

set interfaces ge-5/2/0 gigether-options 802.3ad ae0
set interfaces ge-5/2/1 gigether-options 802.3ad ae0
set interfaces ae0 hierarchical-scheduler maximum-hierarchy-levels 2
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 unit 0 vlan-id 100
set interfaces ae0 unit 1 vlan-id 101
set interfaces ae0 unit 2 vlan-id 102
set interfaces ae0 unit 3 vlan-id 103
set interfaces ae0 unit 4 vlan-id 104
```

3. Configure the interface set.

```
set interfaces interface-set ifset1-ae0 interface ae0 unit 0
set interfaces interface-set ifset1-ae0 interface ae0 unit 1
```

4. Configure class-of-service parameters for the interface sets.

```
set class-of-service interfaces interface-set ifset1-ae0 output-traffic-control-profile
tcp
```



**NOTE:** You also need to configure the parameters of the traffic control profile. For more information, see the Related Documentation section on this page.

5. Configure scheduler mode.

```
set class-of-service interfaces ae0 member-link-scheduler scale
```

### Related Documentation

- [Configuring Traffic Control Profiles for Shared Scheduling and Shaping](#)
- [Hierarchical Schedulers on Aggregated Ethernet Interfaces Overview on page 56](#)
- [Example: Configuring Shared Resources on Ethernet IQ2 Interfaces](#)

- *Configuring Traffic Control Profiles for Shared Scheduling and Shaping*

## Example: Configuring Scheduling Modes on Aggregated Interfaces

---

You can configure class-of-service parameters, such as queuing or shaping parameters on aggregated interfaces, in either link-protect or non-link-protect mode. You can configure these parameters for per-unit schedulers, hierarchical schedulers, or shaping at the physical and logical interface level. You can control the way these parameters are applied by configuring the aggregated interface to operate in **scale** or **replicate** mode.

You can apply these parameters on the following routers:

- MX Series router interfaces on EQ DPCs
- MX Series router interfaces on MICs or MPCs through Junos OS Release 10.2 (non-link-protect mode only)
- M120 or M320 routers
- T Series router interfaces on IQ2 PICs
- PTX Series Packet Transport Routers

You can configure the applied parameters for aggregated interfaces operating in non-link-protected mode. In link-protected mode, only one link in the bundle is active at a time (the other link is a backup link) so schedulers cannot be scaled or replicated. In non-link-protected mode, all the links in the bundle are active and send traffic; however, there is no backup link. If a link fails or is added to the bundle in non-link-protected mode, the links' traffic is redistributed among the active links.

To set the scheduling mode for aggregated interfaces, include the **scale** or **replicate** option of the **member-link-scheduler** statement at the **[edit class-of-service interfaces ae*n*]** hierarchy level, where *n* is the configured number of the interface:

```
[edit class-of-service interfaces aen]  
  member-link-scheduler (replicate | scale);
```

By default, if you do not include the **member-link-scheduler** statement, scheduler parameters are applied to the member links in the **scale** mode (also called "equal division mode").

The aggregated Ethernet interfaces are otherwise configured as usual. For more information on configuring aggregated Ethernet interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

The following examples set **scale** mode on the **ae0** interface and **replicate** mode on the **ae1** interface.

```
[edit class-of-service]  
  interfaces ae0 {  
    member-link-scheduler scale;  
  }  
  
  interfaces ae1 {  
    member-link-scheduler replicate;  
  }
```



```
[edit class-of-service]
interfaces ae1 {
  member-link-scheduler replicate;
}
```



**NOTE:** The `member-link-scheduler` statement only appears for aggregated interfaces. You configure this statement for aggregated interfaces in non-link-protected mode. For more information about link protection modes, see the *Network Interfaces Configuration Guide*.

Aggregated interfaces support both hierarchical and per-unit schedulers.



**NOTE:** The `traffic-control-profiles` statement is not supported for PTX Series Packet Transport Routers.

When interface parameters are using the `scale` option of the `member-link-scheduler` statement, the following parameters under the `[edit class-of-service traffic-control-profiles traffic-control-profile-name]` configuration are scaled on egress when hierarchical schedulers are configured:

- `shaping-rate` (PIR)
- `guaranteed-rate` (CIR)
- `delay-buffer-rate`

When interface parameters are using the `scale` option of the `member-link-scheduler` statement, the following parameters under the `[edit class-of-service schedulers scheduler-name]` configuration are scaled on egress when per-unit schedulers are configured:

- `transmit-rate`
- `buffer-size`



**NOTE:** You cannot apply a hierarchical scheduler at the interface set level for an `ae` interface. (Interface sets cannot be configured under an `ae` interface.)

The following configuration parameters are not supported on `ae` interfaces in non-link-protection mode:

- Input scheduler maps
- Input traffic control profiles
- Input shaping rates

The following configuration conventions are also not supported:

- Scaling of the **input-traffic-control-profile-remaining** statement.
- The **scheduler-map-chassis** statement and the **derived** option for the **ae** interface. Chassis scheduler maps should be applied under the physical interfaces.
- Dynamic and demux interfaces are not supported as part of the **ae** bundle.

Depending on whether the **scale** or **replicate** option is configured, the **member-link-scheduler** statement operates in either scaled mode (also called “equal division mode”) or replicated mode, respectively.

In scaled mode, a VLAN can have multiple flows that can be sent over multiple member links of the **ae** interface. Likewise, a member link can receive traffic from any VLAN in the **ae** bundle. In scaled mode, the physical interface bandwidth is divided equally among all member links of the **ae** bundle.

In scaled mode, the following scheduler parameter values are divided equally among the member links:

- When the parameters are configured using traffic control profiles, then the parameters scaled are the shaping rate, guaranteed rate, and delay buffer rate.
- When the parameters are configured using scheduler maps, then the parameters scaled are the transmit rate and buffer size. Shaping rate is also scaled if you configure it in bits per second (bps). Shaping rate is not scaled if you configure it as a percentage of the available interface bandwidth.

For example, consider an **ae** bundle between routers R1 and R2 consisting of three links. These are **ge-0/0/1**, **ge-0/0/2** and **ge-0/0/3** (**ae0**) on R1; and **ge-1/0/0**, **ge-1/0/1**, and **ge-1/0/2** (**ae2**) on R2. Two logical interfaces (units) are also configured on the **ae0** bundle on R1: **ae0.0** and **ae0.1**.

On **ae0**, traffic control profiles on R1 are configured as follows:

- **ae0** (the physical interface level) has a PIR of 450 Mbps.
- **ae0.0** (VLAN 100 at the logical interface level) has a PIR of 150 Mbps and a CIR of 90 Mbps.
- **ae0.1** (VLAN 200 at the logical interface level) has a PIR of 90 Mbps and a CIR of 60 Mbps.

In scaled mode, the **ae0** PIR is first divided among the member physical interfaces. Because there are three members, each receives  $450 / 3 = 150$  Mbps as a derived value. So the scaled PIR for the members interfaces is 150 Mbps each.

However, there are also two logical interfaces (**ae0.0** and **ae0.1**) and VLANs (100 and 200) on **ae0**. Traffic can leave on any of the three physical interfaces (**ge-0/0/1**, **ge-0/0/2**, or **ge-0/0/3**) in the bundle. Therefore, two derived logical interfaces are added to the member links to represent the two VLANs.

There are now six logical interfaces on the physical interfaces of the links making up the **ae** bundle, one set for VLAN 100 and the other for VLAN 200:

- **ge-0/0/1.0** and **ge-0/0/1.1**
- **ge-0/0/2.0** and **ge-0/0/2.1**
- **ge-0/0/3.0** and **ge-0/0/3.1**


The traffic control profile parameters configured on **ae0.0** are divided across all the underlying logical interfaces (the unit 0s). In the same way, the traffic control profile parameters configured on **ae0.1** are divided across all the underlying logical interfaces (the unit 1s).

Therefore, the derived values of the scaled parameters on the interfaces are:

- For **ge-0/0/1.0** and **ge-0/0/2.0** and **ge-0/0/3.0**, each CIR =  $90 / 3 = 30$  Mbps, and each PIR =  $150 / 3 = 50$  Mbps.
- For **ge-0/0/1.1** and **ge-0/0/2.1** and **ge-0/0/3.1**, each CIR =  $60 / 3 = 20$  Mbps, and each PIR =  $90 / 3 = 30$  Mbps.

The scaled values are shown in [Figure 6 on page 61](#).

**Figure 6: Scaled Mode for Aggregated Ethernet Interfaces**



In scaled mode, when a new member link is added to the bundle, or an existing member link is either removed or fails, then the scaling factor (based on the number of active links) is recomputed and the new scheduler or traffic control profile parameters are reassigned. Only the PIR, CIR, and buffer parameters are recomputed: all other parameters are simply copied at each level.



**NOTE:** In **show class-of-service scheduler-map** commands, values derived in scaled mode instead of explicitly configured are flagged with **&sf\*n** suffix, where **n** indicates the value of the scaling factor.

The following sample shows the output for the scheduler map named **smap-all-abs** with and without a scaling factor:

```

user@host> show class-of-service scheduler-map
Scheduler map: smap-all-abs, Index: 65452
  
```

```
Scheduler: q0_sch_abs, Forwarding class: be, Index: 6775
Transmit rate: 40000000 bps, Rate Limit: none, Buffer size: remainder,
Priority: low
```

```
Excess Priority: unspecified
```

```
Drop profiles:
```

Loss priority	Protocol	Index	Name
Low	any	1	<default-drop-profile>
Medium low	any	1	<default-drop-profile>
Medium high	any	1	<default-drop-profile>
High	any	1	<default-drop-profile>

```
user@host> show class-of-service scheduler-map
```

```
Scheduler map: smap-all-abs, Index: 65452
```

```
Scheduler: q0_sch_abs&***sf**3, Forwarding class: be, Index: 2128
Transmit rate: 13333333 bps, Rate Limit: none, Buffer size: remainder,
Priority: low
```

```
Excess Priority: unspecified
```

```
Drop profiles:
```

Loss priority	Protocol	Index	Name
Low	any	1	<default-drop-profile>
Medium low	any	1	<default-drop-profile>
Medium high	any	1	<default-drop-profile>
High	any	1	<default-drop-profile>



**NOTE:** There can be multiple scheduler maps created with different scaling factors, depending on when the child interfaces come up. For example, if there are only two active children on a parent interface, a new scheduler map with a scaling factor of 2 is created. The scheduler map name is **smap-all-abs&\*\*\*sf\*\*2**.

In replicated mode, in contrast to scaled mode, the configured scheduler parameters are simply replicated, not divided, among all member links of the **ae** bundle.

In replicated mode, the following scheduler parameter values are replicated among the member links and logical interfaces:

- When the parameters are configured using traffic control profiles, then the parameters replicated are the shaping rate, guaranteed rate, and delay buffer rate.
- When the parameters are configured using scheduler maps, then the parameters replicated are the transmit rate and buffer size.

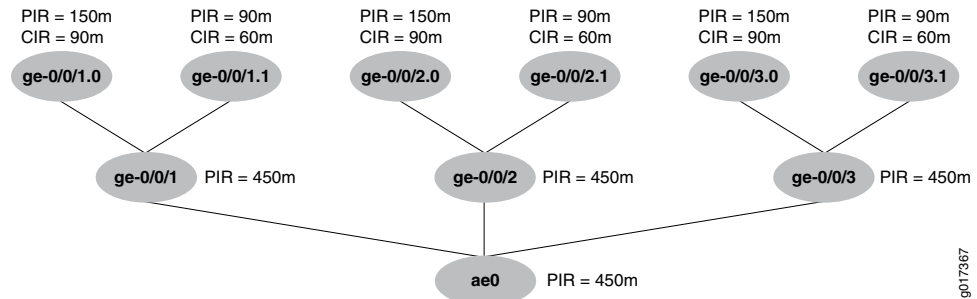
If the scheduler parameters in the example configuration between routers R1 and R2 are applied with the **member-link-scheduler replicate** statement and option, the following parameters are applied:

- The **ae0** PIR is copied among the member physical interfaces. Each receives 450 Mbps as a PIR.
- For each logical interface unit **.0**, the configured PIR and CIR for **ae0.0** is replicated (copied). Each logical interface unit **.0** receives a PIR of 150 Mbps and a CIR of 90 Mbps.

- For each logical interface unit .1, the configured PIR and CIR for **ae0.1** is replicated (copied). Each logical interface unit .1 receives a PIR of 90 Mbps and a CIR of 60 Mbps.

The replicated values are shown in [Figure 7 on page 63](#).

**Figure 7: Replicated Mode for Aggregated Ethernet Interfaces**



In replicated mode, when a new member link is added to the bundle, or an existing member link is either removed or fails, the values are either copied or deleted from the required levels.

#### Related Documentation

- How Schedulers Define Output Queue Properties*
- Default Schedulers Overview*
- Configuring Schedulers*

## Increasing Available Bandwidth on Rich-Queuing MPCs by Bypassing the Queuing Chip

Queuing MPCs contain a queuing chip that enables rich-queuing features such as hierarchical and per-vlan queuing. By default, all traffic passing through an interface on one of these MPCs also passes through the queuing chip, which decreases the available bandwidth of the interface. If you do not require hierarchical or per-vlan queuing on a particular interface of a queuing MPC, you can bypass the queuing chip to increase the available bandwidth.

To bypass the queuing chip on a queuing MPC, you must be running Junos OS Release 14.2 or later. You can bypass the queuing chip on the following line cards:

- MPC1 Q*
- MPC1E Q*
- MPC2 Q*
- MPC2 EQ*
- MPC2E Q*
- MPC2E EQ*
- MPC5E Q (2x100GE + 4x10GE MPC5EQ or 6x40GE + 24x10GE MPC5EQ)*

To bypass the queuing chip on an interface on a queuing MPC:

1. Ensure that neither **per-unit-scheduler** nor **hierarchical-scheduler** is configured on the interface.



**NOTE:** It is not possible to bypass the queuing chip on an interface if per-unit or hierarchical scheduling is configured on that interface.

2. Enable **bypass-queuing-chip** on the interface.

For example:

```
[edit interfaces]
user@router# set interface- name bypass-queuing-chip
```

3. Commit your changes.

```
[edit interfaces]
user@router# show
interface-name {
    bypass-queueing-chip;
}
```

4. Verify your changes.

```
user@router> show interfaces interface-name
Physical interface: interface-name, Enabled, Physical link is Up
  Interface index: 147, SNMP ifIndex: 524
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed:
1000mbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source
filtering: Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues    : 8 supported, 4 maximum usable queues
  Schedulers    : 0, Queuing Chip Bypassed
  Current address: 00:21:59:0f:35:31, Hardware address: 00:21:59:0f:35:31
  Last flapped  : 2014-04-29 14:10:18 PDT (02:27:46 ago)
  Input rate    : 0 bps (0 pps)
  Output rate   : 0 bps (0 pps)
  Active alarms : None
  Active defects: None
  Interface transmit statistics: Disabled
```

**Release History Table**

Release	Description
14.2	To bypass the queuing chip on a queuing MPC, you must be running Junos OS Release 14.2 or later.

**Related  
Documentation**

- *bypass-queuing-chip*





## PART 2

# Hierarchical CoS for Subscriber Management

- [Hierarchical Class of Service for Subscriber Management on page 69](#)
- [Applying CoS to Groups of Subscriber Interfaces on page 95](#)
- [Configuring Hierarchical Scheduling for MPLS Pseudowire Interfaces on page 117](#)
- [Configuring Hierarchical Scheduling for L2TP on page 131](#)
- [Preventing Bandwidth Contention on Subscriber Interfaces on page 141](#)
- [Configuring Targeted Distribution of Subscribers on Aggregated Ethernet Interfaces on page 163](#)
- [Applying CoS Using Parameters Received from RADIUS on page 183](#)



## CHAPTER 3

# Hierarchical Class of Service for Subscriber Management

- [Hierarchical Class of Service for Subscriber Management Overview on page 69](#)
- [Understanding Hierarchical CoS for Subscriber Interfaces on page 70](#)
- [Hardware Requirements for Dynamic Hierarchical CoS on page 77](#)
- [Configuring Static Hierarchical Scheduling in a Dynamic Profile on page 78](#)
- [Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links on page 79](#)
- [Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface on page 80](#)
- [Example: Maintaining a Constant Traffic Flow by Configuring a Static VLAN Interface with a Dynamic Profile for Subscriber Access on page 81](#)

### Hierarchical Class of Service for Subscriber Management Overview

---

The hierarchical class-of-service (HCoS) architecture as supported on fine-grained queuing MPCs is a powerful feature designed to provide a flexible and scalable CoS solution in broadband network gateway (BNG) subscriber access applications where triple-play or business class offerings are enabled through IP CoS.

Hierarchical CoS enables you to apply traffic scheduling and queuing parameters (which can include a delay-buffer bandwidth) and packet transmission scheduling parameters (which can include buffer management parameters) to an individual subscriber interface rather than to all interfaces configured on the port. HCoS enables you to dynamically modify queues when subscribers require services.

The logical interface set construct in a five-level scheduler hierarchy is the key feature that enables HCoS. The interface set feature allows you to group subscribers into aggregate classes with specific guaranteed and peak rates that map to service classes. Service classes ultimately map to how much you can charge for the differentiated service levels.

HCoS can be applied dynamically through the use of dynamic traffic profiles and RADIUS vendor-specific attributes (VSAs).

Dynamic traffic profiles are used to dynamically apply CoS to individual subscribers or groups of subscribers. This enables you, as a service provider, to deploy a BRAS solution

without having to manually provision each customer. In a dynamic traffic profile, variables are used to represent the values for things like shaping rate and drop priority.

Dynamic traffic profiles are used in conjunction with dynamic profiles. Dynamic profiles allow you to dynamically provision IP service definitions by creating a template configuration and having the specific variable values assigned in real time when the subscriber authenticates to the network.



**NOTE:** For a complete list of the Junos OS system variables, see:

[https://www.juniper.net/documentation/en\\_US/junos42/topics/reference/general/junos-defined-variables-table.html](https://www.juniper.net/documentation/en_US/junos42/topics/reference/general/junos-defined-variables-table.html)

---

To learn more about how to use HCoS in conjunction with dynamic traffic control profiles for subscriber management, read the [Day One: Dynamic Subscriber Management](#) book.

In addition, the following topics are very helpful:

- [Hierarchical Class of Service Overview on page 3](#)
- *Subscriber Access Network Overview*
- *CoS for Subscriber Access Overview*
- *Subscriber Management Overview*
- *Class of Service and Subscriber Management Overview*

Before applying dynamic HCoS on your network, you should learn about HCoS, define your needs, plan how you want to implement HCoS, and test the operation in a simulated environment.

**Related  
Documentation**

- [Hierarchical Class of Service Overview on page 3](#)
- [Hierarchical Class of Service Network Scenarios on page 6](#)

---

## Understanding Hierarchical CoS for Subscriber Interfaces

---

Hierarchical CoS enables you to apply traffic scheduling and queuing parameters and packet transmission scheduling parameters to an individual subscriber interface rather than to all interfaces configured on a port. Hierarchical CoS enables you to dynamically modify queues when subscribers require services.

Hierarchical CoS is supported on MX Series routers with either Enhanced Queuing DPCs or queuing MPCs/MICs installed. Beginning with Junos OS Release 16.1, five levels of hierarchy are supported on MPC5E 3D Q line cards.

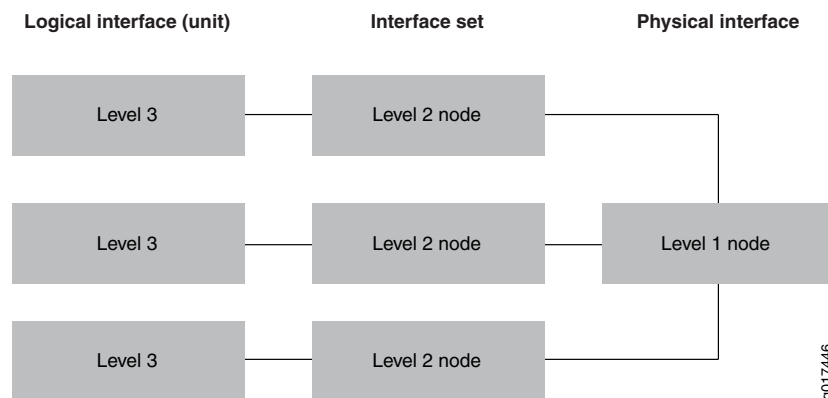
Interfaces support up to a five-level CoS scheduling hierarchy that, when fully configured, consists of the physical interface (level 1), an interface set or underlying interface (level 2), one or more underlying logical interfaces (level 3), one or more session or customer VLANs (level 4), and one or more queues (level 5). Although all CoS scheduling hierarchies are five-level, level 1 is always the physical interface and level 5 is always the queue. Hierarchical scheduling configurations consist of the type of interfaces you configure—for

example, a logical interface or an interface set—and where those interfaces reside in the scheduling hierarchy—level 2, level 3, or level 4. Because many hierarchical scheduling configurations are possible, we use the terms *two-level hierarchical scheduling*, *three-level hierarchical scheduling*, *four-level hierarchical scheduling* in this topic.

## Two-Level Hierarchical Scheduling

Two-level hierarchical scheduling limits the number of hierarchical levels in the scheduling hierarchy to two as shown in [Figure 5 on page 45](#). In this configuration, interface sets are not configured and only the logical interfaces have traffic control profiles (TCPs). Configuring two levels of hierarchy on MPCs that support more levels preserves resources and allows the system to scale higher.

**Figure 8: Two-Level Hierarchical Scheduling**



In a two-level scheduling hierarchy, all logical interfaces and interface sets share a single node; no hierarchical relationship is formed.

You control two-level hierarchical scheduling by setting the **maximum-hierarchy-levels** option under the **[edit interfaces *interface-name* hierarchical-scheduler]** hierarchy to 2:

- If the **maximum-hierarchy-levels** option is not set, then interface sets can be at either level 2 or level 3, depending on whether the member logical interfaces within the interface set have a traffic control profile.
- If any member logical interface has a traffic-control profile, then the interface set is always a level 2 CoS scheduler node.
- If no member logical interface has a traffic-control profile, the interface set is always a level 3 CoS scheduler node.
- If the **maximum-hierarchy-levels** option is set, then the interface set can only be at level 3; it cannot be at level 2. In this case, if you configure a level 2 interface set, you generate Packet Forwarding Engine errors.

[Table 11 on page 72](#) summarizes the interface hierarchy and the CoS scheduler node levels for two-level hierarchical scheduling.

**Table 11: Two-Level Hierarchical Scheduling—Interface Hierarchy Versus Scheduling Nodes**

Level 1	Level 2	Level 3
Physical interface	Logical interface	One or more queues
Physical interface	Interface set	One or more queues
Physical interface	Logical interface	One or more queues

To configure two-level hierarchical scheduling, include the **hierarchical-scheduler** statement at the **[edit interfaces *interface-name*]** hierarchy level and set the **maximum-hierarchy-levels** option to 2.

```
[edit interfaces]
interface-name {
  hierarchical-scheduler {
    maximum-hierarchy-levels 2;
  }
}
```

### Three-Level Hierarchical Scheduling

Three-level hierarchical scheduling is supported only on MX Series routers running MPC/MIC interfaces. Three-level hierarchical scheduling supports up to eight CoS queues. You can configure many different three-level scheduling hierarchies, depending on the location of the interface set or the use of underlying interfaces. In all variations, the physical interface is a level 1 CoS scheduler node and the queues reside at the highest level. Configuring three levels of hierarchy on MPCs that support more levels preserves resources and allows the system to scale higher.



**NOTE:** Three-level hierarchical scheduling is supported only on subscriber interfaces and interface sets running over aggregated Ethernet interfaces on MPC/MIC interfaces in MX Series routers.

When you use three-level hierarchical scheduling, interface sets can reside at either level 3 or level 4. You can also configure an underlying logical interface at level 3 and a logical interface at level 4. [Table 12 on page 72](#) summarizes the most common cases of the interface hierarchy and the CoS scheduler node levels for three-level hierarchical scheduling.

**Table 12: Three-Level Hierarchical Scheduling—Interface Hierarchy Versus CoS Scheduling Node Levels**


Level 1	Level 2	Level 3	Level 4
Physical interface	Interface set	Logical interface	One or more queues
Physical interface	Logical interface	Interface set	One or more queues

**Table 12: Three-Level Hierarchical Scheduling—Interface Hierarchy Versus CoS Scheduling Node Levels** (*continued*)


Level 1	Level 2	Level 3	Level 4
Physical interface	Underlying logical interface	Logical interface	One or more queues

In three-level hierarchical scheduling, the CoS scheduler nodes at level 1, level 2, and level 3 form a hierarchical relationship.

With a three-level hierarchical scheduling, logical interfaces can reside at level 2, or they can reside at level 3 if the logical interface at level 2 is an underlying logical interface. This is shown in [Figure 9 on page 73](#).

**Figure 9: Three-Level Hierarchical Scheduling—Logical Interfaces at Level 3 with Underlying Logical Interfaces at Level 2**

Another possible configuration for three-level hierarchical scheduling is shown in [Figure 10 on page 73](#). In this configuration, the logical interfaces are located at level 2 and the interface sets are located at level 3.

**Figure 10: Three-Level Hierarchical Scheduling—Logical Interfaces at Level 2 with Interface Sets at Level 3**

To configure three-level hierarchical scheduling, include the **implicit-hierarchy** option at the `[edit interfaces interface-name hierarchical-scheduler]` hierarchy level and optionally

set the **maximum-hierarchy-levels** option to 3. (The default value for **maximum-hierarchy-levels** is 3.)

```
[edit interfaces]
interface-name {
  hierarchical-scheduler {
    implicit-hierarchy;
    maximum-hierarchy-levels 3;
  }
}
```


### Interface Hierarchy Versus CoS Hierarchy

An interface hierarchy and a CoS scheduling hierarchy are distinctly different. Interface hierarchy refers to the relationship between the various interfaces—for example, the relationship between logical interfaces and an interface set, the relationship between a logical interface and an underlying logical interface, or the relationship between the physical interface and the logical interface. CoS scheduling hierarchy refers to the hierarchical relationship between the CoS scheduler nodes. In two-level hierarchical scheduling, no hierarchy is formed between the CoS scheduler nodes—the logical interface and interface set share a single level 2 scheduler node. However, when you use the **implicit-hierarchy** option for three-level hierarchical scheduling, the CoS scheduler nodes form a scheduling hierarchy.

Figure 11 on page 74 and Figure 12 on page 75 provide two scenarios for this discussion. Figure 11 on page 74 shows an interface hierarchy where a Gigabit Ethernet interface (ge-1/0/0) is the physical interface. Two logical interfaces (ge-1/0/0.100 and ge-1/0/0.101) are configured on the physical interface:

- Logical interface ge-1/0/0.100 is a member of a PPPoE interface set and a Demux interface set.
- Logical interface ge-1/0/0.101 is a member of a demux interface set.

Figure 11: Logical Interfaces at Level 2 and Interface Sets at Level 3






Each interface set has a dedicated queue. The CoS scheduler nodes at level 1 (physical interface), level 2 (underlying logical interfaces), and level 3 (interface sets) form a scheduling hierarchy.

To configure this scenario, you must include the **implicit-hierarchy** option under the **hierarchical-scheduler** statement on physical interface ge-1/0/0 and configure and apply traffic-control profiles on each interface set and underlying logical interface.

Figure 12 on page 75 shows an interface hierarchy where Gigabit Ethernet interface ge-1/0/0 is the physical interface. Three logical interfaces are configured:

- Two logical interfaces (Pp0.100 and Demux0.100) reside on the underlying logical interface ge-1/0/0.100.
- A third logical interface (Pp0.101) resides on the underlying logical interface ge-1/0/0.101.

**Figure 12: Logical Interfaces at Level 3 and Underlying Logical Interfaces at Level 2**



Each logical interface has a dedicated queue. The CoS scheduler nodes at level 1 (physical interface), level 2 (underlying logical interfaces), and level 3 (logical interfaces) form a scheduling hierarchy.

To configure this scenario, you must include the **implicit-hierarchy** option under the **hierarchical-scheduler** statement on physical interface GE-1/0/0 and configure and apply traffic-control profiles on each logical interface and underlying logical interface.

You can configure many different three-level scheduling hierarchies; Figure 11 on page 74 and Figure 12 on page 75 present just two possible scenarios. Table 12 on page 72 summarizes the possible interface locations and CoS scheduler nodes.

## Four-Level Hierarchical Scheduling

Beginning with Junos OS Release 16.1, four-level hierarchical scheduling is supported on MX Series routers running MPC5E 3D Q interfaces. Four-level hierarchical scheduling

supports up to eight class of service queues. In four-level scheduling hierarchies, the physical interface is a level 1 CoS scheduler node and the queues reside at level 5.



**NOTE:** Four-level hierarchical scheduling is not supported agent circuit identifier (ACI) or aggregated Ethernet (AE) interfaces.

When you use four-level hierarchical scheduling, interface sets reside at either level 2 and logical interfaces reside at levels 3 and 4. [Table 13 on page 76](#) summarizes the most common case of the interface hierarchy and the CoS scheduler node levels for four-level hierarchical scheduling.

**Table 13: Four-Level Hierarchical Scheduling—Interface Hierarchy Versus CoS Scheduling Node Levels**

Level 1	Level 2	Level 3	Level 4	Level 5
Physical interface	Interface set	Customer VLAN (C-VLAN)	Session Logical Interface (ppp or dhcp)	One or more queues

In four-level hierarchical scheduling, the CoS scheduler nodes at level 1, level 2, level 3, and level 4 form a hierarchical relationship.

To configure four-level hierarchical scheduling, include the **implicit-hierarchy** option at the **[edit interfaces *interface-name* hierarchical-scheduler]** hierarchy level and set the **maximum-hierarchy-levels** option to 4.

```
[edit interfaces]
interface-name {
  hierarchical-scheduler {
    implicit-hierarchy;
    maximum-hierarchy-levels 4;
  }
}
```

#### Release History Table

Release	Description
16.1	Beginning with Junos OS Release 16.1, four-level hierarchical scheduling is supported on MX Series routers running MPC5E 3D Q interfaces.

#### Related Documentation

- [Hardware Requirements for Dynamic Hierarchical CoS on page 77](#)
- [Configuring Hierarchical Schedulers for CoS on page 12](#)
- [Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links on page 79](#)
- [Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface on page 80](#)
- [CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 123](#)

- [hierarchical-scheduler \(Subscriber Interfaces on MX Series Routers\) on page 311](#)

## Hardware Requirements for Dynamic Hierarchical CoS

Table 14 on page 77 lists the hardware requirements based on subscriber interface type for hierarchical scheduling in dynamic CoS configurations.

**Table 14: Hardware Required for Dynamic Hierarchical CoS Configurations**

Dynamic CoS Configuration	Subscriber Interface Type	EQ DPCs on MX Series Routers	MPC Q/MIC Modules on MX Series Routers
Hierarchical CoS	Static and dynamic VLANs	Yes	Yes
	Static and dynamic VLANs over aggregated Ethernet	Yes	Yes
	Static or dynamic IP demux interfaces	Yes	Yes
	Static or dynamic IP demux interfaces over aggregated Ethernet	Yes	Yes
	Static or dynamic VLAN demux interfaces	No	Yes
	Static or dynamic VLAN demux interfaces over aggregated Ethernet	No	Yes
	Static PPPoE interfaces	No	Yes
	Dynamic PPPoE interfaces	No	Yes
	Static or dynamic PPPoE interfaces over aggregated Ethernet	No	Yes
	L2TP LAC tunnel over PPP	No	Yes
	L2TP LNS inline service over PPP	No	Yes

### Related Documentation

- [Understanding Hierarchical CoS for Subscriber Interfaces on page 70](#)
- [Guidelines for Configuring Dynamic CoS for Subscriber Access](#)

## Configuring Static Hierarchical Scheduling in a Dynamic Profile

---

You configure static scheduling and queuing in a dynamic profile for subscriber access. To configure CoS in a dynamic profile for subscriber access using static scheduling and queuing parameters:

1. Configure the static CoS parameters in the **[edit class-of-service]** hierarchy.

- a. Enable the hierarchical scheduler for the interface.

See [“Understanding Hierarchical CoS for Subscriber Interfaces” on page 70](#).

- b. Configure the scheduler map and schedulers.

When you configure static scheduling and queuing in a dynamic profile, you reference the scheduler map in the dynamic profile.

See *Configuring Schedulers*.

- c. Configure the drop profiles.

See *Defining Packet Drop Behavior by Configuring RED Drop Profiles*.

- d. Configure the forwarding classes.

See *Configuring a Custom Forwarding Class for Each Queue*.

- e. Configure the rewrite-rules and classifier definitions.

See *Configuring Rewrite Rules* and *Configuring Behavior Aggregate Classifiers*.

See *The Junos OS CoS Components Used to Manage Congestion and Control Service Levels* for information about configuring the remaining CoS parameters.

2. Configure a static or dynamic subscriber interface that can be referenced in the dynamic profile.

3. Configure CoS parameters in a dynamic profile.

- a. Configure the dynamic profile.

See *Configuring a Basic Dynamic Profile*.

- b. Configure traffic shaping and scheduling parameters in the dynamic profile using a traffic-control profile. Reference the scheduler map you configured in the static **[edit class-of-service]** hierarchy.

See *Configuring Static Traffic Shaping and Scheduling Parameters in a Dynamic Profile*.

- c. Apply CoS parameters to a subscriber interface by referencing an interface in the dynamic profile.

See *Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile*.

- 4. To configure default values for subscribers on login, and enable subscribers to replace other CoS parameters when replacing services, configure variables in the dynamic profile.

See “[Configuring Static Default Values for Traffic Scheduling and Shaping](#)” on page 194.

#### Related Documentation

- *Guidelines for Configuring Dynamic CoS for Subscriber Access*
- *CoS for Subscriber Access Overview*
- [Example: Maintaining a Constant Traffic Flow by Configuring a Static VLAN Interface with a Dynamic Profile for Subscriber Access](#) on page 81

## Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links

You can enable hierarchical CoS on a subscriber interface with an underlying aggregated Ethernet interface.

Before you begin, configure the subscriber interface with aggregated Ethernet.

- To configure a VLAN interface over aggregated Ethernet with link protection, see *Configuring a Static or Dynamic VLAN Subscriber Interface over Aggregated Ethernet* and *Configuring Link Protection for Aggregated Ethernet Interfaces*.
- To configure a demux subscriber interface:

For static and dynamic IP demux interfaces, see *Configuring a Static or Dynamic IP Demux Subscriber Interface over Aggregated Ethernet*.

For static and dynamic VLAN demux interfaces, see *Configuring a Static or Dynamic VLAN Demux Subscriber Interface over Aggregated Ethernet*.



**BEST PRACTICE:** Link protection is not required for IP or demux subscriber interfaces. We recommend that you enable targeted distribution on the demux interface to provide accurate hierarchical scheduling for these links. See “[Providing Accurate Scheduling for a Demux Subscriber Interface of Aggregated Ethernet Links](#)” on page 166.

To configure hierarchical CoS on the link aggregation (LAG) bundle:

1. Specify that you want to access the LAG bundle.

```
user@host# edit interfaces aex
```

2. Configure the link aggregation (LAG) bundle with hierarchical scheduler mode.

```
[edit interfaces aex]  
user@host# set hierarchical-scheduler
```

You can then attach static or dynamic traffic shaping and scheduling parameters at the aggregated Ethernet logical interface or its underlying physical interface. See:

- *Configuring Traffic Scheduling and Shaping for Subscriber Access*
- *Configuring Schedulers in a Dynamic Profile for Subscriber Access*
- *Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile*

**Related  
Documentation**

- *Guidelines for Configuring Dynamic CoS for Subscriber Access*
- *Verifying the Scheduling and Shaping Configuration for Subscriber Access*
- *CoS for Subscriber Access Overview*
- [Understanding Hierarchical CoS for Subscriber Interfaces on page 70](#)

---

## Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface

---

You can configure hierarchical CoS on a static PPPoE subscriber interface.

Before you begin:

- Configure the static PPPoE subscriber interface.

See *Configuring PPPoE*.

To configure hierarchical CoS on a static PPPoE subscriber interface:

1. Specify the PPPoE interface that you want to configure.

```
user@host# edit interfaces pppoe-interface-name
```

2. Configure the hierarchical scheduler for the interface.

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

3. (Optional) Group the PPPoE interfaces in an interface set.

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

You can now configure static traffic and scheduling parameters for each traffic-control profile, and attach each traffic-control profile to the PPPoE interface or the PPPoE interface set. For more information, see *Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers*.

- Related Documentation**
- For hardware requirements and configuration guidelines, see *Guidelines for Configuring Dynamic CoS for Subscriber Access*
  - *CoS for PPPoE Subscriber Interfaces Overview*
  - *Verifying the Scheduling and Shaping Configuration for Subscriber Access*
  - [Understanding Hierarchical CoS for Subscriber Interfaces on page 70](#)

## Example: Maintaining a Constant Traffic Flow by Configuring a Static VLAN Interface with a Dynamic Profile for Subscriber Access

---

This example shows how to configure a static VLAN interface with a dynamic profile using static schedulers and CoS parameters for subscriber access to maintain a constant traffic flow. The CoS parameters configure a best-effort data service for subscribers.

- [Requirements on page 81](#)
- [Overview on page 81](#)
- [Configuration on page 82](#)
- [Verification on page 92](#)

### Requirements

Before you begin, be sure that your environment meets the following requirements:

- The interface is hosted on an MX Series router.
- For hierarchical scheduling configurations, hierarchical scheduling is enabled in the static CLI for the interface referenced in the dynamic profile. If not, the dynamic profile fails.
- Only one traffic-control-profile is configured under a dynamic profile.
- The output-traffic-control-profile that binds the traffic-control profile to the interface is defined within the same dynamic profile as the interface.

### Overview

In a dynamic profile, you can configure VLAN subscriber interfaces over the following statically created logical interface types:


- GE—Gigabit Ethernet
- XE—10-Gigabit Ethernet
- AE—Aggregated Ethernet

### Topology

---

We recommend that you configure each subscriber on a statically created VLAN.

[Figure 13 on page 82](#) shows an example of subscriber interfaces on an individual VLAN.

**Figure 13: VLAN Subscriber Interfaces**

You can further separate VLANs on subscriber interfaces by configuring a VLAN interface as the underlying interface for a set of IP demux interfaces.

## Configuration

To configure a static VLAN interface with a dynamic profile for subscriber access, perform these tasks:

- [Configuring a Subscriber Interface with a Static VLAN on page 83](#)
- [Associating the Dynamic Profile with a Statically Created Interface on page 85](#)
- [Configuring the Firewall Filter on page 86](#)
- [Configuring Static Schedulers in a Dynamic Profile on page 87](#)
- [Associating the Scheduler with a Scheduler Map on page 89](#)
- [Configuring and Applying Static Traffic Shaping and Scheduling Parameters in a Dynamic Profile on page 90](#)

### CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```

set interfaces ge-2/2/0
set interfaces ge-2/2/0 hierarchical-scheduler
set interfaces ge-2/2/0 vlan-tagging
set interfaces ge-2/2/0 vlan-tagging unit 100 vlan-id 100
set interfaces ge-2/2/0 vlan-tagging unit 100 vlan-id 100 family inet
set interfaces ge-2/2/0 vlan-tagging unit 100 vlan-id 100 family inet unnumbered-address
  lo0.0 preferred-source-address 10.0.0.1
set dynamic-profiles data-service
set dynamic-profiles data-service interfaces $junos-interface-ifd-name
set dynamic-profiles data-service interfaces $junos-interface-ifd-name unit
  $junos-underlying-interface-unit
set dynamic-profiles data-service interfaces $junos-interface-ifd-name unit
  $junos-underlying-interface-unit family inet
set dynamic-profiles data-service firewall family inet filter filter EF_limit_G=768K
set dynamic-profiles data-service firewall family inet filter filter EF_limit_G=768K term
  EF
set dynamic-profiles data-service firewall family inet filter filter EF_limit_G=768K term
  default
set dynamic-profiles data-service firewall family inet filter filter EF_limit_G=768K term
  EF from forwarding-class EF

```



```

set dynamic-profiles data-service firewall family inet filter filter EF_limit_G=768K term
  EF then policer POL_EF_G=768K
set dynamic-profiles data-service firewall family inet filter filter EF_limit_G=768K term
  default then accept
set dynamic-profiles data-service class-of-service schedulers be-scheduler
set dynamic-profiles data-service class-of-service schedulers be-scheduler buffer-size
  remainder
set dynamic-profiles data-service class-of-service schedulers be-scheduler
  drop-profile-map loss-priority any protocol any
set dynamic-profiles data-service class-of-service schedulers be-scheduler
  drop-profile-map loss-priority any protocol any drop-profile drop3
set dynamic-profiles data-service class-of-service schedulers be-scheduler priority low
user@host# set dynamic-profiles data-service class-of-service schedulers be-scheduler
  transmit-rate percent 40
set dynamic-profiles data-service class-of-service schedulers be-scheduler excess-rate
  percent 90
set dynamic-profiles data-service class-of-service schedulers be-scheduler excess-priority
  high
set dynamic-profiles data-service class-of-service scheduler-maps data-service-map
set dynamic-profiles data-service class-of-service scheduler-maps data-service-map
  forwarding-class best-effort
set dynamic-profiles data-service class-of-service scheduler-maps data-service-map
  forwarding-class best-effort scheduler be-scheduler
set dynamic-profiles data-service class-of-service traffic-control-profiles tcp-data-service
set dynamic-profiles data-service class-of-service traffic-control-profiles tcp-data-service
  scheduler-map data-service-map
set dynamic-profiles data-service class-of-service traffic-control-profiles tcp-data-service
  shaping-rate 50k
set dynamic-profiles data-service class-of-service traffic-control-profiles tcp-data-service
  guaranteed-rate 10k
set dynamic-profiles data-service class-of-service traffic-control-profiles tcp-data-service
  delay-buffer-rate 10k
set dynamic-profiles data-service class-of-service interfaces $junos-interface-ifd-name
  unit $junos-underlying-interface-unit output-traffic-control-profile tcp-data-service

```

### Configuring a Subscriber Interface with a Static VLAN

#### Step-by-Step Procedure

After you configure a static VLAN interface, you can reference it in a dynamic profile.

1. Configure the static VLAN interface.  

```

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

```
2. Enable hierarchical scheduling for the interface.  

```

[edit interfaces ge-2/2/0]
user@host# set hierarchical-scheduler

```
3. Enable VLAN tagging.  

```

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

```

4. Configure the unit and assign a VLAN ID.

```
[edit interfaces ge-2/2/0 vlan-tagging]  
user@host# set unit 100 vlan-id 100
```

5. Define the family address type (inet for IPv4) for the VLAN interface.

```
[edit interfaces ge-2/2/0 vlan-tagging unit 100 vlan-id 100]  
user@host# set family inet
```

6. Enable the physical interface to borrow an IP address from the loopback interface by setting an unnumbered interface address. Configure a secondary IP address on the loopback interface, lo0.0, and configure it as the preferred source address.

```
[edit interfaces ge-2/2/0 vlan-tagging unit 100 vlan-id 100 family inet]  
user@host# set unnumbered-address lo0.0 preferred-source-address 10.0.0.1
```

**Results** Confirm the configuration of the static VLAN interface by entering the **show interfaces** configuration command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]  
user@host# show interfaces  
interfaces {  
  ge-2/2/0 {  
    hierarchical-scheduler;  
    vlan-tagging;  
    unit 100 {  
      vlan-id 100;  
      family inet {  
        unnumbered-address lo0.0 preferred-source-address 10.0.0.1;  
      }  
    }  
  }  
}
```

### Associating the Dynamic Profile with a Statically Created Interface

#### Step-by-Step Procedure

A dynamic profile is a set of characteristics, defined in a type of template, that you can use to provide dynamic subscriber access and services for broadband applications. When configuring the interface at the **[dynamic-profiles *profile-name* interfaces]** hierarchy level for a dynamic profile, you use variables to specify the interface name and the logical unit value. When a DHCP subscriber sends a DHCP request to the interface, the dynamic profile replaces the interface name variable and logical unit name variable with the actual interface name and logical unit number of the interface that received the DHCP request.



**NOTE:** Configuration of the interface name variable and logical interface name variable at the **[edit dynamic-profiles *profile-name* interfaces]** hierarchy level is required for a dynamic profile to function.

1. Create the new dynamic profile for data services for subscribers.

```
[edit]
user@host# set dynamic-profiles data-service
```

2. Define the ***interface-name*** variable 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 data-service]
user@host# set interfaces $junos-interface-ifd-name
```

3. Define the **unit** statement with the internal variable.

- When referencing an existing interface, specify the **\$junos-underlying-interface-unit** variable used by the router to match the unit value of the receiving interface.
- When creating dynamic interfaces, specify the **\$junos-interface-unit** variable used by the router to generate a unit value for the interface.

```
[edit dynamic-profiles data-service interfaces $junos-interface-ifd-name]
user@host# set unit $junos-underlying-interface-unit
```

or

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

4. Define the family address type (inet for IPv4) for the **\$junos-interface-unit** variable.

```
[edit dynamic-profiles data-service interfaces $junos-interface-ifd-name unit
$junos-underlying-interface-unit]
user@host# set family inet
```

**Results** Confirm the configuration of the dynamic profile by entering the **show dynamic-profiles** configuration command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show dynamic-profiles
dynamic-profiles {
  data-service {
    interfaces {
      $junos-interface-ifd-name {
        unit $junos-underlying-interface-unit {
          family inet;
        }
      }
    }
  }
}
```

---

### Configuring the Firewall Filter

---

**Step-by-Step Procedure** To configure a static VLAN interface with a dynamic profile for subscriber access, you can configure a firewall filter to provide enhanced security by blocking packets based on various match criteria, such as subjecting traffic to a policer for rate limiting, assigning the traffic to a class-of-service (CoS) forwarding class for later queuing and packet rewrite operations, or directing traffic to a specific routing instance.

1. Configure the family address type (inet for IPv4) for the firewall filter and specify the filter name.

We recommend that you name the filter something that indicates the filter's purpose. In this example, we use the bandwidth limit settings.

```
[edit dynamic-profiles data-service]
user@host# set firewall family inet filter EF_limit_G=768K
```

2. Specify the term names for the filter. Make each term name unique and represent what its function is. The first term matches traffic that has been classified into the Expedited Forwarding (EF) class, and the second term matches all non-EF traffic.

```
[edit dynamic-profiles data-service firewall family inet filter EF_limit_G=768K]
user@host# set term EF
user@host# set term default
```

3. In each firewall filter term, specify the conditions used to match components of a packet. Configure the first term to match all traffic classified as EF class.

```
[edit dynamic-profiles data-service firewall family inet filter EF_limit_G=768K term EF]
user@host# set from forwarding-class EF
```

4. Specify the actions to take when the packet matches the condition in the first term. Send the EF traffic to the policer named **POL\_EF\_G=768K**.

```
[edit dynamic-profiles data-service firewall family inet filter EF_limit_G=768K term
EF]
user@host# set then policer POL_EF_G=768K
```

- Specify the action to take when the packet matches the condition in the second term. All non-EF packet traffic is accepted.

```
[edit dynamic-profiles data-service firewall family inet filter EF_limit_G=768K term
default]
user@host# set then accept
```

**Results** Confirm the configuration by entering the **show dynamic-profiles data-service firewall** configuration command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show dynamic-profiles data-service firewall
family inet {
  filter EF_limit_G=768K {
    term EF {
      from {
        forwarding-class EF;
      }
      then policer POL_EF_G=768K;
    }
    term default {
      then accept;
    }
  }
}
```

### Configuring Static Schedulers in a Dynamic Profile

**Step-by-Step Procedure** You can configure static scheduling and queuing parameters in a dynamic profile for subscriber access. Schedulers are part of the basic class-of-service (CoS) infrastructure. You must define at least one scheduler per forwarding class. Schedulers indicate a forwarding class's priority, transmit weight, and buffer size, as well as various shaping and rate control mechanisms.

- Specify the best-effort scheduler for which you want to configure parameters.

```
[edit dynamic-profiles data-service class-of-service]
user@host# set schedulers be-scheduler
```



**NOTE:** Set schedulers to the name of the scheduler to be configured or to the Junos OS predefined variable (`$junos-cos-scheduler`) used for dynamic subscriber interfaces. The predefined variable is replaced with the scheduler name obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

2. (Optional) Configure the buffer size to use the remaining buffer available.

This parameter allows you to specify an explicit buffer size, either as a percent of interface speed or as a function of time (specified in microseconds).

```
[edit dynamic-profiles data-service class-of-service schedulers be-scheduler]
user@host# set buffer-size remainder
```

3. (Optional) Configure the drop-profile map to associate one or more drop profiles with a queue.

The default random early detection (RED) drop profile is used when no explicit drop profile mapping is specified. Specify a packet-loss priority (PLP) level of any, and for the specified scheduler to accept any protocol type.

```
[edit dynamic-profiles data-service class-of-service schedulers be-scheduler]
user@host# set drop-profile-map loss-priority any protocol any
```

4. (Optional) Configure the drop profile to map a fill level (fullness of a queue) to a drop probability (probability that a packet is dropped).

```
[edit dynamic-profiles data-service class-of-service schedulers be-scheduler
drop-profile-map loss-priority any protocol any]
user@host# set drop-profile drop3
```

You enable RED by applying a drop profile to a scheduler.

5. (Optional) Configure the queue's scheduler priority to a specific level (low) for guaranteed rate traffic.

```
[edit dynamic-profiles data-service class-of-service schedulers be-scheduler]
user@host# set priority low
```

6. (Optional) Configure the queue's transmit weight [in bits per second (bps)] or as a percentage of transmission capacity.

```
[edit dynamic-profiles data-service class-of-service schedulers be-scheduler]
user@host# set transmit-rate percent 40
```

The transmit rate guarantees the rate for the queue, assuming no priority-based starvation occurs. When you do not specify a transmit weight, or when the transmit rate is reached, the queue can only send excess-rate traffic because that queue's priority is demoted to the excess region. A percentage of zero (0) drops all packets in the queue.

7. (Optional) Configure the queue's weight as either a percentage, or a proportion, for any unused bandwidth traffic to share.

```
[edit dynamic-profiles data-service class-of-service schedulers be-scheduler]
user@host# set excess-rate percent 90
```

Behavior varies based on interface mode, explicit configuration, and whether any other queues have explicit weight configured. By default, excess bandwidth between the guaranteed and shaped rate is shared equally among queues.

8. (Optional) Configure the priority of how excess bandwidth traffic is sent on a scheduler in a dynamic profile.

```
[edit dynamic-profiles data-service class-of-service schedulers be-scheduler]
user@host# set excess-priority high
```

To prevent the queue from sending any excess rate traffic, set to none.

**Results** Confirm the configuration of the scheduler with static values in the dynamic profile by entering the **show dynamic-profiles data-service class-of-service** configuration command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show dynamic-profiles data-service class-of-service
class-of-service {
  schedulers {
    be-scheduler {
      buffer-size remainder;
      drop-profile-map loss-priority any protocol any drop-profile drop3;
      priority low;
      transmit-rate percent 40;
      excess-rate percent 90;
      excess-priority high;
    }
  }
}
```

### Associating the Scheduler with a Scheduler Map

**Step-by-Step Procedure** After you define your schedulers, you must link them to a set of queues on a logical interface using a scheduler map. Applying a scheduler map to an interface places the related set of schedulers and drop profiles into effect.

1. Configure the scheduler map name.

```
[edit dynamic-profiles data-service class-of-service]
user@host# set scheduler-maps data-service-map
```

2. Configure a forwarding class to associate a scheduler with a scheduler map.

```
[edit dynamic-profiles data-service class-of-service scheduler-maps
data-service-map]
user@host# set forwarding-class best-effort
```

3. Associate the scheduler you previously defined (**be-scheduler**) with the scheduler map.

```
[edit dynamic-profiles data-service class-of-service scheduler-maps
data-service-map forwarding-class best-effort]
user@host# set scheduler be-scheduler
```

**Results** Confirm the configuration of the scheduler map by entering the **show dynamic-profiles data-service class-of-service scheduler-maps** configuration command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show dynamic-profiles data-service class-of-service scheduler-maps
scheduler-maps {
  data-service-map {
    forwarding-class best-effort scheduler be-scheduler;
  }
}
```

### Configuring and Applying Static Traffic Shaping and Scheduling Parameters in a Dynamic Profile

---

**Step-by-Step Procedure** Configure static traffic shaping and scheduling parameters in a traffic-control profile. A traffic-control profile is a generic class-of-service (CoS) container that you can apply at all points of a CoS hierarchy to affect the committed information rate (CIR), peak information rate (PIR), and excess bandwidth handling. You can specify the traffic-control profile at the port, logical interface, or logical interface-set level. The traffic-control profile also references the scheduler map.

1. Create the traffic-control profile and assign it a name.

```
[edit dynamic-profiles data-service class-of-service]
user@host# edit traffic-control-profiles tcp-data-service
```

2. Apply the static scheduler map, **data-service-map**, that you previously configured.

```
[edit dynamic-profiles data-service class-of-service traffic-control-profiles
tcp-data-service]
user@host# set scheduler-map data-service-map
```

3. Configure the shaping rate [in bits per second (bps)] to use for the scheduler in the dynamic profile.

```
[edit dynamic-profiles data-service class-of-service traffic-control-profiles
tcp-data-service]
user@host# set shaping-rate 50k
```

The shaping rate places a maximum limit on a queue's transmit capacity. By default, the shaping rate is equal to the interface speed/shaping rate enabling the queue to send at the full rate of the interface.

4. Configure the guaranteed rate [in bits per second (bps)] to use for the scheduler in the dynamic profile.

```
[edit dynamic-profiles data-service class-of-service traffic-control-profiles
tcp-data-service]
user@host# set guaranteed-rate 10k
```



The guaranteed rate is the minimum bandwidth the queue can receive; if excess physical interface bandwidth is available for use, the logical interface can receive more than the guaranteed rate provisioned for the interface, depending on how you choose to manage excess bandwidth and the interface's mode of PIR compared to CIR/PIR.

5. Configure the delay-buffer rate [in bits per second (bps)] based on the delay-buffer calculation.

```
[edit dynamic-profiles data-service class-of-service traffic-control-profiles
tcp-data-service]
user@host# set delay-buffer-rate 10k
```

The delay buffer rate setting at one level of the hierarchy becomes the reference bandwidth used at the next higher level, and the sum of the reference bandwidth cannot exceed the value used at a lower level. If you do not include this statement, the delay-buffer rate is based on the guaranteed rate if one is configured, or on the shaping rate if no guaranteed rate is configured.

6. After you configure the traffic shaping and scheduling CoS parameters in a dynamic profile, you apply them to an interface. The output traffic-control profile enables you to provide traffic scheduling to the interface.

Configure the interface name and logical interface using a variable, and apply the output traffic-control profile to the interface. Specify the previously defined traffic-control profile, **tcp-data-service**.

```
[edit dynamic-profiles data-service class-of-service]
user@host# set interfaces $junos-interface-ifd-name unit
$junos-underlying-interface-unit output-traffic-control-profile tcp-data-service
```

**Results** Confirm the configuration and application of the static traffic shaping and scheduling parameters by entering the **show dynamic-profiles** configuration command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show dynamic-profiles
dynamic-profiles {
  data-service {
    class-of-service {
      interfaces {
        $junos-interface-ifd-name {
          unit $junos-underlying-interface-unit {
            output-traffic-control-profile tcp-data-service;
          }
        }
      }
    }
  }
  traffic-control-profiles {
    tcp-data-service {
      scheduler-map data-service-map;
      shaping-rate 50k;
    }
  }
}
```

```
        guaranteed-rate 10k;
        delay-buffer-rate 10k;
    }
}
}
```

## Verification

Confirm that the configuration is working properly.

- [Verifying Traffic Shaping and Scheduling Profiles for Subscriber Access on page 92](#)
- [Verifying the Mapping of Schedulers for Subscriber Access on page 92](#)

---

### Verifying Traffic Shaping and Scheduling Profiles for Subscriber Access

**Purpose** View the class-of-service (CoS) configurations that are referenced in a dynamic profile for subscriber access.

**Action** `user@host> show class-of-service traffic-control-profile`  
Traffic control profile: tcp-data-service, Index: 57625  
Shaping rate: 50000  
Scheduler map: data-service-map  
Delay Buffer rate: 10000  
Guaranteed rate: 10000

**Meaning** The Shaping rate, Delay Buffer rate, and Guaranteed rate fields indicate rates of 50,000 bps, 10,000 bps, and 10,000 bps, respectively, for the traffic-control profile.

---

### Verifying the Mapping of Schedulers for Subscriber Access

**Purpose** Display the mapping of schedulers to forwarding classes and a summary of scheduler parameters for each entry.

**Action** `user@host> show class-of-service scheduler-map`  
Scheduler map: data-service-map, Index: 84  
  
Scheduler: be-scheduler, Index: 8721, Forwarding class: best-effort  
Transmit rate: 40 percent, Rate Limit: none, Maximum buffer delay: 39 ms,  
Priority: low  
Drop profiles:  
Loss priority      Protocol      Index      Name  
Any                Any                8724      drop3

**Meaning** The Scheduler map field indicates the parameters are for the best-effort scheduler. The Transmit rate field shows 40 percent; the Rate Limit field indicates no limit; and the Drop profiles fields are for drop3.

- Related Documentation**
- *CoS for Subscriber Access Overview*
  - *Guidelines for Configuring Dynamic CoS for Subscriber Access*
  - [Understanding Hierarchical CoS for Subscriber Interfaces on page 70](#)



## CHAPTER 4

# Applying CoS to Groups of Subscriber Interfaces

- [CoS for Interface Sets of Subscribers Overview on page 95](#)
- [Configuring an Interface Set of Subscribers in a Dynamic Profile on page 98](#)
- [Example: Configuring a Dynamic Interface Set of VLAN Subscribers on page 99](#)
- [Example: Configuring a Dynamic Service VLAN Interface Set of Subscribers in a Dynamic Profile on page 112](#)

## CoS for Interface Sets of Subscribers Overview

---

Interface sets enable service providers to group logical interfaces so they can apply CoS parameters to all of the traffic in the group.

Interface sets are beneficial for various scenarios in a subscriber access network. For example, you can use an interface set to configure a local loop with a small number of subscribers. Interface sets are also useful for grouping a large number of subscribers into a particular service class or for defining traffic engineering aggregates for DSLAMs.

- [Guidelines for Configuring Dynamic Interface Sets in a Subscriber Access Network on page 96](#)

## Guidelines for Configuring Dynamic Interface Sets in a Subscriber Access Network

Interface sets enable service providers to group logical interfaces so they can apply CoS parameters to all of the traffic in the group.

Interface sets are beneficial for various scenarios in a subscriber access network. For example, you can use an interface set to configure a local loop with a small number of subscribers. Interface sets are also useful for grouping a large number of subscribers into a particular service class or for defining traffic engineering aggregates for DSLAMs.

When configuring interface sets for subscriber access, keep the following guidelines in mind:

- You can configure interface sets of VLAN demux, PPPoE, or demux interfaces over aggregated Ethernet interfaces.
- An interface can only belong to one interface set. If you try to add the same interface to different interface sets, the commit operation fails.
- You configure the interface set and the traffic scheduling and shaping parameters in a dynamic profile. However, you must apply the traffic-control profile to the interface set in the static **[edit class-of-service]** hierarchy.



**NOTE:** This rule applies to all interface sets except ACI sets.

---

- The **\$junos-interface-set-name** predefined variable is available only for RADIUS Accept messages; change of authorization (CoA) requests are not supported.
- The **\$junos-svlan-interface-set-name** predefined variable locally generates an 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 **physical\_interface\_name - outer\_VLAN\_tag**. For example, an aggregated Ethernet interface “ae0,” with a dual-tagged VLAN interface that has an outer tag of “111,” results in a **\$junos-svlan-interface-set-name** dynamic variable of “ae0-111”. Similarly, a non-aggregated Ethernet interface of ge-1/1/0, with the same dual-tagged VLAN interface that has an outer tag of “111,” results in a **\$junos-svlan-interface-set-name** dynamic variable of “ge-1/1/0-111”.
- The **\$junos-phy-ifd-interface-set-name** predefined variable locally generates an interface set name associated with the underlying physical interface in a dynamic profile. This predefined variable enables you to group all the subscribers on a specific physical interface so that you can apply services to the entire group of subscribers.

Another use case for this predefined variable is to conserve CoS resources in a mixed business and residential topology 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. Otherwise, because the business and residential subscribers share the same interface and business subscribers require three levels of CoS, then three levels are configured for each residential subscriber. That results in an unnecessary level 2 node being consumed for each residential connection, wasting CoS resources.

- The **\$junos-tagged-vlan-interface-set-name** predefined variable locally generates an interface set name used for grouping logical interfaces stacked over logical stacked VLAN demux interfaces for either a 1:1 (dual-tagged; individual client) VLAN or N:1 (single tagged; service) VLAN. The format of the generated variable differs with VLAN type as follows:
  - Dual-tagged (client) VLAN—***physical\_interface\_name - outer\_VLAN\_tag - inner\_VLAN\_tag***. For example, an aggregated Ethernet interface “ae0,” with a dual-tagged VLAN interface that has an outer tag of “111” and an inner tag of “200,” results in a **\$junos-tagged-vlan-interface-set-name** dynamic variable of “ae0-200-111”. Similarly, a non-aggregated Ethernet interface of ge-1/1/0, with the same dual-tagged VLAN interface that has an outer tag of “111” and an inner tag of “200,” results in a **\$junos-tagged-vlan-interface-set-name** dynamic variable of “ge-1/1/0-200-111”.
  - Single tagged (service) VLAN—***physical\_interface\_name - VLAN\_tag***. For example, an aggregated Ethernet interface “ae0,” with an N:1 VLAN using the single tag of “200,” results in a **\$junos-tagged-vlan-interface-set-name** dynamic variable of “ae0-200”. Similarly, a non-aggregated Ethernet interface of ge-1/1/0, with the same N:1 VLAN using the single tag of “200,” results in a **\$junos-tagged-vlan-interface-set-name** dynamic variable of “ge-1/1/0-200”.
- All dynamic demux, dual-tagged VLAN logical interfaces with the same outer VLAN tag and physical interface are assigned to the same interface set and all CoS values provisioned with the dynamic profile are applied to the interfaces that are part of the set.
- The interface set name must be explicitly referenced in the CoS configuration as part of the static configuration outside of the dynamic profile. The CoS configuration is static and the interface set name must be statically referenced.



**NOTE:** This rule applies to all interface sets except ACI sets.

- RADIUS can return an *access-accept* message under certain conditions. A configured RADIUS VSA for the interface set name takes precedence over the locally generated variable on the router. This means that if the interface-set-name VSA is configured on RADIUS, the router continues to use this variable instead of the locally generated value from the dynamic variable.
- Sets of aggregated Ethernet interfaces are supported on MPC/MIC interfaces on MX Series routers only.
- The supported interface stacks for aggregated Ethernet in an interface set include VLAN demux interfaces, IP demux interfaces, and PPPoE logical interfaces over VLAN demux interfaces.
- The link membership list and scheduler mode of the interface set are inherited from the underlying aggregated Ethernet interface over which the interface set is configured.

- When an aggregated Ethernet interface operates in link protection mode, or if the scheduler mode is configured to replicate member links, the scheduling parameters of the interface set are copied to each of the member links.
- If the scheduler mode of the aggregated Ethernet interface is set to scale member links, the scheduling parameters are scaled based on the number of active member links and applied to each of the aggregated interface member links.

**Related  
Documentation**

- [Configuring an Interface Set of Subscribers in a Dynamic Profile on page 98](#)
- [Example: Configuring a Dynamic Service VLAN Interface Set of Subscribers in a Dynamic Profile on page 112](#)

---

## Configuring an Interface Set of Subscribers in a Dynamic Profile

Interface sets enable you to provide hierarchical scheduling to a group of subscriber interfaces.

Before you begin, configure the subscriber interfaces that you intend to include in the interface set.

To configure an interface set of subscriber interfaces:

1. Configure the interface set in the dynamic profile.

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

Replacing the *interface-set-name* variable with the `$junos-interface-set-name`, `$junos-svlan-interface-set-name`, or `$junos-tagged-vlan-interface-set-name` predefined variable. The interface set is created dynamically when the subscriber logs in.

2. Include the interfaces within the dynamic interface-set.

```
[edit dynamic-profiles profile-name interfaces interface-set $junos-interface-set-name]  
user@host# set interface interface-name unit logical-unit-number
```

3. Apply traffic shaping and queuing parameters to the interface set.



**TIP:** You must configure the interface set in the static [edit class-of-service] hierarchy, not in the [edit dynamic-profiles] hierarchy.

```
[edit class-of-service interfaces]  
user@host# edit interface-set interface-set-name  
[edit class-of-service interfaces interface-set interface-set-name]  
user@host# set output-traffic-control-profile profile-name
```

**Related  
Documentation**

- [CoS for Interface Sets of Subscribers Overview on page 95](#)
- [Guidelines for Configuring Dynamic CoS for Subscriber Access](#)



- [CoS for Interface Sets of Subscribers Overview on page 95](#)
- [Example: Configuring a Dynamic Interface Set of VLAN Subscribers on page 99](#)
- [CoS for Aggregated Ethernet Subscriber Interfaces Overview](#)

## Example: Configuring a Dynamic Interface Set of VLAN Subscribers

- [Requirements on page 99](#)
- [Overview on page 99](#)
- [Configuring the Dynamic VLANs on page 99](#)
- [Configuring Dynamic Traffic Scheduling and Shaping on page 101](#)
- [Configuring the Interface Set in the Dynamic Profile on page 105](#)
- [Configuring DHCP Access on page 106](#)
- [Configuring RADIUS Authentication on page 107](#)
- [Verification on page 112](#)

### Requirements

This example uses the following software and hardware components:

- MX Series Router with MPCs

### Overview

In this example, the network administrator groups dynamic VLAN interfaces in an interface set. The interface set is configured in a dynamic profile, and enables hierarchical scheduling for the VLAN interfaces for a multiplay service.

DHCP is used as the access method, and RADIUS is used as the authentication method for the interfaces associated with the interface set.

### Configuring the Dynamic VLANs

**CLI Quick Configuration** To quickly configure the dynamic VLANs, copy the following commands and paste them into the router terminal window:

```
[edit]
edit dynamic-profiles vlan-prof
edit interfaces $junos-interface-ifd-name unit $junos-interface-unit
set vlan-id $junos-vlan-id
set demux-source inet
set family inet unnumbered-address lo0.0 preferred-source-address 203.0.113.32
top
edit interfaces ge-1/0/0
set hierarchical-scheduler
set vlan-tagging
edit auto-configure vlan-ranges dynamic-profile vlan-prof
set ranges any
set accept inet
```

```
top
set interfaces lo0 unit 0 family inet address 203.0.113.32/32
```

---

### Configuring the Dynamic Profile for the Autoconfigured VLANs

---

**Step-by-Step Procedure** In this section, you create a dynamic profile for the VLAN IDs to be automatically assigned when subscribers log in.

To configure the dynamic profile for the VLANs:

1. Configure the dynamic profile.

```
[edit]
user@host#edit dynamic-profile vlan-prof
```

2. Configure the interfaces.

```
[edit dynamic-profiles vlan-prof]
user@host#edit interfaces $junos-interface-ifd-name unit $junos-interface-unit
```

3. Add the VLAN ID variable.

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

4. Configure the demux source as IPv4.

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

5. Configure the family.

```
[edit dynamic-profiles vlan-prof interfaces $junos-interface-ifd-name unit
$junos-interface-unit]
user@host#set family inet unnumbered-address lo0.0 preferred-source-address
203.0.113.32
```

---

### Configuring the VLAN Interfaces

---

**Step-by-Step Procedure** To configure the VLAN interfaces:

1. Create the VLAN interface.

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

2. Enable hierarchical scheduling.

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

3. Configure VLAN tagging.  

```
[edit interfaces ge-1/0/0]  
user@host# set vlan-tagging
```
4. Configure auto-configuration for the dynamic profile.  

```
[edit interfaces ge-1/0/0]  
user@host# edit auto-configure vlan-ranges dynamic-profile vlan-prof
```
5. Configure any VLAN ID range.  

```
[edit interfaces ge-1/0/0 auto-configure vlan-ranges dynamic-profile vlan-prof]  
user@host# set ranges any
```
6. Specify IPv4 traffic for the VLAN.  

```
[edit interfaces ge-1/0/0 auto-configure vlan-ranges dynamic-profile vlan-prof]  
user@host# set accept inet
```

---

### Configuring the Loopback Interface

#### Step-by-Step Procedure

To configure the loopback interface:

1. Create the loopback interface.  

```
[edit]  
user@host# edit interfaces lo0
```
2. Configure the unit and the family.  

```
[edit interfaces lo0]  
user@host# set unit 0 family inet address 203.0.113.32/32
```

### Configuring Dynamic Traffic Scheduling and Shaping

#### CLI Quick Configuration

To quickly configure the traffic scheduling and shaping parameters, copy the following commands and paste them into the router terminal window:

```
[edit]  
edit dynamic-profiles multiplay class-of-service schedulers be_sch  
set transmit-rate percent 12  
set buffer-size percent 12  
set priority low  
up  
edit ef_sch  
set transmit-rate percent 12  
set buffer-size percent 12  
set priority low  
up  
edit af_sch  
set transmit-rate percent 12
```

```
set buffer-size percent 12
set priority low
up
edit nc_sch
set transmit-rate percent 12
set buffer-size percent 12
set priority low
up
edit voice_sch
set transmit-rate percent 12
set buffer-size percent 12
set priority low
up
edit video_sch
set transmit-rate percent 12
set buffer-size percent 12
set priority low
up
edit game_sch
set transmit-rate percent 12
set buffer-size percent 12
set priority low
up
edit data_sch
set transmit-rate percent 12
set buffer-size percent 12
set priority low
up 2
edit scheduler-maps all_smap
set forwarding-class be scheduler be_sch
set forwarding-class ef scheduler ef_sch
set forwarding-class af scheduler af_sch
set forwarding-class nc scheduler nc_sch
set forwarding-class voice scheduler voice_sch
set forwarding-class video scheduler video_sch
set forwarding-class game scheduler game_sch
set forwarding-class data scheduler data_sch
up 2
edit traffic-control-profiles multiplay
set scheduler-map all_smap
set shaping-rate 100m
set guaranteed-rate 20m
```

---

### Configuring the Schedulers in the Dynamic Profile

**Step-by-Step Procedure** In this section, you create a dynamic profile for the multiplay service and configure scheduling and shaping.

To configure the schedulers:

1. Create the **multiplay** dynamic profile.

```
[edit]
user@host# edit dynamic-profiles multiplay class-of-service schedulers
```

2. Configure the best effort scheduler.  

```
[edit dynamic-profiles multiplay class-of-service schedulers]
user@host# edit be_sch
user@host# set transmit-rate percent 12
user@host# set buffer-size percent 12
user@host# set priority low
```
3. Configure the expedited forwarding scheduler.  

```
[edit dynamic-profiles multiplay class-of-service schedulers]
user@host# edit ef_sch
user@host# set transmit-rate percent 12
user@host# set buffer-size percent 12
user@host# set priority low
```
4. Configure the assured forwarding scheduler.  

```
[edit dynamic-profiles multiplay class-of-service schedulers]
user@host# edit af_sch
user@host# set transmit-rate percent 12
user@host# set buffer-size percent 12
user@host# set priority low
```
5. Configure the network control scheduler.  

```
[edit dynamic-profiles multiplay class-of-service schedulers]
user@host# edit nc_sch
user@host# set transmit-rate percent 12
user@host# set buffer-size percent 12
user@host# set priority low
```
6. Configure the voice scheduler.  

```
[edit dynamic-profiles multiplay class-of-service schedulers]
user@host# edit voice_sch
user@host# set transmit-rate percent 12
user@host# set buffer-size percent 12
user@host# set priority low
```
7. Configure the video scheduler.  

```
[edit dynamic-profiles multiplay class-of-service schedulers]
user@host# edit video_sch
user@host# set transmit-rate percent 12
user@host# set buffer-size percent 12
user@host# set priority low
```
8. Configure the gaming scheduler.  

```
[edit dynamic-profiles multiplay class-of-service schedulers]
user@host# edit game_sch
user@host# set transmit-rate percent 12
user@host# set buffer-size percent 12
user@host# set priority low
```

9. Configure the data scheduler.

```
[edit dynamic-profiles multiplay class-of-service schedulers]
user@host# edit data_sch
user@host# set transmit-rate percent 12
user@host# set buffer-size percent 12
user@host# set priority low
```

---

### Configuring the Scheduler Map in the Dynamic Profile

---

#### Step-by-Step Procedure

To configure the scheduler map:

1. Configure the scheduler map for all of the services.

```
[edit dynamic-profiles multiplay class-of-service]
user@host# edit scheduler-maps all_smap
```

2. Configure the forwarding classes for each service in the scheduler map.

```
[edit dynamic-profiles multiplay class-of-service scheduler-maps all_smap]
user@host# set forwarding-class be scheduler be_sch
user@host# set forwarding-class ef scheduler ef_sch
user@host# set forwarding-class af scheduler af_sch
user@host# set forwarding-class nc scheduler nc_sch
user@host# set forwarding-class voice scheduler voice_sch
user@host# set forwarding-class video scheduler video_sch
user@host# set forwarding-class game scheduler game_sch
user@host# set forwarding-class data scheduler data_sch
```

---

### Configuring the Traffic-Control Profile in the Dynamic Profile

---

#### Step-by-Step Procedure

To configure the traffic-control profile the interface set:

1. Configure the traffic-control profile.

```
[edit dynamic-profiles multiplay class-of-service]
user@host# edit traffic control-profiles multiplay
```

2. Configure the scheduler map.

```
[edit dynamic-profiles multiplay class-of-service traffic control-profiles multiplay]
user@host# set scheduler-map all_smap
```

3. Configure the shaping rate.

```
[edit dynamic-profiles multiplay class-of-service traffic control-profiles multiplay]
user@host# set shaping-rate 100m
```

4. Configure the guaranteed rate.

```
[edit dynamic-profiles multiplay class-of-service traffic control-profiles multiplay]
user@host# set guaranteed-rate 20m
```

## Configuring the Interface Set in the Dynamic Profile

**CLI Quick Configuration** To quickly configure the interface set, copy the following commands and paste them into the router terminal window:

```
[edit]
edit dynamic-profiles multiplay
edit interfaces interface-set $junos-interface-set-name
set interface $junos-interface-ifd-name unit $junos-underlying-interface-unit
top
edit class-of-service interfaces interface-set
set output-traffic-control-profile multiplay
```

### Configuring the Interfaces for the Interface Set

**Step-by-Step Procedure** To configure the interface variable for the interface set:

1. Configure the dynamic profile for the interface set.

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

2. Configure the interface using the Junos OS predefined variable.

```
[edit dynamic-profiles multiplay]
user@host#edit interfaces $junos-interface-ifd-name unit
$junos-underlying-interface-unit
```

3. Configure the family.

```
[edit dynamic-profiles multiplay interfaces $junos-interface-set-name unit
$junos-underlying-interface-unit]
user@host#set family inet unnumbered-address lo0.0 preferred-source-address
203.0.113.32
```

### Configuring the Interface Set

**Step-by-Step Procedure** To configure the interface set:

1. Configure the interface set using the Junos OS predefined variable.

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

2. Add the dynamic VLAN interfaces to the interface set.

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

### Applying the Traffic-Control Profile to the Interface Set

---

- Step-by-Step Procedure** You apply the traffic-control profile outside of the dynamic profile in the **[edit class-of-service]** hierarchy.
- To apply the traffic-control profile:
1. Specify the interface set to which you want to apply the traffic-control profile.  

```
[edit class-of-service]  
user@host#edit interfaces interface-set dynamic-set
```
  2. Attach the output traffic-control profile defined in the dynamic profile to the interface set.  

```
[edit class-of-service interfaces]  
user@host#set output-traffic-control-profile multiplay
```

### Configuring DHCP Access

- CLI Quick Configuration** To quickly configure DHCP access, copy the following commands and paste them into the router terminal window:

```
[edit]  
edit system services dhcp-local-server authentication  
set password $ABC123  
set username-include user-prefix multiplay  
up 1  
set dynamic-profile dhcp-vlan-prof aggregate-clients replace  
set group vlans interface ge-1/0/0  
top  
edit access address-assignment pool v4 family inet  
set network 203.0.113.0/16  
set range limited low 203.0.113.10  
set range limited high 203.0.113.250  
set dhcp-attributes maximum-lease-time 84600
```

### Configuring the DHCP Local Server

---

- Step-by-Step Procedure** To configure DHCP access:
1. Configure the DHCP local server.  

```
[edit system]  
user@host# edit services dhcp-local-server authentication
```
  2. Set the password.  

```
[edit system services dhcp-local-server authentication]  
user@host# set password $ABC123
```
  3. Specify that you want to include optional information in the username.



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

4. Attach the dynamic profile with the interface set.

```
[edit system services dhcp-local-server]
user@host# set dynamic-profile dhcp-vlan-prof aggregate-clients replace
```

5. Configure a group for the VLAN interface.

```
[edit system services dhcp-local-server]
user@host# set group vlans interface ge-1/0/0
```

### Configuring Address Assignment Pools

#### Step-by-Step Procedure

To configure address assignment pools:

1. Configure the pool of IPv4 addresses.

```
[edit access]
user@host#edit address-assignment pool v4 family inet
```

2. Configure the family of interfaces in the pool.

```
[edit access address-assignment pool v4]
user@host#set network 203.0.113.0/16
```

3. Configure the upper and lower bounds of the address range.

```
[edit access address-assignment pool v4]
user@host#set range limited low 203.0.113.10
user@host#set range limited high 203.0.113.250
```

4. Configure the maximum length of time in seconds for which a subscriber can request and hold a lease.

```
[edit access address-assignment pool v4]
user@host#set dhcp-attributes maximum-lease-time 84600
```

### Configuring RADIUS Authentication

#### CLI Quick Configuration

To quickly configure RADIUS authentication, copy the following commands and paste them into the router terminal window:

```
[edit]
edit access radius-server 192.51.100.108
set secret $ABC123ABC123ABC123
set timeout 5
set retry 5
up 2
edit profile acc-prof
```

```
set authentication-order radius
set radius authentication-server 192.51.100.108
```

### Configuring RADIUS Access

---

#### Step-by-Step Procedure

To configure RADIUS access:

1. Configure the RADIUS server.  

```
[edit access]
user@host#edit radius-server 192.51.100.108
```
2. Configure the required secret (password) that the local router or switch passes to the RADIUS client.  

```
[edit access radius-server 192.51.100.108]
user@host# set secret $ABC123ABC123ABC123
```
3. Configure the length of time that the local router or switch waits to receive a response from a RADIUS server.  

```
[edit access radius-server 192.51.100.108]
user@host# set timeout 5
```
4. Configure the number of times that the router or switch attempts to contact a RADIUS accounting server.  

```
[edit access radius-server 192.51.100.108]
user@host# set retry 5
```
5. Configure the access profile.  

```
[edit access]
user@host#edit profile acc-prof
```
6. Configure the authentication order.  

```
[edit access profile acc-prof ]
user@host# set authentication-order radius
```
7. Configure the authentication server.  

```
[edit access profile acc-prof]
user@host#set radius authentication-server 192.51.100.108
```

### Results

---

```
dynamic-profiles {
  vlan-prof {
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-interface-unit" {
```

```

        vlan-id "$junos-vlan-id";
        demux-source inet;
        family inet {
            unnumbered-address lo0.0 preferred-source-address 203.0.113.32;
        }
    }
}
}
multiplay {
    class-of-service {
        traffic-control-profiles {
            multiplay {
                scheduler-map all_smap;
                shaping-rate 100m;
                guaranteed-rate 20m;
            }
        }
    }
    interfaces {
        interface-set "$junos-interface-set-name" {
            interface "$junos-interface-ifd-name" {
                unit "$junos-underlying-interface-unit";
            }
        }
        "$junos-interface-ifd-name" {
            unit "$junos-interface-unit" {
                output-traffic-control-profile multiplay;
            }
        }
    }
}
scheduler-maps {
    all_smap {
        forwarding-class be scheduler be_sch;
        forwarding-class ef scheduler ef_sch;
        forwarding-class af scheduler af_sch;
        forwarding-class nc scheduler nc_sch;
        forwarding-class voice scheduler voice_sch;
        forwarding-class video scheduler video_sch;
        forwarding-class game scheduler game_sch;
        forwarding-class data scheduler data_sch;
    }
}
schedulers {
    be_sch {
        transmit-rate percent 12;
        buffer-size percent 12;
        priority low;
    }
    ef_sch {
        transmit-rate percent 12;
        buffer-size percent 12;
        priority low;
    }
    af_sch {
        transmit-rate percent 12;
        buffer-size percent 12;
    }
}

```

```
        priority low;
    }
    nc_sch {
        transmit-rate percent 12;
        buffer-size percent 12;
        priority low;
    }
    voice_sch {
        transmit-rate percent 12;
        buffer-size percent 12;
        priority low;
    }
    video_sch {
        transmit-rate percent 12;
        buffer-size percent 12;
        priority low;
    }
    game_sch {
        transmit-rate percent 12;
        buffer-size percent 12;
        priority low;
    }
    data_sch {
        transmit-rate percent 12;
        buffer-size percent 12;
        priority low;
    }
}
}
}
access {
    radius-server {
        192.51.100.108 {
            secret "$ABC123ABC123ABC123"; ## SECRET-DATA
            timeout 5;
            retry 5;
        }
    }
    profile acc-prof {
        authentication-order radius;
        radius {
            authentication-server 192.51.100.108;
        }
    }
    address-assignment {
        pool v4 {
            family inet {
                network 203.0.113.0/16;
                range limited {
                    low 203.0.113.10;
                    high 203.0.113.250;
                }
                dhcp-attributes {
                    maximum-lease-time 84600;
                }
            }
        }
    }
}
```

```
}  
}  
}  
class-of-service {  
    interfaces {  
        interface-set dynamic-set {  
            output-traffic-control-profile multiplay;  
        }  
    }  
}  
  
interfaces {  
    interface-set "$junos-interface-set-name" {  
        interface "$junos-interface-ifd-name" {  
            unit "$junos-underlying-interface-unit";  
        }  
    }  
    "$junos-interface-ifd-name" {  
        unit "$junos-underlying-interface-unit" {  
            family inet {  
                unnumbered-address lo0.0 preferred-source-address 203.0.113.32;  
            }  
        }  
    }  
}  
}  
}  
}  
  
}   
  
interfaces {  
    ge-1/0/0 {  
        hierarchical-scheduler;  
        vlan-tagging;  
        auto-configure {  
            vlan-ranges {  
                dynamic-profile vlan-prof {  
                    accept inet;  
                    ranges {  
                        any;  
                    }  
                }  
            }  
        }  
    }  
}  
  
lo0 {  
    unit 0 {  
        family inet {  
            address 203.0.113.32/32;  
        }  
    }  
}  
}  
  
system {  
    services {  
        dhcp-local-server {  
            authentication {  
                password $ABC123;  
                username-include {  
                    user-prefix multiplay;  
                }  
            }  
        }  
    }  
}
```

```
    }  
  }  
  dynamic-profile multiplay aggregate-clients replace;  
  group vlans {  
    interface ge-1/0/0.0;  
  }  
}  
}
```

## Verification

To confirm that the configuration is correct, perform these tasks:

- [Verifying the Interfaces that are Included in the Interface Set on page 112](#)
- [Verifying the Traffic Scheduling and Shaping Parameters for the Interface Set on page 112](#)

### Verifying the Interfaces that are Included in the Interface Set

**Purpose** Verify the interfaces included in the interface set.

**Action** user@host> `show interfaces interface-set dynamic-set terse`

### Verifying the Traffic Scheduling and Shaping Parameters for the Interface Set

**Purpose** Verify that the traffic scheduling and shaping parameters are applied properly to an interface included in the interface set.

**Action** user@host> `show class-of-service interface`

**Related Documentation**

- [Understanding Hierarchical CoS for Subscriber Interfaces on page 70](#)
- [Configuring an Interface Set of Subscribers in a Dynamic Profile on page 98](#)

## Example: Configuring a Dynamic Service VLAN Interface Set of Subscribers in a Dynamic Profile

---

Interface sets enable you to provide hierarchical scheduling to a group of subscriber interfaces. In this example, by using the `$junos-svlan-interface-set-name` internal dynamic variable when specifying the interface set name, you can locally generate an interface set name for use by SVLAN interfaces based on the outer tag of the dual-tagged VLAN. The format of the generated variable is *physical\_interface\_name - outer\_VLAN\_tag*.

- [Requirements on page 113](#)
- [Overview on page 113](#)

- [Configuration on page 113](#)
- [Verification on page 115](#)

## Requirements

Before you begin, configure the subscriber interfaces that you intend to include in the interface set. You can find general configuration instructions for the supported dynamic interface configuration in *DHCP Subscriber Interface Overview* and in the following:

- For dynamic VLAN interfaces, see *Configuring a Static or Dynamic VLAN Subscriber Interface over Aggregated Ethernet*.
- For dynamic IP demux interfaces, see *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles* and *Configuring a Static or Dynamic IP Demux Subscriber Interface over Aggregated Ethernet*.
- For dynamic VLAN demux interfaces, see *Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles*.

## Overview

Interface sets enable you to provide hierarchical scheduling to a group of subscriber interfaces. By using the `$junos-svlan-interface-set-name` internal dynamic variable when specifying the interface set name, you can locally generate an interface set name for use by SVLAN interfaces based on the outer tag of the dual-tagged VLAN. The format of the generated variable is *physical\_interface\_name - outer\_VLAN\_tag*.

This example includes the following statements:

- **interface-set**—Configures the name of the scheduler for dynamic CoS. In this example, you use the `$junos-svlan-interface-set-name` variable to obtain the locally generated interface set name for use by SVLAN interfaces based on the outer tag of the dual-tagged VLAN.
- **output-traffic-control-profile**—Applies an output traffic scheduling and shaping profile to the interface set.
- **output-traffic-control-profile-remaining**—Applies an output traffic scheduling and shaping profile for remaining traffic to the interface set.

## Configuration

### CLI Quick Configuration

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

```
[edit]
set dynamic-profiles profile-dhcp-ipdemux interfaces interface-set
  $junos-svlan-interface-set-name interface $junos-interface-ifd-name unit
  $junos-underlying-interface-unit
set dynamic-profiles profile-dhcp-ipdemux interfaces $junos-interface-ifd-name unit
  $junos-underlying-interface-unit
```

```
set class-of-service traffic-control-profiles tcp1 scheduler-map schedMap
set class-of-service traffic-control-profiles tcp1 shaping-rate 50m
set class-of-service traffic-control-profiles tcp1 guaranteed-rate 200k
set class-of-service traffic-control-profiles tcp3 scheduler-map ss1q0q1
set class-of-service traffic-control-profiles tcp3 shaping-rate 20m
set class-of-service traffic-control-profiles tcp3 guaranteed-rate 5m
set class-of-service interfaces interface-set ae0-111 output-traffic-control-profile tcp1
set class-of-service interfaces interface-set ae0-111
output-traffic-control-profile-remaining tcp3
```

### Step-by-Step Procedure

To configure an SVLAN interface set of subscriber interfaces:

1. Access the dynamic profile you want to modify for interface sets.

```
[edit]
user@host# edit dynamic-profiles profile-dhcp-ipdemux
```

2. Access the dynamic profile interface configuration.

```
[edit dynamic-profiles profile-dhcp-ipdemux]
user@host# edit interfaces
```

3. Configure the SVLAN interface set in the dynamic profile.

The interface set is created dynamically when the subscriber logs in.

```
[edit dynamic-profiles profile-dhcp-ipdemux interfaces]
user@host# edit interface-set $junos-svlan-interface-set-name
```

4. Include dynamic IP demux interface creation within the dynamic interface set.

```
[edit dynamic-profiles profile-dhcp-ipdemux interfaces interface-set
$junos-svlan-interface-set-name]
user@host# set interface $junos-interface-ifd-name unit
$junos-underlying-interface-unit
```

5. Access the SVLAN interface set name that you expect `$junos-svlan-interface-set-name` to generate. For example, to specify the expected interface set name for aggregated Ethernet interface ae0 and outer VLAN tag 111, include `ae0-111` for the *interface-set-name* variable.

```
[edit class-of-service interfaces]
user@host# edit interface-set ae0-111
```

6. Apply traffic shaping and queuing parameters to the SVLAN interface set.



**TIP:** You must configure the interface set in the static [edit class-of-service] hierarchy, not in the [edit dynamic-profiles] hierarchy.

---

```
[edit class-of-service interfaces interface-set ae0-111]
user@host# set output-traffic-control-profile tcp1
```



7. Apply traffic shaping and queuing parameters to any remaining traffic on the SVLAN interface set.

```
[edit class-of-service interfaces interface-set ae0-111]
user@host# set output-traffic-control-profile-remaining tcp3
```

## Results

From configuration mode, confirm your configuration by entering the **show dynamic-profiles** command and the **show class-of-service** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show dynamic-profiles
dynamic-profiles {
  profile-dhcp-ipdemux {
    interfaces {
      interface-set "$junos-svlan-interface-set-name" {
        interface "$junos-interface-ifd-name" {
          unit "$junos-underlying-interface-unit";
        }
      }
      "$junos-interface-ifd-name" {
        unit "$junos-underlying-interface-unit";
      }
    }
  }
}

user@host# show class-of-service
class-of-service {
  traffic-control-profiles {
    tcp1 {
      scheduler-map schedMap;
      shaping-rate 50m;
      guaranteed-rate 200k;
    }
    tcp3 {
      inactive: scheduler-map sslq0q1;
      shaping-rate 20m;
      guaranteed-rate 5m;
    }
  }
  interfaces {
    interface-set ae0-111 {
      output-traffic-control-profile tcp1;
      output-traffic-control-profile-remaining tcp3;
    }
  }
}
```

## Verification

To confirm that the configuration is correct, perform these tasks:

### Verifying the Interfaces that are Included in the Interface Set

**Purpose** Verify the interfaces that are included in the interface set.

**Action** user@host> show class-of-service interface-set

### Displaying Information for Active Subscribers

**Purpose** Display information for active subscribers.

**Action** user@host> [show subscribers](#) detail

- Related Documentation**
- *Dynamic Profiles Overview*
  - *Configuring a Basic Dynamic Profile*
  - [Configuring Hierarchical Schedulers for CoS on page 12](#)
  - *Configuring Remaining Common Queues on MIC and MPC Interfaces*

## CHAPTER 5

# Configuring Hierarchical Scheduling for MPLS Pseudowire Interfaces

- [Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 117](#)
- [CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 118](#)
- [CoS Two-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 120](#)
- [Configuring CoS Two-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces on page 122](#)
- [CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 123](#)
- [Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces \(Logical Interfaces over a Transport Logical Interface\) on page 127](#)
- [Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces \(Logical Interfaces over a Pseudowire Interface Set\) on page 128](#)

## Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview

---

Junos OS supports two aspects of CoS for MPLS pseudowire subscriber interfaces. You can apply CoS rewrite rules and behavior aggregate (BA) classifiers to MPLS pseudowire subscriber interfaces. In addition, CoS performs egress hierarchical shaping towards the subscriber on MPLS pseudowire subscriber interfaces.

Hierarchical CoS enables you to apply traffic scheduling and queuing parameters and packet transmission scheduling parameters to an individual subscriber interface rather than to all interfaces configured on the port. Hierarchical CoS is supported on MX Series routers with either EQ DPCs or MPC/MICs installed.

On Juniper Networks MX Series routers, MPC/MIC and EQ DPC interfaces support a four-level CoS scheduling hierarchy that, when fully configured, consists of the physical interface (level 1), the interface set or the underlying interface (level 2), one or more logical interfaces (level 3), and one or more queues (level 4). Although all CoS scheduling hierarchies are four-level, level 1 is always the physical interface and level 4 is always the queue. Hierarchical scheduling configurations consist of the type of interfaces you configure; for example, a logical interface or an interface set and where those interfaces reside in the scheduling hierarchy, either level 2 or level 3. Because many hierarchical

scheduling configurations are possible, we use the terms *two-level hierarchical scheduling* and *three-level hierarchical scheduling* in this discussion.

**Related  
Documentation**

- *Pseudowire Subscriber Logical Interfaces Overview*
- *Configuring a Pseudowire Subscriber Logical Interface*
- [Understanding Hierarchical CoS for Subscriber Interfaces on page 70](#)
- [CoS Two-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 120](#)
- [CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 123](#)
- [CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 118](#)
- [hierarchical-scheduler on page 311](#)

---

## CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces

---

CoS supports two-level and three-level hierarchies for MPLS pseudowire subscriber interfaces.

To configure two-level scheduling, include the **maximum-hierarchy-levels 2** option under the **[edit interfaces *interface-name* hierarchical-scheduler]** statement on the physical interface of the logical tunnel anchor point.

To configure three-level hierarchical scheduling, include the **implicit-hierarchy** option under the **[edit interfaces *interface-name* hierarchical-scheduler]** statement on the physical interface of the logical tunnel anchor point. Use the following guidelines for configuring the **implicit-hierarchy** option:

- If an output traffic-control profile is configured on the pseudowire transport interface and on a pseudowire service interface, the two interfaces form a scheduling hierarchy. The pseudowire transport interface resides in a level 2 scheduler node and the pseudowire service interface resides in a level 3 scheduler node.
- If an output traffic-control profile is configured on the pseudowire services interface but not on a pseudowire transport interface, the pseudowire services interface resides in a level 3 scheduler node.
- If an output traffic-control profile is only configured on the pseudowire transport interface and not on the pseudowire services interface, the pseudowire transport interface resides in a level 3 scheduler node and all pseudowire traffic uses this node.

If the **implicit-hierarchy** option is not set on the logical tunnel anchor point, logical interfaces behave normally with the hierarchical-scheduler mode configured with or without the **hierarchical-scheduler maximum-hierarchy-levels** option under the **[edit interfaces *interface-name* hierarchical-scheduler]** statement. In this case, when you apply a traffic-control profile to the pseudowire and service logical interfaces, they both reside in level 3 scheduler nodes and do not form a scheduling hierarchy, which might not be the desirable behavior. In business edge, where only the pseudowire logical interfaces

need to be shaped, applying the traffic-control profile at just the transport logical interface may be sufficient.

When configuring the logical tunnel physical interface for the maximum hierarchy level, all pseudowire logical interfaces operating on the physical interface use the same hierarchy model. If you want to mix two-level and three-level scheduling hierarchies, you can group the pseudowires together by hierarchy levels and share the same logical tunnel anchor point or you can use three-level scheduling for all pseudowires over the anchor point.

To specify rewrite rules and classifiers on pseudowire interfaces, reference the pseudowire device under the **[edit class-of-service interfaces]** hierarchy level and specify the rewrite rules and classifiers for the pseudowire interfaces.

To control all pseudowire traffic using the same logical tunnel interface, apply CoS policies at the physical interface for the anchor logical tunnel.



**NOTE:** Starting with Junos OS Release 17.3R1, stateful anchor point redundancy support is provided for pseudowire subscriber logical interface by the underlying redundant logical tunnel interface (rlt) in active-backup mode. This redundancy protects the access and the core facing link against anchor PFE (Packet Forwarding Engine) failure. Starting in Junos OS Release 18.1R2, full hierarchical CoS support is provided for stateful anchor point redundancy of pseudowire subscriber logical interfaces. Both transport and services logical interfaces created for the pseudowire subscriber logical interface are stacked on the underlying redundant logical tunnel control logical interface. This logical interface stacking model is used for both redundant and non-redundant pseudowire subscriber logical interfaces. Hierarchical CoS is supported and configured the same on both redundant and non-redundant pseudowire subscriber logical interfaces.

Release History Table

Release	Description
18.1R1	Starting in Junos OS Release 18.1R2, full hierarchical CoS support is provided for stateful anchor point redundancy of pseudowire subscriber logical interfaces.

**Related Documentation**

- [Pseudowire Subscriber Logical Interfaces Overview](#)
- [Configuring a Pseudowire Subscriber Logical Interface](#)
- [Understanding Hierarchical CoS for Subscriber Interfaces on page 70](#)
- [Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 117](#)
- [Configuring CoS Two-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces on page 122](#)
- [Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces \(Logical Interfaces over a Transport Logical Interface\) on page 127](#)

- [Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces \(Logical Interfaces over a Pseudowire Interface Set\)](#) on page 128
- *Anchor Redundancy Pseudowire Subscriber Logical Interfaces Overview*
- [hierarchical-scheduler](#) on page 311

## CoS Two-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces

Two-level hierarchical scheduling limits the number of hierarchical levels in the scheduling hierarchy to two. In a two-level scheduling hierarchy, all logical interfaces and interface sets share a single level 2 node. [Table 15 on page 120](#) summarizes the interface hierarchy and the CoS scheduler node levels for two-level hierarchical scheduling.


**Table 15: Two-Level Hierarchical Scheduling—Interface Hierarchy Versus Scheduling Nodes**

Level 1	Level 2	Level 3	Level 4
Physical interface	—	Pseudowire transport logical interface	One or more queues
Physical interface	—	Interface set	One or more queues
Physical interface	—	Pseudowire service logical interface	One or more queues

You use the two-level hierarchical scheduling when you have many pseudowires but you do not require shaping specific to the subscriber logical interface. For example, when your configuration is one subscriber per pseudowire interface.

[Figure 14 on page 121](#) shows a two-level hierarchical scheduling configuration for the MPLS pseudowires. In this configuration, level 1 is the physical interface used for the logical tunnel anchor node. All of the pseudowire transport interfaces share a single level 2 node. The level 3 nodes are the pseudowire transport logical interfaces (ps0.0, ps1.0, and ps2.0). In this configuration, interface sets are not configured and only the logical interfaces have traffic control profiles.

Figure 14: MPLS Pseudowire Subscriber Interface Two-Level Scheduler Configuration



Two-level hierarchical scheduling has up to eight class of service queues. For this configuration, include the `maximum-hierarchy-levels 2` option under the `[edit interfaces interface-name hierarchical-scheduler]` hierarchy level at the physical interface for the anchor logical tunnel.



**NOTE:** You cannot configure shaping policies on both the pseudowire logical interfaces and the subscriber logical interfaces over the same pseudowire. If a traffic-control profile is configured on a pseudowire logical interface, and CoS policies are configured on the subscriber logical interface over another pseudowire, all of the logical interfaces are at level 3 and act as peers.

#### Related Documentation

- [Pseudowire Subscriber Logical Interfaces Overview](#)
- [Configuring a Pseudowire Subscriber Logical Interface](#)
- [Understanding Hierarchical CoS for Subscriber Interfaces on page 70](#)
- [Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 117](#)
- [CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 123](#)
- [CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 118](#)
- [Configuring CoS Two-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces on page 122](#)
- [hierarchical-scheduler on page 311](#)

## Configuring CoS Two-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces

---

Before configuring CoS parameters for MPLS pseudowire subscriber interfaces, you must first complete these tasks:

1. Configure the pseudowire logical interfaces. See *Configuring a Pseudowire Subscriber Logical Interface*.
2. Configure the pseudowire device count. See *Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router*.
3. Configure the pseudowire device including the logical tunnel anchor point. See *Configuring a Pseudowire Subscriber Logical Interface Device*.
4. Configure the pseudowire transport logical interface. See *Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface*.
5. Configure the pseudowire signaling (either Layer 2 circuit signaling or Layer 2 VPN signaling). See *Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces* or *Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces*.
6. Configure the pseudowire logical interfaces. See *Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface*.

To configure CoS policies on MPLS pseudowire subscriber interfaces using two-level scheduling:

1. Configure the hierarchical scheduler for the physical interface used for the logical tunnel (anchor point). For two-level scheduling the hierarchical scheduler must be set to **maximum-scheduler levels 2**.

```
[edit]
user@host#edit interfaces ps-anchor-device-name
user@host#set hierarchical-scheduler maximum-hierarchy-levels 2
```

2. Specify the traffic-control profile to use on the pseudowire logical interface.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#set output-traffic-control-profile profile-name
```

3. Configure the rewrite rule.

The available rewrite rule types for pseudowire interfaces are **dscp** and **inet-precedence**.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#edit rewrite-rules (dscp | inet-precedence) rewrite-name
user@host#edit forwarding-class class-name
user@host#set loss-priority class-name code-point (alias | bits)
```



## 4. Configure the classifier.

The available classifier types for pseudowire interfaces are **dscp** and **inet-precedence**.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#edit classifiers (dscp | inet-precedence) classifier-name
user@host#edit forwarding-class class-name
user@host#set loss-priority class-name code-points [aliases] [bit-patterns]
```

## 5. Apply the rewrite rule and classifier to the pseudowire interface.

For the *interface\_name* parameter, specify the pseudowire device name.

```
[edit class-of-service interfaces interface_name unit logical-unit-number]
user@host#set rewrite-rule (dscp | inet-precedence) (rewrite-name | default) protocol
protocol-types
user@host#set classifiers (dscp | inet-precedence) (classifier-name | default)
```

#### Related Documentation

- [CoS on Ethernet Pseudowires in Universal Edge Networks Overview](#)
- [Mapping CoS Component Inputs to Outputs](#)
- [Understanding Hierarchical CoS for Subscriber Interfaces on page 70](#)
- [Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 117](#)
- [CoS Two-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 120](#)
- [CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 123](#)
- [CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 118](#)
- [hierarchical-scheduler on page 311](#)

## CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces

In three-level hierarchical scheduling, the CoS scheduler nodes at level 1, level 2, and level 3 form a scheduling hierarchy. You can configure many different three-level scheduling hierarchies, depending on the location of the interface set and the use of underlying interfaces. In all variations, the physical interface on which the logical tunnel resides is a level 1 CoS scheduler node and the queues reside at level 4. Three-level scheduling hierarchies can have up to eight class of service queues.

[Table 16 on page 124](#) summarizes the most common three-level hierarchical scheduling configurations and shows the interface hierarchy and CoS scheduler nodes.

**Table 16: Three-Level Hierarchical Scheduling—Interface Hierarchy Versus CoS Scheduling Node Levels**

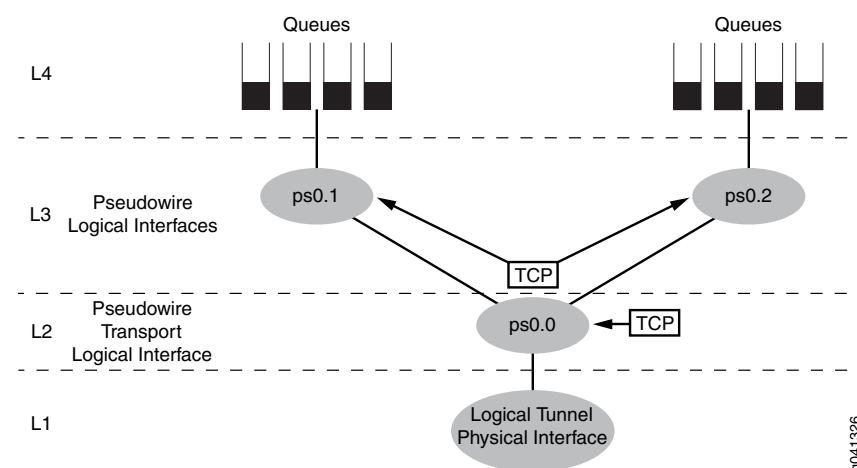
Level 1	Level 2	Level 3	Level 4
Physical interface	Pseudowire interface set	Pseudowire service logical interfaces	One or more queues
Physical interface	Pseudowire transport logical interface	Pseudowire interface set	One or more queues
Physical interface	Pseudowire transport logical interface	Pseudowire service logical interfaces	One or more queues

### Three-Level Scheduling Hierarchy: Pseudowire Logical Interfaces over a Transport Logical Interface

Figure 15 on page 124 shows an MPLS pseudowire three-level scheduling hierarchy that includes two pseudowire service logical interfaces over a pseudowire transport logical interface. This variation uses the following scheduler nodes:

- Level 4—Forwarding class-based queues
- Level 3—Pseudowire service logical interfaces (ps0.1 and ps0.2) for subscriber sessions
- Level 2—Pseudowire transport logical interface (ps0.0)
- Level 1—Common/shared physical interface of the logical tunnel anchor point

You apply the traffic-control profiles at the pseudowire transport logical interfaces (level 2) and the pseudowire service logical interfaces (level 3).

**Figure 15: Three-Level Scheduling Hierarchy Case 1: Pseudowire Service Logical Interfaces over a Transport Logical Interface**

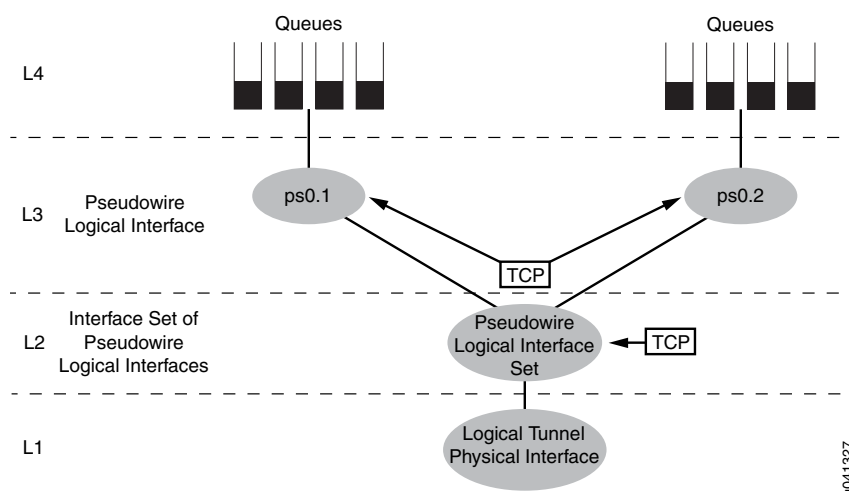
### Three-Level Scheduling Hierarchy : Pseudowire Service Logical Interfaces over a Pseudowire Service Interface Set

Figure 16 on page 125 shows another variation of MPLS pseudowire three-level hierarchical scheduling that includes two pseudowire service logical interfaces over a pseudowire service interface set. This variation uses the following CoS scheduler nodes:

- Level 4—Forwarding class-based queues
- Level 3—Pseudowire service logical interfaces (ps0.1 and ps0.2)
- Level 2—Pseudowire service interface set
- Level 1—Common/shared physical interface of the logical tunnel anchor point

You apply the traffic-control profile at the pseudowire service interfaces (level 3) and at the interface set (level 2). This variation is most useful for subscriber edge deployments.

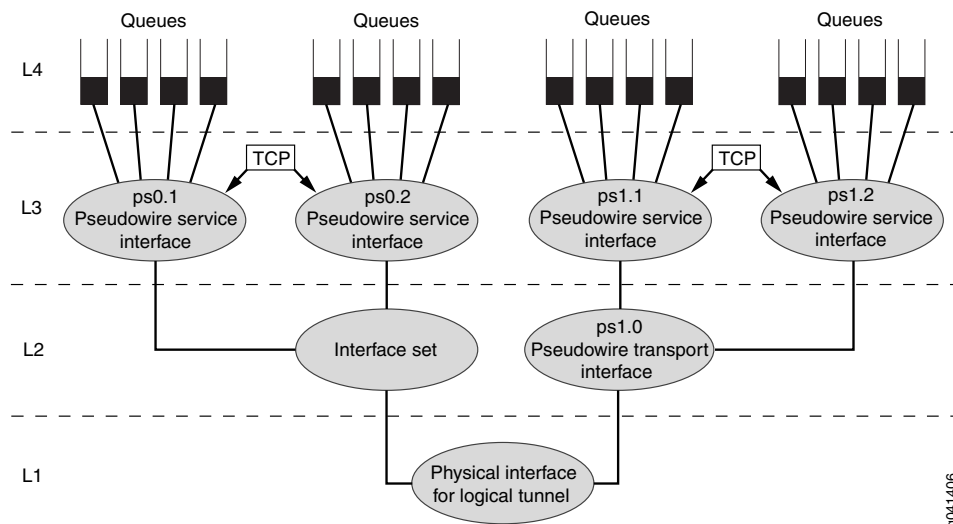
**Figure 16: Three-Level Scheduling Hierarchy Case 2: Pseudowire Service Logical Interfaces over a Pseudowire Service Interface Set**



### Three-Level Scheduling Hierarchy Combined Deployment Scenario

Figure 17 on page 126 shows a deployment scenario that combines the three-level hierarchical scheduling scenarios in Figure 15 on page 124 and Figure 16 on page 125.

**Figure 17: Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces—Deployment Scenario**



9041406

This variation uses the following CoS scheduler nodes:

- Level 4—Forwarding class-based queues
- Level 3—Pseudowire service logical interfaces (ps0.1, ps0.2, ps1.1, and ps1.2)
- Level 2—Service interface set for pseudowire service interfaces (ps0.1 and ps0.2) and transport logical interface (ps1.0) for the pseudowire service logical interfaces (ps1.1 and ps1.2)
- Level 1—Common/shared physical interface of the logical tunnel anchor point

You apply the traffic-control profiles to the interfaces at both level 2 and level 3, as well as the interface set at level 2.

#### Related Documentation

- [Pseudowire Subscriber Logical Interfaces Overview](#)
- [Configuring a Pseudowire Subscriber Logical Interface](#)
- [Understanding Hierarchical CoS for Subscriber Interfaces on page 70](#)
- [Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 117](#)
- [CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 118](#)
- [Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces \(Logical Interfaces over a Transport Logical Interface\) on page 127](#)
- [Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces \(Logical Interfaces over a Pseudowire Interface Set\) on page 128](#)
- [hierarchical-scheduler on page 311](#)

## Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Transport Logical Interface)

Before configuring CoS three-level scheduling on pseudowire logical interfaces over a transport logical interface, you must first complete these tasks:

1. Configure the pseudowire logical interfaces. See *Configuring a Pseudowire Subscriber Logical Interface*.
2. Configure the pseudowire device count. See *Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router*.
3. Configure the pseudowire device including the logical tunnel anchor point. See *Configuring a Pseudowire Subscriber Logical Interface Device*.
4. Configure the pseudowire transport logical interface. See *Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface*.
5. Configure the pseudowire signaling (either Layer 2 circuit signaling or Layer 2 VPN signaling). See *Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces* or *Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces*.
6. Configure the pseudowire logical interfaces. See *Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface*.

Three-level scheduling on pseudowire logical interfaces over a transport logical interface requires you to apply the traffic-control profiles at both the pseudowire logical interface and the pseudowire transport logical interface. To configure CoS policies on three-level scheduling on pseudowire logical interfaces over a transport logical interface:

1. Configure the hierarchical scheduler for the physical interface used for the logical tunnel (anchor point). For three-level scheduling the hierarchical scheduler must be set to **implicit-hierarchy**.

[edit]

user@host#edit interfaces *ps-anchor-device-name*

user@host#set hierarchical-scheduler implicit-hierarchy

2. Specify the traffic-control profile to use on the pseudowire logical interface.

[edit class-of-service]

user@host#edit interfaces *ps ps-device-name*

user@host#edit unit *logical-unit-number*

user@host#set output-traffic-control-profile *profile-name*

3. Specify the traffic-control profile to use on the pseudowire transport logical interface.

[edit class-of-service]

user@host#edit interfaces *ps ps-device-name*

user@host#edit unit *logical-unit-number*

user@host#set output-traffic-control-profile *profile-name*

4. Configure the rewrite rule.

The available rewrite rule types for pseudowire interfaces are **dscp** and **inet-precedence**.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#edit rewrite-rules (dscp | inet-precedence) rewrite-name
user@host#edit forwarding-class class-name
user@host#set loss-priority class-name code-point (alias | bits)
```

5. Configure the classifier.

The available classifier types for pseudowire interfaces are **dscp** and **inet-precedence**.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#edit classifiers (dscp | inet-precedence) classifier-name
user@host#edit forwarding-class class-name
user@host#set loss-priority class-name code-points [aliases] [bit-patterns]
```

6. Apply the rewrite rule and classifier to the pseudowire interfaces.

For the *interface\_name* parameter, specify the pseudowire device name.

```
[edit class-of-service interfaces interface_name unit logical-unit-number]
user@host#set rewrite-rule (dscp | inet-precedence) (rewrite-name | default) protocol
protocol-types
user@host#set classifiers (dscp | inet-precedence) (classifier-name | default)
```

**Related  
Documentation**

- [CoS on Ethernet Pseudowires in Universal Edge Networks Overview](#)
- [Mapping CoS Component Inputs to Outputs](#)
- [Understanding Hierarchical CoS for Subscriber Interfaces on page 70](#)
- [Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 117](#)
- [CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 123](#)
- [CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 118](#)
- [Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces \(Logical Interfaces over a Pseudowire Interface Set\) on page 128](#)
- [hierarchical-scheduler on page 311](#)

---

## Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Pseudowire Interface Set)

---

Before configuring three-level scheduling on pseudowire logical interfaces over a pseudowire logical interface set, you must first complete the following tasks:

1. Configure the pseudowire logical interfaces. See *Configuring a Pseudowire Subscriber Logical Interface*.
2. Configure the pseudowire device count. See *Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router*.
3. Configure the pseudowire device including the logical tunnel anchor point. See *Configuring a Pseudowire Subscriber Logical Interface Device*.
4. Configure the pseudowire transport logical interface. See *Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface*.
5. Configure the pseudowire signaling (either Layer 2 circuit signaling or Layer 2 VPN signaling). See *Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces* or *Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces*.
6. Configure the pseudowire logical interfaces. See *Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface*.

Three-level scheduling on pseudowire logical interfaces over a pseudowire logical interface set requires you to apply the traffic-control profiles at both the pseudowire logical interface and the pseudowire logical interface-set. To configure CoS policies on MPLS pseudowire subscriber interfaces using three-level implicit hierarchical scheduling:

1. Configure the hierarchical scheduler for the physical interface used for the logical tunnel (anchor point). For three-level scheduling the hierarchical scheduler must be set to **implicit-hierarchy**.

```
[edit]
user@host#edit interfaces ps-anchor-device-name
user@host#set hierarchical-scheduler implicit-hierarchy
```

2. Specify the traffic-control profile to use on the pseudowire logical interfaces.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#set output-traffic-control-profile profile-name
```

3. Define a pseudowire logical interface set and configure the traffic-control profile used for the interface set.

```
[edit class-of-service]
user@host#edit interfaces
user@host#edit interface-set interface-set-name
user@host#edit output-traffic-control-profile profile-name
```

4. Group the pseudowire logical interfaces in the pseudowire logical interface set.

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

5. Configure the rewrite rule.

The available rewrite rule types for pseudowire interfaces are **dscp** and **inet-precedence**.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#edit rewrite-rules (dscp | inet-precedence) rewrite-name
user@host#edit forwarding-class class-name
user@host#set loss-priority class-name code-point (alias | bits)
```

6. Configure the classifier.

The available classifier types for pseudowire interfaces are **dscp** and **inet-precedence**.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#edit classifiers (dscp | inet-precedence) classifier-name
user@host#edit forwarding-class class-name
user@host#set loss-priority class-name code-points [aliases] [bit-patterns]
```

7. Apply the rewrite rule and classifier to the pseudowire interfaces.

For the *interface\_name* parameter, specify the ps device name.

```
[edit class-of-service interfaces interface_name unit logical-unit-number]
user@host#set rewrite-rule (dscp | inet-precedence) (rewrite-name | default) protocol
protocol-types
user@host#set classifiers (dscp | inet-precedence) (classifier-name | default)
```

**Related  
Documentation**

- [CoS on Ethernet Pseudowires in Universal Edge Networks Overview](#)
- [Mapping CoS Component Inputs to Outputs](#)
- [Understanding Hierarchical CoS for Subscriber Interfaces on page 70](#)
- [Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 117](#)
- [CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 123](#)
- [CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 118](#)
- [Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces \(Logical Interfaces over a Transport Logical Interface\) on page 127](#)
- [hierarchical-scheduler on page 311](#)



## CHAPTER 6

# Configuring Hierarchical Scheduling for L2TP

- [CoS for L2TP LAC Subscriber Interfaces Overview on page 132](#)
- [Configuring Dynamic CoS for an L2TP LAC Tunnel on page 134](#)
- [CoS for L2TP LNS Inline Services Overview on page 135](#)
- [Configuring an Inline Service Interface for L2TP LNS on page 136](#)
- [Configuring Dynamic CoS for an L2TP LNS Inline Service on page 137](#)


## CoS for L2TP LAC Subscriber Interfaces Overview

You can apply CoS to the Layer 2 Tunnel Protocol (L2TP) access concentrator (LAC) component.

In Layer 2 Tunnel Protocol (L2TP) configurations, IP and L2TP headers are added to packets arriving at a PPP subscriber interface on the L2TP access concentrator (LAC) before being tunneled to the L2TP network server (LNS). You can manage the IP header by configuring classifiers and rewrite-rules that transfer the ToS (Type of Service) value or the 802.1p value from the *inner* IP header to the *outer* IP header of the L2TP packet.

Figure 18 on page 132 shows the classifier and rewrite rules that you can configure from the LAC to the LNS, and from the LNS to the LAC.

**Figure 18: CoS Configuration for L2TP LAC Topology**



- [Traffic from LAC to LNS on page 132](#)
- [LAC Tunnels: Traffic from LNS to LAC on page 133](#)

### Traffic from LAC to LNS

To set the ToS value or the 802.1p value on the inner IP header, you can configure both fixed and behavior aggregate (BA) classifiers for subscribers at Layer 2 or Layer 3 of the network.

Table 17 on page 132 lists the configuration options for applying classifiers to a subscriber interface on an ingress LAC tunnel.

**Table 17: Ingress LAC Tunnel Classifier Options**

Classifier	Subscriber Interface
Fixed	Either of the following: <ul style="list-style-type: none"> <li>• PPP interface</li> <li>• Underlying VLAN interface</li> </ul>

Table 17: Ingress LAC Tunnel Classifier Options (*continued*)

Classifier	Subscriber Interface
Layer 2	Either of the following: <ul style="list-style-type: none"> <li>• PPP interface</li> <li>• Underlying VLAN interface</li> </ul>
Layer 3	Family of PPP interfaces

You cannot configure a Layer 2 and fixed classifier together.

The behavior of the Layer 2 and Layer 3 classifiers depends on the configuration. For example, a Layer 3 classifier for a family of PPP interfaces overrides a Layer 2 classifier configured at the PPP interface, except for the unknown packets and control packets.

If you do not configure a classifier for Layer 2, the system applies the default Layer 3 classifier so that tunneled and terminated subscribers have the same behavior. To prevent unknown packets and control packets from being discarded, the system assigns them to the best-effort forwarding class.

For egress tunnels, you configure rewrite rules at the PPP interface to set the ToS or 802.1p value of the outer IP header. Rewrite rules are applied accordingly to the forwarding class, packet loss priority (PLP), and code point.

## LAC Tunnels: Traffic from LNS to LAC

On a LAC, mapping the inner IP header to the outer IP header of the L2TP packet depends on the classifier and rewrite-rule configurations. For example, [Table 18 on page 133](#) lists the values for the classifier and rewrite rules for a VLAN interface. For assured forwarding, the inner 802.1p value (**ob001**) is classified with the assured-forwarding class and low loss priority at the ingress interface. Based on the assured-forwarding class and low loss priority in the rewrite rule, the ToS value in the outer IP header is set to **ob001**.

Table 18: Sample Result for the Classifier and Rewrite Rules for a VLAN Interface

Inner .1p Value	Forwarding Class	Loss Priority	Code Point	Outer ToS Value
ob000	best-effort	low	000	ob000
ob001	assured-forwarding	low	001	ob001
ob101	expedited-forwarding	low	101	ob101
ob111	network-control	low	11	ob111

**Related Documentation** • [Configuring Dynamic CoS for an L2TP LAC Tunnel on page 134](#)

## Configuring Dynamic CoS for an L2TP LAC Tunnel

---

In L2TP configurations, IP and L2TP headers are added to packets arriving at a PPP subscriber interface on the LAC before being tunneled to the L2TP network server (LNS).

Classifiers and rewrite rules enable you to properly transfer the ToS (Type of Service) value or the 802.1p value from the inner IP header to the outer IP header of the L2TP packet.

Before you begin, configure the L2TP LAC. See *Configuring an L2TP LAC*.

To manage the IP header values for a LAC tunnel:

1. Configure the classifier for the inner tunnel.
  - a. Define the fixed or behavior aggregate (BA) classifier.
    - To configure a fixed classifier:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
user@host# set forwarding-class class-name
```
    - To configure a BA classifier:

```
[edit class-of-service]
user@host# set classifiers (ieee-802.1 | inet-precedence) classifier-name
forwarding-class class-name loss-priority level code-points [ aliases ] [
bit-patterns]
```
  - b. Apply the classifier to the Layer 2 interface or Layer 3 interface. For Layer 2, you can apply the classifier at the PPP interface or an underlying VLAN interface. For Layer 3, you can apply classifiers to a family of PPP interfaces.
    - To apply the classifier for the IEEE 802.1p value:

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name
unit logical-unit-number classifiers]
user@host# set ieee-802.1 (classifier-name | default) vlan-tag (inner | outer)
```
    - To apply the classifier for the ToS value:

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name
unit logical-unit-number classifiers]
user@host# set inet-precedence (classifier-name | default)
```
2. Configure the rewrite rule for the egress tunnel.
  - a. Configure the rewrite rule with the forwarding class and the loss priority value.

```
[edit class-of-service]
user@host# set rewrite-rules (ieee-802.1 | inet-precedence) rewrite-name
forwarding-class class-name loss-priority level code-point (alias | bits)
```
  - b. Apply the rewrite rule to the PPP interface for which the L2TP tunnel is configured.
    - To apply the rewrite-rule for the IEEE 802.1p value:

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name
unit logical-unit-number rewrite-rules]
```

```
user@host# set ieee-802.1 (rewrite-name | default) vlan-tag (outer |
outer-and-inner)
```

- To apply the rewrite rule for the ToS value:

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name
unit logical-unit-number rewrite-rules]
user@host# set inet-precedence (rewrite-name | default)
```

#### Related Documentation

- [Guidelines for Configuring Dynamic CoS for Subscriber Access](#)
- [CoS for L2TP LAC Subscriber Interfaces Overview on page 132](#)

## CoS for L2TP LNS Inline Services Overview

You can apply hierarchical scheduling and per-session shaping to Layer 2 Tunnel Protocol (L2TP) network server (LNS) inline services using a static or dynamic CoS configuration.

This feature is supported on MIC and MPC interfaces on MX240, MX480, and MX960 routers.

- [Guidelines for Applying CoS to the LNS on page 135](#)
- [Hardware Requirements for Inline Services on the LNS on page 136](#)

### Guidelines for Applying CoS to the LNS

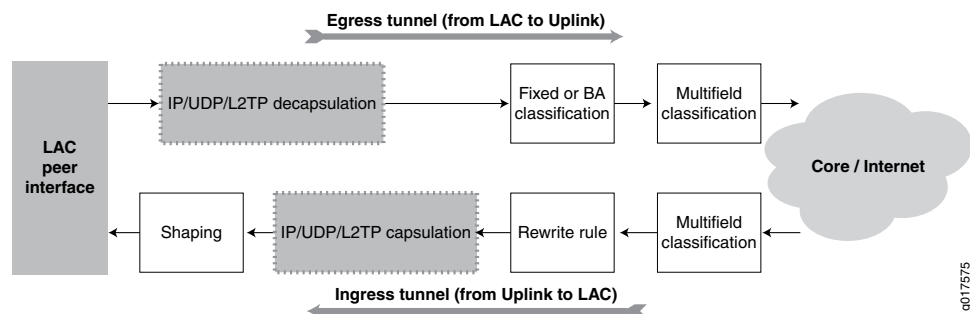
In L2TP configurations, IP, UDP, and L2TP headers are added to packets arriving at a PPP subscriber interface on the L2TP access concentrator (LAC) before being tunneled to the LNS.

When a service interface is configured for an L2TP LNS session, it has an *inner* IP header and an outer IP header. You can configure CoS for an LNS session that corresponds to the inner IP header only. The *outer* IP header is used for L2TP tunnel processing only.

However, we recommend that you configure classifiers and rewrite-rules to transfer the ToS (type of service) value from the inner IP header to the outer IP header of the L2TP packet.

[Figure 19 on page 135](#) shows the classifier and rewrite rules that you can configure on an LNS inline service.

**Figure 19: Processing of CoS Parameters in an L2TP LNS Inline Service**



By default, the shaping calculation on the service interface includes the L2TP encapsulation. If necessary, you can configure additional adjustments for downstream ATM traffic from the LAC or differences in Layer 2 protocols.

## Hardware Requirements for Inline Services on the LNS

Hierarchical scheduling for L2TP LNS inline services is supported on MIC and MPC interfaces only. The services that you can configure depend on the hardware combination. [Table 19 on page 136](#) lists the supported inline services and peer interfaces for each MIC and MPC combination.

**Table 19: Hardware Requirements for L2TP LNS Inline Services**

MPC Module	Inline Service Support—With Per-Session Shaping	Inline Service Support—Without Per-Session Shaping
MPC2E-3D-NG	No	Yes
MPC2E-3D-NG-Q	Yes	Yes
MX80		
MPC-3D-16XGE-SFPP	No	No

- Related Documentation**
- [Configuring Static CoS for an L2TP LNS Inline Service](#)
  - [Configuring Dynamic CoS for an L2TP LNS Inline Service on page 137](#)

## Configuring an Inline Service Interface for L2TP LNS

The inline service interface is a virtual physical service interface that resides on the Packet Forwarding Engine. This si interface, referred to as an *anchor* interface, makes it possible to provide L2TP services without a special services PIC. The inline service interface is supported only by MPCs on MX Series routers. Four inline service interfaces are configurable per MPC-occupied chassis slot.

You can maximize the number of sessions that can be shaped in one service interface by setting the maximum number of hierarchy levels to two. In this case, each LNS session consumes one L3 node in the scheduler hierarchy for shaping.

If you do not specify the number of levels (two is the only option), then the number of LNS sessions that can be shaped on the service interface is limited to the number of L2 nodes, or 4096 sessions. Additional sessions still come up, but they are not shaped.

To configure an inline service interface:

1. Access the service interface.
 

```
[edit interfaces]
user@host# edit si-slot/pic/port
```

2. (Optional; for per-session shaping only) Enable the inline service interface for hierarchical schedulers and limit the number of scheduler levels to two.

```
[edit interfaces si-slot/pic/port]
user@host# set hierarchical-scheduler maximum-hierarchy-levels 2
```

3. (Optional; for per-session shaping only) Configure services encapsulation for inline service interface.

```
[edit interfaces si-slot/pic/port]
user@host# set encapsulation generic-services
```

4. Configure the IPv4 family on the reserved unit 0 logical interface.

```
[edit interfaces si-slot/pic/port]
user@host# set unit 0 family inet
```

#### Related Documentation

- [Configuring an L2TP LNS with Inline Service Interfaces](#)

## Configuring Dynamic CoS for an L2TP LNS Inline Service

You can configure hierarchical scheduling for an L2TP LNS inline service and manage the IP header values using rewrite rules and classifiers.

Before you begin, configure the L2TP LNS inline service interface. See *Configuring an L2TP LNS with Inline Service Interfaces*.

To configure CoS for an L2TP LNS inline service in a dynamic profile:

1. Configure the hierarchical scheduler for the service interface (si) interface.

```
[edit interfaces si-fpc/port/pic ]
user@host# set hierarchical-scheduler maximum-hierarchy-levels 2
```



**BEST PRACTICE:** To enable Level 3 nodes in the LNS scheduler hierarchy and to provide better scaling, we recommend that you also specify a maximum of two hierarchy levels.

2. Configure the LNS to reflect the IP ToS value in the inner IP header to the outer IP header.

```
[edit services l2tp tunnel-group name]
user@host# set tos-reflect
```

3. Configure the classifier for egress traffic from the LAC.
  - a. Define the fixed or behavior aggregate (BA) classifier.

- To configure a fixed classifier:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

```
user@host# set forwarding-class class-name
```

- To configure a BA classifier:

```
[edit class-of-service]
user@host# set classifiers (dscp | dscp-ipv6 | inet-precedence) classifier-name
forwarding-class class-name loss-priority level code-points [ aliases ] [
bit-patterns]
```

- Apply the classifier to the service interface.

- To apply the classifier for the DSCP or DSCP IPv6 value:

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name
unit logical-unit-number classifiers]
user@host# set dscp (classifier-name | default)
user@host# set dscp-ipv6 (classifier-name | default)
```

- To apply the classifier for the ToS value:

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name
unit logical-unit-number classifiers]
user@host# set inet-precedence (classifier-name | default)
```

- Configure and apply a rewrite-rule to ingress traffic to the LAC:

- Configure the rewrite rule with the forwarding class and the loss priority value.

```
[edit class-of-service]
user@host# set rewrite-rules (dscp | dscp-ipv6 | inet-precedence) rewrite-name
forwarding-class class-name loss-priority level code-point (alias | bits)
```

- Apply the rewrite rule to the service interface.

- To apply the rewrite rule for the DSCP or DSCP IPv6 value:

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name
unit logical-unit-number rewrite-rules]
user@host# set dscp (rewrite-name | default)
user@host# set dscp-ipv6 (rewrite-name | default)
```

- To apply the rewrite rule for the ToS value:

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name
unit logical-unit-number rewrite-rules]
user@host# set inet-precedence (rewrite-name | default)
```

- (Optional) Configure additional adjustments for downstream ATM traffic.

By default, the shaping calculation on the service interface includes the L2TP encapsulation. If necessary, you can configure additional adjustments for downstream ATM traffic from the LAC or differences in Layer 2 protocols.

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
user@host# set overhead-accounting (frame-mode | cell-mode |
$junos-cos-shaping-mode) <bytes (byte-value | $junos-cos-byte-adjust)
```

- Apply the traffic-control profile.

```
[edit dynamic-profiles profile-name class-of-service interfaces
$junos-interface-ifd-name unit $junos-interface-unit]
```



```
user@host# set output-traffic-control-profile profile-name
```

**Related  
Documentation**

- *Guidelines for Configuring Dynamic CoS for Subscriber Access*
- [CoS for L2TP LNS Inline Services Overview on page 135](#)
- *Example: Configuring an L2TP LNS*
- *Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates*



## CHAPTER 7

# Preventing Bandwidth Contention on Subscriber Interfaces

- [Hierarchical CoS Shaping-Rate Adjustments Overview on page 141](#)
- [Shaping Rate Adjustments for Subscriber Local Loops Overview on page 143](#)
- [Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops on page 144](#)
- [Configuring the Minimum Adjusted Shaping Rate on Scheduler Nodes for Subscribers on page 145](#)
- [Configuring Shaping-Rate Adjustments on Queues on page 146](#)
- [Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 148](#)
- [Disabling Shaping-Rate Adjustments for Subscriber Local Loops on page 154](#)
- [Disabling Hierarchical Bandwidth Adjustment for Subscriber Interfaces with Reverse-OIF Mapping on page 154](#)
- [Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops on page 155](#)
- [Verifying the Configuration of Shaping-Rate Adjustments for Subscriber Local Loops on page 158](#)
- [Verifying the Configuration of ANCP for Shaping-Rate Adjustments on page 158](#)
- [Using Hierarchical CoS to Adjust Shaping Rates Based on Multicast Traffic on page 159](#)

## Hierarchical CoS Shaping-Rate Adjustments Overview

---

This overview describes how MX Series 3D Universal Edge Routers installed in a subscriber access network can adjust hierarchical class-of-service (CoS) parameters to prevent bandwidth contention at subscriber interfaces.

Hierarchical CoS is supported only for subscriber interfaces on Enhanced Queueing (EQ) DPCs or MPC interfaces operating in hierarchical scheduler mode.

The characteristics of voice, data, and video applications vary widely in their requirements for traffic throughput, bandwidth management, delay and jitter tolerance, and buffer depth. To prevent bandwidth contention at subscriber interfaces, you can configure applications such as ANCP and Multicast to perform real-time adjustments to the shaping rate configured for subscriber interfaces for residential gateways. Enabling shaping-rate

adjustments on the router can prevent bandwidth contention at the interface from causing degradation of the subscriber's voice, data, or video services.

## Types of Shaping-Rate Adjustments

The ANCP application supports *absolute* adjustments to a specific shaping-rate value. You can configure ANCP to communicate the subscriber local loop speed to the MX Series router, which in turn throttles traffic destined to the associated subscriber interface so that it matches the subscriber local loop speed. ANCP acquires subscriber line rate information from DSLAMs and then communicates this data transmission rate for use with CoS.

The OIF mapping and reverse OIF mapping multicast applications support *delta* adjustments that increase or decrease the current shaping rate by a certain value. The system adjusts traffic destined to the subscriber using reverse OIF mapping enabled on a specified multicast interface. Reverse OIF mapping is used to determine the subscriber VLAN interface and the multicast traffic bandwidth on the interface.

## Levels of Shaping-Rate Adjustments

Both absolute and delta adjustments are made to a subscriber's aggregate shaping rate on a level 3 scheduler node.

Adjustments that occur on the scheduler node can also impact the shaping rates for all queues. This adjustment can be undesirable for service providers who want to provide a premium level of service on specific queues.

For delta-based adjustments by multicast applications, you can control the distribution of shaping rates among queues by assigning the percentage of adjustment allowed for each queue. In addition, you can set a minimum adjusted shaping rate for each queue.

Figure 20 on page 142 shows a sample multicast network with shaping rates adjusted at the scheduler node level. The shaping rate is reduced by 4 Mbps (from 41 Mbps to 37 Mbps) at the scheduler node for subscriber interface 1, which reduces the rates of both the best effort and video on demand (VoD) service queues.

**Figure 20: Scheduler Node and Queues with Adjusted Shaping Rates**





Figure 21 on page 143 shows the same network with queue-based adjustments enabled for the best-effort queue on subscriber 1. The shaping rate of the best-effort queue is reduced by 4 Mbps (from 5 Mbps to 1 Mbps). The VoD service queue is not affected.

Figure 21: Queue with Adjusted Shaping Rate



#### Related Documentation

- [Configuring the Minimum Adjusted Shaping Rate on Scheduler Nodes for Subscribers on page 145](#)
- [Configuring Shaping-Rate Adjustments on Queues on page 146](#)
- [Shaping Rate Adjustments for Subscriber Local Loops Overview on page 143](#)
- [Disabling Hierarchical Bandwidth Adjustment for Subscriber Interfaces with Reverse-OIF Mapping on page 154](#)
- [Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops on page 155](#)

## Shaping Rate Adjustments for Subscriber Local Loops Overview

This overview describes how an MX Series 3D Universal Edge Router installed as an edge router can adjust hierarchical CoS policy for subscriber interfaces for subscriber local loops. You can configure the router to throttle the traffic sent to subscriber local loops so that the traffic does not exceed the current data transmission rate of those lines. This feature ensures that changes to subscriber local loop speeds do not cause bandwidth contention at the subscriber's residential gateway.

In a typical subscriber access network, traffic destined to a subscriber is delivered from the access network, through an edge router, to a DSLAM. The DSLAM multiplexes subscriber traffic through a DSL, also known as a *local loop*, to the subscriber's residential gateway. When line noise or cross talk in a subcarrier causes the error rate on a DSL to exceed a certain threshold, the DSLAM can adapt itself by lowering the data transmission rate to that carrier device. A lower data transmission rate is less susceptible to induced errors.

You can configure an MX Series router to adjust the configured shaping rates on scheduler nodes for subscriber interfaces that represent subscriber local loops. Whenever a DSLAM resynchronizes a subscriber local loop speed, the router adjusts the configured shaping rate for that line so that the aggregate egress traffic to those subscribers is shaped to the local loop speed before the traffic reaches the DSLAM. Unless the maximum amount of bandwidth allocated to the subscriber interface on the router is throttled to the local loop speed, bandwidth contention can occur at the subscriber's residential gateway, which can cause the DSLAM to drop packets. This type of shaping-rate adjustment requires the topology discovery and traffic-monitoring features of the Access Node Control Protocol (ANCP).

You can enable ANCP to communicate the subscriber local loop speed to CoS, which in turn throttles traffic destined to the associated subscriber interface so that it matches the subscriber local loop speed. The ANCP agent acquires unadjusted (net) subscriber line rate information from DSLAMs and then communicates this data transmission rate for use with CoS. You can also configure percentage and byte adjustments that the ANCP agent can make to the received net data rate for frame-mode DSL types before communicating the adjusted rate and overhead to CoS.

**Related  
Documentation**

- [Hierarchical CoS Shaping-Rate Adjustments Overview on page 141](#)
- [Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops on page 144](#)
- [Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 148](#)
- [Disabling Shaping-Rate Adjustments for Subscriber Local Loops on page 154](#)
- [Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops on page 155](#)
- For more information about the ANCP protocol, see the *ANCP and the ANCP Agent Overview*.

---

## Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops

These guidelines apply to configuring an MX Series 3D Universal Edge Router installed as an edge router to adjust the configured shaping rates on scheduler nodes for subscriber interfaces that represent subscriber local loops. This shaping-rate feature uses the topology discovery and traffic-monitoring features of ANCP.

When you enhance hierarchical CoS policy by configuring ANCP-driven shaping-rate adjustments, consider the following guidelines:

- Shaping-rate adjustments are supported only for subscriber local loops that terminate at DSLAMs that you have configured as ANCP neighbors of the MX Series router.
- Shaping-rate adjustments are supported only for scheduler nodes for which you have configured an initial shaping rate by including the **shaping-rate** statement in a traffic-control profile applied to the scheduler node. Specify the initial shaping rate as a peak rate, in bits per second (bps), and not as a percentage. Other methods of configuring a shaping rate are not supported with this feature.
- Shaping-rate adjustments are supported only for scheduler nodes that are static logical interface sets that you have configured to operate at Level 3 of the scheduler hierarchy on the router. If an interface set is configured with a logical interface (such as unit 0) and queue, then the interface set is an internal scheduler node (as opposed to a root node or a leaf node) at Level 2 of the hierarchy. However, if there are no traffic-control profiles are configured on logical interfaces in an interface set, then the interface set is an internal scheduler node at Level 3 of the hierarchy.
- Shaping-rate adjustments are supported only for subscriber interfaces over physical interfaces that you have configured to operate in hierarchical scheduler mode.

- After shaping-rate adjustments are enabled and the router has performed shaping-rate adjustments on a scheduler node, you can configure a new shaping rate by including the **shaping-rate** statement in a traffic-control profile and then applying that profile to that scheduler node. However, this new shaping-rate value does not immediately result in shaping traffic at the new rate. The scheduler node continues to be shaped at rate set by ANCP. Only when the ANCP shaping-rate adjustment feature is disabled is the scheduler node shaped at the newly configured shaping-rate.
- The Layer 2 Tunneling Protocol (L2TP) is often used to carry traffic securely between an L2TP Network Server (LNS) and an L2TP Access Concentrator (LAC). The QoS adjustment feature supports the shaping overhead options that you can use to add a specified number of bytes to the actual packet length when determining shaped session packet length. ANCP shaping-rate adjustments are not supported for ingress traffic, only for egress traffic. To configure the number of bytes to add to the packet at the egress side of the tunnel, include the **egress-shaping-overhead** and **mode** statements at the **[edit chassis fpc slot-number pic pic-number traffic-manager]** hierarchy level. Use the shaping overhead options if you need to account for encapsulation overhead.

For more information about the ANCP protocol, see the *ANCP and the ANCP Agent Overview*.

#### Related Documentation

- [Hierarchical CoS Shaping-Rate Adjustments Overview on page 141](#)
- [Shaping Rate Adjustments for Subscriber Local Loops Overview on page 143](#)
- [Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 148](#)
- [Disabling Shaping-Rate Adjustments for Subscriber Local Loops on page 154](#)
- [Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops on page 155](#)

## Configuring the Minimum Adjusted Shaping Rate on Scheduler Nodes for Subscribers

- [Overview on page 145](#)
- [Configuring a Static Minimum Adjusted Shaping Rate on Scheduler Nodes on page 146](#)
- [Configuring a Dynamic Minimum Adjusted Shaping Rate on Scheduler Nodes on page 146](#)

### Overview

Absolute adjustments and delta adjustments are performed at the scheduler node level. You can configure a minimum adjusted shaping rate at the scheduler node level using static or dynamic CoS parameters.

This feature is supported for adjustments performed by the ANCP and multicast applications.



**BEST PRACTICE:** For multicast traffic, you can configure a minimum adjusted shaping rate at the queue level. We recommend that you configure the minimum adjusted value at the scheduler node or the queue, but not both.

When you configure a minimum adjusted value for a node and for a scheduler that is referenced by a scheduler map in the same traffic-control-profile, the system uses the minimum value from the scheduler.

---

## Configuring a Static Minimum Adjusted Shaping Rate on Scheduler Nodes

To apply a minimum adjusted shaping rate for a scheduler node:

- Configure the **adjust-minimum** statement for the static traffic-control profile.

```
[edit class-of-service traffic-control-profiles profile-name]  
user@host# set adjust-minimum rate
```

- See Also**
- [Verifying the Scheduling and Shaping Configuration for Subscriber Access](#)
  - [Configuring a Dynamic Shaping-Rate Adjustment for Queues on page 148](#)
  - [Hierarchical CoS Shaping-Rate Adjustments Overview on page 141](#)
  - [Configuring Traffic Scheduling and Shaping for Subscriber Access](#)

## Configuring a Dynamic Minimum Adjusted Shaping Rate on Scheduler Nodes

To apply a minimum adjusted shaping rate for a scheduler node:

- Configure the **adjust-minimum** statement for the dynamic traffic-control profile.

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]  
user@host# set adjust-minimum rate
```

- See Also**
- [Verifying the Scheduling and Shaping Configuration for Subscriber Access](#)
  - [Configuring a Dynamic Shaping-Rate Adjustment for Queues on page 148](#)
  - [Hierarchical CoS Shaping-Rate Adjustments Overview on page 141](#)
  - [Configuring Traffic Scheduling and Shaping for Subscriber Access](#)

- Related Documentation**
- [Verifying the Scheduling and Shaping Configuration for Subscriber Access](#)
  - [Configuring Shaping-Rate Adjustments on Queues on page 146](#)
  - [Hierarchical CoS Shaping-Rate Adjustments Overview on page 141](#)

---

## Configuring Shaping-Rate Adjustments on Queues

- [Overview on page 147](#)
- [Configuring a Static Shaping-Rate Adjustment for Queues on page 147](#)
- [Configuring a Dynamic Shaping-Rate Adjustment for Queues on page 148](#)



## Overview

By default, the multicast application adjusts the shaping rates at the scheduler node level. This adjustment also impacts the shaping rates for all queues, which can be undesirable for service providers who want to provide a premium level of service on specific queues.

For multicast applications, you can control the distribution of shaping rates among queues by assigning the percentage of adjustment allowed for each queue. In addition, you can set a minimum adjusted shaping rate for each queue.

This feature is supported for adjustments performed by the multicast application.



**BEST PRACTICE:** We recommend that you configure the minimum adjusted value at the scheduler node or the queue, but not both.

When you configure a minimum adjusted value for a node and for a scheduler that is referenced by a scheduler map in the same traffic-control-profile, the system uses the minimum value from the scheduler.

## Configuring a Static Shaping-Rate Adjustment for Queues

To configure adjustment parameters for a queue:

1. Configure the percentage of adjustment for the shaping rate.

```
[edit class-of-service schedulers scheduler-name]
user@host# set adjust-percent percentage
```

2. Configure the minimum adjusted value for the shaping rate.

Do one of the following:

- Configure the minimum adjusted value for the queue.

```
[edit class-of-service schedulers scheduler-name]
user@host# set adjust-minimum rate
```

- Configure the minimum adjusted value for the node.

```
[edit class-of-service traffic-control-profile profile-name]
user@host# set adjust-minimum rate
```



**BEST PRACTICE:** Ensure that the minimum adjusted value that you configure does not exceed the shaping rate and is not lower than the configured transmit rate.

**See Also** • *Verifying the Scheduling and Shaping Configuration for Subscriber Access*

- [Hierarchical CoS Shaping-Rate Adjustments Overview on page 141](#)
- [Configuring Traffic Scheduling and Shaping for Subscriber Access](#)

## Configuring a Dynamic Shaping-Rate Adjustment for Queues

To configure adjustment parameters for a queue in a dynamic profile:

1. Configure the percentage of adjustment for the shaping rate.

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]  
user@host# set adjust-percent percentage
```

2. Configure the minimum adjusted value for the shaping rate.

Do one of the following:

- Configure the minimum adjusted value for the queue.

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]  
user@host# set adjust-minimum (rate | $junos-cos-adjust-minimum)
```

- Configure the minimum adjusted value for the node.

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profile  
  profile-name]  
user@host# set adjust-minimum rate
```



**BEST PRACTICE:** Ensure that the minimum adjusted value that you configure does not exceed the shaping rate and is not lower than the configured transmit rate.

- See Also**
- [Verifying the Scheduling and Shaping Configuration for Subscriber Access](#)
  - [Configuring a Dynamic Minimum Adjusted Shaping Rate on Scheduler Nodes on page 146](#)
  - [Hierarchical CoS Shaping-Rate Adjustments Overview on page 141](#)
  - [Configuring Traffic Scheduling and Shaping for Subscriber Access](#)

- Related Documentation**
- [Verifying the Scheduling and Shaping Configuration for Subscriber Access](#)
  - [Configuring the Minimum Adjusted Shaping Rate on Scheduler Nodes for Subscribers on page 145](#)
  - [Hierarchical CoS Shaping-Rate Adjustments Overview on page 141](#)

---

## Enabling Shaping-Rate Adjustments for Subscriber Local Loops

You can enhance a CoS implementation by enabling an MX Series 3D Universal Edge Router to adjust the hierarchical CoS policy shaping rate configured for static interface

sets that consist of two or more VLANs and represent subscriber local loops. Whenever the digital subscriber line access multiplexer (DSLAM) resynchronizes its data transmission rate to a digital subscriber line (DSL), the router adjusts the shaping rate for the associated subscriber interface so that the maximum bandwidth allocation cannot exceed the current data rate for the associated subscriber local loop. This feature ensures that data transmission rate adjustments by the DSLAM do not cause bandwidth contention at the subscriber's residential gateway.

This topic includes the following tasks:

- [Configuring Static Logical Interface Sets to Serve as CoS Hierarchical Scheduler Nodes for Subscriber Loops on page 149](#)
- [Configuring the Logical Interfaces That Compose the Static Logical Interface Sets on page 150](#)
- [Configuring Hierarchical CoS on the Static Logical Interface Sets That Serve as Hierarchical Scheduler Nodes for Subscriber Local Loops on page 150](#)
- [Configuring ANCP Functionality That Supports and Drives Shaping-Rate Adjustments for Subscriber Local Loops on page 152](#)

## Configuring Static Logical Interface Sets to Serve as CoS Hierarchical Scheduler Nodes for Subscriber Loops

To configure a logical interface set, begin by including the **interface-set** statement with the *interface-set-name* option at the **[edit interfaces]** hierarchy level.

An interface set is composed of two or more logical interfaces on the same physical interface. Each logical interface in an interface set corresponds to an individual subscriber service, such as voice, video, or data. To specify either a list of logical unit numbers or the single outer VLAN tag used to identify the logical interfaces that compose the interface set, include statements at the **[edit interfaces interface-set *interface-set-name*]** hierarchy level:

- For an interface set composed of a list of logical interfaces identified by an inner VLAN tag on Ethernet frames (called the customer VLAN, or C-VLAN, tag), you must specify each logical interface by including the **unit** statement with the ***logical-unit-number*** option.

```
[edit]
interfaces {
  interface-set interface-set-name {
    interface ethernet-interface-name { # EQ DPC port
      unit logical-unit-number;
      unit logical-unit-number;
      ...
    }
    ...
  }
}
```

- For an interface set composed of a set of VLANs grouped at the DSLAM and identified by the same service VLAN (S-VLAN) tag), you must specify the S-VLAN tag as the outer VLAN tag for each VLAN by including the **vlan-tags-outer** statement with the **vlan-tag** option.

```
[edit]
interfaces {
  interface-set interface-set-name {
    interface ethernet-interface-name { # EQ DPC port
      vlan-tags-outer vlan-tag; # Identify the DSLAM
    }
    ...
  }
}
```

For more information, see “Configuring Hierarchical Schedulers for CoS” on page 12.

## Configuring the Logical Interfaces That Compose the Static Logical Interface Sets

Each underlying physical interface must be configured to operate in hierarchical scheduler mode and to support stacked VLAN tagging on all logical interfaces. To configure, include the **hierarchical-scheduler** statement and the **stacked-vlan-tagging** statement at the **[edit interfaces *ethernet-interface-name*]** hierarchy level.

To associate the individual logical interfaces of an interface set with specific subscriber services provided by the subscriber local loop, bind an S-VLAN tag and a C-VLAN tag to each logical interface that belongs to a scheduler node that represents a subscriber local loop. Ethernet frames sent from the logical interfaces contain an outer VLAN tag that identifies a DSLAM and an inner VLAN tag that identifies a subscriber port on the DSLAM. To configure, include the **vlan-tags** statement at each logical interface:

```
[edit]
interfaces {
  ethernet-interface-name { # EQ DPC port underlying an interface set
    hierarchical-scheduler;
    stacked-vlan-tagging; # Support 802.1Q VLAN dual-tagged frames
    unit logical-unit-number { # Bind S-VLAN and C-VLAN tags to logical interface
      vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
    }
    ...
  }
}
```

For more information, see *802.1Q VLANs Overview*.

## Configuring Hierarchical CoS on the Static Logical Interface Sets That Serve as Hierarchical Scheduler Nodes for Subscriber Local Loops

To configure hierarchical CoS on the static logical interface set that serves as the hierarchical scheduler node for a subscriber local loop:

1. For each scheduler node that represents a subscriber local loop, configure an initial shaping rate.



**NOTE:** The CoS shaping-rate feature is supported only for scheduler nodes with a configured shaping rate. The initial shaping rate must be configured by applying a traffic-control profile that includes the **shaping-rate** statement. Specify the initial shaping rate as a peak rate, in bits per second (bps), and not as a percentage. Other methods of configuring a shaping rate are not supported with this feature.

- To enable traffic heading downstream (from the router to the DSLAM) to be gathered into an interface set, include the **interface-set** statement and define the logical interface set name as the **interface-set-name** option at the **[edit class-of-service interfaces]** hierarchy level.
- To apply output traffic scheduling and shaping parameters at the logical interface set level (rather than at the logical unit level), include the **output-traffic-control-profile** statement and specify the name of a traffic-control profile as the **profile-name** option at the **[edit class-of-service interfaces interface-set interface-set-name]** hierarchy level.

To configure, include the following statements:

```
interfaces { # Configure interface-specific CoS for incoming packets
  interface-set interface-set-name { # Configure a hierarchical scheduler
    output-traffic-control-profile tc-profile-name; # Level 3 scheduler node
  }
  ...
}
traffic-control-profiles { # Define traffic-control profiles
  tc-profile-name { # Specify a scheduler map and traffic-shaping parameters
    scheduler-map map-name;
    shaping-rate rate; # This is the "configured shaping rate"
    guaranteed-rate (percent percentage | rate);
    delay-buffer-rate (percent percentage | rate);
  }
  ...
}
```

You can include the statements at the following hierarchy levels:

- **[edit class-of-service]**
  - **[edit *dynamic-profiles* *profile-name* class-of-service]**
2. Configure the scheduler maps referenced in the traffic-control profiles applied to the interface sets, the schedulers referenced in those scheduler maps, and the drop profiles referenced in those schedulers.
    - A scheduler map establishes the traffic output queues (forwarding classes) for a scheduler node and associates each queue with a specific scheduler map.
    - A scheduler defines queue properties (transmit rate, buffer size, priority, and drop profile) that specify how traffic is treated in the output queue.

- A drop profile specifies how aggressively the MX Series router drops packets that are managed by a particular scheduler by defining either a segmented or interpolated graph that maps output queue fullness to packet drop probability.

To configure, include the statements at the static **[edit class-of-service]** hierarchy level:

```
[edit]
class-of-service {
  scheduler-maps { # Assign queuing characteristics to output queues
    map-name { # Map output queues to
      forwarding-class class-name scheduler scheduler-name;
      forwarding-class class-name scheduler scheduler-name;
      ...
    }
    ...
  }
  schedulers { # Define queuing characteristics
    scheduler-name { # Specify queuing and buffer management
      transmit-rate transmit-rate-option;
      buffer-size buffer-size-option;
      priority priority-level;
      drop-profile-map loss-priority loss-priority-option protocol any drop-profile
        drop-profile-name;
      ...
    }
  }
  drop-profiles { # Define random early detection (RED) for the delay buffer
    drop-profile-name { # Specify how to drop packets from an output queue
      drop-profile-name { # Map a queue fullness to a drop probability
        fill-level percentage drop-probability percentage; # Option 1: segmented
        fill-level percentage drop-probability percentage;
        ...
      }
      interpolate { # Option 2: interpolated
        drop-probability [ values ];
        fill-level [ values ];
      }
    }
    ...
  }
}
```

For more information about configuring scheduler maps, schedulers, and drop profiles, see *Mapping CoS Component Inputs to Outputs*.

## Configuring ANCP Functionality That Supports and Drives Shaping-Rate Adjustments for Subscriber Local Loops

To configure the Access Node Control Protocol (ANCP) functionality that supports and drives the shaping-rate adjustments for subscriber local loops:

- Enable the ANCP agent to monitor subscriber local loop rates at the DSLAMs and communicate this information to CoS.
- For frame-mode DSL types, optionally configure adjustments that are made to the net data rates, the frame overhead, or both before the ANCP agent reports the values to CoS. Rates are adjusted by a percentage. Bytes are added to or subtracted from the overhead per frame.
- Configure each DSLAM as an ANCP neighbor of the router so that TCP connections can be established between the router and each DSLAM.
- Identify the subscriber interface sets whose traffic is monitored and shaped by the ANCP agent, and associate those interface sets with the corresponding identifiers configured on the access node (DSLAM) to uniquely identify the subscriber local loops within the access network.

The ANCP agent uses this information to build a mapping of subscribers to subscriber interfaces. When the ANCP agent receives port management messages from a DSLAM or other access node, it uses the access identifier contained in the message to determine which hierarchical scheduler node corresponds to the subscriber.

To configure, include statements at the **[edit protocols ancp]** hierarchy level:

```
[edit]
protocols {
  ancp {
    qos-adjust; # Enable ANCP to monitor and adjust CoS shaping rates
    other-bytes bytes; # Specify number of bytes to adjust OTHER access technology
      rate
    other-overhead-adjust percentage; # Specify percentage by which to adjust OTHER
      access technology rate
    sdsl-bytes bytes; # Specify number of bytes to adjust SDSL rate
    sdsl-overhead-adjust percentage; # Specify percentage by which to adjust SDSL
      rate
    vdsl-bytes bytes; # Specify number of bytes to adjust VDSL rate
    vdsl-overhead-adjust percentage; # Specify percentage by which to adjust VDSL
      rate
    vdsl2-bytes bytes; # Specify number of bytes to adjust VDSL2 rate
    vdsl2-overhead-adjust percentage; # Specify percentage by which to adjust VDSL2
      rate
  }
  neighbor ip-address; # Configure each DSLAM as an ANCP neighbor
  ...
  interfaces { # Identify subscribers for which ANCP can adjust shaping rates
    interface-set {
      interface-set-name {
        access-identifier identifier-string; # DSLAM ID for the local loop
      }
    }
    ...
  }
  ...
}
```

- Related Documentation**
- For hardware requirements and configuration guidelines, see [Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops on page 144](#)
  - [Shaping Rate Adjustments for Subscriber Local Loops Overview on page 143](#)
  - *Traffic Rate Reporting and Adjustment by the ANCP Agent*
  - *Configuring the ANCP Agent to Report Traffic Rates to CoS*
  - [Verifying the Configuration of ANCP for Shaping-Rate Adjustments on page 158](#)
  - [Verifying the Configuration of Shaping-Rate Adjustments for Subscriber Local Loops on page 158](#)
  - [Disabling Shaping-Rate Adjustments for Subscriber Local Loops on page 154](#)
  - [Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops on page 155](#)

---

## Disabling Shaping-Rate Adjustments for Subscriber Local Loops

To disable hierarchical CoS shaping-rate adjustments for subscriber local loops:

- Disable hierarchical CoS traffic-shaping adjustment by ANCP:

```
[edit protocols ancp]  
user@host# delete qos-adjust
```

Traffic-shaping parameters for all subscriber local loops revert to their current configured values.

- Related Documentation**
- For hardware requirements and configuration guidelines, see [Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops on page 144](#)
  - [Shaping Rate Adjustments for Subscriber Local Loops Overview on page 143](#)
  - [Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 148](#)
  - [Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops on page 155](#)

---

## Disabling Hierarchical Bandwidth Adjustment for Subscriber Interfaces with Reverse-OIF Mapping

You can disable hierarchical bandwidth adjustment for all subscriber interfaces with reverse OIF mapping enabled on a specified multicast interface. Reverse OIF mapping is used to determine the subscriber VLAN interface and the multicast traffic bandwidth on the interface.

To disable hierarchical bandwidth adjustment:

1. Specify that you want to access the subscriber interfaces with reverse-OIF mapping enabled.



```
[edit routing-instances routing-instance routing-options multicast interface
interface-name]
```

```
user@host# edit reverse-oif-mapping
```

2. Disable hierarchical bandwidth adjustment for all subscriber interfaces on the interface.

```
user@host# set no-qos-adjust
```

- Related Documentation**
- [Hierarchical CoS Shaping-Rate Adjustments Overview on page 141](#)
  - [Example: Configuring Multicast with Subscriber VLANs](#)

## Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops

This example shows how you can enable shaping-rate adjustments for static logical interface sets that represent subscriber local loops:

1. Configure static logical interface sets to serve as CoS hierarchical scheduler nodes for subscriber local loops.

This example uses a single scheduler node that represents two subscriber local loops. The scheduler node is a static logical interface composed of two logical interfaces. The underlying physical interface is port 0 on a Gigabit Ethernet EQ DPC in slot 4, PIC 0:

```
[edit]
interfaces {
  interface-set ifset-of-logical-interfaces {
    interface ge-4/0/0 {
      unit 1;
      unit 2;
    }
  }
  ge-4/0/0 {
    description "access interface ge-4/0/0";
    hierarchical-scheduler;
    stacked-vlan-tagging;
    unit 1 {
      description "DSL type ADSL1 = 0x01";
      proxy-arp;
      vlan-tags outer 1 inner 1; # S-VLAN tag is '1' and C-VLAN tag is '1'
      family inet { # Specify a secondary loopback address
        unnumbered-address lo0.0 preferred-source-address 192.168.7.3;
      }
    }
    unit 2 {
      description "DSL type ADSL1 = 0x01";
      proxy-arp;
      vlan-tags outer 1 inner 2; # S-VLAN tag is '1' and C-VLAN tag is '2'
      family inet { # Specify a secondary loopback address
        unnumbered-address lo0.0 preferred-source-address 192.168.7.4;
      }
    }
  }
}
```

```
    }  
  }
```

2. Begin configuring hierarchical CoS on the static logical interface set that serves as the hierarchical scheduler node for the group of subscriber local loops.

```
[edit]  
class-of-service {  
  interfaces {  
    interface-set ifset-of-logical-interfaces {  
      output-traffic-control-profile tcp-premium-with-4-queues;  
    }  
  }  
}
```

3. Configure the traffic-control profiles that can be applied to the scheduler node:

```
[edit]  
class-of-service {  
  traffic-control-profiles {  
    tcp-basic-rate { # Specify a scheduler map and traffic controls  
      shaping-rate 10m;  
    }  
    tcp-premium-with-4-queues { # Specify a scheduler map and traffic controls  
      scheduler-map smap-premium-4q;  
      shaping-rate 20m;  
      guaranteed-rate 10m;  
      delay-buffer-rate 5m;  
    }  
  }  
}
```

In this example, the **tcp-premium-with-4-queues** traffic-control profile is applied to the interface set. The other profile provides a lower shaping rate and no guaranteed rate.

4. Configure the scheduler map **smap-premium-4q** that is referenced in the traffic-control profile for the scheduler node:

```
[edit]  
class-of-service {  
  scheduler-maps { # Define the queues that comprise each scheduler node  
    smap-premium-4q { # Map each queue in the scheduler node to a scheduler  
      forwarding-class be scheduler be_sch;  
      forwarding-class af scheduler af_sch;  
      forwarding-class ef scheduler ef_sch;  
      forwarding-class nc scheduler nc_sch;  
    }  
  }  
}
```

5. Configure the four schedulers (referenced in the scheduler map) that define the four output queues for the scheduler node:

```
[edit]
class-of-service {
  schedulers { # Define scheduling characteristics of each queue
    be_sch { # Transmit rate and buffer management parameters
      transmit-rate percent 10;
      buffer-size remainder;
      priority low;
    }
    ef_sch { # Transmit rate and buffer management parameters
      ...
    }
    af_sch { # Transmit rate and buffer management parameters
      ...
    }
    nc_sch { # Transmit rate and buffer management parameters
      ...
    }
  }
}
```

6. Enable ANCP to communicate with the DSLAM to adjust the CoS shaping rate for the scheduler node.

You must enable the ANCP feature for performing CoS traffic shaping adjustments, configure the DSLAM as an ANCP neighbor, and specify the DSLAM-assigned identifier for the subscriber local loop represented by the scheduler node. Optionally specify byte or percentage adjustments for frame-mode DSL types.

```
[edit]
protocols {
  ancp {
    qos-adjust; # Enable ANCP to adjust CoS shaping rates and specify rate adjustments
    sds1-bytes 30;
    sds1-overhead-adjust 87;
    vds1-bytes 20;
    vds1-overhead-adjust 95;
    vds12-bytes 20;
    vds12-overhead-adjust 87;
  }
  neighbor 10.2.3.4; # Configure the DSLAM as an ANCP neighbor
  interfaces { # Identify subscribers for which ANCP can adjust shaping rates
    interface-set {
      ifset-of-logical-interfaces {
        access-identifier "dslam port 2/3"; # DSLAM ID for the local loop
      }
    }
  }
}
```



**NOTE:** If ANCP is not yet enabled, the process starts when you commit a configuration that contains the protocols ancp stanza.

7. You can display the configured shaping rate and the adjusted shaping rate for each logical interface set configured for hierarchical CoS, issue the **show class-of-service interface-set** operational command.

**Related Documentation**

- [Hierarchical CoS Shaping-Rate Adjustments Overview on page 141](#)
- [Shaping Rate Adjustments for Subscriber Local Loops Overview on page 143](#)
- [Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops on page 144](#)
- [Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 148](#)

---

## Verifying the Configuration of Shaping-Rate Adjustments for Subscriber Local Loops

**Purpose** Display the configured shaping rate and the adjusted shaping rate for each logical interface set configured for hierarchical CoS.



**NOTE:** After shaping-rate adjustments are enabled and the router has performed shaping-rate adjustments on a scheduler node, you can configure a new shaping rate by including the `shaping-rate` statement in a traffic-control profile and then applying that profile to that scheduler node. However, this new shaping-rate value does not immediately result in shaping traffic at the new rate. The scheduler node continues to be shaped at rate set by ANCP. Only when the ANCP shaping-rate adjustment feature is disabled is the scheduler node shaped at the newly configured shaping-rate.

**Action** Issue the **show class-of-service interface-set** operational command.

**Related Documentation**

- [Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 148](#)

---

## Verifying the Configuration of ANCP for Shaping-Rate Adjustments

**Purpose** Use to display or clear information about the ANCP configuration for shaping-rate adjustments.

- Action**
- To display ANCP neighbor information, issue the ***show ancp neighbor*** operational command.
  - To clear ANCP neighbors, issue the ***clear ancp neighbor*** operational command.
  - To display ANCP subscriber information, issue the ***show ancp subscriber*** operational command.
  - To display ANCP class-of-service information, issue the ***show ancp cos*** operational command.

If ANCP is not yet enabled, the process starts when you commit a configuration that contains the ***protocols ancp*** stanza.

- Related Documentation**
- *ANCP and the ANCP Agent Overview*
  - *Configuring the ANCP Agent*

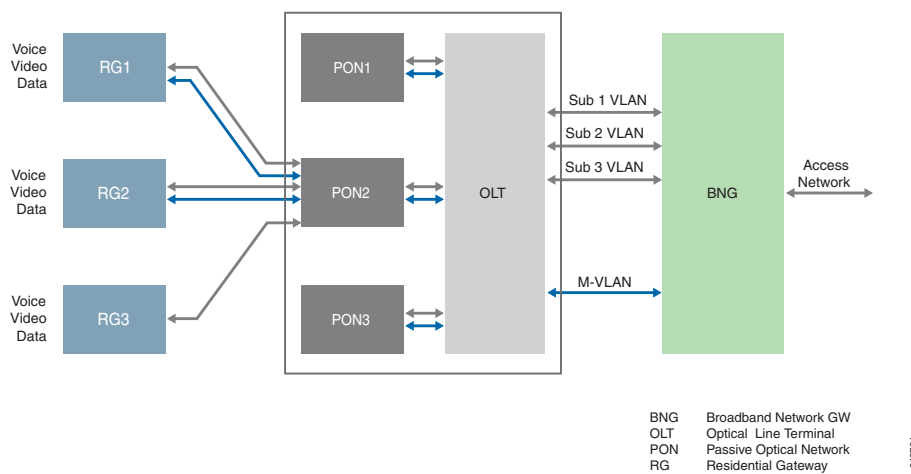
---

## Using Hierarchical CoS to Adjust Shaping Rates Based on Multicast Traffic

For service providers that are using interface sets to deliver services such as voice and data and multicast VLANs (M-VLANs) to deliver broadcast television, you can set up CoS so that when a subscriber begins receiving multicast traffic, the shaping rate of the subscriber interface is adjusted to account for the multicast traffic. You can also set up the class of service (CoS) multicast adjustment to be propagated from the subscriber interface to the interface set, which is the parent in the scheduler hierarchy. This feature prevents oversubscription of the multicast replicator, such as a PON, which can result in dropped traffic and service disruption.

For broadcast television, instead of transporting separate video streams from the source to each subscriber receiving the same stream, the broadband network gateway (BNG) uses M-VLANs to send one stream to the access node. The access node, such as an Optical Line Terminal (OLT) or a DSLAM, replicates the video stream for each subscriber that is currently watching a particular television channel. In this scenario, M-VLANs are used to stream television and interface sets are used to manage traffic to the access node. An interface set contains an interface for each subscriber that is attached to the access node. See [Figure 22 on page 160](#) for a typical broadcast television network topology.

Figure 22: Typical Broadcast Television Network Topology.



When a subscriber begins watching a television channel, the BNG detects that the subscriber has joined a multicast group and has begun receiving traffic from an M-VLAN. The BNG adjusts traffic shaping on the subscriber interface to account for the bandwidth that the multicast traffic is using. For example, if a subscriber begins watching a television channel, the BNG reduces the bandwidth on the subscriber interface to account for the bandwidth being used by the multicast traffic. If you set up CoS to propagate the reduced bandwidth of the subscriber interface to the interface set, the interface set's shaping rate is also reduced.


The BNG uses CoS adjustments based on three-level hierarchical CoS scheduling to implement this feature. Level 2 in the scheduler hierarchy is the interface set, and level 3 is the logical interface. To use this feature, you configure traffic-control profiles for logical interfaces (level 3) to adjust the parent interface set (level 2).

You can set up CoS adjustments to the parent interface set to happen once for each broadcast television channel being streamed over the M-VLAN by including the [qos-adjust-hierarchical interface-set](#) at the `[edit dynamic-profiles dynamic-profile-name access-cac interface $junos-interface-name]` hierarchy. For example:

```
user@host# set dynamic-profiles dynamic-profile-name access-cac interface
$junos-interface-name qos-adjust-hierarchical interface-set
```

To illustrate the benefit of this feature, consider two scenarios from [Figure 23 on page 161](#), which shows a possible subscriber scheduler hierarchy with M-VLANs.

Figure 23: Typical Subscriber Scheduler Hierarchy with M-VLANs



### Scenario 1

Each of 32 subscribers on PON2 are consuming as much data as they can. Each subscriber receives  $(2.4 \text{ Gbps} / 32 \text{ subscribers}) = 75\text{Mbps}$ . Subscriber 1 on PON2 starts watching a 120Mbps multicast video stream. Each subscriber receives  $(2.4\text{Gbps} - 120\text{Mbps}) / 32 \text{ subscribers} = 71.25\text{Mbps}$  of data. Subscriber 1 receives 120Mbps of video, and 71.25Mbps of data. Then Subscriber 2 on PON2 starts watching a *different* 120Mbps multicast video stream. Each subscriber then receives  $(2.4\text{Gbps} - 240\text{Mbps}) / 32 \text{ subscribers} = 67.5 \text{ Mbps}$  of data. Subscribers 1 and 2 receive 120Mbps of video and 67.5Mbps of data.

Final QoS adjustments:

- Subscriber 1's shaping-rate is reduced by 120Mbps to 880Mbps.
- Subscriber 2's shaping-rate is reduced by 120Mbps to 880Mbps.
- PON2's shaping-rate is reduced by 240Mbps to 2.160Gbps, leaving all 32 subscribers to equally share 2.160Gbps – which is 67.5Mbps each.

### Scenario 2

Each of 32 subscribers are consuming as much data as they can. Each subscriber receives  $(2.4 \text{ Gbps} / 32 \text{ subscribers}) = 75\text{Mbps}$ . Subscriber 1 on PON2 starts watching a 120Mbps multicast video stream. Each subscriber then receives  $(2.4\text{Gbps} - 120\text{Mbps}) / 32 \text{ subscribers} = 71.25\text{Mbps}$  of data. Subscriber 1 gets 120Mbps of video, and 71.25Mbps of data. Subscriber 2 on PON2 then starts watching the *same* 120Mbps multicast video

stream. This is completely serviced. Since this is the same video stream as Subscriber 1 is watching, and the stream is being replicated at the PON, only one subscriber adjustment is needed. No additional PON adjustment is needed due to subscriber 2 watching the same 120Mbps multicast video stream.

Final QoS adjustments:

- Subscriber 1's shaping-rate is reduced by 120Mbps to 880Mbps.
- Subscriber 2's shaping-rate is reduced by 120Mbps to 880Mbps.
- PON2's shaping-rate is reduced by 240Mbps to 2.280Gbps, leaving all 32 subscribers to equally share 2.280Gbps – which is 71.25Mbps each.

When any one of the two subscribers receiving the 120Mbps video stream unsubscribes, the adjustment for that subscriber is reverted. The other subscriber's adjustment remains, as does the PON's adjustment. When the remaining subscriber unsubscribes the adjustment for that subscriber is reverted, as is the PON's adjustment.

- Related Documentation**
- [qos-adjust-hierarchical on page 362](#)
  - [show access-cac interface-set on page 432](#)



## CHAPTER 8

# Configuring Targeted Distribution of Subscribers on Aggregated Ethernet Interfaces

- [Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 163](#)
- [Providing Accurate Scheduling for a Demux Subscriber Interface of Aggregated Ethernet Links on page 166](#)
- [Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces on page 167](#)
- [Configuring Link and Module Redundancy for Demux Subscribers in an Aggregated Ethernet Interface on page 168](#)
- [Configuring Rebalancing of Demux Subscribers in an Aggregated Ethernet Interface on page 168](#)
- [Example: Separating Targeted Multicast Traffic for Demux Subscribers on Aggregated Ethernet Interfaces on page 170](#)
- [Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 181](#)
- [Configuring the Distribution Type for PPPoE Subscribers on Aggregated Ethernet Interfaces on page 181](#)
- [Verifying the Distribution of PPPoE Subscribers in an Aggregated Ethernet Interface on page 182](#)

## Distribution of Demux Subscribers in an Aggregated Ethernet Interface

---

This topic describes the distribution options available for demux subscriber interfaces over aggregated Ethernet.

### Distribution Models

By default, the system supports hash-based distribution for all subscriber interface types in an aggregated Ethernet bundle configured without link protection. In this model, traffic for a logical interface can be distributed over multiple links in the bundle. This model is desirable when there are many flows through the logical interface and you need to load balance those flows.

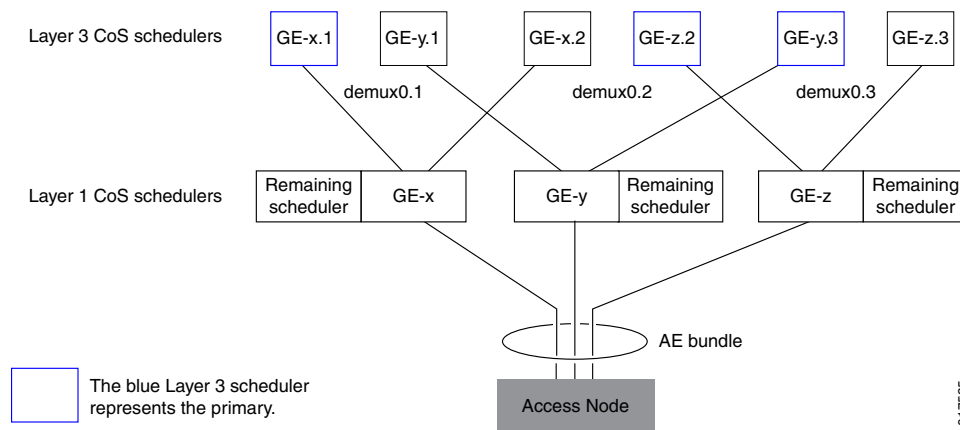
Note that if the distribution flows are not even, egress CoS scheduling can be inaccurate. In addition, scheduler resources are required on every link of the aggregated Ethernet interface. For example, if subscriber traffic is allocated 10 MB for a triple-play service over four links in a bundle, each of the links could receive 2.5 MB of traffic. High-density services such as video could be limited by the bandwidth on one of the links.

Targeted distribution enables you to target the egress traffic for an IP or VLAN demux subscriber on a single member link, using a single scheduler resource. To achieve load balancing over the member links, the system distributes the subscriber interfaces equally among the links. This enables the subscriber that is allocated 10 MB to be accurately scheduled as the traffic flows through.

### Sample Targeted Distribution Topology

Figure 24 on page 164 displays a sample targeted distribution of subscriber traffic across links in an aggregated Ethernet interface. A primary and backup link is allocated for each subscriber.

Figure 24: Targeted Subscriber Links



For example, if link **GE-x** went down, subscriber 1 can begin forwarding over the backup, which is link **GE-y**. When link **GE-y** comes back up, subscriber 1 switches back to its primary link, **GE-x**.

In the event that both **GE-x** and **GE-y** go down, subscriber 3 starts forwarding through its backup, **GE-z**. Subscriber 1 will have lost its primary and backup links, and will also begin forwarding out the **GE-z** link. A new level 3 scheduler is assigned for this subscriber on link **GE-z**. If there is a momentary lapse between the time that a new scheduler is allocated and forwarding switches to **GE-z**, the traffic will be forwarding through to the remaining scheduler. Subscriber 2 continues to forward through its primary link, **GE-z**.

### Redundancy and Redistribution Mechanisms

Two types of redundancy are available in the targeted distribution model: link redundancy and module redundancy.

By default, an aggregated Ethernet interface is enabled with link redundancy. Backup links for a subscriber are chosen based on the link with the least number of subscribers, which provides redundancy if a link fails.

The module redundancy option enables you to provide redundancy if a module or a link fails. Backup links for a subscriber are chosen on a different DPC or MPC from the primary link, based on the link with the least number of subscribers among the links on different modules. You can enable this for the aggregated Ethernet interface.

When links are removed, affected subscribers are redistributed among the active remaining backup links. When links are added to the system, no automatic redistribution occurs. New subscribers are assigned to the links with the fewest subscribers (which are typically the new links).

## Considerations and Best Practices

Keep the following guidelines in mind when configuring targeted distribution for demux subscribers:

- You can manage subscribers with both hash-based and targeted distribution models in the same network. For example, you can allocate subscribers with interface types such as PPPoE with hash-based distribution, and enable demux subscribers with targeted distribution.
- We recommend that you configure module redundancy to protect against module failures. When module redundancy is enabled, you can ensure an even distribution of subscribers if you allocate no more than 50 percent of the links on a single DPC or MPC.
- During normal network operations, the system maintains an even balance of subscribers among the links in a bundle, even as subscribers log in and out. However, if the distribution of a bundle becomes uneven (for example, when a link goes down and new subscribers are logging in), you can perform a manual rebalance of the bundle. In addition, you can configure periodic rebalancing of the bundle with a specific time interval.
- When you anticipate that a link will be down for an extended time, and you want to ensure that backup links are provisioned for all subscribers, we recommend that you remove the failed link from the bundle. This forces the affected subscribers to redistribute to other links.
- We recommend that you apply a remaining traffic-control profile to the logical interface to ensure that minimal scheduling parameters are applied to the remaining subscriber traffic. This provides scheduling for subscribers that do not have schedulers allocated because they have not been configured or they have been over-provisioned, or because of scheduler transitions on multiple link failures.
- If you perform a cold restart on the router when it is forwarding active subscribers, the subscriber interfaces with targeted distribution are assigned to the first links that become available when the system is initializing so forwarding can begin. To rebalance the system following a cold restart, perform a manual rebalance of the bundle. In addition, we recommend that you configure Graceful Routing Engine switchover (GRES) on the router to enable nonstop forwarding during switchover, and avoid performing cold restarts.

- To ensure appropriate and predictable targeted distribution, you must configure chassis network services to use **enhanced-ip** mode.
- Unless specifically separated, multicast traffic egresses in parallel with unicast traffic, sharing the CoS hierarchy and aggregated Ethernet flow distribution.

**Related Documentation**

- [Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces on page 167](#)
- [Configuring Link and Module Redundancy for Demux Subscribers in an Aggregated Ethernet Interface on page 168](#)
- [Configuring Rebalancing of Demux Subscribers in an Aggregated Ethernet Interface on page 168](#)
- [Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet Overview](#)

## Providing Accurate Scheduling for a Demux Subscriber Interface of Aggregated Ethernet Links

---

Unlike VLAN subscriber interfaces, enabling link protection is not required for configuring hierarchical CoS on demux interfaces. Instead, we recommend that you enable targeted distribution on the demux interface to provide accurate scheduling for the aggregated Ethernet links.

Before you begin, configure the subscriber interface with aggregated Ethernet:

- For static and dynamic IP demux interfaces, see *Configuring a Static or Dynamic IP Demux Subscriber Interface over Aggregated Ethernet*.
- For static and dynamic VLAN demux interfaces, see *Configuring a Static or Dynamic VLAN Demux Subscriber Interface over Aggregated Ethernet*.

To provide accurate scheduling for a demux subscriber interface of aggregated Ethernet links:

1. Enable targeted distribution for the demux interface.

See “[Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces](#)” on page 167.

2. Enable hierarchical scheduling on the link aggregation bundle.

See “[Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links](#)” on page 79.

3. (Optional) Enable module redundancy to ensure that CoS resources are provisioned for the aggregated Ethernet links if a module or a link fails. By default, link redundancy is supported.

See “[Configuring Link and Module Redundancy for Demux Subscribers in an Aggregated Ethernet Interface](#)” on page 168.

4. (Optional) Configure rebalancing periodically or manually for the subscribers. See [“Configuring Rebalancing of Demux Subscribers in an Aggregated Ethernet Interface” on page 168](#).
5. Attach static or dynamic traffic shaping and scheduling parameters at the aggregated Ethernet logical interface or its underlying physical interface. See:
  - *Configuring Traffic Scheduling and Shaping for Subscriber Access*
  - *Configuring Schedulers in a Dynamic Profile for Subscriber Access*
  - *Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile*
  - *Applying Minimal Shaping and Scheduling to Remaining Subscriber Traffic*

**Related  
Documentation**

- *Guidelines for Configuring Dynamic CoS for Subscriber Access*
- [Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 181](#)

---

## Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces

---

By default, the system supports hash-based distribution of subscriber traffic in aggregated Ethernet bundles. You can configure the system to target the egress traffic for a subscriber on a single member link, using a single scheduler resource. The system distributes the subscriber interfaces equally among the member links.

To configure targeted distribution:

1. Edit the chassis hierarchy level.

```
[edit]
user@host#edit chassis
```

2. Enable chassis network services for **enhanced-ip** mode.

```
[edit chassis]
user@host#set network-services enhanced-ip
```

3. Access the logical interface.

```
[edit]
user@host#edit interfaces demux0 unit logical-unit-number
```

4. Enable targeted distribution for the interface.

```
[edit interfaces demux0 unit logical-unit-number]
user@host#set targeted-distribution
```

**Related  
Documentation**

- [Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 181](#)

- [Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 163](#)

## Configuring Link and Module Redundancy for Demux Subscribers in an Aggregated Ethernet Interface

---

By default, an aggregated Ethernet bundle with targeted distribution is enabled with link redundancy. Backup links for a subscriber are chosen based on the link with the fewest subscribers, which provides redundancy if a link fails.

We recommend that you configure the module redundancy option to provide redundancy if a module or a link fails. Backup links for a subscriber are chosen on a different DPC or MPC from the primary link, based on the link with the fewest subscribers among the links on different modules.

To configure module redundancy for an aggregated Ethernet bundle:

1. Access the aggregated Ethernet bundle for which you want to configure module redundancy.

```
edit
user@host# edit interfaces aex aggregated-ether-options
```

2. Enable module redundancy for the bundle.

```
[edit interfaces aex aggregated-ether-options]
user@host# logical-interface-fpc-redundancy
```

### Related Documentation

- [Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces on page 167](#)
- [Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 163](#)

## Configuring Rebalancing of Demux Subscribers in an Aggregated Ethernet Interface

---

In a targeted distribution model, the system allocates demux subscriber interfaces equally among the member links in the aggregated Ethernet interface. When links are removed, affected subscribers are redistributed among the active remaining backup links. When links are added to the system, no automatic redistribution occurs. New subscribers are assigned to the links with the fewest subscribers (which are typically the new links).

During normal network operations, the system maintains an even balance of traffic among the links in a bundle, even as subscribers log in and out. However, if the distribution of a bundle becomes uneven (for example, when a link goes down for a period of time

and new subscribers are logging in), you can perform a manual rebalance of the bundle. In addition, you can configure periodic rebalancing of the bundle with a specific interval.

- [Configuring Periodic Rebalancing of Subscribers in an Aggregated Ethernet Interface on page 169](#)
- [Configuring Manual Rebalancing of Subscribers on an Aggregated Ethernet Interface on page 169](#)

## Configuring Periodic Rebalancing of Subscribers in an Aggregated Ethernet Interface

If subscribers are frequently logging in and logging out of your network, you can configure the system to periodically rebalance the links based on a specific time and interval.

To configure periodic rebalancing:

1. Access the aggregated Ethernet interface for which you want to configure periodic rebalancing.

```
edit
user@host# edit interfaces aenumber aggregated-ether-options
```

2. Configure the rebalancing parameters for the interface, including the time and the interval between rebalancing actions.

```
[edit interfaces aenumber aggregated-ether-options]
user@host# rebalance-periodic time hour:minute <interval hours>
```

- See Also**
- [Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 181](#)
  - [Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces on page 167](#)
  - [Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 163](#)

## Configuring Manual Rebalancing of Subscribers on an Aggregated Ethernet Interface

To manually rebalance the subscribers among the links in an aggregated Ethernet bundle with targeted distribution:

- Issue the **request interface rebalance** command:

```
user@host# request interface rebalance interface <interface-name>
```

- See Also**
- [Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 181](#)
  - [Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces on page 167](#)
  - [Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 163](#)

- Related Documentation**
- [Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 181](#)
  - [Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces on page 167](#)
  - [Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 163](#)

## Example: Separating Targeted Multicast Traffic for Demux Subscribers on Aggregated Ethernet Interfaces

---

This example shows how to separate targeted multicast traffic from targeted unicast traffic and send that multicast traffic to a different interface through the use of OIF maps.

- [Requirements on page 170](#)
- [Overview on page 170](#)
- [Configuration on page 171](#)
- [Verification on page 176](#)

### Requirements

Before configuring this example, make sure to configure the distribution type for the interface. See [“Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces” on page 167](#) for instructions.

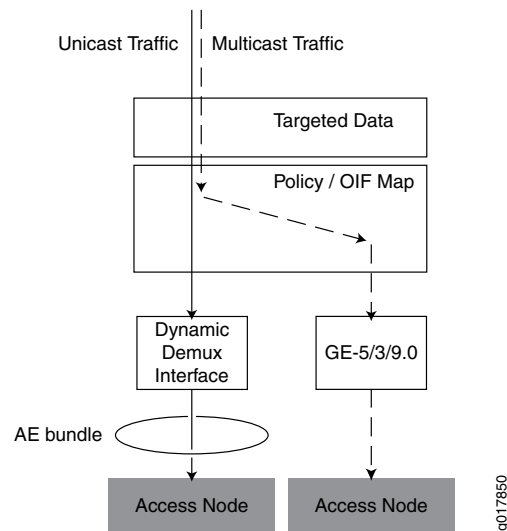
### Overview

In this example, targeted traffic distribution is already configured on the router. Dynamically created interfaces each carry their unicast traffic but all multicast traffic is sent to the GE-5/3/9.0 interface.

[Figure 25 on page 171](#) shows the sample network.



Figure 25: Multicast Traffic Separation Using OIF Mapping



## Configuration

- [Configure an OIF Map Policy on page 172](#)
- [Configure a DHCP VLAN Dynamic Profile on page 173](#)
- [Configure a VLAN Demux Dynamic Profile on page 174](#)

### CLI Quick Configuration

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

```
set policy-options policy-statement OIF-v4-all term oif539 from route-filter 192.168.20.0/4
  orlonger
set policy-options policy-statement OIF-v4-all term oif539 then map-to-interface
  ge-5/3/9.0
set policy-options policy-statement OIF-v4-all term oif539 then accept
set dynamic-profiles dhcp-vlan-prof interfaces "$junos-interface-ifd-name" unit
  "$junos-underlying-interface-unit" family inet unnumbered-address lo0.0
set dynamic-profiles dhcp-vlan-prof interfaces "$junos-interface-ifd-name" unit
  "$junos-underlying-interface-unit" family inet unnumbered-address preferred-source-address 10.20.0.2
set dynamic-profiles demux-vlan-prof interfaces demux0 unit "$junos-interface-un it"
  vlan-id "$junos-vlan-id"
set dynamic-profiles demux-vlan-prof interfaces demux0 unit "$junos-interface-un it"
  demux-options underlying-interface "$junos-interface-ifd-name"
set dynamic-profiles demux-vlan-prof interfaces demux0 unit "$junos-interface-un it"
  targeted-distribution
set dynamic-profiles demux-vlan-prof interfaces demux0 unit "$junos-interface-un it"
  family inet unnumbered-address lo0.0
set dynamic-profiles demux-vlan-prof interfaces demux0 unit "$junos-interface-un it"
  family inet unnumbered-address preferred-source-address 10.20.0.2
set dynamic-profiles demux-vlan-prof protocols igmp interface "$junos-interface- name"
  version 2
```

```
set dynamic-profiles demux-vlan-prof protocols igmp interface "$junos-interface- name"  
promiscuous-mode  
set dynamic-profiles demux-vlan-prof protocols igmp interface "$junos-interface- name"  
passive allow-receive  
set dynamic-profiles demux-vlan-prof protocols igmp interface "$junos-interface- name"  
passive send-group-query  
set dynamic-profiles demux-vlan-prof protocols igmp interface "$junos-interface- name"  
oif-map OIF-v4-all
```

---

### Configure an OIF Map Policy

---

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy.

To configure the OIF map:

1. Access the router policy options:

```
[edit]  
user@host#edit policy-options
```

2. Edit a policy statement.

```
[edit policy-options]  
user@host edit policy-statement OIF-v4-all
```

3. Create a term for mapping incoming multicast traffic to a specific interface.

```
[edit policy-options OIF-v4-all]  
user@host edit term oif539
```

4. Define the match condition for the term. In this case, the term matches any route prefix of 192.168.20.0/4 or longer (all multicast traffic).

```
[edit policy-options OIF-v4-all term oif539]  
user@host set from route-filter 192.168.20.0/4 orlonger
```

5. Define the action for the term. In this case, when a match occurs, the term accepts the traffic and maps it to interface GE-5/3/9.0.

```
[edit policy-options OIF-v4-all term oif539]  
user@host set then map-to-interface ge-5/3/9.0  
user@host set then accept
```

#### Results

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

```
[edit]  
user@host# show policy-options  
policy-statement OIF-v4-all {  
  term oif539 {
```

```

    from {
        route-filter 192.168.20.0/4 orlonger;
    }
    then {
        map-to-interface ge-5/3/9.0;
        accept;
    }
}
}

```

### Configure a DHCP VLAN Dynamic Profile

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy.

To configure a DHCP VLAN dynamic profile for client access:

1. Create a dynamic VLAN demux profile.

```

[edit]
user@host#edit dynamic-profiles dhcp-vlan-prof

```

2. Edit the dynamic profile interface.

```

[edit dynamic-profiles dhcp-vlan-prof]
user@host edit interfaces $junos-ifd-name

```

3. Edit the interface unit dynamic variable.

```

[edit dynamic-profiles demux-vlan-prof interfaces $junos-ifd-name]
user@host edit unit $junos-underlying-interface-unit

```

4. Edit the interface family.

```

[edit dynamic-profiles demux-vlan-prof interfaces $junos-ifd-name unit
 $junos-underlying-interface-unit]
user@host edit family inet

```

5. Define the loopback address.

```

[edit dynamic-profiles demux-vlan-prof interfaces $junos-ifd-name unit
 $junos-underlying-interface-unit ]
user@host set unnumbered-address lo0.0 preferred-source-address 10.20.0.2

```

**Results** Confirm your configuration by issuing the **show dynamic-profiles** command. If the output for the dhcp-vlan-prof dynamic profile does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

[edit]
user@host# show dynamic-profiles
dhcp-vlan-prof {
    interfaces {

```

```
"$junos-interface-ifd-name" {  
  unit "$junos-underlying-interface-unit" {  
    family inet {  
      unnumbered-address lo0.0 preferred-source-address 10.20.0.2;  
    }  
  }  
}  
}
```

---

### Configure a VLAN Demux Dynamic Profile

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy.

To configure the OIF map:

1. Create a dynamic VLAN demux profile.

```
[edit]  
user@host#edit dynamic-profiles demux-vlan-prof
```

2. Edit the dynamic profile demux0 interface.

```
[edit dynamic-profiles demux-vlan-prof]  
user@host edit interfaces demux0
```

3. Edit the interface unit dynamic variable.

```
[edit dynamic-profiles demux-vlan-prof interfaces demux0]  
user@host edit unit $junos-interface-unit
```

4. Specify the VLAN ID dynamic variable.

```
[edit dynamic-profiles demux-vlan-prof interfaces demux0 unit  
"$junos-interface-unit"]  
user@host set vlan-id $junos-vlan-id
```

5. Access the demux options.

```
[edit dynamic-profiles demux-vlan-prof interfaces demux0 unit  
"$junos-interface-unit"]  
user@host edit demux-options
```

6. Define the demux underlying interface.

```
[edit dynamic-profiles demux-vlan-prof interfaces demux0 unit  
"$junos-interface-unit" demux-options]  
user@host set underlying-interface $junos-interface-ifd-name
```

7. Specify that dynamically created VLANs are using targeted distribution.  

```
[edit dynamic-profiles demux-vlan-prof interfaces demux0 unit  
"$junos-interface-unit"]  
user@host set targeted-distribution
```
8. Edit the interface family.  

```
[edit dynamic-profiles demux-vlan-prof interfaces demux0 unit  
"$junos-interface-unit"]  
user@host edit family inet
```
9. Define the loopback address.  

```
[edit dynamic-profiles demux-vlan-prof interfaces demux0 unit  
"$junos-interface-unit" family inet]  
user@host set unnumbered-address lo0.0 preferred-source-address 10.20.0.2
```
10. Edit the dynamic profile IGMP protocol.  

```
[edit dynamic-profiles demux-vlan-prof]  
user@host edit protocols igmp
```
11. Enable IGMP on dynamically created interfaces.  

```
[edit dynamic-profiles demux-vlan-prof protocols igmp]  
user@host edit interface $junos-interface-name
```
12. Specify the IGMP version that you want dynamically created interfaces to use.  

```
[edit dynamic-profiles demux-vlan-prof protocols igmp interface  
$junos-interface-name]  
user@host set version 2
```
13. Specify the OIF map that you want dynamically created IGMP interfaces to use.  

```
[edit dynamic-profiles demux-vlan-prof protocols igmp interface  
$junos-interface-name]  
user@host set oif-map OIF-v4-all
```
14. Specify that IGMP selectively sends and receives control traffic such as IGMP reports, queries, and leaves.  

```
[edit dynamic-profiles demux-vlan-prof protocols igmp interface  
$junos-interface-name]  
user@host set passive allow-receive send-group-query
```
15. Specify that the interface accepts IGMP reports from hosts on any subnetwork.  

```
[edit dynamic-profiles demux-vlan-prof protocols igmp interface  
$junos-interface-name]  
user@host set promiscuous-mode
```

**Results** Confirm your configuration by issuing the **show dynamic-profiles** commands. If the output for the dhcp-vlan-prof dynamic profile does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show dynamic-profiles
demux-vlan-prof {
  interfaces {
    demux0 {
      unit "$junos-interface-unit" {
        vlan-id "$junos-vlan-id";
        demux-options {
          underlying-interface "$junos-interface-ifd-name";
        }
        targeted-distribution;
        family inet {
          unnumbered-address lo0.0 preferred-source-address 10.20.0.2;
        }
      }
    }
  }
  protocols {
    igmp {
      interface "$junos-interface-name" {
        version 2;
        promiscuous-mode;
        passive allow-receive send-group-query;
        oif-map OIF-v4-all;
      }
    }
  }
}
...
```

## Verification

Confirm that the configuration is working properly.

- [Locate the Multicast Group Member on page 176](#)
- [Ensure the Targeting Aggregated Ethernet Interface for the Subscriber is Functional on page 177](#)
- [View the Packets for the Targeted Interface on page 178](#)

---

### Locate the Multicast Group Member

**Purpose** Locate the dynamic interface and ensure that it is associated with the appropriate IGMP group.

**Action** user@host>show igmp group

```
Interface: demux0.1073741824, Groups: 1
  Group: 192.0.2.1
    Source: 0.0.0.0
    Last reported by: 10.20.0.10
    Timeout: 52 Type: Dynamic
Interface: local, Groups: 2
  Group: 192.0.2.2
    Source: 0.0.0.0
    Last reported by: Local
    Timeout: 0 Type: Dynamic
  Group: 192.0.2.22
    Source: 0.0.0.0
    Last reported by: Local
    Timeout: 0 Type: Dynamic
```

**Meaning** The first **Interface** field shows the dynamically created demux interface, **demux0.1073741824**, and the **Group** field immediately below the first **Interface** field shows the group, **192.0.2.1**, to which the subscriber belongs.

#### Ensure the Targeting Aggregated Ethernet Interface for the Subscriber is Functional

---

**Purpose** Use the dynamic subscriber interface value to ensure that the targeting aggregated interface is functional.

**Action** user@host>show interfaces demux0.1073741824 extensive

```
Logical interface demux0.1073741824 (Index 810) (SNMP ifIndex 1613)
(Generation 170)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2
  Demux:
    Underlying interface: ae0 (Index 708)
  Link:
    ge-1/0/0
    ge-5/3/7
  Targeting summary:
    ge-1/0/0, backup, Physical link is Up
    ge-5/3/7, primary, Physical link is Up
  Traffic statistics:
    Input bytes :                862
    Output bytes :               3160
    Input packets:                3
    Output packets:              30
  Local statistics:
    Input bytes :                862
    Output bytes :               3160
    Input packets:                3
    Output packets:              30
  Transit statistics:
    Input bytes :                0          0 bps
    Output bytes :                0          0 bps
    Input packets:                0          0 pps
    Output packets:               0          0 pps
  Protocol inet, MTU: 1500, Generation: 212, Route table: 0
  Flags: Sendbcst-pkt-to-re, Unnumbered
  Donor interface: lo0.0 (Index 802)
  Preferred source address: 10.20.0.2
```

**Meaning** The **Targeting summary** field shows that the primary interface, **ge-5/3/7**, is up.

---

#### [View the Packets for the Targeted Interface](#)

**Purpose** Verify that packet traffic sent to targeted interface GE-5/3/9 consists only of multicast packets.



**Action** user@host>show interfaces ge-5/3/9 extensive

```

Physical interface: ge-5/3/9, Enabled, Physical link is Up
Interface index: 704, SNMP ifIndex: 1605, Generation: 197
Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Disabled, Auto-negotiation: Enabled, Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Schedulers    : 0
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:21:59:ab:85:2a, Hardware address: 00:21:59:ab:85:2a
Last flapped   : 2012-09-26 17:32:24 EDT (6d 20:44 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes :          97857650          1320 bps
  Output bytes :              0          0 bps
  Input packets:         889615          1 pps
  Output packets:              0        889620 pps
IPv6 transit statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:              0
Dropped traffic statistics due to STP State:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:              0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

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

Queue number:        Mapped forwarding classes
  0                  best-effort
  1                  expedited-forwarding
  2                  assured-forwarding
  3                  network-control
Active alarms : None
Active defects : None
MAC statistics:
  Total octets          Receive          Transmit
  Total packets         0          113871616
  Unicast packets       0          889620
  Multicast packets     0              0

```

```

Broadcast packets                0                0
Multicast packets              0              889620
CRC/Align errors                 0                0
FIFO errors                      0                0
MAC control frames               0                0
MAC pause frames                 0                0
Oversized frames                 0
Jabber frames                    0
Fragment frames                  0
VLAN tagged frames               0
Code violations                  0
Total errors                     0                0
Filter statistics:
  Input packet count              0
  Input packet rejects            0
  Input DA rejects                0
  Input SA rejects                0
  Output packet count              0              889620
  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, Remote fault: OK
  Local resolution:
    Flow control: None, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 0 (0x00)
CoS information:
  Direction : Output
  CoS transmit queue              Bandwidth              Buffer Priority  Limit
                                %          bps          %          usec
  0 best-effort                  95          950000000    95           0        low    none
  3 network-control              5           500000000     5           0        low    none
Interface transmit statistics: Disabled

Logical interface ge-5/3/9.0 (Index 818) (SNMP ifIndex 1597) (Generation 149)
Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
Traffic statistics:
  Input bytes :                0
  Output bytes :              97857650
  Input packets:                0
  Output packets:              889620
Local statistics:
  Input bytes :                0
  Output bytes :                0
  Input packets:                0
  Output packets:                0
Transit statistics:
  Input bytes :                0              0 bps
  Output bytes :              97857650      1320 bps
  Input packets:                0              0 pps
  Output packets:              889615       1 pps
Protocol aenet, AE bundle: ae4.0, Generation: 180, Route table: 0

```

**Meaning** The MAC statistics **Unicast packet** field shows that the interface is not transmitting any unicast packet traffic and the **Multicast packet** field shows that the total number of packets being transmitted from the interface are multicast packets.

**Related Documentation**

- [Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces on page 167](#)

## Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface

**Purpose** View the distribution status of subscribers that are targeted to links in an aggregated Ethernet interface.

**Action**

- To display a summary of the distribution of links on the demux interface:  
user@host> **show interfaces demux0 extensive**
- To display the targeted distribution on a specific aggregated Ethernet interface:  
user@host> **show interfaces targeting aex**

**Related Documentation**

- [Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces on page 167](#)
- [Configuring Rebalancing of Demux Subscribers in an Aggregated Ethernet Interface on page 168](#)

## Configuring the Distribution Type for PPPoE Subscribers on Aggregated Ethernet Interfaces

By default, the system supports hash-based distribution of subscriber traffic in aggregated Ethernet bundles. You can configure the system to target the egress traffic for a subscriber on a single member link, using a single scheduler resource. The system distributes the subscriber interfaces equally among the member links.

To configure targeted distribution:

1. Edit the chassis hierarchy level.

```
[edit]  
user@host#edit chassis
```

2. Enable chassis network services for **enhanced-ip** mode.

```
[edit chassis]  
user@host#set network-services enhanced-ip
```

3. Access the logical interface.

```
[edit]
```

```
user@host#edit interfaces pp0 unit logical-unit-number
```

4. Enable targeted distribution for the interface.

```
[edit interfaces pp0 unit logical-unit-number]  
user@host#set targeted-distribution
```

- Related Documentation**
- *CoS for PPPoE Subscriber Interfaces Overview*
  - [Verifying the Distribution of PPPoE Subscribers in an Aggregated Ethernet Interface on page 182](#)

---

## Verifying the Distribution of PPPoE Subscribers in an Aggregated Ethernet Interface

**Purpose** View the distribution status of subscribers that are targeted to links in an aggregated Ethernet interface.

- Action**
- To display a summary of the distribution of links on the demux interface:  

```
user@host> show interfaces pp0 extensive
```
  - To display the targeted distribution on a specific aggregated Ethernet interface:  

```
user@host> show interfaces targeting aex
```

- Related Documentation**
- *CoS for PPPoE Subscriber Interfaces Overview*
  - [Configuring the Distribution Type for PPPoE Subscribers on Aggregated Ethernet Interfaces on page 181](#)

## CHAPTER 9

# Applying CoS Using Parameters Received from RADIUS

- [Subscriber Interfaces That Provide Initial CoS Parameters Dynamically Obtained from RADIUS on page 183](#)
- [Changing CoS Services Overview on page 187](#)
- [CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview on page 190](#)
- [Guidelines for Configuring CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions on page 192](#)
- [Configuring Initial CoS Parameters Dynamically Obtained from RADIUS on page 193](#)
- [Configuring Static Default Values for Traffic Scheduling and Shaping on page 194](#)
- [Applying CoS Traffic-Shaping Attributes to Dynamic Interface Sets and Member Subscriber Sessions on page 195](#)
- [CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets on page 198](#)
- [Example: Configuring Dynamic Hierarchical Scheduling for Subscribers on page 203](#)

## Subscriber Interfaces That Provide Initial CoS Parameters Dynamically Obtained from RADIUS

---

You can configure interface-specific CoS parameters that the router obtains when subscribers log in at appropriately configured static or dynamic subscriber interfaces. This feature is supported only for interfaces on Enhanced Queuing Dense Port Concentrators (EQ DPCs) in MX Series 3D Universal Edge Routers.

To configure a dynamic profile to provide initial CoS Services, make sure you understand the following concepts:

- [Dynamic Configuration of Initial CoS in Access Profiles on page 184](#)
- [Predefined Variables for Dynamic Configuration of Initial Traffic Shaping on page 184](#)
- [Predefined Variables for Dynamic Configuration of Initial Scheduling and Queuing on page 185](#)

## Dynamic Configuration of Initial CoS in Access Profiles

When a router interface receives a join message from a DHCP subscriber, the Junos OS applies the values configured in the dynamic profile associated with that router interface. A dynamic profile that is activated through its association with a subscriber interface is known as an *access dynamic profile*. You can associate a dynamic profile with a subscriber interface on the router by including statements at the **[edit dynamic-profiles profile-name class-of-service interfaces]** hierarchy level.

The Junos OS supports predefined variables for obtaining CoS parameters from the RADIUS authentication server. When a client authenticates over a router interface associated with the access dynamic profile, the router replaces the predefined variables with interface-specific values obtained from the RADIUS server.



**NOTE:** To associate dynamically configured initial CoS features with a subscriber interface, reference *Junos OS predefined variables*—and not *user-defined variables*—in an access dynamic profile for that interface.

## Predefined Variables for Dynamic Configuration of Initial Traffic Shaping

You can configure an access dynamic profile that provides initial traffic-shaping parameters when a subscriber logs in. The Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.

If you define the Juniper Networks authentication and authorization VSA for CoS traffic-shaping parameter values (attribute number 26–108) on the RADIUS authentication server, the RADIUS server includes the values in RADIUS Access-Accept messages it sends to the router when a subscriber successfully authenticates over the interface.

To provide an initial scheduler map name and traffic shaping parameters obtained from the RADIUS authentication server when a subscriber logs in, reference the Junos OS predefined variables for CoS listed in [Table 20 on page 184](#) in an access dynamic profile associated with the subscriber interface.

**Table 20: CoS Predefined Variables for Scheduler Map and Traffic Shaping**

Variable	Description
\$junos-cos-scheduler-map	Scheduler-map name to be dynamically configured in a traffic-control profile in the access dynamic profile when a subscriber logs in.  <b>NOTE:</b> The scheduler map referenced by the <b>scheduler-map</b> statement can be defined dynamically (at the <b>[edit dynamic-profiles profile-name class-of-service scheduler-maps]</b> hierarchy level) or statically (at the <b>[edit class-of-service scheduler-maps]</b> hierarchy level).
\$junos-cos-shaping-rate	Shaping rate to be dynamically configured in a traffic-control profile in the access dynamic profile when a subscriber logs in. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.

Table 20: CoS Predefined Variables for Scheduler Map and Traffic Shaping (*continued*)

Variable	Description
\$junos-cos-guaranteed-rate	Guaranteed rate to be dynamically configured in a traffic-control profile in the access dynamic profile when a subscriber logs in. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.
\$junos-cos-delay-buffer-rate	Delay-buffer rate to be dynamically configured in a traffic-control profile in the access dynamic profile when a subscriber logs in. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.

### Predefined Variables for Dynamic Configuration of Initial Scheduling and Queuing

You can configure an access dynamic profile that provides initial traffic-shaping parameters when a subscriber logs in. The Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.

If you define the Juniper Networks authentication and authorization VSA for CoS scheduling and queuing parameter values (attribute number 26–146) on the RADIUS authentication server, the RADIUS server includes the values in RADIUS Access-Accept messages it sends to the router when a subscriber successfully authenticates over the interface.

To provide an initial scheduler name and scheduler and queuing parameters obtained from the RADIUS authentication server when a subscriber logs in, reference the Junos OS predefined variables listed in [Table 21 on page 185](#) in an access dynamic profile associated with the subscriber interface.

Table 21: CoS Predefined Variables for Scheduling and Queuing

Variable	Description
\$junos-cos-scheduler	Name of a scheduler to be dynamically configured in the access dynamic profile. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.
\$junos-cos-scheduler-transmit-rate	Transmit rate to be dynamically configured for the scheduler in the access dynamic profile. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.
\$junos-cos-scheduler-bs	Buffer size, as a percentage of total buffer, to be dynamically configured for the scheduler in the access dynamic profile. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.

Table 21: CoS Predefined Variables for Scheduling and Queuing (*continued*)

Variable	Description
\$junos-cos-scheduler-pri	<p>Packet-scheduling priority value to be dynamically configured for the scheduler in the access dynamic profile. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.</p>
\$junos-cos-scheduler-dropfile-low	<p>Name of the drop profile for RED for loss-priority level <b>low</b> to be dynamically configured for the scheduler in the access dynamic profile. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.</p> <p><b>NOTE:</b> The drop profile must be configured statically (at the <b>[edit class-of-service schedulers <i>scheduler-name</i> drop-profiles]</b> hierarchy level) for loss-priority <b>low</b>.</p>
\$junos-cos-scheduler-dropfile-medium-low	<p>Name of the drop profile for RED for loss-priority level <b>medium-low</b> to be dynamically configured for the scheduler in the access dynamic profile. The Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.</p> <p><b>NOTE:</b> The drop profile must be configured statically (at the <b>[edit class-of-service schedulers <i>scheduler-name</i> drop-profiles]</b> hierarchy level).</p>
\$junos-cos-scheduler-dropfile-medium-high	<p>Name of the drop profile for RED for loss-priority level <b>medium-high</b> to be dynamically configured for the scheduler in the access dynamic profile. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.</p> <p><b>NOTE:</b> The drop profile must be configured statically (at the <b>[edit class-of-service schedulers <i>scheduler-name</i> drop-profiles]</b> hierarchy level).</p>
\$junos-cos-scheduler-dropfile-high	<p>Name of the drop profile for RED for loss-priority level <b>high</b> to be dynamically configured for the scheduler in the access dynamic profile. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.</p> <p><b>NOTE:</b> The drop profile must be configured statically (at the <b>[edit class-of-service schedulers <i>scheduler-name</i> drop-profiles]</b> hierarchy level).</p>
\$junos-cos-scheduler-dropfile-any	<p>Name of the drop profile for RED for loss-priority level <b>any</b> to be dynamically configured for the scheduler in the access dynamic profile. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.</p> <p><b>NOTE:</b> The drop profile must be configured statically (at the <b>[edit class-of-service schedulers <i>scheduler-name</i> drop-profiles]</b> hierarchy level).</p>



**Related Documentation**

- *Subscriber Activation and Service Management in an Access Network*
- *Dynamic Profiles Overview*
- *Dynamic Variables Overview*
- *Junos OS Predefined Variables*
- [Configuring Initial CoS Parameters Dynamically Obtained from RADIUS on page 193](#)
- *Example: Configuring Initial CoS Parameters Dynamically Obtained from RADIUS*

---

## Changing CoS Services Overview

This topic describes how to provide CoS when subscribers dynamically upgrade or downgrade services in an access environment.

You can configure your network with an *access profile* that provides all subscribers with default CoS parameters when they log in. For example, all subscribers can receive a basic data service. By configuring the access profile with Junos OS predefined variables for RADIUS-provided CoS parameters, you also enable the service to be activated for those subscribers at login.

To enable subscribers to activate a service or upgrade to different services through RADIUS change-of-authorization (CoA) messages after login, configure a *service profile* that includes user-defined variables.

### Types of CoS Variables Used in a Service Profile

You can configure variables for the following CoS parameters in a service profile:

- Shaping rate
- Delay buffer rate
- Guaranteed rate
- Scheduler map

For each CoS parameter, you must associate a RADIUS vendor ID. For each vendor ID, you must assign an attribute number and a tag. The tag is used to differentiate between values for different CoS variables when you specify the same attribute number for those variables. These values are matched with the values supplied by RADIUS during subscriber authentication. All of the values in the dynamic profile must be defined in RADIUS or none of the values are passed.

Optionally, you can configure default values for each parameter. Configuring default values is beneficial if you do not configure RADIUS to enable service changes. During service changes, RADIUS takes precedence over the default value that is configured.

## Static and Dynamic CoS Configurations

Depending on how you configure CoS parameters in the access and service profiles, certain CoS parameters are replaced or merged when subscribers change or activate new services.

Static configuration is when you configure the scheduler map and schedulers in the static **[edit class-of-service]** hierarchy and reference the scheduler map in the dynamic profile. Dynamic configuration is when you configure the scheduler map and schedulers within the dynamic profile.

The CoS configuration also depends on whether you have enabled multiple subscribers on the same logical interface using the **aggregate-clients** statements in the dynamic profile referenced by DHCP. When you specify the **aggregate-clients replace** statement, the scheduler map names are replaced. In both cases, if the length of the scheduler map name exceeds 128 characters, subscribers cannot log in. When you specify the **aggregate-clients merge** statement, the scheduler map names specified in the dynamic profile are appended.



**BEST PRACTICE:** To improve CoS performance in IPv4, IPv6, and dual-stack networks, we recommend that you use the **aggregate-clients replace** statement rather than the **aggregate-clients merge** statement.

## Scenarios for Static and Dynamic Configuration of CoS Parameters

Table 22 on page 188 lists the scenarios for static and dynamic configuration of CoS parameters in access profiles and service profiles at subscriber login. The table also lists the behavior for each configuration for service activation and service modification using RADIUS CoA messages.

Table 22: CoS Services and Variables

Scenario	Static CoS Configuration (Single Subscriber)	Dynamic CoS Configuration (Single Subscriber)	Dynamic CoS Configuration (Multiple Subscribers Enabled on a Logical Interface with the <b>aggregate-clients merge</b> Statement)	Dynamic CoS Configuration (Multiple Subscribers Enabled on a Logical Interface with the <b>aggregate-clients replace</b> Statement)
Subscriber login	<ul style="list-style-type: none"> <li>Configure RADIUS values or default values for all parameters in access profile</li> <li>Configure scheduler map in <b>edit class-of-service</b> hierarchy and reference in access profile</li> </ul>	<ul style="list-style-type: none"> <li>Configure RADIUS values or default values for all parameters in access profile</li> <li>Configure scheduler map and schedulers in access profile</li> </ul>	<ul style="list-style-type: none"> <li>Configure RADIUS values or default values for all parameters in access profile</li> <li>Configure scheduler map and schedulers in access profile</li> </ul>	<ul style="list-style-type: none"> <li>Configure RADIUS values or default values for all parameters in access profile</li> <li>Configure scheduler map and schedulers in access profile</li> </ul>

Table 22: CoS Services and Variables (*continued*)

Scenario	Static CoS Configuration (Single Subscriber)	Dynamic CoS Configuration (Single Subscriber)	Dynamic CoS Configuration (Multiple Subscribers Enabled on a Logical Interface with the aggregate-clients merge Statement)	Dynamic CoS Configuration (Multiple Subscribers Enabled on a Logical Interface with the aggregate-clients replace Statement)
RADIUS CoA for service or variable change	Replaces the following parameters: <ul style="list-style-type: none"> <li>• Delay buffer rate</li> <li>• Guaranteed rate</li> <li>• Scheduler map</li> <li>• Shaping rate</li> </ul>	Replaces the following parameters: <ul style="list-style-type: none"> <li>• Delay buffer rate</li> <li>• Guaranteed rate</li> <li>• Shaping rate</li> <li>• Scheduler map</li> </ul>	Combines the values of the following parameters to their maximum scalar value: <ul style="list-style-type: none"> <li>• Delay buffer rate</li> <li>• Guaranteed rate</li> <li>• Shaping rate</li> </ul> Appends the scheduler map parameter	Replaces the following parameters: <ul style="list-style-type: none"> <li>• Delay buffer rate</li> <li>• Guaranteed rate</li> <li>• Shaping rate</li> <li>• Scheduler map</li> </ul>
RADIUS CoA for service activation	Does not merge queues  <b>NOTE:</b> In this case, use a similar configuration to the access profile, including the same name for the traffic-control-profile. During service activation, this configuration replaces the original configuration in the access profile.	Merge queues if the queue specified in the service profile is not already in use for the subscriber  <b>NOTE:</b> Do not instantiate a CoA request using a service dynamic profile that is already in use on the same logical interface.	Merge queues if the queue specified in the service profile is not already in use for the subscriber  <b>NOTE:</b> Do not instantiate a CoA request using a service dynamic profile that is already in use on the same logical interface.	Merge queues if the queue specified in the service profile is not already in use for the subscriber  <b>NOTE:</b> Do not instantiate a CoA request using a service dynamic profile that is already in use on the same logical interface.

**Related Documentation**

- [Configuring Static Hierarchical Scheduling in a Dynamic Profile on page 78](#)
- [Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview](#)
- [RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework](#)
- [Guidelines for Configuring Dynamic CoS for Subscriber Access](#)

## CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview

---

To control bandwidth at a household level in a subscriber access network, you can apply RADIUS dynamic class of service (CoS) traffic-shaping attributes to a dynamic interface set and its member subscriber sessions when the subscriber sessions are authenticated. (The dynamic interface set itself does not go through the authentication process.)

A *household* is represented by either a dynamic interface set or a dynamic agent-circuit-identifier (ACI) interface set from which the subscriber sessions originate. For this feature, dynamic interface sets and dynamic ACI interface sets are mapped to Level 2 of the Junos OS CoS scheduler hierarchy, which enables you to use CoS traffic-shaping to shape the bandwidth at the household (interface set) level.

The *subscriber sessions*, also referred to as *subscriber interfaces* or *client sessions*, can be dynamic VLAN, PPPoE, or IP demultiplexing (IP demux) subscriber interfaces. The subscriber interfaces are mapped to Level 3 of the Junos OS CoS scheduler hierarchy.

- [Supported Network Configurations on page 190](#)
- [Traffic-Control Profiles in Subscriber Interface Dynamic Profiles on page 190](#)
- [CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets and Member Subscriber Sessions on page 191](#)

### Supported Network Configurations

Applying RADIUS dynamic CoS traffic-shaping attributes to a dynamic interface set and its member subscriber sessions is supported for the following network configurations:

- Dynamic IP demux subscriber interfaces (for DHCP subscribers) over either a dynamic interface set or a dynamic ACI interface set
- Dynamic PPPoE subscriber interfaces over either a dynamic interface set or a dynamic ACI interface set

### Traffic-Control Profiles in Subscriber Interface Dynamic Profiles

To apply dynamic CoS traffic-shaping attributes to a dynamic interface set and its member subscriber sessions, you must define and attach the traffic-control profiles for *both* the dynamic interface set and the dynamic subscriber sessions within the dynamic profile for the subscriber interface.

At the `[edit dynamic-profiles profile-name class-of-service traffic-control-profiles]` hierarchy level in the dynamic profile, configure both of the following:

- Traffic-control profile for the dynamic VLAN, PPPoE, or IP demux subscriber interfaces
- Traffic-control profile for the dynamic interface set or dynamic ACI interface set to which the subscriber interfaces belong

RADIUS tag values for the Junos OS CoS traffic shaping predefined variables used in both traffic-control profiles must be in the 100s range, as described in [“CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets” on page 198](#).

At the `[edit dynamic-profiles profile-name interfaces]` hierarchy level in the dynamic profile, use the `output-traffic-control-profile` statement to apply the traffic-control profiles to the dynamic subscriber interface and the dynamic interface set or dynamic ACI interface set.

## CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets and Member Subscriber Sessions

The set of `$junos-cos-parameter` predefined dynamic variables has been duplicated and assigned a RADIUS tag value in the 100s range for use with this feature. The RADIUS tag value is the only difference between the existing CoS traffic-shaping predefined dynamic variables and the predefined dynamic variables that you must use with this feature.

Both RADIUS instances of the `$junos-cos-parameter` predefined dynamic variables are available, but you must use the dynamic variables with tag values in the 100s range to apply CoS traffic-shaping attributes to both the dynamic interface set and member subscriber sessions in a subscriber interface dynamic profile.

For example, the existing `$junos-cos-shaping-rate` predefined variable is assigned RADIUS vendor ID 4874, attribute number 108, and tag value 2. To apply CoS traffic-shaping attributes to the dynamic interface set and its member subscriber sessions, you must instead use the `$junos-cos-shaping-rate` predefined variable that is assigned RADIUS vendor ID 4874, attribute number 108, and tag value 102.



**NOTE:** Do not configure a combination of `$junos-cos-parameter` predefined dynamic variables with RADIUS tag values in the 100s range and `$junos-cos-parameter` predefined dynamic variables with tag values not in the 100s range in the same traffic-control profile. If you do so, the subscriber authentication process fails.

### Related Documentation

- [Guidelines for Configuring CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions on page 192](#)
- [Applying CoS Traffic-Shaping Attributes to Dynamic Interface Sets and Member Subscriber Sessions on page 195](#)
- [CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets on page 198](#)

## Guidelines for Configuring CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions

---

Observe the following guidelines when you apply dynamic CoS traffic-shaping attributes to a dynamic interface set or a dynamic ACI interface set and its member subscriber sessions. For complete information about the Junos OS CoS traffic-shaping predefined dynamic variables and RADIUS tag values used with this feature, see [“CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets” on page 198](#).

- This feature is supported only for dynamically configured and instantiated subscriber interfaces.
- Do not configure a combination of **\$junos-cos-parameter** predefined dynamic variables with RADIUS tag values in the 100s range and **\$junos-cos-parameter** predefined dynamic variables with tag values not in the 100s range in the same traffic-control profile. If you do so, the subscriber authentication process fails.
- Use the **\$junos-cos-adjust-minimum** predefined variable (tag 109) only in traffic-control profiles for dynamic subscriber interfaces. Using this variable in a traffic-control profile for a dynamic interface set or dynamic ACI interface set has no effect.
- Do not configure the **\$junos-cos-excess-rate-high** predefined variable (tag 110) when the **\$junos-cos-excess-rate** predefined variable (tag 105) is configured, and vice-versa.
- Do not configure the **\$junos-cos-excess-rate-low** predefined variable (tag 111) when the **\$junos-cos-excess-rate** predefined variable (tag 105) is configured, and vice-versa.
- Do not configure the **\$junos-cos-byte-adjust-frame** predefined variable (tag 114) when the **\$junos-cos-byte-adjust** predefined variable (tag 108) is configured, and vice-versa.
- Do not configure the **\$junos-cos-byte-adjust-cell** predefined variable (tag 115) when the **\$junos-cos-byte-adjust** predefined variable (tag 108) is configured, and vice-versa.
- Use the per-priority **\$junos-cos-shaping-rate-parameter** predefined variables (tags 116 through 125) only in traffic-control profiles for dynamic interface sets or dynamic ACI interface sets. Using these variables in traffic-control profiles for a dynamic logical subscriber interface causes the subscriber session to fail.

### Related Documentation

- [Applying CoS Traffic-Shaping Attributes to Dynamic Interface Sets and Member Subscriber Sessions on page 195](#)
- [CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets on page 198](#)
- [CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview on page 190](#)

## Configuring Initial CoS Parameters Dynamically Obtained from RADIUS

You can configure a subscriber interface so that subscribers receive initial CoS parameters that the router obtains from the RADIUS authentication server when subscribers log in using that logical interface on the router.

1. Configure external RADIUS server VSAs with values that you expect subscribers to log in with.
  - To configure a RADIUS authentication server to include CoS traffic-shaping parameters in authentication grants on certain subscriber interfaces, configure Juniper Networks VSA 26–108.
  - To configure a RADIUS authentication server to include CoS scheduling and queuing parameters in authentication grants a certain subscriber interfaces, configure Juniper Networks VSA 28–146.

*See [Configuring Router or Switch Interaction with RADIUS Servers](#) and [Configuring RADIUS Server Parameters for Subscriber Access](#).*

2. Configure a subscriber interface that supports hierarchical CoS.

3. Associate a traffic-control profile with the interface.

*See [Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile](#).*

4. Configuring initial traffic-shaping parameters to be obtained from RADIUS.

*See [Configuring Dynamic Traffic Shaping and Scheduling Parameters in a Dynamic Profile](#).*

5. Configure forwarding classes and scheduler maps statically.

*See [Configuring a Custom Forwarding Class for Each Queue](#) and [Configuring Scheduler Maps](#).*

6. Configure a scheduler to specify initial scheduling and queuing parameters to be dynamically obtained from RADIUS when a subscriber logs in.

*See [Configuring Dynamic Schedulers with Variables in a Dynamic Profile](#).*

### Related Documentation

- [Subscriber Interfaces That Provide Initial CoS Parameters Dynamically Obtained from RADIUS on page 183](#)
- [Example: Configuring Initial CoS Parameters Dynamically Obtained from RADIUS](#)
- [Guidelines for Configuring Dynamic CoS for Subscriber Access](#)
- [Subscriber Activation and Service Management in an Access Network](#)

- *Juniper Networks VSAs Supported by the AAA Service Framework*
- *Dynamic Profiles Overview*
- *Dynamic Variables Overview*
- *Junos OS Predefined Variables*

---

## Configuring Static Default Values for Traffic Scheduling and Shaping

To provide subscribers with default values for CoS parameters, configure user-defined variables for CoS parameters and assign static default values to the variables. If you have configured values to be supplied by a RADIUS CoA, subscribers receive the default value when deactivating a service.

To configure user-defined variables with default values for CoS in a dynamic profile:

1. Specify that you want to configure variables in the dynamic profile.

```
[edit dynamic-profiles residential-silver variables]
```

2. Configure a default value for the shaping rate.

```
[edit dynamic-profiles residential-silver variables]  
user@host# set srate default-value 5m
```

3. Configure a default value for the guaranteed rate.

```
[edit dynamic-profiles residential-silver variables]  
user@host# set grate default-value 5m
```

4. Configure a default value for the delay buffer rate.

```
[edit dynamic-profiles residential-silver variables]  
user@host# set dbrate default-value 10m
```

5. Configure a default value for the scheduler map.

```
[edit dynamic-profiles residential-silver variables]  
user@host# set smap default-value triple-play
```

6. Configure the variables for the CoS parameters in the traffic-control profile.

Either the shaping rate or the guaranteed rate is required in the traffic-control profile.

- a. Access the traffic-control profile in the dynamic profile.

```
user@host# edit dynamic-profiles residential-silver class-of-service  
traffic-control-profiles tcp1
```

- b. Configure the scheduler map variable.

```
[edit dynamic-profiles residential-silver class-of-service traffic-control-profiles  
tcp1]
```



```
user@host# set scheduler-map "$smap"
```

- c. Configure the shaping rate variable.

```
[edit dynamic-profiles residential-silver class-of-service traffic-control-profiles
tcp1]
user@host# set shaping-rate "$srate"
```

- d. Configure the guaranteed rate variable.

```
[edit dynamic-profiles residential-silver class-of-service traffic-control-profiles
tcp1]
user@host# set guaranteed-rate "$grate"
```

- e. Configure the delay buffer rate variable.

```
[edit dynamic-profiles residential-silver class-of-service traffic-control-profiles
tcp1]
user@host# set delay-buffer-rate "$dbrate"
```

#### Related Documentation

- [Guidelines for Configuring Dynamic CoS for Subscriber Access](#)
- [Changing CoS Services Overview on page 187](#)

## Applying CoS Traffic-Shaping Attributes to Dynamic Interface Sets and Member Subscriber Sessions

To control bandwidth at a household level in a subscriber access network, you can apply RADIUS dynamic class of service (CoS) traffic-shaping attributes to a dynamic interface set or agent-circuit-identifier (ACI) interface set and its member subscriber sessions when the member sessions are authenticated. The dynamic interface set or ACI interface set represents the *household* from which the subscriber sessions originate. The *subscriber sessions*, also referred to as *client sessions* or *subscriber interfaces*, can be dynamic VLAN, PPPoE, or IP demultiplexing (IP demux, for DHCP) subscriber interfaces.

To apply RADIUS dynamic CoS traffic-shaping attributes to both the dynamic interface set and its member subscriber sessions, you must configure two traffic-control profiles in the dynamic profile for the subscriber interface: one traffic-control profile for the “parent” dynamic interface set, and a second traffic-control profile for the dynamic subscriber interfaces. RADIUS tag values for the Junos OS CoS traffic shaping predefined variables used in both traffic-control profiles must be in the 100s range.

Before you begin:

- Create a dynamic profile that defines the VLAN, PPPoE, or IP demux logical subscriber interface.

See the following topics:

- *Configuring a Basic Dynamic Profile*
- *Configuring a Dynamic Profile Used to Create Single-Tag VLANs*
- *Configuring a Dynamic Profile Used to Create Stacked VLANs*
- *Configuring Dynamic PPPoE Subscriber Interfaces*
- *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*

To apply dynamic CoS traffic-shaping attributes to a dynamic ACI or non-ACI interface set and its member subscriber sessions in a dynamic profile for the subscriber interface:

1. Configure two traffic-control profiles at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control profiles]** hierarchy level:
  - Traffic-control profile for the VLAN, PPPoE, or IP demux dynamic subscriber interfaces
  - Traffic-control profile for the dynamic interface set or dynamic ACI interface set to which the subscriber interfaces belong
2. In the traffic-control profiles configured for the dynamic interface set and the subscriber interfaces, reference Junos OS CoS traffic-shaping predefined variables with RADIUS tag values in the 100s range.

See [“CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets” on page 198](#) for a complete list of the Junos OS predefined variables and RADIUS tag values that you must use in the traffic-control profiles for the dynamic subscriber interfaces and the dynamic interface set.

3. At the **[edit dynamic-profiles *profile-name* interfaces]** hierarchy level, use the **output-traffic-control-profile** statement to apply the traffic-control profiles to the dynamic subscriber interface and the dynamic interface set or dynamic ACI interface set.

#### **Example: Dynamic PPPoE Subscriber Interface over Dynamic ACI Interface Set**

The following example shows a dynamic profile named `pppoe-subscriber` that configures a dynamic PPPoE (**pp0**) subscriber interface over a dynamic ACI interface set.

The **traffic-control-profiles** stanza defines two traffic-control profiles: `tcp-pppoe-session` for the dynamic PPPoE subscriber interface, and `tcp-parent-aci-set` for the dynamic “parent” ACI interface set. The **\$junos-cos-shaping-rate** predefined variable included in each of these traffic-control profiles is assigned RADIUS vendor ID 4874, attribute number 108, and tag value 102. The **\$junos-cos-shaping-mode** variable is assigned RADIUS vendor ID 4874, attribute number 108, and tag value 107.

The **interfaces** stanza applies output traffic-control profile tcp-pppoe-session to the dynamic PPPoE (**pp0**) subscriber interface, and output traffic-control profile tcp-parent-aci-set to the dynamic ACI interface set.

```
[edit dynamic-profiles]
pppoe-subscriber {
  interfaces {
    interface-set "$junos-interface-set-name" {
      interface pp0 {
        unit "$junos-interface-unit";
      }
    }
    pp0 {
      unit "$junos-interface-unit" {
        ppp-options {
          pap;
        }
        pppoe-options {
          underlying-interface "$junos-underlying-interface";
          server;
        }
        no-keepalives;
        family inet {
          unnumbered-address lo0.0;
        }
      }
    }
  }
}
class-of-service {
  traffic-control-profiles {
    tcp-pppoe-session {
      scheduler-map smap-1;
      shaping-rate $junos-cos-shaping-rate;
      overhead-accounting $junos-cos-shaping-mode frame-mode-bytes -4
        cell-mode-bytes 12;
    }
    tcp-parent-aci-set {
      shaping-rate $junos-cos-shaping-rate;
      overhead-accounting $junos-cos-shaping-mode frame-mode-bytes -4
        cell-mode-bytes 12;
    }
  }
  interfaces {
    pp0 {
      unit "$junos-interface-unit" {
        output-traffic-control-profile tcp-pppoe-session;
      }
    }
    interface-set $junos-interface-set-name {
      output-traffic-control-profile tcp-parent-aci-set;
    }
  }
}
```

**Related  
Documentation**

- [CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets on page 198](#)
- [CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview on page 190](#)
- [Guidelines for Configuring CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions on page 192](#)

---

## CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets

---

To control bandwidth at a household level in a subscriber access network, you can apply RADIUS CoS traffic-shaping attributes to a dynamic interface set and its member subscriber sessions when the member sessions are authenticated. The dynamic interface set, which represents the household level in a subscriber access network, can be either a dynamic agent-circuit-identifier (ACI) interface set or a non-ACI-based dynamic interface set. The subscriber sessions belonging to the interface set can be dynamic VLAN, DHCP, or PPPoE subscriber interfaces.

To apply RADIUS CoS traffic-shaping attributes to both the dynamic interface set and its member subscriber sessions, you must configure two traffic-control profiles in the dynamic profile for the subscriber interface: one traffic-control profile for the “parent” dynamic interface set, and a second traffic-control profile for the dynamic subscriber interfaces. RADIUS tag values for the Junos OS CoS traffic-shaping predefined variables used in these traffic-control-profiles must be in the 100s range, as described in [Table 23 on page 199](#).

To accommodate this feature, the set of existing **\$junos-cos-parameter** predefined dynamic variables for traffic shaping have been duplicated and assigned a tag value in the 100s range, as listed in [Table 23 on page 199](#). The tag value is the only difference between the existing predefined dynamic variables and the predefined dynamic variables that you must use with this feature.

For example, the existing **\$junos-cos-shaping-rate** predefined variable is assigned RADIUS vendor ID 4874, attribute number 108, and tag value 2. To apply RADIUS CoS traffic-shaping attributes to the dynamic interface set and its member subscriber sessions, you must instead use the **\$junos-cos-shaping-rate** predefined variable that is assigned RADIUS vendor ID 4874, attribute number 108, and tag value 102.

[Table 23 on page 199](#) describes the Junos OS predefined dynamic variables and RADIUS tag values that you can use in a dynamic profile to apply RADIUS CoS traffic-shaping attributes to the dynamic interface set and its member subscriber sessions. The table lists the predefined dynamic variables in ascending order by tag value.



**NOTE:** All of the predefined variables listed in [Table 23 on page 199](#) use RADIUS vendor ID 4874 and RADIUS attribute value 108.

---

**Table 23: Junos OS CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets**

Predefined Variable	RADIUS Tag Value	Description
\$junos-cos-scheduler-map	101	Scheduler-map name configured in a traffic-control profile in a dynamic profile.
\$junos-cos-shaping-rate	102	Shaping rate configured in a traffic-control profile in a dynamic profile. Represents the maximum bandwidth of a CoS scheduler node.
\$junos-cos-guaranteed-rate	103	Guaranteed rate configured in a traffic-control profile in a dynamic profile. Represents the minimum bandwidth of a CoS scheduler node.
\$junos-cos-delay-buffer-rate	104	Delay-buffer rate configured in a traffic-control profile in a dynamic profile.
\$junos-cos-excess-rate	105	Excess rate configured in a traffic-control profile in a dynamic profile; scheduler weighting when operating in the excess region between the guaranteed rate and the shaping rate.  <b>NOTE:</b> Do not configure the \$junos-cos-excess-rate variable when either the \$junos-cos-excess-rate-high variable or the \$junos-cos-excess-rate-low variable is configured.
\$junos-cos-traffic-control-profile	106	Traffic-control profile configured in a dynamic profile for subscriber access.
\$junos-cos-shaping-mode	107	Overhead-accounting mode configured in a traffic-control profile in a dynamic profile to shape downstream ATM traffic based on either frames (frame-mode) or cells (cell-mode).

Table 23: Junos OS CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets (*continued*)

Predefined Variable	RADIUS Tag Value	Description
\$junos-cos-byte-adjust	108	<p>Byte adjustment value for the cell or frame shaping mode configured in a traffic-control profile in a dynamic profile.</p> <p><b>NOTE:</b> Do not configure the \$junos-cos-byte-adjust variable when either the \$junos-cos-byte-adjust-frame variable or the \$junos-cos-byte-adjust-cell variable is configured.</p>
\$junos-cos-adjust-minimum	109	<p>Minimum adjusted shaping rate configured in a traffic-control profile for a dynamic subscriber interface. Specifying this variable in a traffic-control profile for a dynamic interface set has no effect.</p>
\$junos-cos-excess-rate-high	110	<p>Shaping rate configured for excess high-priority traffic in a traffic-control profile in a dynamic profile.</p> <p><b>NOTE:</b> Do not configure the \$junos-cos-excess-rate-high variable when the \$junos-cos-excess-rate variable is configured.</p>
\$junos-cos-excess-rate-low	111	<p>Shaping rate configured for excess low-priority traffic in a traffic-control profile in a dynamic profile.</p> <p><b>NOTE:</b> Do not configure the \$junos-cos-excess-rate-low variable when the \$junos-cos-excess-rate variable is configured.</p>
\$junos-cos-shaping-rate-burst	112	<p>Burst size for the shaping rate configured in a traffic-control profile in a dynamic profile.</p>
\$junos-cos-guaranteed-rate-burst	113	<p>Burst size for the guaranteed rate configured in a traffic-control profile in a dynamic profile.</p>

Table 23: Junos OS CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets (*continued*)

Predefined Variable	RADIUS Tag Value	Description
\$junos-cos-byte-adjust-frame	114	Overhead bytes when downstream ATM traffic is in frame-mode.  <b>NOTE:</b> Do not configure the \$junos-cos-byte-adjust-frame variable when the \$junos-cos-byte-adjust variable is configured.
\$junos-cos-byte-adjust-cell	115	Overhead bytes when downstream ATM traffic is in cell-mode.  <b>NOTE:</b> Do not configure the \$junos-cos-byte-adjust-cell variable when the \$junos-cos-byte-adjust variable is configured.
\$junos-cos-shaping-rate-priority-high	116	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	117	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-medium	118	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.

**Table 23: Junos OS CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets** *(continued)*

Predefined Variable	RADIUS Tag Value	Description
\$junos-cos-shaping-rate-priority-medium-burst	119	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-shaping-rate-priority-low	120	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	121	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-excess-high	122	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.
\$junos-cos-shaping-rate-excess-high-burst	123	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.



**Table 23: Junos OS CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets** (*continued*)

Predefined Variable	RADIUS Tag Value	Description
\$junos-cos-shaping-rate-excess-low	124	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.
\$junos-cos-shaping-rate-excess-low-burst	125	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.

**Related Documentation**

- [Applying CoS Traffic-Shaping Attributes to Dynamic Interface Sets and Member Subscriber Sessions on page 195](#)
- [CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview on page 190](#)
- [Guidelines for Configuring CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions on page 192](#)
- [Junos OS Predefined Variables](#)

**Example: Configuring Dynamic Hierarchical Scheduling for Subscribers**

In this example, subscribers are provided with a data and voice service defined in an access profile when they initially log in. The RADIUS administrator supplies the initial values on the RADIUS server, and the service activation is performed at subscriber login.

After the initial login, the subscriber adds an assured forwarding service that is not defined in the original access profile. A service profile is used to configure the schedulers and a RADIUS CoA activates the service. The queues defined for the schedulers in the initial scheduler map and the new scheduler map are merged.

In addition, the values for the initial data and voice service are upgraded by the RADIUS administrator through a separate RADIUS CoA message.

To configure the initial service and enable the activation through a RADIUS CoA:

1. Configure the access profile for the service activation.
  - a. Configure the VLAN interface for the access profile.

```
[edit]
dynamic-profiles access-profile {
  interfaces {
    $junos-interface-ifd-name {
      unit $junos-underlying-interface-unit {
        family inet;
      }
    }
  }
}
```

- b. Configure the class of service parameters in the access profile. In this example, you configure Junos OS predefined variables that provide the initial scheduler name and scheduler parameters obtained from the RADIUS authentication server when the subscriber logs in.

Include the configurations for the interfaces, schedulers, and the scheduler maps.

```
[edit]
dynamic-profiles access-profile {
  class-of-service {
    traffic-control-profiles {
      tcp1 {
        scheduler-map $junos-cos-scheduler-map;
        shaping-rate $junos-cos-shaping-rate;
        guaranteed-rate $junos-cos-guaranteed-rate;
        delay-buffer-rate $junos-cos-delay-buffer-rate;
      }
    }
    interfaces {
      $junos-interface-ifd-name {
        unit "$junos-underlying-interface-unit" {
          classifiers {
            ieee-802.1 l2_classifier;
          }
          rewrite-rules {
            ieee-802.1 l2_rewrite;
          }
          output-traffic-control-profile tcp1;
        }
      }
    }
    schedulers {
      $junos-cos-scheduler {
        buffer-size percent $junos-cos-scheduler-bs;
        priority $junos-cos-scheduler-pri;
        transmit-rate percent $junos-cos-scheduler-tx;
        drop-profile-map loss-priority low protocol any $junos-cos-scheduler-low;
        drop-profile-map loss-priority medium-low protocol any
          $junos-cos-scheduler-medium-low;
        drop-profile-map loss-priority medium-high protocol any
          $junos-cos-scheduler-medium-high;
        drop-profile-map loss-priority high protocol any $junos-cos-scheduler-high;
      }
    }
  }
}
```

```

scheduler-maps {
  data_voice_smap {
    forwarding-class be scheduler be_sch;
    forwarding-class ef scheduler ef_sch;
  }
}

```

Table 24 on page 205 lists the initial values defined by the RADIUS administrator for the scheduler map and shaping rates.

**Table 24: Initial Scheduler Map and Shaping Values at Subscriber Login**

Predefined Variable	RADIUS Tag	Value
\$junos-cos-scheduler-map	T01	data_voice_smap
\$junos-cos-shaping-rate	T02	6m
\$junos-cos-guaranteed-rate	T03	4m
\$junos-cos-delay-buffer-rate	T04	4m

Table 25 on page 205 lists the initial values defined by the RADIUS administrator for the voice (expedited forwarding) scheduler.

**Table 25: Initial CoS Values for the Voice Scheduler at Subscriber Login**

Predefined Variable	Tag	Value
\$junos-cos-scheduler	—	ef_sch
\$junos-cos-scheduler-tx	T01	10
\$junos-cos-scheduler-bs	T02	10
\$junos-cos-scheduler-pri	T03	medium-high
\$junos-cos-scheduler-dropfile-low	T04	d3
\$junos-cos-scheduler-dropfile-medium-low	T05	d2
\$junos-cos-scheduler-dropfile-medium-high	T06	d1
\$junos-cos-scheduler-dropfile-high	T07	d0

Table 26 on page 206 lists the initial values defined by the RADIUS administrator for the data (best effort) scheduler.

Table 26: Initial CoS Values for the Data Scheduler at Subscriber Login

Predefined Variable	Tag	Value
\$junos-cos-scheduler	—	be_sch
\$junos-cos-scheduler-tx	T01	10
\$junos-cos-scheduler-bs	T02	10
\$junos-cos-scheduler-pri	T03	low
\$junos-cos-scheduler-dropfile-low	T04	d0
\$junos-cos-scheduler-dropfile-medium-low	T05	d1
\$junos-cos-scheduler-dropfile-medium-high	T06	d2
\$junos-cos-scheduler-dropfile-high	T07	d3

2. Configure the classifiers, drop profiles, forwarding classes, and rewrite rules in the static **[edit class-of-service]** hierarchy.

```
[edit]
class-of-service {
  classifiers {
    dscp dscp_classifier {
      forwarding-class be {
        loss-priority low code-points 000000;
      }
      forwarding-class af {
        loss-priority medium-low code-points 000001;
      }
    }
    ieee-802.112_classifier {
      forwarding-class be {
        loss-priority medium-low code-points 000;
      }
      forwarding-class ef {
        loss-priority medium-low code-points 100;
      }
      forwarding-class af {
        loss-priority medium-low code-points 010;
      }
    }
  }
  drop-profiles {
    d0 {
      fill-level 25 drop-probability 100;
      fill-level 0 drop-probability 0;
    }
    d1 {
```

```

        fill-level 50 drop-probability 100;
        fill-level 0 drop-probability 0;
    }
    d2 {
        fill-level 75 drop-probability 100;
        fill-level 0 drop-probability 0;
    }
    d3 {
        fill-level 0 drop-probability 0;
        fill-level 100 drop-probability 100;
    }
}
forwarding-classes {
    queue 0 be;
    queue 1 ef;
    queue 2 af;
    queue 3 nc;
}
interfaces {
    ge-1/2/9 {
        shaping-rate 100m;
    }
}
rewrite-rules {
    ieee-802.1 l2_rewrite {
        forwarding-class be {
            loss-priority medium-low code-point 000;
        }
        forwarding-class ef {
            loss-priority medium-low code-point 001;
        }
        forwarding-class af {
            loss-priority medium-low code-point 100;
        }
    }
    dscp l2_rewrite {
        forwarding-class be {
            loss-priority medium-low code-points 000;
        }
        forwarding-class ef {
            loss-priority medium-low code-points 001;
        }
        forwarding-class af {
            loss-priority medium-low code-points 001;
        }
    }
}
}

```

3. Configure the service profile enable RADIUS to activate the video service after login. The video service corresponds to assured forwarding PHB.

In this example, you configure Junos OS predefined variables that provide the initial scheduler name and scheduler parameters obtained from the RADIUS authentication server when the subscriber logs in.

[edit]

```

dynamic-profiles service-af {
  variables {
    af_fc default-value video;
    af_sch default-value af_sch;
    sch-drop-any default-value all;
    sch-pri-2 default-value strict-high;
    sch-bs-2 default-value 40;
    sch-tx-2 default-value 3m;
    smap default-value any
  }
  class-of-service {
    scheduler-maps {
      "$smap" {
        forwarding-class "$af_fc" scheduler "$af_sch";
      }
    }
    schedulers {
      "$af_sch" {
        transmit-rate percent "$sch-tx-2";
        buffer-size percent "$sch-bs-2";
        priority "$sch-pri-2";
        drop-profile-map loss-priority any protocol any drop-profile "$sch-drop-any";
      }
    }
  }
}

```

After the three services are activated, subscribers receive upgraded values for the data and voice service when RADIUS sends a change of authorization (CoA). In this case, the CoS parameters are replaced, because multiple subscribers were not enabled on the logical interface.

[Table 27 on page 208](#) lists the upgraded values defined by the RADIUS administrator.

**Table 27: Upgraded CoS Values for the Video Service**

Variable	RADIUS Tag	Value
junos-cos-scheduler-map	T01	data_voice_smap
junos-cos-shaping-rate	T02	14m
junos-cos-guaranteed-rate	T03	13m
junos-cos-delay-buffer-rate	T04	12m

[Table 28 on page 209](#) lists the values defined by the RADIUS administrator for the video (assured forwarding) scheduler.

**Table 28: Upgraded CoS Values for the Video Scheduler**

Predefined Variable	Tag	Value
\$junos-cos-scheduler	—	af_sch
\$junos-cos-scheduler-tx	T01	10
\$junos-cos-scheduler-bs	T02	10
\$junos-cos-scheduler-pri	T03	medium
\$junos-cos-scheduler-dropfile-low	T04	d3
\$junos-cos-scheduler-dropfile-medium-low	T05	d2
\$junos-cos-scheduler-dropfile-medium-high	T06	d1
\$junos-cos-scheduler-dropfile-high	T07	d0

[Table 29 on page 209](#) lists the values defined by the RADIUS administrator for the expedited forwarding scheduler in the CoA message. The values are the same as the initial service.

**Table 29: Initial CoS Values for the Expedited Forwarding Scheduler at Subscriber Login**

Predefined Variable	Tag	Value
\$junos-cos-scheduler	—	ef_sch
\$junos-cos-scheduler-tx	T01	10
\$junos-cos-scheduler-bs	T02	10
\$junos-cos-scheduler-pri	T03	medium-high
\$junos-cos-scheduler-dropfile-low	T04	d3
\$junos-cos-scheduler-dropfile-medium-low	T05	d2
\$junos-cos-scheduler-dropfile-medium-high	T06	d1
\$junos-cos-scheduler-dropfile-high	T07	d0

[Table 30 on page 210](#) lists the values defined by the RADIUS administrator for the best effort scheduler in the CoA message. The values are the same as the initial service.

**Table 30: Initial CoS Values for the Best Effort Scheduler at Subscriber Login**

Predefined Variable	Tag	Value
\$junos-cos-scheduler	—	be_sch
\$junos-cos-scheduler-tx	T01	10
\$junos-cos-scheduler-bs	T02	10
\$junos-cos-scheduler-pri	T03	low
\$junos-cos-scheduler-dropfile-low	T04	d0
\$junos-cos-scheduler-dropfile-medium-low	T05	d1
\$junos-cos-scheduler-dropfile-medium-high	T06	d2
\$junos-cos-scheduler-dropfile-high	T07	d3

**Related Documentation**

- [Changing CoS Services Overview on page 187](#)
- [Guidelines for Configuring Dynamic CoS for Subscriber Access](#)
- [Understanding Hierarchical CoS for Subscriber Interfaces on page 70](#)



## PART 3

# Configuration Statements and Operational Commands

- Configuration Statements on page 213
- Operational Commands on page 431



## CHAPTER 10

# Configuration Statements

- [accept](#) on page 218
- [address-assignment](#) (Address-Assignment Pools) on page 220
- [authentication](#) (DHCP Local Server) on page 222
- [authentication](#) (DHCP Relay Agent) on page 223
- [authentication-order](#) on page 224
- [authentication-server](#) on page 225
- [auto-configure](#) on page 226
- [adjust-minimum](#) on page 227
- [adjust-percent](#) on page 228
- [atm-service](#) on page 229
- [buffer-size](#) (Schedulers) on page 230
- [cbr](#) on page 232
- [classifiers](#) (Definition) on page 233
- [classifiers](#) (Logical Interface) on page 234
- [classifiers](#) (Physical Interface) on page 235
- [code-points](#) (CoS) on page 236
- [delay-buffer-rate](#) on page 237
- [demux0](#) (Dynamic Interface) on page 238
- [demux-options](#) (Dynamic Interface) on page 239
- [demux-source](#) (Dynamic IP Demux Interface) on page 240
- [demux-source](#) (Dynamic Underlying Interface) on page 241
- [demux-source](#) (Underlying Interface) on page 242
- [dhcp-attributes](#) (Address-Assignment Pools) on page 243
- [dhcp-local-server](#) on page 245
- [dhcp-relay](#) on page 252
- [drop-probability](#) (Interpolated Value) on page 264
- [drop-probability](#) (Percentage) on page 265
- [drop-profile](#) (Schedulers) on page 266



- [drop-profile-map \(Schedulers\) on page 267](#)
- [drop-profiles on page 268](#)
- [dscp \(CoS Classifiers\) on page 269](#)
- [dscp \(Rewrite Rules\) on page 270](#)
- [dscp-ipv6 \(CoS Rewrite Rules\) on page 272](#)
- [dynamic-profile \(DHCP Local Server\) on page 273](#)
- [dynamic-profile \(VLAN\) on page 274](#)
- [dynamic-profiles on page 275](#)
- [egress-shaping-overhead on page 284](#)
- [excess-bandwidth-share on page 286](#)
- [excess-priority on page 287](#)
- [excess-rate \(Dynamic Traffic Shaping\) on page 288](#)
- [excess-rate-high \(Dynamic Traffic Shaping\) on page 289](#)
- [excess-rate-low \(Dynamic Traffic Shaping\) on page 290](#)
- [family on page 291](#)
- [family \(Address-Assignment Pools\) on page 296](#)
- [family \(Dynamic Demux Interface\) on page 297](#)
- [family \(Dynamic PPPoE\) on page 298](#)
- [family \(Dynamic Standard Interface\) on page 299](#)
- [fill-level \(Drop Profiles\) on page 301](#)
- [fill-level \(Interpolated Value\) on page 302](#)
- [filter \(Configuring\) on page 303](#)
- [flexible-vlan-tagging on page 304](#)
- [forwarding-classes \(Class-of-Service\) on page 305](#)
- [group \(DHCP Local Server\) on page 306](#)
- [guaranteed-rate on page 309](#)
- [hierarchical-scheduler on page 310](#)
- [hierarchical-scheduler \(Subscriber Interfaces on MX Series Routers\) on page 311](#)
- [ieee-802.1ad on page 312](#)
- [inet-precedence \(CoS Classifiers\) on page 313](#)
- [input-excess-bandwidth-share on page 314](#)
- [input-scheduler-map on page 315](#)
- [input-shaping-rate \(Logical Interface\) on page 316](#)
- [input-shaping-rate \(Physical Interface\) on page 317](#)
- [input-traffic-control-profile on page 318](#)
- [input-traffic-control-profile-remaining on page 319](#)
- [interface \(DHCP Local Server\) on page 320](#)

- [interface-set \(Ethernet Interfaces\)](#) on page 321
- [interface-set \(Hierarchical Schedulers\)](#) on page 322
- [interfaces](#) on page 323
- [interfaces \(CoS\)](#) on page 324
- [interfaces \(Static and Dynamic Subscribers\)](#) on page 326
- [internal-node](#) on page 330
- [interpolate](#) on page 331
- [loss-priority \(BA Classifiers\)](#) on page 332
- [loss-priority \(Scheduler Drop Profiles\)](#) on page 333
- [max-queues](#) on page 334
- [max-queues-per-interface](#) on page 336
- [maximum-lease-time](#) on page 337
- [member-link-scheduler](#) on page 338
- [mode \(Layer 2 Tunneling Protocol Shaping\)](#) on page 339
- [output-traffic-control-profile](#) on page 340
- [output-traffic-control-profile-remaining](#) on page 342
- [overhead-accounting](#) on page 343
- [pap \(Dynamic PPP\)](#) on page 344
- [peak-rate](#) on page 345
- [per-unit-scheduler](#) on page 346
- [pool \(Address-Assignment Pools\)](#) on page 348
- [pppoe-options \(Dynamic PPPoE\)](#) on page 349
- [ppp-options \(Dynamic PPP\)](#) on page 350
- [priority \(Schedulers\)](#) on page 352
- [profile \(Access\)](#) on page 353
- [protocols](#) on page 358
- [proxy-arp](#) on page 361
- [qos-adjust-hierarchical](#) on page 362
- [queue \(Global Queues\)](#) on page 363
- [radius \(Access Profile\)](#) on page 364
- [radius-server](#) on page 367
- [range \(Address-Assignment Pools\)](#) on page 368
- [ranges \(Dynamic VLAN\)](#) on page 369
- [rewrite-rules \(Definition\)](#) on page 370
- [rewrite-rules \(Interfaces\)](#) on page 371
- [routing-instances \(Multiple Routing Entities\)](#) on page 373
- [scheduler \(Scheduler Map\)](#) on page 374

- [scheduler-map \(Interfaces and Traffic-Control Profiles\) on page 375](#)
- [scheduler-maps \(For Most Interface Types\) on page 376](#)
- [schedulers \(CoS\) on page 377](#)
- [secret on page 378](#)
- [server \(Dynamic PPPoE\) on page 379](#)
- [server-group on page 380](#)
- [shaping-rate \(Applying to an Interface\) on page 381](#)
- [shaping-rate \(Schedulers\) on page 383](#)
- [shaping-rate \(Oversubscribing an Interface\) on page 385](#)
- [shaping-rate-excess-high on page 387](#)
- [shaping-rate-excess-low on page 388](#)
- [shaping-rate-priority-high on page 389](#)
- [shaping-rate-priority-low on page 390](#)
- [shaping-rate-priority-medium on page 391](#)
- [shared-bandwidth-policer \(Configuring\) on page 392](#)
- [shared-instance on page 393](#)
- [shared-scheduler on page 394](#)
- [simple-filter \(Applying to an Interface\) on page 395](#)
- [simple-filter on page 396](#)
- [site-identifier \(VPLS\) on page 397](#)
- [stacked-vlan-tagging on page 398](#)
- [system on page 398](#)
- [term \(Simple Filter\) on page 399](#)
- [three-color-policer \(Applying\) on page 400](#)
- [three-color-policer \(Configuring\) on page 401](#)
- [traffic-control-profiles on page 403](#)
- [traffic-manager on page 404](#)
- [transmit-rate \(Schedulers\) on page 408](#)
- [transparent on page 410](#)
- [underlying-interface \(demux0\) on page 411](#)
- [underlying-interface \(Dynamic PPPoE\) on page 412](#)
- [unit on page 413](#)
- [unit \(Dynamic Demux Interface\) on page 414](#)
- [unit \(Dynamic PPPoE\) on page 416](#)
- [unnumbered-address \(Dynamic PPPoE\) on page 418](#)
- [unnumbered-address \(Dynamic Profiles\) on page 419](#)
- [unnumbered-address \(Ethernet\) on page 421](#)

- [user-prefix \(DHCP Relay Agent\) on page 422](#)
- [vlan-id \(Dynamic Profiles\) on page 423](#)
- [vlan-id \(VLAN ID to Be Bound to a Logical Interface\) on page 424](#)
- [vlan-model on page 425](#)
- [vlan-ranges on page 426](#)
- [vlan-tag on page 427](#)
- [vlan-tags on page 428](#)
- [vlan-tags \(Stacked VLAN Tags\) on page 429](#)

## 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 dynamic-profile <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*
  - *Configuring Subscriber Packet Types to Trigger VLAN Authentication*

## 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 {
        active-drain;
        family family {
            dhcp-attributes {
                protocol-specific attributes;
            }
            excluded-address ip-address;
            excluded-range name low minimum-value high maximum-value;
            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;
        linked-pool-aggregation;
    }
}

```

**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. See [CLI Explorer](#).

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

**Related Documentation**

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

## authentication (DHCP Local Server)


---

<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-description (device-interface   logical-interface);     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 <a href="#">dhcp-local-server</a> dual-stack-group <i>dual-stack-group-name</i>], [edit system services dhcp-local-server dhcpv6], [edit system services dhcp-local-server dhcpv6 <a href="#">group</a> <i>group-name</i>], [edit system services dhcp-local-server <a href="#">group</a> <i>group-name</i>], [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. See <a href="#">CLI Explorer</a>.</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-description (device-interface   logical-interface);     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;     <b>user-prefix</b> <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 group <i>group-name</i>], [edit forwarding-options dhcp-relay dual-stack-group <i>dual-stack-group-name</i>], [edit forwarding-options dhcp-relay group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   forwarding-options dhcp-relay ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</pre>
<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. See <a href="#">CLI Explorer</a>.</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 252</a></li> <li>• <a href="#">Using External AAA Authentication Services with DHCP</a></li> </ul>

## authentication-order

<b>Syntax</b>	<code>authentication-order [ <i>authentication-methods</i> ];</code>
<b>Hierarchy Level</b>	[edit access <i>profile profile-name</i> ]
<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>none</b> option added in Junos OS Release 11.2.</p> <p><b>nasreq</b> option added in Junos OS Release 16.1.</p>
<b>Description</b>	<p>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.</p> <p>A given subscriber does not undergo both authentication and authorization as separate steps. When both <b>authentication-order</b> and <b>authorization-order</b> are specified, DHCP subscribers honor the configured authorization order, all other subscribers use the configured authentication-order.</p>
<b>Options</b>	<p><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:</p> <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 [edit access <i>profile profile-name client client-name</i>] hierarchy level.</li> <li>• <b>radius</b>—Verify the client using RADIUS authentication services.</li> </ul>
	<div>  <p><b>NOTE:</b> Subscriber access management does not support the password option, and authentication fails when no method (<b>none</b>) is specified.</p> </div>
	<b>Default:</b> password
<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>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>	<code>authentication-server [ <i>ip-address</i> ];</code>
<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;
    }

```

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

**Release Information** Statement introduced in Junos OS Release 9.5.

**Description** Enable the configuration of dynamic, auto-sensed VLANs.

The remaining statements are explained separately.

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

**Related Documentation**

- *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*

## adjust-minimum

**Syntax** `adjust-minimum rate;`

**Hierarchy Level** [edit class-of-service [schedulers](#) *scheduler-name*],  
[edit class-of-service [traffic-control-profiles](#) *traffic-control-profile-name*]

**Release Information** Statement introduced in Junos OS Release 11.4.

**Description** For adjustments performed by the ANCP or multicast applications on EQ DPC, MIC, or MPC interfaces, specify the minimum shaping rate for an adjusted scheduler node. The node is associated with a traffic-control profile.

For adjustments performed by the multicast application on MIC or MPC interfaces, specify the minimum shaping rate for an adjusted queue. The queue is associated with a scheduler.

**Options** *rate*—Minimum shaping rate for a node or a queue, in Mbps

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

**Related Documentation**


- [Configuring the Minimum Adjusted Shaping Rate on Scheduler Nodes for Subscribers on page 145](#)

## adjust-percent

---

<b>Syntax</b>	adjust-percent <i>percentage</i> ;
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">schedulers</a> <i>scheduler-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	For a MIC or MPC interface, determine the percentage of adjustment for the shaping rate of a queue.
<b>Options</b>	<b><i>percentage</i></b> —Percentage of the shaping rate to adjust. <b>Range:</b> 0 through 100 percent
<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 Shaping-Rate Adjustments on Queues on page 146</a></li></ul>

## atm-service

<b>Syntax</b>	<code>atm-service (cbr   rtvbr  nrtvbr);</code>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">traffic-control-profiles</a> <i>traffic-control-profile-name</i> ] [edit firewall atm-policer <i>atm-policer-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	(MX Series routers) Configure the ATM service category on ATM MICs to define bandwidth shaping and utilization. Shaping is based on the ATM service category.
<b>Default</b>	If the ATM service category is not specified, bandwidth utilization is unlimited.
<b>Options</b>	<p><b>cbr</b>—Use the constant bit rate.</p> <p><b>nrtvbr</b>—Use the non real-time variable bit rate.</p> <p><b>rtvbr</b>—Use the real-time variable bit rate.</p>
<div>  <p><b>NOTE:</b> (MX Series with MPCs and ATM MICs with SFP) To configure up to OC12 CBR bandwidth speed per virtual circuit (VC) on an ATM MIC with SFP (MIC-3D-8OC3-2OC12-ATM), specify <b>cbr</b> as the ATM service category.</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>show class-of-service traffic-control-profile</i></li> </ul>

## buffer-size (Schedulers)

---

**Syntax** `buffer-size (percent percentage | remainder | shared | temporal microseconds);`

**Hierarchy Level** `[edit class-of-service schedulers scheduler-name]`

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.  
Statement introduced in Junos OS Release 12.2 for ACX Series Routers.  
**shared** option introduced in Junos OS Release 18.1 for PTX Series Packet Transport Routers.

**Description** Specify buffer size.



**NOTE:** On PTX Series Packet Transport Routers, buffer-size cannot be configured on rate-limited queues.

**Default** If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent, respectively.

**Options** **percent *percentage***—Buffer size as a percentage of the total buffer.  
**Range:** 0 through 100



**NOTE:** For the routers with channelized OC12/STM4 IQE PIC with SFP (PB-4CHOC12-STM4-IQE-SFP) and channelized OC48/STM16 IQE PIC with SFP (PB-1CHOC48-STM16-IQE-SFP), the minimum buffer allocated to any queue is 18,432 bytes. If a queue is configured to have a buffer size less than 18K, the queue retains a buffer size of 18,432 bytes.

**remainder**—Remaining buffer available.

**shared**—On PTX Series routers, set a queue's buffer to be up to 100 percent of the interface's buffer. This option allows the queue's buffer to grow as large as 100 percent of the interface's buffer if and only if it is the only active queue for the interface.

**temporal *microseconds***—Buffer size as a temporal value. The queuing algorithm starts dropping packets when it queues more than a computed number of bytes. This maximum is computed by multiplying the logical interface speed by the configured temporal value.

**Range:** The ranges vary by platform as follows:

- For SRX Series Services Gateways: 1 through 2,000,000 microseconds.
- For vSRX instances: 1 through 32,000,000 microseconds.

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

**Related Documentation**


- *Managing Congestion on the Egress Interface by Configuring the Scheduler Buffer Size*
- *Buffer Size Temporal Value Ranges by Router Type*

## cbr

---

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

## classifiers (Definition)

<b>Syntax</b>	<pre> classifiers {     type classifier-name {         import (classifier-name   default);         forwarding-class class-name {             loss-priority level code-points [ aliases ] [ bit-patterns ];         }     } } </pre>
<b>Hierarchy Level</b>	[edit class-of-service], [edit class-of-service routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>ieee-802.1ad</b> option introduced in Junos OS Release 9.2.
<b>Description</b>	Define a CoS behavior aggregate (BA) classifier for classifying packets. You can associate the classifier with a forwarding class or code-point mapping, and import a default classifier or one that is previously defined.
<div style="display: flex; align-items: center;">  <div> <p><b>NOTE:</b> The [edit class-of-service routing-instances <i>routing-instance-name</i>] hierarchy level and the <b>dscp-ipv6</b> and <b>ieee-802.1ad</b> classifier types are not supported on ACX Series routers.</p> </div> </div>	
<b>Options</b>	<p><b>classifier-name</b>—Name of the aggregate behavior classifier.</p> <p><b>type</b>—Traffic type: <b>dscp</b>, <b>dscp-ipv6</b>, <b>exp</b>, <b>ieee-802.1</b>, <b>ieee-802.1ad</b>, <b>inet-precedence</b>.</p>
<b>Required Privilege Level</b>	<p><b>interface</b>—To view this statement in the configuration.</p> <p><b>interface-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic</i></li> </ul>

## classifiers (Logical Interface)

---

**Syntax**    classifiers {  
              type (classifier-name | default) family (mpls | inet);  
              no-default;  
          }

**Hierarchy Level**    [edit class-of-service interfaces *interface-name* **unit** *logical-unit-number*]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 14.2 for PTX Series Packet Transport Routers.  
**no-default** option added for MX Series devices only in Junos OS Release 16.1.

**Description**    Apply a CoS aggregate behavior classifier to a logical interface. You can apply a default classifier or one that is previously defined.

On MX Series devices, if you do not explicitly apply a classifier configuration to the interface, the default classifier is applied to the interface. Apply the **no-default** option to disable the application of any default classifier to the routing instance.

**Options**    *classifier-name*—Name of the aggregate behavior classifier.

*type*—Traffic type.

**Values:** dscp, dscp-ipv6, exp, ieee-802.1, inet-precedence



**NOTE:** You can only specify a family for the dscp and dscp-ipv6 types.

---

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

**Related Documentation**    • *Default DSCP and DSCP IPv6 Classifiers*  
                                  • *Applying Behavior Aggregate Classifiers to Logical Interfaces*




## classifiers (Physical Interface)


<b>Syntax</b>	<pre> classifiers {     type (classifier-name   default) ; } </pre>
<b>Hierarchy Level</b>	[edit class-of-service interfaces <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.
<b>Description</b>	Apply a CoS aggregate behavior classifier to a physical interface. You can apply a default classifier or one that is previously defined.
<b>Options</b>	<p><b><i>classifier-name</i></b>—Name of the aggregate behavior classifier.</p> <p><b><i>type</i></b>—Traffic type.</p> <p><b>Values:</b> <i>dscp</i>, <i>ieee-802.1</i>, and <i>inet-precedence</i></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>dscp</i></li> <li>• <i>inet-precedence</i></li> <li>• <i>ieee-802.1</i></li> </ul>

## code-points (CoS)

---

<b>Syntax</b>	<code>code-points ([ <i>aliases</i> ]   [ <i>bit-patterns</i> ] );</code>
<b>Hierarchy Level</b>	[edit class-of-service classifiers <i>type classifier-name</i> forwarding-class <i>class-name</i> loss-priority <i>level</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.2 for SRX Series devices. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 12.1X44 for the SRX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. Statement introduced in Junos OS Release 14.2 for PTX Series Packet Transport Routers.
<b>Description</b>	Specify one or more DSCP code-point aliases or bit sets to apply to a forwarding class..  <div> <b>NOTE:</b> OCX Series switches do not support MPLS, and therefore, do not support EXP code points or code point aliases.</div>
<b>Options</b>	<i>aliases</i> —Name of the DSCP alias.  <i>bit-patterns</i> —Value of the code-point bits, in six-bit binary form.
<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>Understanding Interfaces</i></li><li>• <i>Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic</i></li><li>• <i>Example: Configuring Behavior Aggregate Classifiers</i></li><li>• <i>Example: Configuring BA Classifiers on Transparent Mode Security Devices</i></li></ul>

## delay-buffer-rate

<b>Syntax</b>	delay-buffer-rate (percent <i>percentage</i>   <i>rate</i> );
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">traffic-control-profiles</a> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	For Gigabit Ethernet IQ, Channelized IQ PICs, and FRF.15 and FRF.16 LSQ interfaces only, base the delay-buffer calculation on a delay-buffer rate.
<b>Default</b>	If you do not include this statement, the delay-buffer calculation is based on the guaranteed rate if one is configured, or the shaping rate if no guaranteed rate is configured. For more information, see <i>Oversubscribing Interface Bandwidth</i> .
<b>Options</b>	<p><b>percent</b><i>percentage</i>—For LSQ interfaces, delay-buffer rate as a percentage of the available interface bandwidth.</p> <p><b>Range:</b> 1 through 100 percent</p> <p><b>rate</b>—For IQ and IQ2 interfaces, delay-buffer rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 1000 through 6,400,000,000,000 bps</p>
<div>  <p><b>NOTE:</b> Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Beginning with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.</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>Oversubscribing Interface Bandwidth</i></li> <li>• <i>Providing a Guaranteed Minimum Rate</i></li> <li>• <i>Configuring Traffic Control Profiles for Shared Scheduling and Shaping</i></li> <li>• <a href="#">output-traffic-control-profile on page 340</a></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. See [CLI Explorer](#).

**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)

<b>Syntax</b>	<pre>demux-options {   underlying-interface interface-name }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">dynamic-profiles</a> profile-name <a href="#">interfaces</a> demux0 interface-name <a href="#">unit</a> logical-unit-number]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3.
<b>Description</b>	<p>Configure logical demultiplexing (demux) interface options in a dynamic profile.</p> <p>The remaining statement is 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 Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles</i></li> <li>• <i>Demultiplexing Interface Overview</i></li> </ul>

## demux-source (Dynamic IP Demux Interface)

---

<b>Syntax</b>	<code>demux-source {     source-address; }</code>
<b>Hierarchy Level</b>	[edit <a href="#">dynamic-profiles profile-name</a> <a href="#">interfaces demux0</a> unit <i>logical-unit-number</i> <a href="#">family family</a> ]
<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)

**Syntax** `demux-source family;`

**Hierarchy Level** [edit `dynamic-profiles interfaces interface-name` unit `logical-unit-number`]

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

**Description** Configure the logical demultiplexing (demux) source family type on the IP demux underlying interface within a dynamic profile.



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


**Options** *family*—Protocol family:

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

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

## demux-source (Underlying Interface)

---

<b>Syntax</b>	<code>demux-source <i>family</i>;</code>
<b>Hierarchy Level</b>	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> interfaces</code> <code>    <i>interface-name</i> unit <i>logical-unit-number</i>]</code>
<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.
<div> <b>NOTE:</b> The IP demux interface feature currently supports only Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet underlying interfaces.</div>	
<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.
<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)

**Syntax** `dhcp-attributes {`  
     `boot-file filename;`  
     `boot-server (address | hostname);`  
     `dns-server [ ipv6-address ];`  
     `domain-name domain-name;`  
     `exclude-prefix-len exclude-prefix-length;`  
     `grace-period seconds;`  
     `maximum-lease-time seconds;`  
     `name-server [ server-list ];`  
     `netbios-node-type node-type;`  
     `option {`  
         `[ (id-number option-type option-value)`  
         `(id-number array option-type option-value) ];`  
     `}`  
     `option-match {`  
         `option-82 {`  
             `circuit-id value range named-range;`  
             `remote-id value range named-range;`  
         `}`  
     `}`  
     `preferred-lifetime seconds;`  
     `router [ router-address ];`  
     `server-identifier ip4-address;`  
     `sip-server-address [ ipv6-address ];`  
     `sip-server-domain-name domain-name;`  
     `t1-percentage percentage;`  
     `t1-renewal-time;`  
     `t2-percentage percentage;`  
     `t2-rebinding-time;`  
     `tftp-server address;`  
     `valid-lifetime seconds;`  
     `wins-server [ servers ];`  
`}`

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

**Release Information** Statement introduced in Junos OS Release 9.0.  
 Statement introduced in Junos OS Release 12.3 for EX Series switches.  
**exclude-prefix-len** statement introduced in Junos OS Release 17.3 for MX Series.

**Description** Configure DHCP attributes for the protocol family in a specific address pool. The attributes determine options and behaviors for the DHCP clients.

The remaining statements are explained separately.

**Options** **exclude-prefix-len** *exclude-prefix-length*—Specify the length of the IPv6 prefix to be excluded from the delegated prefix.  
**Range:** 1 through 128

<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-description (device-interface | logical-interface);
                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;
                overrides {
                    asymmetric-lease-time seconds;
                    asymmetric-prefix-lease-time seconds;
                    interface-client-limit number;
                    multi-address-embedded-option-response;
                    process-inform {
                        pool pool-name;
                    }
                    protocol-attributes attribute-set-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;  
    }  
    layer2-liveness-detection {  
        max-consecutive-retries number;  
        transmit-interval interval;  
    }  
}  
}  
overrides {  
    asymmetric-lease-time seconds;  
    asymmetric-prefix-lease-time seconds;  
    delegated-pool;  
    interface-client-limit number;  
    multi-address-embedded-option-response;  
    process-inform {  
        pool pool-name;  
    }  
    protocol-attributes attribute-set-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;  
        }  
        layer2-liveness-detection {
```

```

        max-consecutive-retries number;
        transmit-interval interval;
    }
}
overrides {
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    delegated-pool;
    include-option-82 {
        forcerenew;
        nak;
    }
    interface-client-limit number;
    multi-address-embedded-option-response;
    process-inform {
        pool pool-name;
    }
    protocol-attributes attribute-set-name;
    rapid-commit;
}
reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    support-option-pd-exclude;
    timeout timeout-value;
    token token-value;
    trigger {
        radius-disconnect;
    }
}
reauthenticate (<lease-renewal> <remote-id-mismatch >);
requested-ip-network-match subnet-mask;
route-suppression;
server-duid-type type;
service-profile dynamic-profile-name;
}
dual-stack-group name {
    access-profile access-profile;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description (device-interface | logical-interface);
            interface-name ;
            logical-system-name;
            mac-address;
            relay-agent-interface-id;
            relay-agent-remote-id;
            routing-instance-name;
            user-prefix user-prefix-string;
        }
    }
}

```

```
classification-key {
  circuit-id circuit-id;
  mac-address mac-address;
  remote-id remote-id;
}
dual-stack-interface-client-limit number;
dynamic-profile profile-name {
  aggregate-clients (merge | replace);
  use-primary primary-profile-name;
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    layer2-liveness-detection {
      max-consecutive-retries number;
      transmit-interval interval;
    }
  }
}
on-demand-address-allocation;
protocol-master (inet | inet6);
reauthenticate (<lease-renewal> <remote-id-mismatch>);
service-profile service-profile;
}
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;
    overrides {
      asymmetric-lease-time seconds;
      client-discover-match (option60-and-option82 | incoming-interface);
      include-option-82 {
        forcerenew;
        nak;
      }
      interface-client-limit number;
      process-inform {
        pool pool-name;
      }
      protocol-attributes attribute-set-name;
    }
    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;
  }
  layer2-liveness-detection {
    max-consecutive-retries number;
    transmit-interval interval;
  }
}
}
overrides {
  asymmetric-lease-time seconds;
  client-discover-match (option60-and-option82 | incoming-interface);
  include-option-82 {
    forcerenew;
    nak;
  }
  interface-client-limit number;
  process-inform {
    pool pool-name;
  }
  protocol-attributes attribute-set-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;
    }
    layer2-liveness-detection {
        max-consecutive-retries number;
        transmit-interval interval;
    }
}
on-demand-address-allocation;
overrides {
    asymmetric-lease-time seconds;
    client-discover-match <option60-and-option82 | incoming-interface>;
    interface-client-limit number;
    process-inform {
        pool pool-name;
    }
    protocol-attributes attribute-set-name;
}
pool-match-order {
    external-authority;
    ip-address-first;
    option-82;
}
protocol-master;
reauthenticate (<lease-renewal> <remote-id-mismatch >);
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 to 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, therefore, 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 **dhcpv6** 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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

**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-description (device-interface | logical-interface);
                interface-name;
                logical-system-name;
                mac-address;
                option-60;
                option-82 <circuit-id> <remote-id>;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        bulk-leasequery {
            attempts number-of-attempts;
            timeout seconds;
            trigger automatic;
        }
        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-description (device-interface | logical-interface);
                interface-name interface-name;
                logical-system-name;
                mac-address mac-address;
                relay-agent-interface-id;
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        bulk-leasequery {
            attempts number-of-attempts;
            timeout seconds;
            trigger automatic;
        }
        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;
}
forward-snooped-clients (all-interfaces | configured-interfaces |
    non-configured-interfaces);
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;
            interface-description (device-interface | logical-interface);
            interface-name interface-name;
            logical-system-name;
            mac-address mac-address;
            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;
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    exclude;
    overrides {
        allow-snooped-clients;
        asymmetric-lease-time seconds;
        asymmetric-prefix-lease-time seconds;
        client-negotiation-match incoming-interface;
        delay-authentication;
        delete-binding-on-renegotiation;
        dual-stack dual-stack-group-name;
    }
}

```

```
    interface-client-limit number;  
    no-allow-snooped-clients;  
    no-bind-on-request;  
    relay-source interface-name;  
    send-release-on-delete;  
  }  
  service-profile dynamic-profile-name;  
  trace;  
  upto upto-interface-name;  
}  
}  
lease-time-validation {  
  lease-time-threshold seconds;  
  violation-action action;  
}  
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;  
    }  
    layer2-liveness-detection {  
      max-consecutive-retries number;  
      transmit-interval interval;  
    }  
  }  
}  
overrides {  
  allow-snooped-clients;  
  asymmetric-lease-time seconds;  
  asymmetric-prefix-lease-time seconds;  
  client-negotiation-match incoming-interface;  
  delay-authentication;  
  delete-binding-on-renegotiation;  
  dual-stack dual-stack-group-name;  
  interface-client-limit number;  
  no-allow-snooped-clients;  
  no-bind-on-request;  
  relay-source interface-name;  
  send-release-on-delete;  
}  
relay-agent-interface-id {  
  include-irb-and-l2;
```

```

    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
}
relay-agent-remote-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
    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;
        relay-server-group relay-server-group;
    }
}
remote-id-mismatch disconnect;
route-suppression;
service-profile dynamic-profile-name;
}
leasequery {
    attempts number-of-attempts;
    timeout seconds;
}
lease-time-validation {
    lease-time-threshold seconds;
    violation-action action;
}
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;
}
layer2-liveness-detection {
    max-consecutive-retries number;
    transmit-interval interval;
}
route-suppression;
service-profile dynamic-profile-name;
}
}
no-snoop;
overrides {
    allow-snooped-clients;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-negotiation-match incoming-interface;
    delay-authentication;
    delete-binding-on-renegotiation;
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    relay-source interface-name;
    send-release-on-delete;
}
relay-agent-interface-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
}
relay-agent-remote-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
}
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;
        relay-server-group relay-server-group;
    }
}
relay-option-vendor-specific{
    host-name;
    location;
    remote-id-mismatch disconnect;
    route-suppression;
    server-group {
        server-group-name {
            server-ip-address;
        }
    }
    server-response-time seconds;
    service-profile dynamic-profile-name;
}
dual-stack-group dual-stack-group-name {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description (device-interface | logical-interface);
            interface-name;
            logical-system-name;
            mac-address;
            relay-agent-interface-id;
            relay-agent-remote-id;
            routing-instance-name;
            user-prefix user-prefix-string;
        }
    }
    classification-key {
        circuit-id circuit-id;
        mac-address mac-address;
        remote-id remote-id;
    }
    dual-stack-interface-client-limit number;
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            layer2-liveness-detection {
                max-consecutive-retries number;
                transmit-interval interval;
            }
        }
    }
}

```

```
    }
  }
  protocol-master;
  relay-agent-interface-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
  }
  relay-agent-remote-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
  }
  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-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 {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-description (device-interface | logical-interface);
      interface-name interface-name;
      logical-system-name;
      mac-address;
      option-60;
      option-82 [circuit-id] [remote-id];
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
}
dynamic-profile profile-name {
  aggregate-clients (merge | replace);
  use-primary primary-profile-name;
}
```



```

forward-only {
  logical-system <current | default | logical-system-name>;
  routing-instance <current | default | routing-instance-name>;
}
forward-only {
  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 {
  allow-no-end-option;
  allow-snooped-clients;
  always-write-giaddr;
  always-write-option-82;
  asymmetric-lease-time seconds;
  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;
  relay-source
  replace-ip-source-with;
  send-release-on-delete;
  trust-option-82;
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;

```

```
}
overrides {
  allow-no-end-option
  allow-snooped-clients;
  always-write-giaddr;
  always-write-option-82;
  asymmetric-lease-time seconds;
  asymmetric-prefix-lease-time seconds;
  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;
  relay-source
  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
}
remote-id-mismatch disconnect;
route-suppression:
service-profile dynamic-profile-name;
```

```

}
leasequery {
  attempts number-of-attempts;
  timeout seconds;
}
lease-time-validation {
  lease-time-threshold seconds;
  violation-action action;
}
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;
    }
    layer2-liveness-detection {
      max-consecutive-retries number;
      transmit-interval interval;
    }
  }
}
no-snoop;
overrides {
  allow-no-end-option
  allow-snooped-clients;
  always-write-giaddr;
  always-write-option-82;
  asymmetric-lease-time seconds;
  asymmetric-prefix-lease-time seconds;
  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;
  relay-source
  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
}
}
remote-id-mismatch disconnect;
route-suppression:
server-group {
  server-group-name {
    server-ip-address;
  }
}
server-response-time seconds;
service-profile dynamic-profile-name;
```

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

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

**Description** Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the router or switch to enable the router (or switch) to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.

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

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

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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

**Related Documentation**

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

## drop-probability (Interpolated Value)

---

<b>Syntax</b>	<code>drop-probability [values];</code>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">drop-profiles</a> <i>profile-name</i> <a href="#">interpolate</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS 11.4 for EX Series switches.
<b>Description</b>	<p>Define values for interpolated drop probabilities. The maximum number of drop probability values supported per drop profile is based on the line card.</p> <p>On EX Series switches, this statement is supported only on the EX9200 switch, EX8200 standalone switches, and EX8200 Virtual Chassis.</p>
<b>Options</b>	<p><b>percentage</b>—The probability (expressed in percentage) for a packet to be dropped from the queue.</p> <p><b>Range:</b> 0 through 100</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>Managing Congestion Using RED Drop Profiles and Packet Loss Priorities</i></li><li>• <i>Defining Packet Drop Behavior by Configuring RED Drop Profiles</i></li></ul>

---

## drop-probability (Percentage)

---

<b>Syntax</b>	<code>drop-probability <i>percentage</i>;</code>
<b>Hierarchy Level</b>	[edit class-of-service drop-profiles <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define drop probability percentages. The maximum number of drop probability values supported per drop profile is based on the line card.
<b>Options</b>	<b><i>percentage</i></b> —Probability that a packet is dropped, expressed as a percentage. A value of 0 means that a packet is never dropped, and a value of 100 means that all packets are dropped. <b>Range:</b> 0 through 100 percent
<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>Managing Congestion Using RED Drop Profiles and Packet Loss Priorities</i></li><li>• <i>Defining Packet Drop Behavior by Configuring RED Drop Profiles</i></li></ul>

## drop-profile (Schedulers)

---

<b>Syntax</b>	<code>drop-profile <i>profile-name</i>;</code>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">schedulers</a> <i>scheduler-name</i> <a href="#">drop-profile-map</a> <a href="#">loss-priority</a> (any   low   medium-low   medium-high   high) protocol (any   non-tcp   tcp)]
<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. Statement introduced in Junos OS Release 12.2 for ACX Series Routers.
<b>Description</b>	Define drop profiles for RED. When a packet arrives, RED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the arriving packet.
<b>Options</b>	<i>profile-name</i> —Name of the drop 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>Determining Packet Drop Behavior by Configuring Drop Profile Maps for Schedulers</i></li><li>• <i>Managing Congestion Using RED Drop Profiles and Packet Loss Priorities</i></li></ul>



## drop-profile-map (Schedulers)

<b>Syntax</b>	drop-profile-map <b>loss-priority</b> (any   low   medium-low   medium-high   high) protocol (any   non-tcp   tcp) <b>drop-profile (Schedulers)</b> <i>profile-name</i> ;
<b>Hierarchy Level</b>	[edit class-of-service <b>schedulers</b> <i>scheduler-name</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. Statement introduced in Junos OS Release 12.2 for ACX Series Routers.
<b>Description</b>	Define the loss-priority value for a drop profile.  The remaining statements are explained separately. See <a href="#">CLI Explorer</a> .
<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>Default Schedulers Overview</i></li> <li>• <i>Determining Packet Drop Behavior by Configuring Drop Profile Maps for Schedulers</i></li> </ul>

## drop-profiles

---

<b>Syntax</b>	<pre>drop-profiles {   profile-name {     fill-level percentage drop-probability percentage;     interpolate {       drop-probability [values];       fill-level [values]     }   } }</pre>
<b>Hierarchy Level</b>	[edit class-of-service]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS 11.4 for EX Series switches.
<b>Description</b>	<p>Define drop profiles for RED.</p> <p>For a packet to be dropped, it must match the drop profile. When a packet arrives, RED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the packet.</p>
<b>Options</b>	<p><i>profile-name</i>—Name of the drop profile.</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>Defining Packet Drop Behavior by Configuring RED Drop Profiles</i></li></ul>

## dscp (CoS Classifiers)

<b>Syntax</b>	<pre> dscp classifier-name {     import (classifier-name   default);     forwarding-class class-name {         loss-priority level ] {             code-points [ aliases ] [ bit-patterns;         }     } } </pre>
<b>Hierarchy Level</b>	[edit class-of-service <i>classifiers</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define the diffserv code point (DSCP) mapping that is applied to the packets.
<b>Options</b>	<p><b>classifier-name</b>—Name of the classifier.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>
<b>Required Privilege Level</b>	<p>interfaces—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>Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic</i></li> <li>• <i>Applying DSCP Classifiers to MPLS Traffic</i></li> </ul>

## dscp (Rewrite Rules)

---

<b>Syntax</b>	<code>dscp (rewrite-name   default) protocol (inet-both   inet-outer   mpls);</code>
<b>Hierarchy Level</b>	<code>[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>For IPv4 traffic, apply a Differentiated Services (DiffServ) code point (DSCP) rewrite rule.</p> <p>Logical interfaces do not support multiple <b>dscp</b> rewrite rules for the same protocol.</p> <p>DSCP and DSCP IPv6 rewrite rules are supported on M Series and T Series routers when non-queuing PICs are installed, but are disabled when queuing PICs are installed with the following exceptions:</p> <ul style="list-style-type: none"><li>• On M320 routers, DSCP rewrite is supported on IQ, IQ2, IQE, and IQ2E PICs when used with the Enhanced III FPC.</li><li>• On M120 routers, DSCP rewrite is supported on IQ, IQ2, IQE, and IQ2E PICs.</li></ul> <p>DSCP and DCSP IPv6 rewrite rules are supported on MIC and MPC interfaces on MX Series routers.</p> <p>DSCP rewrite rules are not supported on T Series routers when IQ, IQ2, IQE, IQ2E, SONET/SDH OC48/STM16 IQE, or PD-5-10XGE-SFPP PICs are installed.</p>
<b>Options</b>	<p><b>rewrite-name</b>—Name of a <b>rewrite-rules</b> mapping configured at the <code>[edit class-of-service rewrite-rules dscp]</code> hierarchy level.</p> <p><b>default</b>—The default mapping.</p> <p><b>protocol inet-both</b>—For gr- interfaces (GRE tunnels) on MPCs, rewrite the DSCP CoS value to both the inner and outer header for Unicast/Multicast IPv4 traffic. The first six bits of the CoS value are rewritten and the final two bits are taken from the incoming CoS value.</p> <p><b>protocol inet-outer</b>—For gr- interfaces on MPCs, rewrite the DSCP CoS value to the outer header for Unicast/Multicast IPv4 traffic. The first six bits of the CoS value are rewritten and the final two bits are taken from the incoming CoS value.</p> <p><b>protocol mpls</b>—(Optional for ingress MPLS tunnel nodes) For interfaces on MX Series routers or hosted on Enhanced III FPCs in M120 or M320 routers only, rewrite the MPLS EXP bits in the MPLS header independently of the IPv4 DSCP value for IPv4 packets entering an MPLS tunnel.</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>

- Related Documentation**
- *Configuring Rewrite Rules*
  - *Applying Rewrite Rules to Output Logical Interfaces*
  - *protocol (Rewrite Rules)*
  - *Rewriting MPLS and IPv4 Packet Headers*
  - [rewrite-rules \(Definition\) on page 370](#)

## dscp-ipv6 (CoS Rewrite Rules)

---

<b>Syntax</b>	<code>dscp-ipv6 (rewrite-name   &lt;default&gt;) protocol mpls;</code>
<b>Hierarchy Level</b>	<code>[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Support for <b>protocol mpls</b> option introduced in Junos OS Release 10.4R2.
<b>Description</b>	<p>For IPv6 traffic, apply a DSCP rewrite rule.</p> <p>Logical interfaces do not support multiple <b>dscp-ipv6</b> rewrite rules for the same protocol.</p> <p>DSCP and DSCP IPv6 rewrite rules are supported on M Series and T Series routers when non-queuing PICs are installed, but are disabled when queuing PICs are installed with the following exceptions:</p> <ul style="list-style-type: none"><li>• On M320 routers, DSCP rewrite is supported on IQ, IQ2, IQE, and IQ2E PICs when used with the Enhanced III FPC.</li><li>• On M120 routers, DSCP rewrite is supported on IQ, IQ2, IQE, and IQ2E PICs.</li></ul> <p>DSCP and DCSP IPv6 rewrite rules are supported on MIC and MPC interfaces on MX Series routers.</p> <p>DSCP rewrite rules are not supported on T Series routers when IQ, IQ2, IQE, IQ2E, SONET/SDH OC48/STM16 IQE, or PD-5-10XGE-SFPP PICs are installed.</p>
<b>Options</b>	<p><b>rewrite-name</b>—Name of a <b>rewrite-rules</b> mapping configured at the <code>[edit class-of-service rewrite-rules dscp-ipv6]</code> hierarchy level.</p> <p><b>default</b>—Default mapping.</p> <p><b>protocol mpls</b>—(Optional for ingress MPLS tunnel nodes) For interfaces on MX Series routers or hosted on Enhanced III FPCs in M120 or M320 routers only, rewrite the MPLS EXP bits in the MPLS header independently of the IPv6 DSCP value for IPv6 packets entering an MPLS tunnel.</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 Rewrite Rules</a></li><li>• <a href="#">protocol</a></li><li>• <a href="#">Setting IPv6 DSCP and MPLS EXP Values Independently</a></li><li>• <a href="#">Configuring DSCP Values for IPv6 Packets Entering the MPLS Tunnel</a></li><li>• <a href="#">Applying Rewrite Rules to Output Logical Interfaces</a></li></ul>

- [rewrite-rules \(Definition\)](#) on page 370

## dynamic-profile (DHCP Local Server)

<b>Syntax</b>	dynamic-profile <i>profile-name</i> { aggregate-clients (merge   replace); use-primary <i>primary-profile-name</i> ; }
<b>Hierarchy Level</b>	[edit system services <a href="#">dhcp-local-server</a> ], [edit system services <a href="#">dhcp-local-server</a> dual-stack-group <i>dual-stack-group-name</i> ], [edit system services dhcp-local-server dhcpv6], [edit system services dhcp-local-server dhcpv6 <a href="#">group</a> <i>group-name</i> ], [edit system services dhcp-local-server dhcpv6 group <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 group <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> ...]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2. Statement introduced in Junos OS Release 12.3R2 for EX Series switches. Options <b>aggregate-clients</b> and <b>use-primary</b> introduced in Junos OS Release 9.3. Support at the [edit ... <a href="#">interface</a> ] hierarchy levels introduced in Junos OS Release 11.2.
<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>	<b><i>profile-name</i></b> —Name of the dynamic profile.  The remaining statements are explained separately. See <a href="#">CLI Explorer</a> .
<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>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 (VLAN)

---

<b>Syntax</b>	<pre>dynamic-profile <i>profile-name</i> {     <b>accept</b> (any   dhcp-v4   dhcp-v6   inet   inet6   pppoe);     accept-out-of-band <i>protocol</i>;     access-profile <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);
                    report-ingress-shaping-rate bps;
                    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;
    }
}
}
}
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 {
  actual-transit-statistics;
  auto-configure {
    agent-circuit-identifier {
      dynamic-profile profile-name;
    }
    line-identity {
      include {
        accept-no-ids;
        circuit-id;
        remote-id;
      }
      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;
    }
    pcef pcef-profile-name {
        activate rule-name | activate-all;
    }
}
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;
    }
}
host-prefix-only;
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-limit limit;
      group-policy;
      group-threshold value;
      immediate-leave;
      log-interval seconds;
      no-accounting;
      oif-map;
      passive;
      promiscuous-mode;
      ssm-map ssm-map-name;
      ssm-map-policy ssm-map-policy-name
      static {
        group group {
          source source;
        }
      }
      version version;
    }
  }
  mld {
    interface interface-name {
      (accounting | no-accounting);
      disable;
      group-limit limit;
      group-policy;
      group-threshold value;
      immediate-leave;
      log-interval seconds;
      oif-map;
      passive;
      ssm-map ssm-map-name;
      ssm-map-policy ssm-map-policy-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;
    }
}
}
services {
    captive-portal-content-delivery {
        rule name {
            match-direction (input | input-output | output);
            term name {
                from {
                    applications application-name {
                        application-protocol type;
                        destination-port port-type;
                        protocol ip-protocol-type;
                        source-port port-type;
                    }
                    destination-address name <except>;
                    destination-address-range low minimum-value high maximum-value <except>;
                    destination-prefix-list name <except>;
                }
                then {
                    accept;
                    redirect url;
                    rewrite destination-address address <destination-port port-number>;
                    syslog;
                }
            }
        }
    }
}
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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

**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-shaping-overhead

**Syntax** `egress-shaping-overhead number;`

**Hierarchy Level** `[edit chassis fpc slot-number pic pic-number traffic-manager],`  
`[edit chassis lcc number fpc slot-number pic pic-number traffic-manager]`

**Release Information** Statement introduced in Junos OS Release 8.3.

**Description** Number of bytes to add to packet to determine shaped session packet length.



**NOTE:** On M Series and T Series routers with Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs and Enhanced IQ2 (IQ2E) PICs and on MX Series routers with Dense Port Concentrators (DPCs) only, to account for egress shaping overhead bytes added to output traffic on the line card, you must use the `egress-policer-overhead` statement to explicitly configure corresponding egress policing overhead for Layer 2 policers, MAC policers, or queue rate limits applied to output traffic on the line card.



**NOTE:** For MIC and MPC interfaces on MX Series routers, by default the value of `egress-shaping-overhead` is configured to 20, which means that the number of class-of-service (CoS) shaping overhead bytes to be added to the packets is 20. The interfaces on DPCs in MX Series routers, the default value is zero. For interfaces on PICs other than the 10-port 10-Gigabit Oversubscribed Ethernet (OSE) Type 4, you should configure `egress-shaping-overhead` to a minimum of 20 bytes to add a shaping overhead of 20 bytes to the packets.



**NOTE:** When you change the `egress-shaping-overhead` value, on M Series, T Series, and MX104 routers the PIC on which it is changed is restarted. On other MX Series routers the DPC/MPC on which it is changed is restarted.

**Options** *number*—When traffic management (queuing and scheduling) is configured on the egress side, the number of CoS shaping overhead bytes to add to the packets on the egress interface.

**Range:**

- –63 through 192.
- –62 through 192 for vSRX.



**NOTE:** The L2 headers (DA/SA + VLAN tags) are automatically a part of the shaping calculation.

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

**Related Documentation**

- *egress-policer-overhead*
- *Configuring CoS for L2TP Tunnels on ATM Interfaces*
- *ingress-shaping-overhead*
- [mode \(Layer 2 Tunneling Protocol Shaping\) on page 339](#), *ingress-shaping-overhead*
- [traffic-manager on page 404](#)


## excess-bandwidth-share

---

<b>Syntax</b>	<code>excess-bandwidth-share (proportional <i>value</i>   equal);</code>
<b>Hierarchy Level</b>	<code>[edit class-of-service <a href="#">interfaces interface-set interface-set-name</a>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5.
<b>Description</b>	<p>Specify the method of sharing excess bandwidth in a hierarchical scheduler environment. With hierarchical schedulers, you can provide shaping and scheduling at the service VLAN level as well as other levels, such as the physical interface. You can also group a set of logical interfaces and then apply scheduling and shaping parameters to the logical interface set.</p> <p>To configure CoS hierarchical schedulers, you must enable hierarchical scheduling by including the <a href="#">hierarchical-scheduler</a> statement at the <code>[edit interfaces]</code> hierarchy for the physical interface. If you do not include this statement, the interfaces on the MX Series router cannot use hierarchical interfaces.</p> <p>The Enhanced Queuing DPC supports the following hierarchical scheduler characteristics:</p> <ul style="list-style-type: none"><li>• Shaping at the physical interface level</li><li>• Shaping and scheduling at the service VLAN interface set level</li><li>• Shaping and scheduling at the customer VLAN logical interface level</li><li>• Scheduling at the queue level</li></ul>
<b>Options</b>	<p><b>equal</b>—Share excess bandwidth equally among the configured interfaces.</p> <p><b>proportional <i>value</i></b>—(Default) Share excess bandwidth proportionally according to the specified value. In this mode, the excess bandwidth is shared at the ratio of the logical interface shaping rates.</p> <p><b>Default:</b> 32.64 Mbps</p> <p>This example sets the excess bandwidth sharing for an Enhanced Queuing DPC interface proportionally at a rate of 100 Mbps and a shaping rate of 80 Mbps applied to the interface through the output-traffic-control profile for scheduling and shaping:</p> <pre>[edit class-of-service interfaces interface-set example-interface-set] user@host# set excess-bandwidth-share proportional 100m user@host# set output-traffic-control-profile PIR-80Mbps</pre>
<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 Hierarchical Schedulers for CoS on page 12</a></li><li>• <a href="#">Configuring Interface Sets</a></li></ul>


- *Enhanced Queuing DPC CoS Properties*
- *Configuring MDRR on Enhanced Queuing DPCs*

## excess-priority


<b>Syntax</b>	excess-priority [ low   medium-low   medium-high   high   none];
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">schedulers</a> <i>scheduler-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3. Option <b>none</b> introduced in Junos OS Release 11.4.
<b>Description</b>	Determine the priority of excess bandwidth traffic on a scheduler.
	<div>  <p><b>NOTE:</b> For Link Services IQ (LSQ) PICs or Multiservices PIC (MS-PICs), the <b>excess-priority</b> statement is allowed for consistency, but ignored. If an explicit priority is not configured for these interfaces, a default low priority is used. This default priority is also used in the excess region.</p> </div>
<b>Options</b>	<p><b>low</b>—Excess traffic for this scheduler has low priority.</p> <p><b>medium-low</b>—Excess traffic for this scheduler has medium-low priority.</p> <p><b>medium-high</b>—Excess traffic for this scheduler has medium-high priority.</p> <p><b>high</b>—Excess traffic for this scheduler has high priority.</p> <p><b>none</b>—System does not demote the priority of guaranteed traffic when the bandwidth exceeds the shaping rate or the guaranteed rate.</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 Excess Bandwidth Sharing on IQE PICs</i></li> <li>• <i>Bandwidth Sharing on Nonqueuing Packet Forwarding Engines Overview</i></li> <li>• <i>Managing Excess Bandwidth Distribution on Static Interfaces on MICs and MPCs</i></li> </ul>

## excess-rate (Dynamic Traffic Shaping)

---


<b>Syntax</b>	excess-rate (percent <i>percentage</i>   \$junos-cos-excess-rate)   proportion <i>value</i> ;
<b>Hierarchy Level</b>	[edit <a href="#">dynamic-profiles</a> <i>profile-name</i> class-of-service traffic-control-profiles <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.2.
<b>Description</b>	For an MPC interface, determine the percentage or proportion of excess bandwidth traffic to share for all priorities of traffic.
<b>Options</b>	<p><b><i>percentage</i></b>—Percentage of the excess bandwidth to share. <b>Range:</b> 0 through 100 percent</p> <p><b><i>value</i></b>—Proportion of the excess bandwidth to share. <b>Range:</b> 0 through 1000</p>
	<hr/> <div> <b>NOTE:</b> The proportion of excess bandwidth on MPC2-3D MPCs can be configured with increments of 1 from 0 through 1000. All other MPCs should be configured with increments of 10 from 0 through 1000.</div> <hr/>
	<p><b>\$junos-cos-excess-rate</b>—Variable for the excess rate that is specified for the logical interface. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.</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>Guidelines for Configuring Dynamic CoS for Subscriber Access</i></li><li>• <i>Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MPC Interfaces</i></li><li>• <i>output-traffic-control-profile</i></li></ul>

## excess-rate-high (Dynamic Traffic Shaping)

<b>Syntax</b>	<code>excess-rate-high ((percent <i>percentage</i>   \$junos-cos-excess-rate-high)   proportion <i>value</i>);</code>
<b>Hierarchy Level</b>	[edit <a href="#">dynamic-profiles</a> <i>profile-name</i> class-of-service traffic-control-profiles <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	For an MPC/MIC interface, determine the percentage of excess bandwidth for high-priority traffic to share.
<b>Options</b>	<p><b><i>percentage</i></b>—Percentage of the excess bandwidth to share.  <b>Range:</b> 0 through 100 percent</p> <p><b><i>value</i></b>—Proportion of the excess bandwidth to share.  <b>Range:</b> 0 through 1000</p>
<div>  <p><b>NOTE:</b> The proportion of excess bandwidth on MPC2-3D MPCs can be configured with increments of 1 from 0 through 1000. All other MPCs should be configured with increments of 10 from 0 through 1000.</p> </div>	
<p><b>\$junos-cos-excess-rate-high</b>—Variable for the excess rate that is specified for high-priority traffic on the logical interface. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.</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>Guidelines for Configuring Dynamic CoS for Subscriber Access</i></li> <li><i>Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MPC Interfaces</i></li> <li><i>output-traffic-control-profile</i></li> </ul>

## excess-rate-low (Dynamic Traffic Shaping)

---

Syntax	excess-rate-low ((percent <i>percentage</i>   \$junos-cos-excess-rate-low)   proportion <i>value</i> );
Hierarchy Level	[edit <a href="#">dynamic-profiles</a> <i>profile-name</i> class-of-service traffic-control-profiles <i>profile-name</i> ]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	For an MPC/MIC interface, determine the percentage of excess bandwidth for low-priority traffic to share.
Options	<p><b>percentage</b>—Percentage of the excess bandwidth to share. <b>Range:</b> 0 through 100 percent</p> <p><b>value</b>—Proportion of the excess bandwidth to share. <b>Range:</b> 0 through 1000</p>
<hr/> <div> <b>NOTE:</b> The proportion of excess bandwidth on MPC2-3D MPCs can be configured with increments of 1 from 0 through 1000. All other MPCs should be configured with increments of 10 from 0 through 1000.</div> <hr/>	
<p><b>\$junos-cos-excess-rate-low</b>—Variable for the excess rate that is specified for low-priority traffic on the logical interface. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.</p> <hr/>	
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><i>Guidelines for Configuring Dynamic CoS for Subscriber Access</i></li><li><i>Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MPC Interfaces</i></li><li><i>output-traffic-control-profile</i></li></ul>



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

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

**Description** Configure protocol family information for the logical interface.



**NOTE:** Not all subordinate 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*.

The remaining statements are explained separately. See [CLI Explorer](#).

<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]`  
                   `}`  
                   `excluded-address ip-address;`  
                   `excluded-range name low minimum-value high maximum-value;`  
                   `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. See [CLI Explorer](#).

**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. See [CLI Explorer](#).

**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. See [CLI Explorer](#).

**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. See [CLI Explorer](#).

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

**Related Documentation**

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

## fill-level (Drop Profiles)

---

<b>Syntax</b>	fill-level <i>percentage</i> ;
<b>Hierarchy Level</b>	[edit class-of-service <b>drop-profiles</b> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS 11.4 for EX Series switches.
<b>Description</b>	When configuring RED, map the fullness of a queue to a drop probability.
<b>Options</b>	<p><b>percentage</b>—How full the queue is, expressed as a percentage. You configure the <b>fill-level</b> and <b>drop-probability</b> statements in pairs. To specify multiple fill levels, include multiple <b>fill-level</b> and <b>drop-probability</b> statements. The values you assign to each statement pair must increase relative to the previous pair's values. This is shown in the discrete graph in <i>Managing Congestion Using RED Drop Profiles and Packet Loss Priorities</i>.</p> <p><b>Range:</b> 0 through 100 percent</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>Managing Congestion Using RED Drop Profiles and Packet Loss Priorities</i></li> <li>• <i>Defining Packet Drop Behavior by Configuring RED Drop Profiles</i></li> </ul>

## fill-level (Interpolated Value)

---

<b>Syntax</b>	fill-level [ <i>values</i> ];
<b>Hierarchy Level</b>	[edit class-of-service <b>drop-profiles</b> <i>profile-name</i> interpolate]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS 11.4 for EX Series switches.
<b>Description</b>	Define up to 64 values for interpolating queue fill level.  On EX Series switches, this statement is supported only on EX8200 standalone switches and EX8200 Virtual Chassis.
<b>Options</b>	<b>values</b> —Data points for mapping queue fill percentage. <b>Range:</b> 0 through 100 <b>Default:</b> In the default tail drop profile, when the fill level is 0 percent, the drop probability is 0 percent. When the fill level is 100 percent, the drop probability is 100 percent.
<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>Managing Congestion Using RED Drop Profiles and Packet Loss Priorities</i></li><li>• <i>Defining Packet Drop Behavior by Configuring RED Drop Profiles.</i></li></ul>

## filter (Configuring)


<b>Syntax</b>	<pre>filter <i>filter-name</i> {     accounting-profile <i>name</i>;     enhanced-mode;     fast-lookup-filter;     filter-list-template;     interface-shared;     interface-specific;     physical-interface-filter;     promote gre-key;     term <i>term-name</i> {         ... term configuration ...     } }</pre>
<b>Hierarchy Level</b>	<p>[edit <a href="#">dynamic-profiles</a> <i>profile-name</i> firewall family <i>family-name</i>],          [edit firewall family <i>family-name</i>],          [edit logical-systems <i>logical-system-name</i> firewall family <i>family-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.          Logical systems support introduced in Junos OS Release 9.3.  <b>physical-interface-filter</b> statement introduced in Junos OS Release 9.6.          Support for the <b>interface-shared</b> statement introduced in Junos OS Release 12.2.          Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Configure firewall filters.
<b>Options</b>	<p><b><i>filter-name</i></b>—Name that identifies the filter. This must be a non-reserved string of not more than 64 characters. To include spaces in the name, enclose it in quotation marks (" "). Firewall filter names are restricted from having the form <b>__.*__</b> (beginning and ending with underscores) or <b>__.*</b> (beginning with an underscore).</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>
<b>Required Privilege Level</b>	<p>firewall—To view this statement in the configuration.          firewall-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Guidelines for Configuring Firewall Filters</i></li> <li>• <i>Guidelines for Applying Standard Firewall Filters</i></li> <li>• <i>Configuring Multifield Classifiers</i></li> <li>• <i>Using Multifield Classifiers to Set Packet Loss Priority</i></li> <li>• <a href="#">simple-filter on page 396</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>	Statement introduced in Junos OS Release 8.1. Support for aggregated Ethernet added in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers. Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.
<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.</p> <p>This statement is supported on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series and QFX Series switches.</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>Enabling VLAN Tagging</i></li><li>• <i>Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers</i></li><li>• <i>Configuring Double-Tagged VLANs on Layer 3 Logical Interfaces</i></li></ul>

## forwarding-classes (Class-of-Service)

<b>Syntax</b>	<pre>forwarding-classes {   class queue-num <i>queue-number</i> priority (high   low);   <i>queue</i> <i>queue-number</i> <i>class-name</i> priority (high   low) [ policing-priority (premium   normal) ]; }</pre>
<b>Hierarchy Level</b>	[edit class-of-service]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>policing-priority</b> option introduced in Junos OS Release 9.5.</p> <p>Statement introduced on PTX Series Packet Transport Routers in Junos OS Release 12.1.</p>
<b>Description</b>	<p>Associate the forwarding class with a queue name and number. For M320, MX Series, T Series routers and EX Series switches only, you can configure fabric priority queuing by including the <b>priority</b> statement. For Enhanced IQ PICs, you can include the <b>policing-priority</b> option.</p>
<div style="display: flex; align-items: center;">  <div> <p><b>NOTE:</b> The <b>priority</b> and <b>policing-priority</b> options are not supported on PTX Series Packet Transport Routers.</p> </div> </div>	
<p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</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 Custom Forwarding Class for Each Queue</i></li> <li>• <i>Forwarding Classes and Fabric Priority Queues</i></li> <li>• <i>Configuring Layer 2 Policers on IQE PICs</i></li> <li>• <i>Classifying Packets by Egress Interface</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;
                interface-description (device-interface | logical-interface);
                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 {
                asymmetric-lease-time seconds;
                asymmetric-prefix-lease-time seconds;
                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;
}
layer2-liveness-detection {
    max-consecutive-retries number;
    transmit-interval interval;
}
}
}
overrides {
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    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;
    }
    protocol-attributes attribute-set-name;
    rapid-commit;
}
reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
        radius-disconnect;
    }
}
}
route-suppression;
service-profile dynamic-profile-name;
}

```

**Hierarchy Level** [edit system services [dhcp-local-server](#)],  
 [edit system services [dhcp-local-server](#) dhcpv6],  
 [edit logical-systems *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 12.1 for EX Series switches.

**Description** Configure a group of interfaces that have a common configuration, such as authentication parameters. A group must contain at least one interface.

**Options** *group-name*—Name of the group.


The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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

**Related Documentation**

- *Extended DHCP Local Server Overview*
- *Grouping Interfaces with Common DHCP Configurations*
- *Using External AAA Authentication Services with DHCP*
- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*
- *DHCP Liveness Detection Using ARP and Neighbor Discovery Packets*


## guaranteed-rate

<b>Syntax</b>	<code>guaranteed-rate (percent <i>percentage</i>   <i>rate</i>) &lt;burst-size <i>bytes</i>&gt;;</code>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">traffic-control-profiles</a> <i>profile-name</i> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Option <b>burst-size</b> introduced for Enhanced Queuing (EQ) DPC interfaces in Junos OS Release 9.4.</p> <p>Option <b>burst-size</b> introduced for MIC and MPC interfaces in Junos OS Release 11.4.</p> <p>Option <b>burst-size</b> introduced for IQ2 and IQ2E interfaces in Junos OS Release 12.3</p>
<b>Description</b>	For Gigabit Ethernet IQ, Channelized IQ PICs, Multiservices and Services PICs FRF.16 LSQ interfaces, and EQ DPCs only, configure a guaranteed minimum rate. You can also configure an optional burst size for a logical interface on EQ DPCs and on IQ2 and IQ2E PICs. This can help to ensure that higher priority services do not starve lower priority services.
<b>Default</b>	If you do not include this statement and you do not include the <b>delay-buffer-rate</b> statement, the logical interface receives a minimal delay-buffer rate and minimal bandwidth equal to 2 MTU-sized packets.
<b>Options</b>	<p><b>percent <i>percentage</i></b>—For LSQ interfaces, guaranteed rate as a percentage of the available interface bandwidth.</p> <p><b>Range:</b> 1 through 100 percent</p> <p><b><i>rate</i></b>—For IQ and IQ2 interfaces, guaranteed rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 1000 through 6,400,000,000,000 bps</p>
	<p> <b>NOTE:</b> Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Beginning with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.</p>
	<b>burst-size <i>bytes</i></b> —(Optional) Maximum burst size, in bytes.
<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>Providing a Guaranteed Minimum Rate</i></li> <li><i>Configuring Traffic Control Profiles for Shared Scheduling and Shaping</i></li> </ul>

- [output-traffic-control-profile on page 340](#)

## hierarchical-scheduler

---

<b>Syntax</b>	hierarchical-scheduler;
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">interfaces</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5.
<b>Description</b>	On MX Series, M Series, and T Series routers with IQ2E PIC, enables the use of hierarchical schedulers.
	<div> <b>NOTE:</b> To enable hierarchical scheduling on MX80 and MX104 routers, configure the <code>hierarchical-scheduler</code> statement at each member physical interface level of a particular aggregated Ethernet interface as well as at that aggregated Ethernet interface level. On other routing platforms, it is enough if you include this statement at the aggregated Ethernet interface level.</div>
<b>Default</b>	If you do not include this statement, the interfaces on the MX Series router cannot use hierarchical 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>• <a href="#">Configuring Hierarchical Schedulers for CoS on page 12</a></li><li>• <a href="#">Understanding Hierarchical CoS for Subscriber Interfaces on page 70</a></li><li>• <a href="#">hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) on page 311</a></li></ul>

## hierarchical-scheduler (Subscriber Interfaces on MX Series Routers)

<b>Syntax</b>	<pre> hierarchical-scheduler {     implicit-hierarchy;     maximum-hierarchy-levels <i>number</i>; } </pre>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.1.</p> <p><b>implicit-hierarchy</b> option added in Junos OS Release 13.1.</p> <p>Support on GRE tunnel interfaces configured on physical interfaces on MICs or MPCs in MX Series routers added in Junos OS Release 13.3.</p> <p>Support for up to four hierarchy levels added in Junos OS Release 16.1.</p>
<b>Description</b>	<p>Configure hierarchical scheduling options on the interface.</p> <p>The statement is supported on the following interfaces:</p> <ul style="list-style-type: none"> <li>• MIC and MPC interfaces in MX Series routers</li> <li>• GRE tunnel interfaces configured on physical interfaces hosted on MIC or MPC line cards in MX Series routers</li> </ul> <p>To enable hierarchical scheduling on MX Series routers, configure the <b>hierarchical-scheduler</b> statement at each member physical interface level of a particular aggregated Ethernet interface as well as at that aggregated Ethernet interface level. On other routing platforms, it is enough if you include this statement at the aggregated Ethernet interface level.</p>
<b>Options</b>	<p><b>implicit-hierarchy</b>—Configure four-level hierarchical scheduling. When you include the <b>implicit-hierarchy</b> option, a hierarchical relationship is formed between the CoS scheduler nodes at level 1, level 2, level 3, and level 4. The <b>implicit-hierarchy</b> option is supported only on MPC/MIC subscriber interfaces and interface sets on MX Series routers.</p> <p><b>maximum-hierarchy-levels <i>number</i></b>—Specify the maximum number of hierarchical scheduling levels allowed for node scaling, from 2 through 4 levels. The default number of levels is 3. The <b>maximum-hierarchy-levels</b> option is supported on MPC/MIC or EQ DPC subscriber interfaces and interface sets on MX Series routers.</p> <ul style="list-style-type: none"> <li>• If you set <b>maximum-hierarchy-levels</b> to 2, interface sets are not allowed. In this case, if you configure a level 2 interface set, you generate Packet Forwarding Engine errors.</li> <li>• If you do not include the <b>maximum-hierarchy-levels</b> option, keeping the default number of hierarchy levels at 3, interface sets can be at either level 2 or level 3, depending on whether the member logical interfaces within the interface set have a traffic control profile. If any member logical interface has a traffic control profile, then the interface set is a level 2 CoS scheduler node. If no member logical interface has a traffic control profile, the interface set is at level 3.</li> </ul>

<b>Required Privilege Level</b>	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Hierarchical CoS for Subscriber Interfaces on page 70</a></li><li>• <a href="#">Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links on page 79</a></li><li>• <a href="#">Configuring Hierarchical Schedulers for CoS on page 12</a></li><li>• <a href="#">Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface on page 80</a></li><li>• <a href="#">Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 117</a></li></ul>

---

## ieee-802.1ad

<b>Syntax</b>	ieee-802.1ad ( <i>rewrite-name</i>   default) <b>vlan-tag</b> (outer   outer-and-inner);
<b>Hierarchy Level</b>	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>rewrite-rules</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2.
<b>Description</b>	Apply a IEEE-802.1ad rewrite rule.
<b>Options</b>	<p><b>rewrite-name</b>—Name of a <b>rewrite-rules</b> mapping configured at the [edit class-of-service <b>rewrite-rules</b> <b>ieee-802.1ad</b>] hierarchy level.</p> <p><b>default</b>—The default rewrite bit mapping.</p> <p><b>vlan-tag</b>—The rewrite rule is applied to the <b>outer</b> or <b>outer-and-inner</b> VLAN tag.</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 Rewrite Rules</a></li><li>• <a href="#">dscp (Rewrite Rules) on page 270</a></li><li>• <a href="#">dscp-ipv6 (CoS Rewrite Rules) on page 272</a></li><li>• <i>exp</i></li><li>• <i>exp-push-push-push</i></li><li>• <i>exp-swap-push-push</i></li><li>• <i>ieee-802.1 (Rewrite Rules on Logical Interface)</i></li><li>• <i>inet-precedence (CoS Rewrite Rules)</i></li><li>• <a href="#">rewrite-rules (Definition) on page 370</a></li></ul>

## inet-precedence (CoS Classifiers)

<b>Syntax</b>	<pre>inet-precedence (classifier-name   default) {     import (classifier-name   default);     forwarding-class class-name {         loss-priority level code-points [ aliases ] [ bit-patterns ]     } }</pre>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">classifiers</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Apply an IPv4 classifier.
<b>Options</b>	<p><b>default</b>—The default mapping. By default, IP precedence rewrite rules alter the first three bits on the type of service (ToS) byte while leaving the last three bits unchanged.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</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>Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic</i></li> </ul>


## input-excess-bandwidth-share

---

<b>Syntax</b>	input-excess-bandwidth-share (proportional <i>value</i>   equal);
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">interfaces</a> <i>interface-name</i> ], [edit class-of-service <a href="#">interfaces</a> <a href="#">interface-set</a> <i>interface-set-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0.
<b>Description</b>	Determines the method of sharing excess bandwidth on the ingress interface in a hierarchical scheduler environment. If you do not include this statement, the node shares excess bandwidth proportionally at 32.64 Mbps.
<b>Options</b>	<b>proportional <i>value</i></b> —(Default) Share ingress excess bandwidth proportionally (default value is 32.64 Mbps).  <b>equal</b> —Share ingress excess bandwidth equally.
<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 Ingress Hierarchical CoS on Enhanced Queuing DPCs on page 35</a></li></ul>




## input-scheduler-map


<b>Syntax</b>	<code>input-scheduler-map <i>map-name</i>;</code>
<b>Hierarchy Level</b>	<code>[edit class-of-service interfaces <i>interface-name</i>],</code> <code>[edit class-of-service interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	<p>Associate a scheduler map with a physical or logical input interface. The <b>input-scheduler-map</b> and <b>input-traffic-control-profile</b> statements are mutually exclusive at the same hierarchy level.</p> <p><b>input-scheduler-map</b> is supported on the following Ethernet interfaces:</p> <ul style="list-style-type: none"> <li>• IQ2 and IQ2E PICs</li> <li>• DPCs and MPCs that support Enhanced Queuing (Q/EQ)</li> <li>• MX80 with support for per-VLAN queuing</li> </ul>
	<p> <b>NOTE:</b></p> <p>For an Enhanced Queuing (EQ) DPC on an MX Series router, CoS queuing and scheduling are enabled on the egress side but disabled on the ingress side by default. To enable ingress CoS on the EQ DPC, you must configure the <b>traffic-manager</b> statement with ingress-and-egress mode:</p> <pre>[edit chassis fpc <i>slot-number</i> pic <i>pic-number</i>] traffic-manager mode ingress-and-egress;</pre>
<b>Options</b>	<p><b>map-name</b>—Name of scheduler map that you define at the <code>[edit class-of-service scheduler-maps]</code> hierarchy level.</p> <p><b>default</b>—The default scheduler mapping.</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 Input Scheduler on an Interface</a></li> <li>• <a href="#">Configuring Ingress Hierarchical CoS on Enhanced Queuing DPCs on page 35</a></li> <li>• <a href="#">input-traffic-control-profile on page 318</a></li> </ul>

## input-shaping-rate (Logical Interface)



---

<b>Syntax</b>	input-shaping-rate (percent <i>percentage</i>   <i>rate</i> );
<b>Hierarchy Level</b>	[edit class-of-service interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	For Gigabit Ethernet IQ2, Enhanced Queuing DPC, MIC, and MPC interfaces, configure input traffic shaping by specifying the amount of bandwidth to be allocated to the logical interface. You can configure hierarchical shaping, meaning you can apply an input shaping rate to both the physical interface and the logical interface.
<b>Default</b>	If you do not include this statement, logical interfaces share a default scheduler. This scheduler has a committed information rate (CIR) that equals 0. (The CIR is the guaranteed rate.) The default scheduler has a peak information rate (PIR) that equals the physical interface shaping rate.
<b>Options</b>	<p><b>percent <i>percentage</i></b>—Shaping rate as a percentage of the available interface bandwidth. <b>Range:</b> 0 through 100 percent</p> <p><b><i>rate</i></b>—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000). <b>Range:</b> 1000 through 6,400,000,000,000 bps</p>
	<div> <b>NOTE:</b> Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Beginning with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.</div>
<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 Ingress Hierarchical CoS on MIC and MPC Interfaces on page 41</a></li><li>• <a href="#">Configuring Input Shaping Rates for Both Physical and Logical Interfaces</a></li><li>• <a href="#">Configuring Ingress Hierarchical CoS on Enhanced Queuing DPCs on page 35</a></li><li>• <a href="#">input-traffic-control-profile on page 318</a></li></ul>

## input-shaping-rate (Physical Interface)

<b>Syntax</b>	<code>input-shaping-rate rate;</code>
<b>Hierarchy Level</b>	[edit class-of-service interfaces <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	For Gigabit Ethernet IQ2, Enhanced Queuing DPC, MIC, and MPC interfaces, configure input traffic shaping by specifying the amount of bandwidth to be allocated to the physical interface. You can configure hierarchical shaping, meaning you can apply an input shaping rate to both the physical interface and the logical interface.
<b>Options</b>	<p><b>rate</b>—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 1000 through 6,400,000,000,000 bps</p>
<div>  <p><b>NOTE:</b> Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Beginning with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.</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>• <a href="#">Configuring Input Shaping Rates for Both Physical and Logical Interfaces</a></li> <li>• <a href="#">Configuring Ingress Hierarchical CoS on Enhanced Queuing DPCs on page 35</a></li> <li>• <a href="#">input-traffic-control-profile on page 318</a></li> </ul>

## input-traffic-control-profile

<b>Syntax</b>	<code>input-traffic-control-profile <i>profile-name</i> <b>shared-instance</b> <i>instance-name</i>;</code>
<b>Hierarchy Level</b>	<code>[edit class-of-service interfaces <i>interface-name</i>],</code> <code>[edit class-of-service interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	For Gigabit Ethernet IQ2 and IQ2E PIC, Enhanced Queuing DPC, MIC, and MPC interfaces, apply an input traffic scheduling and shaping profile to the logical interface. The <b>input-traffic-control-profile</b> and <b>input-scheduler-map</b> statements are mutually exclusive at the same hierarchy level.
<div>  <p><b>NOTE:</b> The <b>shared-instance</b> statement applies only to Gigabit Ethernet IQ2 and IQ2E PICs.</p> </div>	
<div>  <p><b>NOTE:</b></p> <p>For an Enhanced Queuing (EQ) DPC on an MX Series router, CoS queuing and scheduling are enabled on the egress side but disabled on the ingress side by default. To enable ingress CoS on the EQ DPC, you must configure the <b>traffic-manager</b> statement with <b>ingress-and-egress</b> mode:</p> <pre>[edit chassis fpc <i>slot-number</i> pic <i>pic-number</i>] traffic-manager mode ingress-and-egress;</pre> </div>	
<b>Options</b>	<p><b><i>profile-name</i></b>—Name of the traffic-control profile to be applied to this interface.</p> <p><b><i>instance-name</i></b>—Name of the shared scheduler and shaper to be applied to this 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>• <a href="#">Configuring Traffic Control Profiles for Shared Scheduling and Shaping</a></li> <li>• <a href="#">Configuring Ingress Hierarchical CoS on Enhanced Queuing DPCs on page 35</a></li> <li>• <a href="#">input-shaping-rate (Logical Interface) on page 316</a></li> <li>• <a href="#">input-scheduler-map on page 315</a></li> <li>• <a href="#">traffic-control-profiles on page 403</a></li> </ul>

## input-traffic-control-profile-remaining

---

<b>Syntax</b>	<code>input-traffic-control-profile-remaining <i>profile-name</i>;</code>
<b>Hierarchy Level</b>	[edit class-of-service interfaces <i>interface-name</i> ], [edit class-of-service interfaces <i>interface-name</i> <b>interface-set</b> <i>interface-set-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0.
<b>Description</b>	For Enhanced Queuing DPC, MICs, or MPC interfaces on MX Series routers, or for IQ2E PICs interfaces on M Series and T Series router, apply an input traffic scheduling and shaping profile for the remaining traffic to the logical interface or interface set.
<b>Options</b>	<b><i>profile-name</i></b> —Name of the traffic-control profile for the remaining traffic to be applied to this interface or interface set.
<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 Ingress Hierarchical CoS on Enhanced Queuing DPCs on page 35</a></li> <li>• <a href="#">input-traffic-control-profile on page 318</a></li> </ul>

## interface (DHCP Local Server)

**Syntax**    `interface interface-name {  
               access-profile profile-name;  
               exclude;  
               overrides {  
                   asymmetric-lease-time seconds;  
                   asymmetric-prefix-lease-time seconds;  
                   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 supports only 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. See [CLI Explorer](#).

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

**Related Documentation**

- *Extended DHCP Local Server Overview*
- *Grouping Interfaces with Common DHCP Configurations*
- *Using External AAA Authentication Services with DHCP*

## interface-set (Ethernet Interfaces)

**Syntax**

```
interface-set interface-set-name {
    interface ethernet-interface-name {
        (unit unit-number | vlan-tags-outer vlan-tag);
    }
}
```

**Hierarchy Level** [edit interfaces]

**Release Information** Statement introduced in Junos OS Release 8.5.

**Description** The set of interfaces used to configure hierarchical CoS schedulers on Ethernet interfaces on the MX Series router and IQ2E PIC on M Series and T Series routers.

The remaining statements are described separately.

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

**Related Documentation**

- [interface-set \(Hierarchical Schedulers\) on page 322](#)

## interface-set (Hierarchical Schedulers)

---

<b>Syntax</b>	<pre>interface-set <i>interface-set-name</i> {     excess-bandwidth-share (proportional <i>value</i>   equal);     internal-node;     output-traffic-control-profile <i>profile-name</i>;     output-traffic-control-profile-remaining <i>profile-name</i>; }</pre>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">interfaces</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5.
<b>Description</b>	For Enhanced Queuing DPC, MIC, or MPC interfaces on MX Series routers, or for IQ2E PIC interfaces on M Series routers, configure hierarchical schedulers for an interface set.
<b>Options</b>	<p><i>interface-set-name</i>—Name of the interface set.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</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 Interface Sets</a></li><li>• <a href="#">Configuring Hierarchical Schedulers for CoS on page 12</a></li></ul>



## interfaces

<b>List of Syntax</b>	<a href="#">Syntax (QFX Series) on page 323</a> <a href="#">Syntax (EX Series, MX Series and T Series) on page 323</a>
<b>Syntax (QFX Series)</b>	<pre>interfaces <i>interface-name</i> {     no-mac-learning; }</pre>
<b>Syntax (EX Series, MX Series and T Series)</b>	<pre>interfaces { ... }</pre>
<b>QFX Series</b>	[edit ethernet-switching-options]
<b>EX Series, MX Series and T Series</b>	[edit]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Configure settings for interfaces that have been assigned to family <b>ethernet-switching</b> .
<b>Default</b>	The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.
<b>Options</b>	<p><b><i>interface-name</i></b> —Name of an interface that is configured for family <b>ethernet-switching</b>.</p> <p>The remaining statement is explained separately. See <a href="#">CLI Explorer</a>.</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> <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>Physical Interface Configuration Statements Overview</i></li> <li>• <i>Configuring Aggregated Ethernet Link Protection</i></li> </ul>

## interfaces (CoS)

```
Syntax interfaces {
    interface-name {
        classifiers{
            dscp(classifier-name | default) {
            }
            ieee-802.1 (classifier-name | default) vlan-tag (inner | outer | classifier-name);
            inet-precedence (rewrite-name | default);
        }
        input-scheduler-map map-name;
        input-shaping-rate rate;
        irb {
            unit logical-unit-number {
                classifiers {
                    type (classifier-name | default);
                    no-default;
                }
                rewrite-rules {
                    dscp (rewrite-name | default);
                    dscp-ipv6 (rewrite-name | default);
                    exp (rewrite-name | default) protocol protocol-types;
                    ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                    inet-precedence (rewrite-name | default);
                }
            }
        }
        member-link-scheduler (replicate | scale);
        rewrite-rules {
            dscp (rewrite-name | default);
            ieee-802.1 (rewrite-name | default) vlan-tag (outer);
            inet-precedence (rewrite-name | default);
        }
        scheduler-map map-name;
        scheduler-map-chassis map-name;
        shaping-rate rate;
        unit logical-unit-number {
            classifiers {
                type (classifier-name | default) family (mpls | inet);
            }
            forwarding-class class-name;
            fragmentation-map map-name;
            input-shaping-rate (percent percentage | rate);
            input-traffic-control-profile profile-name shared-instance instance-name;
            output-traffic-control-profile profile-name shared-instance instance-name;
            per-session-scheduler;
            policy-map policy-map-name;
            rewrite-rules {
                dscp (rewrite-name | default);
                dscp-ipv6 (rewrite-name | default);
                exp (rewrite-name | default) protocol protocol-types;
                exp-push-push-push default;
                exp-swap-push-push default;
            }
        }
    }
}
```

```

    ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
    inet-precedence (rewrite-name | default);
  }
  scheduler-map map-name;
  shaping-rate rate;
  translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 | to-exp-from-exp
    | to-inet-precedence-from-inet-precedence) table-name;
}
}
interface-set interface-set-name {
  excess-bandwidth-share;
  internal-node;
  output-traffic-control-profile profile-name;
  output-traffic-control-profile-remaining profile-name;
}
}

```

<b>Hierarchy Level</b>	[edit class-of-service]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Interface-set level added in Junos OS Release 8.5.
<b>Description</b>	Configure interface-specific CoS properties for incoming packets.
<b>Options</b>	The remaining statements are explained separately. See <a href="#">CLI Explorer</a> .
<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>Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic</i></li> <li>• <i>Configuring Rewrite Rules</i></li> </ul>

## interfaces (Static and Dynamic Subscribers)

---

```
Syntax  interfaces {
        interface-name {
            unit logical-unit-number {
                actual-transit-statistics;
                auto-configure {
                    agent-circuit-identifier {
                        dynamic-profile profile-name;
                    }
                    line-identity {
                        include {
                            accept-no-ids;
                            circuit-id;
                            remote-id;
                        }
                        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;
    }
}
host-prefix-only;
ppp-options {
    chap;
    pap;
}
proxy-arp;
service {
    pcef pcef-profile-name {
        activate rule-name | activate-all;
    }
}
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;  
            }  
            ignore-magic-number-mismatch;  
            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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

**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*

---

## internal-node

---

**Syntax** internal-node;

**Hierarchy Level** [edit class-of-service [interfaces interface-set](#) *interface-set-name*]

**Release Information** Statement introduced in Junos OS Release 8.5.

**Description** The statement is used to raise the interface set without children to the same level as the other configured interface sets with children, allowing them to compete for the same set of resources.

**Default** If you do not include this statement, the node is internal only if its children have a traffic control profile configured.

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

**Related Documentation**

- *Configuring Internal Scheduler Nodes*



## interpolate

---


<b>Syntax</b>	<pre> interpolate {   drop-probability [values];   fill-level [values]; } </pre>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">drop-profiles</a> <i>profile-name</i> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced before Junos OS 11.4 for EX Series switches.</p>
<b>Description</b>	<p>Specify values for interpolating relationship between queue fill level and drop probability.</p> <p>On EX Series switches, this statement is supported only on EX8200 standalone switches and EX8200 Virtual Chassis.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</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>• See <i>Defining Packet Drop Behavior by Configuring RED Drop Profiles</i>.</li> </ul>

## loss-priority (BA Classifiers)

---

<b>Syntax</b>	<code>loss-priority <i>level</i>;</code>
<b>Hierarchy Level</b>	[edit class-of-service classifiers <i>type classifier-name</i> forwarding-class <i>class-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 14.2 for PTX Series Packet Transport Routers.
<b>Description</b>	Specify packet loss priority value for a specific set of code-point aliases and bit patterns.
<b>Options</b>	<i>level</i> can be one of the following: <ul style="list-style-type: none"><li>• <b>high</b>—Packet has high loss priority.</li><li>• <b>medium-high</b>—Packet has medium-high loss priority.</li><li>• <b>medium-low</b>—Packet has medium-low loss priority.</li><li>• <b>low</b>—Packet has low loss priority.</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>Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic</i></li><li>• <i>Configuring and Applying Tricolor Marking Policers</i></li></ul>

## loss-priority (Scheduler Drop Profiles)

<b>Syntax</b>	loss-priority (any   high   low   medium-high   medium-low);
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">schedulers</a> <i>scheduler-name</i> <a href="#">drop-profile-map</a> ]
<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. Statement introduced in Junos OS Release 12.2 for ACX Series Routers.
<b>Description</b>	Specify a loss priority to which to apply a drop profile. The drop profile map sets the drop profile for a specific PLP and protocol type. The inputs for the map are the PLP designation and the protocol type. The output is the drop profile.
<b>Options</b>	<p><b>any</b>—The drop profile applies to packets with any PLP.</p> <p><b>high</b>—The drop profile applies to packets with high PLP.</p> <p><b>low</b>—The drop profile applies to packets with low PLP.</p> <p><b>medium-high</b>—The drop profile applies to packets with medium-high PLP.</p> <p><b>medium-low</b>—The drop profile applies to packets with medium-low PLP.</p>
<div>  <p><b>NOTE:</b> On ACX Series Routers, if you configure the <i>protocol</i> as <i>tcp</i>, then the loss-priority (any   high   low   medium-high) values are supported. If you configure the <i>protocol</i> with either <i>non-tcp</i> or <i>any</i> option, then irrespective of the loss-priority value, only one drop-profile can be specified.</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>Default Schedulers Overview</i></li> <li>• <i>Determining Packet Drop Behavior by Configuring Drop Profile Maps for Schedulers</i></li> <li>• <i>Configuring Schedulers for Priority Scheduling</i></li> <li>• <i>Configuring and Applying Tricolor Marking Policers</i></li> <li>• <i>protocol (Schedulers)</i></li> </ul>

## max-queues

---

<b>Syntax</b>	<code>max-queues <i>queues-per-line-card</i>;</code>
<b>Hierarchy Level</b>	[edit chassis fpc <i>slot-number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2. Support for MPC7E, MPC8E, and MPC9E from Junos OS Release 15.1F4.
<b>Description</b>	<p>Configure the maximum number of queues allowed per MPC1 Q, MPC2 Q, MPC2 EQ, MPC7E, MPC8E, and MPC9E line cards in an MX240, MX480, MX960, MX2010, or MX2020 router or per 2-port or 4-port 10-Gigabit Ethernet fixed MIC with XFP in an MX5, MX10, MX40, or MX80 modular chassis router.</p> <p>Reducing the number of queues allowed in a hierarchical scheduling environment in turn reduces the degree of jitter in the queues.</p>



**CAUTION:** Committing a change to the maximum number of queues on a line card causes a reset of the line card.

---

<b>Options</b>	<p><b><i>queues-per-line-card</i></b>—Maximum number of queues allowed for the line card. Only the following keywords are valid: <b>8k, 16k, 32k, 64k, 128k, 256k, 512k, or 1M..</b></p> <ul style="list-style-type: none"><li>Built-in 10-Gigabit Ethernet MICs and MPC1 Q line cards support up to 128 K queues. You can use this statement to configure the single Packet Forwarding Engine to support a lower maximum number of queues.</li><li>MPC2 Q and MPC2 EQ line cards support up to 256 K queues. You can use this statement to configure the two Packet Forwarding Engines to support a lower maximum number of queues.</li></ul> <p>If you configure a keyword for a value that exceeds the number of queues supported by the line card hardware, the system uses the maximum number of queues supported by the line card.</p> <p>If the <b>max-queues</b> statement is <i>not</i> configured on MPC7E, MPC8E, and MPC9E, which is the default mode, the MPC starts with a message similar to the following:</p> <p><b>FPC 0 supports only port based queuing. A license is required for per-VLAN and hierarchical features.</b></p> <p>If the <b>max-queues</b> statement is configured on MPC7E, MPC8E, and MPC9E and the value is less than or equal to 32,000, the MPC starts with a message similar to the following:</p>
----------------	--

FPC 0 supports port based queuing and is configured in 16384 queue mode. A limited per-VLAN queuing license is required for per VLAN and hierarchical queuing features.

If the **max-queues** statement is configured on MPC7E, MPC8E, and MPC9E and the value is greater than 32,000, the MPC starts with a message similar to the following:

FPC 0 supports port based queuing and is configured in 524288 queue mode. A full scale per-VLAN queuing license is required for per VLAN and hierarchical queuing features.

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

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Jitter Reduction in Hierarchical CoS Queues on page 48</a></li><li>• <i>MPC1 Q</i></li><li>• <i>MPC2 Q</i></li><li>• <i>MPC2 EQ</i></li><li>• <i>MIC/MPC Compatibility</i></li><li>• <i>10-Gigabit Ethernet MICs with XFP</i></li></ul>
------------------------------	---

## max-queues-per-interface

---

<b>Syntax</b>	max-queues-per-interface (8   4);
<b>Hierarchy Level</b>	[edit chassis fpc <i>slot-number</i> pic <i>pic-number</i> ], [edit chassis lcc <i>number</i> fpc <i>slot-number</i> pic <i>pic-number</i> ] (Routing Matrix)
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Support for TX Matrix and TX Matrix Plus added in Junos OS Release 9.6. On MIC or MPC interfaces on MX Series routers, configure eight egress queues.
<b>Description</b>	On IQ, MPC, and DPC interfaces on M120, T320, T640, T1600, TX Matrix, and TX Matrix Plus routers, or on MIC or MPC interfaces on MX Series routers, configure eight egress queues.
<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 Junos OS to Support Eight Queues on IQ Interfaces for T Series and M320 Routers</i></li><li>• <i>Configuring Up to 16 Custom Forwarding Classes</i></li><li>• <i>Enabling Eight Queues on ATM Interfaces</i></li><li>• <i>Configuring the Maximum Number of Queues for Trio MPC/MIC Interfaces</i></li></ul>

## maximum-lease-time

---

<b>Syntax</b>	<code>maximum-lease-time seconds;</code>
<b>Hierarchy Level</b>	[edit access address-assignment pool <i>pool-name</i> <b>family</b> (inet   inet6) <b>dhcp-attributes</b> ], [edit access protocol-attributes <i>attribute-set-name</i> ]
<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>	<b>seconds</b> —Maximum number of seconds the lease can be held. <b>Range:</b> 30 through 4,294,967,295 seconds <b>Default:</b> 86,400 (24 hours)
<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> <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>

## member-link-scheduler

---


<b>Syntax</b>	member-link-scheduler (replicate   scale);
<b>Hierarchy Level</b>	[edit class-of-service interfaces], [edit logical-systems <i>logical-system-name</i> class-of-service interfaces <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> class-of-service interfaces <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6.
<b>Description</b>	Determines whether scheduler parameters for aggregated interface member links are applied in a replicated or scaled manner.
<b>Default</b>	By default, scheduler parameters are scaled (in “equal division mode”) among aggregated interface member links.
<b>Options</b>	<b>replicate</b> —Scheduler parameters are copied to each level of the aggregated interface member links.  <b>scale</b> —Scheduler parameters are scaled based on number of member links and applied each level of the aggregated interface member links.
<b>Required Privilege Level</b>	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Hierarchical Schedulers for CoS on page 12</a></li></ul>



## mode (Layer 2 Tunneling Protocol Shaping)

<b>Syntax</b>	<code>mode <i>traffic-manager-mode</i>;</code>
<b>Hierarchy Level</b>	[edit chassis fpc <i>slot-number</i> pic <i>pic-number</i> <b>traffic-manager</b> ], [edit chassis lcc <i>number</i> fpc <i>slot-number</i> pic <i>pic-number</i> <b>traffic-manager</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3.
<b>Description</b>	Enable shaping on an L2TP session.
<b>Options</b>	<p><b>traffic-manager-mode</b>—Configure CoS traffic manager mode of operation on this interface. This option has the following suboptions:</p> <p><b>egress-only</b>—Enable CoS queuing and scheduling on the egress side for the PIC that houses the interface. This is the default mode for an Enhanced Queuing (EQ) DPC on MX Series routers.</p> <p>If ingress packet drops are observed at a high rate for an IQ2 or IQ2E PIC, configure the <b>traffic-manager</b> statement to work in the <b>egress-only</b> mode.</p> <p><b>ingress-and-egress</b>—Enable CoS queuing and scheduling on both the egress and ingress sides for the PIC. This is the default mode for IQ2 and IQ2E PICs on M Series and T Series routers.</p> <p>Junos OS does not support <b>ingress-and-egress</b> mode on label-switched interfaces (LSIs) configured with VPLS.</p> <p>For EQ DPCs, you must configure the <b>traffic-manager</b> statement with <b>ingress-and-egress</b> mode to enable ingress CoS on the EQ DPC. EQ DPCs have 250 ms of buffering, with only egress queuing (default mode). When <b>ingress-and-egress</b> is configured, the buffer is partitioned as 50 ms for the ingress direction and 200 ms for the egress direction.</p> <p><b>session-shaping</b>—(M10i and M120 routers only) Configure the IQ2 PIC mode for session-aware traffic shaping to enable L2TP session shaping.</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 CoS for L2TP Tunnels on ATM Interfaces</i></li> <li>• <a href="#">egress-shaping-overhead on page 284</a></li> <li>• <i>ingress-shaping-overhead</i></li> <li>• <a href="#">traffic-manager on page 404</a></li> </ul>


## output-traffic-control-profile

<b>Syntax</b>	<code>output-traffic-control-profile <i>profile-name</i> <b>shared-instance</b> <i>instance-name</i>;</code>
<b>Hierarchy Level</b>	<code>[edit class-of-service interfaces <i>interface-name</i> ],</code> <code>[edit class-of-service interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i>],</code> <code>[edit class-of-service interfaces <i>interface-name</i> <b>interface-set</b> <i>interface-set-name</i>]</code>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.6.</p> <p><b>interface-set</b> option added for Enhanced Queuing DPCs on MX Series routers in Junos OS Release 8.5.</p> <p><b>interface-set</b> option added for MIC and MPC interfaces on MX Series routers in Junos OS Release 10.2.</p> <p>Support on GRE tunnel interfaces configured on physical and logical interfaces on MICs or MPCs in MX Series routers added in Junos OS Release 13.3.</p>
<b>Description</b>	<p>Apply the specified CoS traffic control profile (traffic scheduling and shaping configuration objects) to the output traffic at the physical interface, logical interface, or interface set.</p> <p>The statement is supported on the following interfaces:</p> <ul style="list-style-type: none"> <li>• Channelized IQ PIC interfaces</li> <li>• Gigabit Ethernet IQ, Gigabit Ethernet IQ2, and IQ2E PIC interfaces</li> <li>• Link services IQ (LSQ) interfaces on Multiservices and Services PICs</li> <li>• Enhanced Queuing DPC, MIC, and MPC interfaces on MX Series routers</li> <li>• GRE tunnel interfaces configured on physical or logical interfaces hosted on MIC or MPC line cards in MX Series routers.</li> </ul>
	<div>  <p><b>NOTE:</b> Interface sets (sets of interfaces used to configure hierarchical CoS schedulers on supported Ethernet interfaces) are not supported on GRE tunnel interfaces.</p> </div>
	<p>The <b>shared-instance</b> statement is supported on Gigabit Ethernet IQ2 PICs only.</p>
<b>Options</b>	<p><b><i>profile-name</i></b>—Name of the traffic-control profile to be applied to this interface.</p> <p><b><i>shared-instance</i><i>instance-name</i></b>—Name of the shared scheduler and shaper to be applied to this interface.</p>
<b>Required Privilege Level</b>	<p><b>interface</b>—To view this statement in the configuration.</p> <p><b>interface-control</b>—To add this statement to the configuration.</p>

- Related Documentation**
- *Oversubscribing Interface Bandwidth*
  - *Configuring Traffic Control Profiles for Shared Scheduling and Shaping*
  - [Configuring Hierarchical Schedulers for CoS on page 12](#) (Enhanced Queuing DPC, MIC, and MPC interfaces on MX Series routers)
  - *Configuring Interface Sets* (Enhanced Queuing DPC, MIC, and MPC interfaces on MX Series routers)
  - [output-traffic-control-profile-remaining on page 342](#)
  - [traffic-control-profiles on page 403](#)

## output-traffic-control-profile-remaining

---

<b>Syntax</b>	<code>output-traffic-control-profile-remaining <i>profile-name</i>;</code>
<b>Hierarchy Level</b>	[edit class-of-service interfaces <i>interface-name</i> ], [edit class-of-service interfaces <i>interface-name</i> <b>interface-set</b> <i>interface-set-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5. Support on GRE tunnel interfaces configured on physical interfaces on MICs or MPCs in MX Series routers added in Junos OS Release 13.3.
<b>Description</b>	<p>Apply the specified traffic control profile (traffic scheduling and shaping configuration objects) to the remaining output traffic at the physical interface or interface set. The remaining traffic is transmitted by the default interface or interface set.</p> <p>This statement is supported on the following interfaces:</p> <ul style="list-style-type: none"><li>• IQ2E PIC interfaces on M Series and T Series routers</li><li>• Enhanced Queuing DPC, MICs, and MPC interfaces on MX Series routers</li><li>• GRE tunnel interfaces configured on physical interfaces hosted on MIC or MPC line cards in MX Series routers.</li></ul> <div> <b>NOTE:</b> Interface sets (sets of interfaces used to configure hierarchical CoS schedulers on supported Ethernet interfaces) are not supported on GRE tunnel interfaces.</div> <p>You can map the TCP to the interface or interface set by using the <b>output-traffic-control-profile-remaining</b> statement to explicitly configure the queues of the default interface or interface set scheduler that transmits the remaining traffic.</p>
<b>Options</b>	<b><i>profile-name</i></b> —Name of the traffic-control profile for remaining traffic to be applied to this interface or interface set.
<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 Hierarchical Schedulers for CoS on page 12</a></li><li>• <a href="#">Configuring Remaining Common Queues on MIC and MPC Interfaces</a></li><li>• <a href="#">output-traffic-control-profile on page 340</a></li></ul>

## overhead-accounting

---


<b>Syntax</b>	overhead-accounting { bytes <i>bytes</i> ; cell-mode cell-mode-bytes <i>cell-mode-bytes</i> ; frame-mode frame-mode-bytes <i>frame-mode-bytes</i> ; }
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">traffic-control-profiles</a> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.2.
<b>Description</b>	Configure the mode to shape downstream ATM traffic based on either frames or cells.
<b>Default</b>	The default is <b>frame-mode</b> .
<b>Options</b>	The remaining statements are explained separately. See <a href="#">CLI Explorer</a> .
<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 Static Shaping Parameters to Account for Overhead in Downstream Traffic Rates</i></li> <li>• <i>Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates</i></li> <li>• <a href="#">egress-shaping-overhead on page 284</a></li> </ul>

## pap (Dynamic PPP)

---

<b>Syntax</b>	<code>pap;</code>
<b>Hierarchy Level</b>	[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit "\$junos-interface-unit" <a href="#">ppp-options</a> ], [edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit" <a href="#">ppp-options</a> ]
<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" <a href="#">ppp-options</a> ] 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>

## peak-rate

<b>Syntax</b>	<code>peak-rate rate;</code>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">traffic-control-profiles</a> <i>traffic-control-profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	(MX Series routers) Define ATM peak cell rate on ATM MICs in cells per second by entering a decimal number followed by the abbreviation c; where 1 cps = 384 bps.
<b>Options</b>	<p><i>rate</i>—ATM peak rate in cells per second.</p> <p><b>Range:</b> 61 cps through 353,206 cps.</p> <p><b>Range:</b> (MX Series with MPCs and ATM MICs with SFP) 61 cps through 1,412,829 cps</p>
	<div>  <p><b>NOTE:</b> Beginning with Junos OS Release 14.2, to configure OC12 CBR bandwidth speed per virtual circuit (VC) on an ATM MIC with SFP (MIC-3D-8OC3-2OC12-ATM), specify up to 1,412,829 cps as the ATM peak cell rate.</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><code>show class-of-service traffic-control-profile</code></li> </ul>

## per-unit-scheduler

**Syntax** per-unit-scheduler;

**Hierarchy Level** [edit interfaces *interface-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 13.2 on 16x10GE MPC and MPC3E line cards.  
Statement introduced in Junos OS Release 13.2 on PTX Series Packet Transport Routers.  
Statement introduced in Junos OS Release 13.3 on MPC4E line cards.  
Statement introduced in Junos OS Release 15.1 on MPC6E line cards.

**Description** For Channelized OC3 IQ, Channelized OC12 IQ, Channelized STM1 IQ, Channelized T3 IQ, Channelized E1 IQ, E3 IQ, link services IQ interfaces (lsq-), Gigabit Ethernet IQ, Gigabit Ethernet IQ2 and IQ2-E, and 10-, 40-, and 100-Gigabit Ethernet interfaces (including the 16x10GE MPC), enable the association of scheduler maps with logical interfaces.



**CAUTION:** Turning on per-unit scheduling causes the interface to reinitialize, which means all logical interfaces (units) on the interface are deleted and recreated.



**NOTE:** To enable per-unit scheduling on MX80 and MX104 routers, configure the per-unit-scheduler statement at each member physical interface level of a particular aggregated Ethernet interface as well as at that aggregated Ethernet interface level. On other routing platforms, it is enough if you include this statement at the aggregated Ethernet interface level.



**NOTE:** Per-unit scheduling is not supported on T1 interfaces configured on the Channelized OC12 IQ PIC.



**NOTE:** On Gigabit Ethernet IQ2 and IQ2-E PICs without the per-unit-scheduler statement, the entire PIC supports 4071 VLANs and the user can configure all the VLANs on the same port.

On Gigabit Ethernet IQ2 and IQ2-E PICs with the per-unit-scheduler statement, the entire PIC supports  $1024 - 2 * \text{number of ports}$  (1024 minus two times the number of ports), because each port is allocated two default schedulers.



When including the **per-unit-scheduler** statement, you must also include the **vlan-tagging** statement or the **flexible-vlan-tagging** statement (to apply scheduling to VLANs) or the **encapsulation frame-relay** statement (to apply scheduling to DLCIs) at the **[edit interfaces interface-name]** 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>• <i>Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs</i></li><li>• <i>vlan-tagging</i></li><li>• <a href="#">flexible-vlan-tagging on page 304</a></li><li>• <i>Example: Applying Scheduling and Shaping to VLANs</i></li><li>• <i>Configuring Virtual LAN Queuing and Shaping on PTX Series Routers</i></li></ul>

## pool (Address-Assignment Pools)

**Syntax** `pool pool-name {`  
     `active-drain;`  
     `family family {`  
         `dhcp-attributes {`  
             `[ protocol-specific attributes ]`  
         `}`  
         `excluded-address ip-address;`  
         `excluded-range name low minimum-value high maximum-value;`  
         `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;`  
     `linked-pool-aggregation;`  
`}`

**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. See [CLI Explorer](#).

**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](#)

---

## pppoe-options (Dynamic PPPoE)

---

<b>Syntax</b>	<pre>pppoe-options {     <a href="#">underlying-interface</a> <i>interface-name</i>;     <a href="#">server</a>; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">dynamic-profiles</a> <i>profile-name</i> <a href="#">interfaces</a> pp0 <a href="#">unit</a> "\$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. See <a href="#">CLI Explorer</a>.</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>

## ppp-options (Dynamic PPP)

<b>Syntax</b>	<pre>ppp-options {   aaa-options <i>aaa-options-name</i>;   authentication [ <i>authentication-protocols</i> ];   chap {     challenge-length minimum <i>minimum-length</i> maximum <i>maximum-length</i>;   }   ignore-magic-number-mismatch;   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 <b>unit</b> "\$junos-interface-unit"], [edit dynamic-profiles <i>profile-name</i> interfaces "\$junos-interface-ifd-name" <b>unit</b> "\$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" <b>unit</b> "\$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. See <a href="#">CLI Explorer</a> .



### NOTE:

PPP options can also be configured in a group profile with the `ppp-options` (L2TP) statement. The following behavior determines the interaction between the PPP options configured in a group profile and the PPP options configured in a dynamic profile:

- When PPP options are configured only in the group profile, the group profile options are applied to the subscriber.
- When PPP options are configured in both a group profile and a dynamic profile, the dynamic profile configuration takes complete precedence over the group profile when the dynamic profile includes one or more of the PPP options that can be configured in the group profile. Complete precedence means that there is no merging of options between the profiles. The group profile is applied to the subscriber only when the dynamic profile does not include any PPP option available in the group profile.


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

**Related Documentation**

- *Dynamic Profiles Overview*
- *Configuring a PPPoE Dynamic Profile*
- *Configuring Dynamic Authentication for PPP Subscribers*
- *Attaching Dynamic Profiles to Static PPP Subscriber Interfaces*
- *Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface*

## priority (Schedulers)

---

<b>Syntax</b>	<code>priority <i>priority-level</i>;</code>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">schedulers</a> <i>scheduler-name</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. Statement introduced in Junos OS Release 12.2 for ACX Series Routers.
<b>Description</b>	Specify the packet-scheduling priority value.
<b>Options</b>	<p><i>priority-level</i> can be one of the following:</p> <ul style="list-style-type: none"><li>• <b>low</b>—Scheduler has low priority.</li><li>• <b>medium-low</b>—Scheduler has medium-low priority.</li><li>• <b>medium-high</b>—Scheduler has medium-high priority.</li><li>• <b>high</b>—Scheduler has high priority. Assigning high priority to a queue prevents the queue from being underserved.</li><li>• <b>strict-high</b>—Scheduler has strictly high priority. Configure a <b>high</b> priority queue with unlimited transmission bandwidth available to it. As long as it has traffic to send, the <b>strict-high</b> priority queue receives precedence over <b>low</b>, <b>medium-low</b>, and <b>medium-high</b> priority queues, but not <b>high</b> priority queues. You can configure <b>strict-high</b> priority on only one queue per interface.</li></ul>
	<div> <b>NOTE:</b> The <b>strict-high</b> priority level is the only priority level supported on ACX Series Routers. However, multiple strict-high priority queues can be configured per interface on ACX Series Routers.</div>
<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 Schedulers for Priority Scheduling</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;
        }
        accounting-order (radius | [accounting-order-data-list]);
        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;
                service-profile profile-name(parameter)&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;  
local {  
    flat-file-profile profile-name;  
}  
preauthentication-order preauthentication-method;  
provisioning-order (gx-plus | jsr | pcrf);  
radius {  
    accounting-server [ ip-address ];  
    attributes {  
        exclude {  
            attribute-name packet-type;  
            standard-attribute number {  
                packet-type [ access-request | accounting-off | accounting-on | accounting-start  
                    | accounting-stop ];  
            }  
            vendor-id id-number {  
                vendor-attribute vsa-number {  
                    packet-type [ access-request | accounting-off | accounting-on | accounting-start  
                        | accounting-stop ];  
                }  
            }  
        }  
    }  
    ignore {  
        dynamic-iflset-name;  
        framed-ip-netmask;  
        idle-timeout;  
        input-filter;  
        logical-system:routing-instance;  
        output-filter;  
        session-timeout;  
        standard-attribute number;  
        vendor-id id-number {  
            vendor-attribute vsa-number;  
        }  
    }  
}  
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-channel;
    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 {
        statistics (time | volume-time);
        update-interval minutes;
    }
    accounting-order (activation-protocol | local | 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** *profile-name*—Name of the profile.

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.

The remaining statements are explained separately. See [CLI Explorer](#).

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

**Related Documentation**

- *Configuring the PPP Authentication Protocol*
- *Configuring Access Profiles for L2TP or PPP Parameters*
- *Configuring L2TP Properties for a Client-Specific Profile*
- *Configuring an L2TP Access Profile on the LNS*
- *Configuring an L2TP LNS with Inline Service Interfaces*
- *Configuring PPP Properties for a Client-Specific Profile*
- *Configuring Service Accounting with JSRC*
- *Configuring Service Accounting in Local Flat Files*
- *AAA Service Framework Overview*

## 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.


**Required Privilege Level**

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

**Related Documentation**

- *Example: Configuring Multiple Routing Instances of OSPF*

## 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 on Switches (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>

## qos-adjust-hierarchical

---

<b>Syntax</b>	qos-adjust-hierarchical [interface-set];
<b>Hierarchy Level</b>	[edit dynamic-profiles <i>dynamic-profile-name</i> access-cac interface <i>\$junos-interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 17.2R1.
<b>Description</b>	Specify that multicast QoS adjustments received on the subscriber contribute to hierarchical QoS adjustments.
<b>Options</b>	<b>interface-set</b> —(Optional) Specify that the immediate parent interface-set receive QoS adjustments.



**NOTE:** For any multicast-enabled interface, only directly stacked interface-set parents are adjusted. The following scenario is not supported: multicast-enabled client interface is stacked over a VLAN interface and the VLAN interface in turn is stacked over and interface-set.

---

<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 the Minimum Adjusted Shaping Rate on Scheduler Nodes for Subscribers on page 145</a></li></ul>



## queue (Global Queues)

<b>Syntax</b>	<code>queue <i>queue-number</i> <i>class-name</i>;</code>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">forwarding-classes</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 14.2 for PTX Series Packet Transport Routers.
<b>Description</b>	<p>Specify the output transmission queue to which to map all input from an associated forwarding class.</p> <p>On M120, M320, MX Series, T Series routers and on EX Series switches, this statement enables you to configure up to eight forwarding classes with one-to-one mapping to output queues. If you want to configure up to 16 forwarding classes with multiple forwarding classes mapped to single output queues, include the <b>class</b> statement instead of the <b>queue</b> statement at the [edit class-of-service forwarding-classes] hierarchy level.</p>
<b>Options</b>	<p><b><i>class-name</i></b>—Name of forwarding class.</p> <p><b><i>queue-number</i></b>—Output queue number.</p> <p><b>Range:</b> 0 through 7. For M Series routers and some T Series router PICs, 0 through 3.</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 Custom Forwarding Class for Each Queue</i></li> <li>• <i>class (Forwarding Classes)</i></li> </ul>

## radius (Access Profile)

---

```
Syntax  radius {
    accounting-server [ ip-address ];
    attributes {
        exclude
            attribute-name packet-type;
        standard-attribute number {
            packet-type [ access-request | accounting-off | accounting-on | accounting-start |
                accounting-stop ];
        }
        vendor-id id-number {
            vendor-attribute vsa-number {
                packet-type [ access-request | accounting-off | accounting-on | accounting-start
                    | accounting-stop ];
            }
        }
    }
    ignore {
        dynamic-iflset-name;
        framed-ip-netmask;
        idle-timeout;
        input-filter;
        logical-system-routing-instance;
        output-filter;
        session-timeout;
        standard-attribute number;
        vendor-id id-number {
            vendor-attribute vsa-number;
        }
    }
}

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-channel;
        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. See [CLI Explorer](#).

**Required Privilege** admin—To view this statement in the configuration.  
**Level** 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 port-number;     accounting-retry number;     accounting-timeout seconds;     dynamic-request-port     max-outstanding-requests value;     port port-number;     preauthentication-port port-number;     preauthentication-secret password;     retry attempts;     routing-instance routing-instance-name;     secret password;     source-address source-address;     timeout seconds; }</pre>
<b>Hierarchy Level</b>	<p>[edit access],</p> <p>[edit access <b>profile</b> profile-name]</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. See <a href="#">CLI Explorer</a>.</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>• <i>show network-access aaa statistics</i></li> </ul>

- *clear network-access aaa statistics*

## 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</b> <i>upper-limit</i>—Upper limit of an address range or IPv6 prefix range.</p> <p><b>low</b> <i>lower-limit</i>—Lower limit of an address range or IPv6 prefix range.</p> <p><b>prefix-length</b> <i>prefix-length</i>—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>	<p><b>admin</b>—To view this statement in the configuration.</p> <p><b>admin-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring a Named Address Range for Dynamic Address Assignment</i></li><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>• <i>Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs</i></li></ul>

## rewrite-rules (Definition)

---

<b>Syntax</b>	<pre>rewrite-rules {     type <i>rewrite-name</i> {         import (<i>rewrite-name</i>   default);         forwarding-class <i>class-name</i> {             loss-priority <i>level</i> code-point [ <i>aliases</i> ] [ <i>bit-patterns</i> ];         }     } }</pre>
<b>Hierarchy Level</b>	[edit class-of-service]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>ieee-802.1ad</b> option introduced in Junos OS Release 9.2.
<b>Description</b>	Specify a rewrite-rules mapping for the traffic that passes through all queues on the interface.
<b>Options</b>	<p><b><i>rewrite-name</i></b>—Name of a <b>rewrite-rules</b> mapping.</p> <p><b><i>type</i></b>—Traffic type.</p> <p><b>Values:</b> <b>dscp</b>, <b>dscp-ipv6</b>, <b>exp</b>, <b>ieee-802.1</b>, <b>ieee-802.1ad</b>, <b>inet-precedence</b>, <b>inet6-precedence</b></p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>
<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>Configuring Rewrite Rules</i></li></ul>



## rewrite-rules (Interfaces)

**Syntax** `rewrite-rules {  
     dscp (rewrite-name | default) protocol (inet-both | inet-outer | mpls);  
     dscp-ipv6 (rewrite-name | default) protocol mpls;  
     exp (rewrite-name | default) protocol protocol-types;  
     exp-push-push-push default;  
     exp-swap-push-push default;  
     ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);  
     ieee-802.1ad (rewrite-name | default) vlan-tag (outer | outer-and-inner);  
     inet-precedence (rewrite-name | default) protocol (inet-both | inet-outer | mpls);  
 }`

**Hierarchy Level** `[edit class-of-service interfaces interface-name],`  
`[edit class-of-service interfaces interface-name unit logical-unit-number]`

**Release Information** Statement introduced before Junos OS Release 7.4.

**Description** Associate a rewrite-rules configuration or default mapping with a specific interface.

The `[edit class-of-service interfaces interface-name]` hierarchy level is not supported on M Series routers.

The `[edit class-of-service interfaces interface-name unit logical-unit-number]` hierarchy level is not supported on ACX Series routers.

Integrated Bridging and Routing (IRB) interfaces are used to tie together Layer 2 switched and Layer 3 routed domains on MX routers. MX routers support classifiers and rewrite rules on the IRB interface at the `[edit class-of-service interfaces irb unit logical-unit-number]` level of the hierarchy. All types of classifiers and rewrite rules are allowed, including IEEE 802.1p.



**NOTE:** The IRB classifiers and rewrite rules are used only for *routed* packets; in other words, it is for traffic that originated in the Layer 2 domain and is then routed through IRB into the Layer 3 domain, or vice versa. Only IEEE classifiers and IEEE rewrite rules are allowed for pure Layer 2 interfaces within a bridge domain.

On an MX Series router and on an EX Series switch, **exp-push-push-push**, **exp-swap-push-push**, and **frame-relay-de** are not supported on an integrated routing and bridging (IRB) interface.

On an ACX Series router, only the outer tag is supported for **dscp**, **inet-precedence**, and **ieee802.1**.

On M Series routers only, if you include the **control-word** statement at the `[edit protocols l2circuit neighbor address interface interface-name]` hierarchy level, the software cannot rewrite MPLS EXP bits.

For IQ PICs, you can configure only one IEEE 802.1 rewrite rule on a physical port. All logical ports (units) on that physical port should apply the same IEEE 802.1 rewrite rule.

On M320 and T Series routers (except for T4000 routers with Type 5 FPCs), for a single interface, you cannot enable a rewrite rule on a subset of forwarding classes. You must assign a rewrite rule to either none of the forwarding classes or all of the forwarding classes. When you assign a rewrite rule to a subset of forwarding classes, the commit does not fail, and the subset of forwarding classes works as expected. However, the forwarding classes to which the rewrite rule is not assigned are rewritten to all zeros.

For example, if you configure a Differentiated Services code point (DSCP) rewrite rule, the bits in the forwarding classes to which you do not assign the rewrite rule are rewritten to 000000. If you configure an IP precedence rewrite rule, the bits in the forwarding classes to which you do not assign the rewrite rule are rewritten to 000.

**Options**    ***rewrite-name***—Name of a **rewrite-rules** mapping configured at the **[edit class-of-service rewrite-rules]** hierarchy level.

**default**—The default mapping.

The remaining statements are explained separately. See [CLI Explorer](#).

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

**Related Documentation**

- *Configuring Rewrite Rules*
- [rewrite-rules \(Definition\) on page 370](#)
- *Applying Rewrite Rules to Output Logical Interfaces*

## routing-instances (Multiple Routing Entities)

<b>Syntax</b>	<code>routing-instances <i>routing-instance-name</i> { ... }</code>
<b>Hierarchy Level</b>	<code>[edit],</code> <code>[edit logical-systems <i>logical-system-name</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>remote-vtep-v6-list</b> statement introduced in Junos OS Release 17.3 for MX Series routers with MPC and MIC interfaces.
<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 <code>my-instance.inet.0</code>. All routes for <b>my-instance</b> are installed into <code>my-instance.inet.0</code>.</p> <p>Routes are installed into the default routing instance <code>inet.0</code> 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>	<p><b><i>routing-instance-name</i></b>—Name of the routing instance. This must be a non-reserved string of not more than 128 characters.</p> <p><b>remote-vtep-list</b>—Configure static remote VXLAN tunnel endpoints.</p> <p><b>remote-vtep-v6-list</b>—Configure static IPv6 remote VXLAN tunnel endpoints.</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>Example: Configuring Interprovider Layer 3 VPN Option A</i></li><li>• <i>Example: Configuring Interprovider Layer 3 VPN Option B</i></li><li>• <i>Example: Configuring Interprovider Layer 3 VPN Option C</i></li></ul>

---

## scheduler (Scheduler Map)

---

<b>Syntax</b>	<code>scheduler <i>scheduler-name</i>;</code>
<b>Hierarchy Level</b>	[edit class-of-service scheduler-maps <i>map-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 Routers.
<b>Description</b>	Associate a scheduler with a scheduler map.
<b>Options</b>	<i>scheduler-name</i> —Name of the scheduler configuration block.
<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 Schedulers</i></li></ul>

## scheduler-map (Interfaces and Traffic-Control Profiles)

<b>Syntax</b>	<code>scheduler-map <i>map-name</i>;</code>
<b>Hierarchy Level</b>	<p>[edit class-of-service interfaces <i>interface-name</i>],          [edit class-of-service interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i>],          [edit class-of-service <b>traffic-control-profiles</b>]</p>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>For Gigabit Ethernet IQ, Channelized IQ PICs, and FRF.15 and FRF.16 LSQ interfaces only, associate a scheduler map name with an interface or with a traffic-control profile.</p> <p>For channelized OC12 intelligent queuing (IQ), channelized T3 IQ, channelized E1 IQ, and Gigabit Ethernet IQ interfaces only, you can associate a scheduler map name with a logical interface.</p>
<b>Options</b>	<b><i>map-name</i></b> —Name of the scheduler map.
<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 Schedulers</i></li> <li>• <i>Oversubscribing Interface Bandwidth</i></li> <li>• <a href="#">output-traffic-control-profile on page 340</a></li> </ul>

## scheduler-maps (For Most Interface Types)

---

**Syntax** scheduler-maps {  
    *map-name* {  
        forwarding-class *class-name* scheduler *scheduler-name*;  
    }  
}

**Hierarchy Level** [edit class-of-service]

**Release Information** Statement introduced before Junos OS Release 7.4.

**Description** Specify a scheduler map name and associate it with the scheduler configuration and forwarding class.

**Options** *map-name*—Name of the scheduler map.

The remaining statements are explained separately. See [CLI Explorer](#).

See *Configuring Schedulers* .

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

## schedulers (CoS)

<b>Syntax</b>	<pre> schedulers {   scheduler-name {     adjust-minimum rate;     adjust-percent percentage;     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   medium-low   medium-high   high   none];     excess-rate (percent percentage   proportion value);     priority priority-level;     shaping-rate (percent percentage   rate);     transmit-rate (percent percentage   rate   remainder) &lt;exact   rate-limit&gt;;   } } </pre>
<b>Hierarchy Level</b>	[edit class-of-service]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1X48 for PTX Series routers.</p>
<b>Description</b>	Specify the scheduler name and parameter values.
<b>Options</b>	<p><b>scheduler-name</b>—Name of the scheduler to be configured.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</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>How Schedulers Define Output Queue Properties</i></li> <li>• <i>Default Schedulers Overview</i></li> <li>• <i>Configuring Schedulers</i></li> <li>• <i>Configuring a Scheduler</i></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>	server;
<b>Hierarchy Level</b>	[edit <a href="#">dynamic-profiles</a> <i>profile-name</i> <a href="#">interfaces</a> pp0 <a href="#">unit</a> "\$junos-interface-unit" <a href="#">pppoe-options</a> ]
<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

---

Syntax	<pre>server-group {     server-group-name {         server-ip-address;     } }</pre>
Hierarchy Level	<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>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Support at the <b>[edit ... dhcpv6]</b> hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Specify the 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>
Options	<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>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">dhcp-relay on page 252</a></li><li>• <i>Extended DHCP Relay Agent Overview</i></li><li>• <i>Configuring Named Server Groups</i></li></ul>

## shaping-rate (Applying to an Interface)

<b>Syntax</b>	<code>shaping-rate rate;</code>
<b>Hierarchy Level</b>	[edit class-of-service interfaces <i>interface-name</i> ], [edit class-of-service interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. [edit class-of-service interfaces <i>interface-name</i> ] hierarchy level added in Junos OS Release 7.5. Statement introduced in Junos OS Release 13.2 on PTX Series Packet Transport Routers. Statement introduced in Junos OS Release 17.3 on PTX10008 Routers.
<b>Description</b>	For logical interfaces on which you configure packet scheduling, configure traffic shaping by specifying the amount of bandwidth to be allocated to the logical interface. Applying a shaping rate can help ensure that higher-priority services do not starve lower-priority services.  For physical interfaces, configure traffic shaping based on the rate-limited bandwidth of the total interface bandwidth.  Logical and physical interface traffic shaping rates are mutually exclusive. This means you can include the <b>shaping-rate</b> statement at the [edit class-of-service interfaces <i>interface-name</i> ] hierarchy level or at the [edit class-of-service interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ] hierarchy level, but not at both.



**NOTE:** For MX Series routers and for EX Series switches, the shaping rate value for the physical interface at the [edit class-of-service interfaces *interface-name*] hierarchy level must be a minimum of 160 Kbps. If the value is less than the sum of the logical interface guaranteed rates, you cannot apply the shaping rate to a physical interface.

For PTX Series routers, the shaping rate value for the physical interface at the [edit class-of-service interfaces *interface-name*] hierarchy level must be a minimum of 1 Gbps and an incremental granularity of 0.1 percent of the physical interface speed after that (for example, 10 Mbps increments on a 10 Gbps interface).

For T4000 routers with Type 5 FPCs, the shaping rate value for the physical interface must be a minimum of 292 Kbps. The maximum value of **shaping-rate** is limited by the maximum transmission rate of the interface.

Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the **shaping-rate** statement at the [edit class-of-service **traffic-control-profiles**] hierarchy level. With this configuration approach, you can

independently control the delay-buffer rate, as described in *Oversubscribing Interface Bandwidth*.

For FRF.15 and FRF.16 bundles on link services interfaces, only shaping rates based on percentage are supported.

**Default** If you do not include this statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level, the default logical interface bandwidth is the average of unused bandwidth for the number of logical interfaces that require default bandwidth treatment. If you do not include this statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level, the default physical interface bandwidth is the average of unused bandwidth for the number of physical interfaces that require default bandwidth treatment.

**Options** **rate**—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).  
**Range:** 1000 through 6,400,000,000,000 bps.



**NOTE:** Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Beginning with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.



**NOTE:** For all MX Series and EX Series interfaces, the rate can be from 65,535 to 6,400,000,000,000 bps.

For all PTX Series interfaces, the rate can be from 1,000,000,000 to 160,000,000,000 bps in increments of 0.1 percent of the interface speed.

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

**Related Documentation**

- *Applying Scheduler Maps Overview*
- *Configuring Virtual LAN Queuing and Shaping on PTX Series Routers*

## shaping-rate (Schedulers)

**Syntax** `shaping-rate (percent percentage | rate) <burst-size bytes>;`

**Hierarchy Level** `[edit class-of-service schedulers scheduler-name]`

**Release Information** Statement introduced before Junos OS Release 7.4.  
The **burst-size** option added for MIC and MPC interfaces on MX Series routers in Junos OS Release 11.4.  
Statement introduced in Junos OS Release 12.2 for ACX Series Routers.

**Description** Define a limit on excess bandwidth usage for a forwarding class/queue.

The **transmit-rate** statement at the `[edit class-of-service schedulers scheduler-name]` hierarchy level configures the minimum bandwidth allocated to a queue. The transmission bandwidth can be configured as an exact value or allowed to exceed the configured rate if additional bandwidth is available from other queues.

Configure the shaping rate as an absolute maximum usage and not the additional usage beyond the configured transmit rate.

**Default** If you do not include this statement, the default shaping rate is 100 percent, which is the same as no shaping at all.

**Options** **percent *percentage***—Shaping rate as a percentage of the available interface bandwidth.  
**Range:** 0 through 100 percent



**NOTE:** If you configure a shaping rate as a percent in a scheduler, the effective shaping rate is calculated based on the following hierarchy:

1. Logical interface shaping rate, if configured
2. Physical interface shaping rate, if configured
3. Physical interface bandwidth

With SRX300, SRX320, SRX340, SRX345, SRX550m, SRX1500, and vSRX2.0 devices, you can configure both logical interface shaping rates and physical interface shaping rates on the same physical interface. On all other models, you can only configure one or the other on a particular physical interface.

**rate**—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

**Range:** 3200 through 6,400,000,000,000 bps



**NOTE:** Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Beginning with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.

**burst-size bytes**—Maximum burst size, in bytes. The burst value determines the number of rate credits that can accrue when the queue or scheduler node is held in the inactive round robin.

**Range:** 0 through 1,000,000,000

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

**Related Documentation**    • *Applying Scheduler Maps Overview*

## shaping-rate (Oversubscribing an Interface)

<b>Syntax</b>	<code>shaping-rate (percent <i>percentage</i>   <i>rate</i>) &lt;burst-size <i>bytes</i>&gt;;</code>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">traffic-control-profiles</a> <i>profile-name</i> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Option <b>burst-size</b> introduced for Enhanced Queuing (EQ) DPC interfaces on MX Series routers in Junos OS Release 9.4.</p> <p>Option <b>burst-size</b> option introduced for MIC and MPC interfaces on MX Series routers in Junos OS Release 11.4.</p> <p>Option <b>burst-size</b> introduced for IQ2 and IQ2E interfaces in Junos OS Release 12.3.</p> <p>Statement introduced for PTX Series Packet Transport Routers in Junos OS Release 16.1. PTX Series Packet Transport Routers do not support the <b>burst-size</b> option or defining the <b>shaping-rate</b> as a percentage.</p>
<b>Description</b>	<p>For Gigabit Ethernet IQ, Channelized IQ PIC, FRF.15 and FRF.16 LSQ interfaces, and for EQ DPC, MIC, and MPC interfaces on MX Series routers, configure a shaping rate for a logical interface. You can also configure an optional burst size for a logical interface on EQ DPC interfaces and on IQ2 and IQ2E PIC interfaces. This can help to ensure that higher-priority services do not starve lower-priority services.</p> <p>For physical interfaces on T4000 router interfaces on Type 5 FPCs and on PTX Series routers, configure traffic shaping rate.</p> <p>The sum of the shaping rates for all logical interfaces on the physical interface can exceed the physical interface bandwidth. This practice is known as oversubscription of the peak information rate (PIR).</p>
<b>Default</b>	The default behavior depends on various factors. For more information, see <i>Oversubscribing Interface Bandwidth</i> .
<b>Options</b>	<p><b>percent <i>percentage</i></b>—For LSQ interfaces, shaping rate as a percentage of the available interface bandwidth.</p> <p><b>Range:</b> 1 through 100 percent</p> <p><b><i>rate</i></b>—For IQ and IQ2 interfaces, and T4000 routers with Type 5 FPCs, peak shaping rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b></p> <ul style="list-style-type: none"> <li>IQ and IQ2 interfaces—1000 through 6,400,000,000,000 bps</li> </ul>



**NOTE:** Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Beginning with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.

- T4000 routers with Type 5 FPCs—The shaping rate value for the physical interface must be a minimum of 292 Kbps. The maximum value of **shaping-rate** is limited by the maximum transmission rate of the interface.

**burst-size bytes**—(Optional) Maximum burst size, in bytes.


**Range:** 0 through 1,000,000,000

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

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Traffic Control Profiles for Shared Scheduling and Shaping</i></li><li>• <i>Oversubscribing Interface Bandwidth</i></li><li>• <a href="#">output-traffic-control-profile on page 340</a></li></ul>
------------------------------	---




## shaping-rate-excess-high


<b>Syntax</b>	<code>shaping-rate-excess-high <i>rate</i> [ <i>burst-size bytes</i> ];</code>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">traffic-control-profiles</a> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	For MIC and MPC interfaces on MX Series routers, configure a shaping rate and optional burst size for high-priority excess traffic. This can help to make sure higher priority services do not starve lower priority services.
<b>Default</b>	If you do not include this statement, the default shaping rate for this priority is determined by the <b>shaping-rate</b> statement in the traffic control profile.
<b>Options</b>	<p><b>rate</b>—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 1000 through 6,400,000,000,000</p>
<div>  <p><b>NOTE:</b> Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Beginning with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.</p> </div>	
<b>Default:</b>	None
	<p><b>burst-size <i>bytes</i></b>—Maximum burst size, in bytes.</p> <p><b>Range:</b> 0 through 1,000,000,000</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>Per-Priority Shaping on MIC and MPC Interfaces Overview</i></li> <li>• <i>Oversubscribing Interface Bandwidth</i></li> <li>• <i>Configuring Traffic Control Profiles for Shared Scheduling and Shaping</i></li> <li>• <a href="#">shaping-rate-excess-low on page 388</a></li> <li>• <a href="#">shaping-rate-priority-high on page 389</a></li> <li>• <a href="#">shaping-rate-priority-low on page 390</a></li> <li>• <a href="#">shaping-rate-priority-medium on page 391</a></li> </ul>

## shaping-rate-excess-low

---


<b>Syntax</b>	<code>shaping-rate-excess-low rate [ burst-size bytes ];</code>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">traffic-control-profiles</a> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	For MIC and MPC interfaces on MX Series routers, configure a shaping rate and optional burst size for low-priority excess traffic. This can help to make sure higher priority services do not starve lower priority services.
<b>Default</b>	If you do not include this statement, the default shaping rate for this priority is determined by the <b>shaping-rate</b> statement in the traffic control profile.
<b>Options</b>	<p><b>rate</b>—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 1000 through 6,400,000,000,000</p> <hr/>
	<div> <b>NOTE:</b> Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Beginning with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.</div> <hr/>
	<p><b>Default:</b> None</p> <p><b>burst-size bytes</b>—Maximum burst size, in bytes.</p> <p><b>Range:</b> 0 through 1,000,000,000</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>Per-Priority Shaping on MIC and MPC Interfaces Overview</i></li><li>• <i>Oversubscribing Interface Bandwidth</i></li><li>• <i>Configuring Traffic Control Profiles for Shared Scheduling and Shaping</i></li><li>• <a href="#">shaping-rate-excess-high on page 387</a></li><li>• <a href="#">shaping-rate-priority-high on page 389</a></li><li>• <a href="#">shaping-rate-priority-low on page 390</a></li><li>• <a href="#">shaping-rate-priority-medium on page 391</a></li></ul>

## shaping-rate-priority-high


<b>Syntax</b>	<code>shaping-rate-priority-high rate [ burst-size bytes ];</code>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">traffic-control-profiles</a> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	For MIC and MPC interfaces on MX Series routers, configure a shaping rate and optional burst size for high priority traffic. This can help to make sure higher priority services do not starve lower priority services.
<b>Default</b>	If you do not include this statement, the default shaping rate for this priority is determined by the <b>shaping-rate</b> statement in the traffic control profile.
<b>Options</b>	<p><b>rate</b>—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 1000 through 6,400,000,000,000</p>
	<p> <b>NOTE:</b> Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Beginning with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.</p>
	<p><b>Default:</b> None</p> <p><b>burst-size bytes</b>—Maximum burst size, in bytes.</p> <p><b>Range:</b> 0 through 1,000,000,000</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>Per-Priority Shaping on MIC and MPC Interfaces Overview</i></li> <li>• <i>Oversubscribing Interface Bandwidth</i></li> <li>• <i>Configuring Traffic Control Profiles for Shared Scheduling and Shaping</i></li> <li>• <a href="#">shaping-rate-excess-high on page 387</a></li> <li>• <a href="#">shaping-rate-excess-low on page 388</a></li> <li>• <a href="#">shaping-rate-priority-low on page 390</a></li> <li>• <a href="#">shaping-rate-priority-medium on page 391</a></li> </ul>

## shaping-rate-priority-low

---


<b>Syntax</b>	<code>shaping-rate-priority-low rate [ burst-size bytes ];</code>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">traffic-control-profiles</a> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	For MIC and MPC interfaces on MX Series routers, configure a shaping rate and optional burst size for low priority traffic. This can help to make sure higher priority services do not starve lower priority services.
<b>Default</b>	If you do not include this statement, the default shaping rate for this priority is determined by the <b>shaping-rate</b> statement in the traffic control profile.
<b>Options</b>	<p><b>rate</b>—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 1000 through 6,400,000,000,000</p> <hr/>
	<div> <b>NOTE:</b> Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Beginning with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.</div> <hr/>
	<p><b>Default:</b> None</p> <p><b>burst-size bytes</b>—Maximum burst size, in bytes.</p> <p><b>Range:</b> 0 through 1,000,000,000</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>Per-Priority Shaping on MIC and MPC Interfaces Overview</i></li><li>• <i>Oversubscribing Interface Bandwidth</i></li><li>• <i>Configuring Traffic Control Profiles for Shared Scheduling and Shaping</i></li><li>• <a href="#">shaping-rate-excess-high on page 387</a></li><li>• <a href="#">shaping-rate-excess-low on page 388</a></li><li>• <a href="#">shaping-rate-priority-high on page 389</a></li><li>• <a href="#">shaping-rate-priority-medium on page 391</a></li></ul>

## shaping-rate-priority-medium

<b>Syntax</b>	<code>shaping-rate-priority-medium rate [ burst-size bytes ];</code>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">traffic-control-profiles</a> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	For MIC and MPC interfaces on MX Series routers, configure a shaping rate and optional burst size for medium priority traffic. This can help to make sure higher priority services do not starve lower priority services.
<b>Default</b>	If you do not include this statement, the default shaping rate for this priority is determined by the <b>shaping-rate</b> statement in the traffic control profile.
<b>Options</b>	<p><b>rate</b>—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 1000 through 6,400,000,000,000</p>
<div>  <p><b>NOTE:</b> Through Junos OS Release 13.3, the upper limit is 160,000,000,000 bps. Beginning with Junos OS Release 14.1, the upper limit is 6,400,000,000,000 bps.</p> </div>	
<b>Default:</b>	None
<b>burst-size bytes</b>	Maximum burst size, in bytes.
<b>Range:</b>	0 through 1,000,000,000
<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>Per-Priority Shaping on MIC and MPC Interfaces Overview</i></li> <li>• <i>Oversubscribing Interface Bandwidth</i></li> <li>• <i>Configuring Traffic Control Profiles for Shared Scheduling and Shaping</i></li> <li>• <a href="#">shaping-rate-excess-high on page 387</a></li> <li>• <a href="#">shaping-rate-excess-low on page 388</a></li> <li>• <a href="#">shaping-rate-priority-high on page 389</a></li> <li>• <a href="#">shaping-rate-priority-low on page 390</a></li> </ul>

## shared-bandwidth-policer (Configuring)

---

<b>Syntax</b>	shared-bandwidth-policer;
<b>Hierarchy Level</b>	[edit firewall policer <i>policer-name</i> ], [edit firewall <b>three-color-policer</b> <i>policer-name</i> ], [edit firewall hierarchical-policer <i>policer-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2. Support for MX Series MPC and MIC interfaces added in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Policer instances share bandwidth. This enables configuration of interface-specific policers applied on an aggregated Ethernet bundle or an aggregated SONET bundle to match the effective bandwidth and burst-size to user-configured values. This feature is supported on the following platforms: T Series routers, M120, M10i, M7i (CFEB-E only), M320 (SFPC only), MX240, MX480, and MX960 with DPC, MIC, and MPC interfaces and EX Series switches.
<div> <b>NOTE:</b> This statement is not supported on T4000 Type 5 FPCs.</div>	
<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>Policer Support for Aggregated Ethernet Interfaces Overview</i></li></ul>

## shared-instance

---

<b>Syntax</b>	<code>shared-instance <i>instance-name</i>;</code>
<b>Hierarchy Level</b>	[edit class-of-service interfaces <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">input-traffic-control-profile</a> ], [edit class-of-service interfaces <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">output-traffic-control-profile</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	For Gigabit Ethernet IQ2 and IQ2E PICs only, apply a shared traffic scheduling and shaping profile to the logical interface.
<b>Options</b>	<i>instance-name</i> —Name of the shared scheduler and shaper to be applied to this interface
<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 Shaping on 10-Gigabit Ethernet IQ2 PICs</a></li> <li>• <a href="#">traffic-control-profiles on page 403</a></li> </ul>

## shared-scheduler

---

<b>Syntax</b>	shared-scheduler;
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	<p>For Gigabit Ethernet IQ2 PICs only, enable shared schedulers and shapers on this interface. This statement and the <b>per-unit-scheduler</b> statement are mutually exclusive. Even so, you can configure one logical interface for each shared instance. This effectively provides the functionality of per-unit scheduling.</p> <p>For Gigabit Ethernet IQ2 and Ethernet Enhanced IQ2 (IQ2E) PICs on M320 routers, enable shared schedulers on aggregated Ethernet interfaces in link protection mode.</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 Traffic Control Profiles for Shared Scheduling and Shaping</i></li><li>• <i>Configuring Shaping on 10-Gigabit Ethernet IQ2 PICs</i></li><li>• <i>Configuring Shared Scheduling on Aggregated Ethernet Interfaces</i></li><li>• <a href="#">traffic-control-profiles on page 403</a></li></ul>




## simple-filter (Applying to an Interface)

---

<b>Syntax</b>	simple-filter { input <i>filter-name</i> ; }
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	Apply a simple filter to an interface. You can apply simple filters to the family <b>inet</b> only, and only in the input direction.
<b>Options</b>	<b>input <i>filter-name</i></b> —Name of one filter to evaluate when packets are received on the interface.
<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 Multifield Classifiers</i></li> <li>• <i>filter (Applying to an Interface)</i></li> </ul>

## simple-filter

<b>Syntax</b>	<pre> simple-filter <i>filter-name</i> {     term <i>term-name</i> {         from {             <i>match-conditions</i>;         }         then {             forwarding-class <i>class-name</i>;             loss-priority (high   low   medium);         }     } } </pre>
<b>Hierarchy Level</b>	[edit firewall family inet], [edit logical-systems <i>logical-system-name</i> firewall family inet]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Logical systems support introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Configure simple filters.
<b>Options</b>	<p><b><i>filter-name</i></b>—Name that identifies the simple filter. The name must be a non-reserved string of not more than 64 characters. No special characters are restricted. To include spaces in the name, enclose them in quotation marks (" ").</p> <p><b><i>from</i></b>—Match packet fields to values. If the <b><i>from</i></b> option is not included, all packets are considered to match and the actions and action modifiers in the <b><i>then</i></b> statement are taken.</p> <p><b><i>match-conditions</i></b>—One or more conditions to use to make a match.</p> <p><b><i>term term-name</i></b>—Define a simple-filter term. The name that identifies the term can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include spaces in the name, enclose them in quotation marks (" ").</p> <p><b><i>then</i></b>—Actions to take on matching packets. If the <b><i>then</i></b> option is not included and a packet matches all the conditions in the <b><i>from</i></b> statement, the packet is accepted.</p>
	<div>  <p><b>NOTE:</b> Only <b><i>forwarding-class</i></b> and <b><i>loss-priority</i></b> are valid actions in a simple filter configuration.</p> </div>
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.

- Related Documentation**
- [simple-filter \(Applying to an Interface\) on page 395](#)
  - *Simple Filter Overview*
  - *How Simple Filters Evaluate Packets*
  - *Guidelines for Configuring Simple Filters*
  - *Guidelines for Applying Simple Filters*

## site-identifier (VPLS)

<b>Syntax</b>	<code>site-identifier <i>identifier</i>;</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> ], [edit routing-instances <i>routing-instance-name</i> protocols vpls 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>	<i>identifier</i> —Specify the numerical identifier for the local VPLS site. The identifier must be an unsigned 16-bit number greater than zero.
<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 the VPLS Site Name and Site Identifier</i></li> </ul>

## 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></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>

## term (Simple Filter)

<b>Syntax</b>	<pre>term <i>term-name</i> {     from {         <i>match-conditions</i>;     }     then {         forwarding-class <i>class-name</i>;         loss-priority (high   low   medium);     } }</pre>
<b>Hierarchy Level</b>	[edit firewall family inet <a href="#">simple-filter</a> <i>filter-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	Define a simple filter term.
<b>Options</b>	<p><b>from</b>—Match packet fields to values. If the <b>from</b> option is not included, all packets are considered to match and the actions and action modifiers in the <b>then</b> statement are taken.</p> <p><b>match-conditions</b>—One or more conditions to use to make a match.</p> <p><b>term-name</b>—Name that identifies the term. The name can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").</p> <p><b>then</b>—Actions to take on matching packets. If the <b>then</b> option is not included and a packet matches all the conditions in the <b>from</b> statement, the packet is accepted. For CoS, only the actions listed are allowed. These statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Multifield Classification</i></li> <li>• <i>Simple Filter Overview</i></li> <li>• <i>Firewall Filter Match Conditions for IPv4 Traffic</i></li> <li>• <i>Firewall Filter Match Conditions for IPv6 Traffic</i></li> </ul>

## three-color-policer (Applying)

---

<b>Syntax</b>	<pre>three-color-policer {     (single-rate   two-rate) <i>policer-name</i>; }</pre>
<b>Hierarchy Level</b>	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> then] [edit logical-systems <i>logical-system-name</i> firewall family <i>family-name</i> <b>filter</b> <i>filter-name</i> term <i>term-name</i> then]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. <b>single-rate</b> statement added in Junos OS Release 8.2. Logical systems support introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Apply a tricolor marking policer.
<b>Options</b>	<b>single-rate</b> —Named tricolor policer is a single-rate policer.  <b>two-rate</b> —Named tricolor policer is a two-rate policer.  <b><i>policer-name</i></b> —Name of a tricolor policer.
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring and Applying Tricolor Marking Policers</i></li><li>• <i>Firewall Filter Nonterminating Actions</i></li><li>• <i>Three-Color Policer Configuration Overview</i></li></ul>

## three-color-policer (Configuring)

<b>Syntax</b>	<pre> three-color-policer <i>policer-name</i>   <i>uid</i> {     action {         loss-priority high then discard;     }     filter-specific;     logical-interface-policer;     physical-interface-policer;     shared-bandwidth-policer;     single-rate {         (color-aware   color-blind);         committed-burst-size <i>bytes</i>;         committed-information-rate <i>bps</i>;         excess-burst-size <i>bytes</i>;     }     two-rate {         (color-aware   color-blind);         committed-burst-size <i>bytes</i>;         committed-information-rate <i>bps</i>;         peak-burst-size <i>bytes</i>;         peak-information-rate <i>bps</i>;     } } </pre>
<b>Hierarchy Level</b>	<p>[edit <a href="#">dynamic-profiles</a> <i>profile-name</i> firewall],  [edit firewall],  [edit logical-systems <i>logical-system-name</i> firewall]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.  The <b>action</b> and <b>single-rate</b> statements added in Junos OS Release 8.2.  Logical systems support introduced in Junos OS Release 9.3.  Support at the [edit <b>dynamic-profiles ... firewall</b>] hierarchy level introduced in Junos OS Release 11.4.</p>
<b>Description</b>	<p>Configure a three-color policer.</p>
<b>Options</b>	<p><b><i>policer-name</i></b>—Name of the three-color policer. Reference this name when you apply the policer to an interface.</p> <p><b><i>uid</i></b>—When you configure a policer at the [edit <b>dynamic-profiles</b>] hierarchy level, you must assign a variable UID as the policer name.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</p>
<b>Required Privilege Level</b>	<p>firewall—To view this statement in the configuration.  firewall-control—To add this statement to the configuration.</p>

**Related  
Documentation**

- *Configuring and Applying Tricolor Marking Policers*
- *Three-Color Policer Configuration Guidelines*
- *Basic Single-Rate Three-Color Policers*
- *Basic Two-Rate Three-Color Policers*
- *Two-Color and Three-Color Logical Interface Policers*
- *Two-Color and Three-Color Physical Interface Policers*
- *Two-Color and Three-Color Policers at Layer 2*



## traffic-control-profiles

<b>Syntax</b>	<pre> traffic-control-profiles <i>profile-name</i> {   adjust-minimum <i>rate</i>;   atm-service (cbr   rtvbr   nrtvbr);   delay-buffer-rate (percent <i>percentage</i>   <i>rate</i>);   excess-rate (percent <i>percentage</i>   proportion <i>value</i> );   excess-rate-high (percent <i>percentage</i>   proportion <i>value</i>);   excess-rate-low (percent <i>percentage</i>   proportion <i>value</i>);   guaranteed-rate (percent <i>percentage</i>   <i>rate</i>) &lt;burst-size <i>bytes</i>&gt;;   max-burst-size <i>cells</i>;   overhead-accounting (frame-mode   cell-mode   frame-mode-bytes   cell-mode-bytes)     &lt;bytes (<i>byte-value</i>)&gt;;   peak-rate <i>rate</i>;   scheduler-map <i>map-name</i>;   shaping-rate (percent <i>percentage</i>   <i>rate</i>) &lt;burst-size <i>bytes</i>&gt;;   shaping-rate-excess-high <i>rate</i> [ burst-size <i>bytes</i> ];   shaping-rate-excess-low <i>rate</i> [ burst-size <i>bytes</i> ];   shaping-rate-priority-high <i>rate</i> [ burst-size <i>bytes</i> ];   shaping-rate-priority-low <i>rate</i> [ burst-size <i>bytes</i> ];   shaping-rate-priority-medium <i>rate</i> [ burst-size <i>bytes</i> ];   strict-priority-scheduler;   sustained-rate <i>rate</i>; } </pre>
<b>Hierarchy Level</b>	[edit class-of-service]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	For Gigabit Ethernet IQ, Channelized IQ PICs, FRF.15 and FRF.16 LSQ interfaces, Enhanced Queuing (EQ) DPCs, and PTX Series routers only, configure traffic shaping and scheduling profiles. For Enhanced EQ PICs, EQ DPCs, and PTX Series routers only, you can include the <b>excess-rate</b> statement.
<b>Options</b>	<p><b><i>profile-name</i></b>—Name of the traffic-control profile.</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</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>Oversubscribing Interface Bandwidth</i></li> <li><i>Understanding Scheduling on PTX Series Routers</i></li> </ul>

## traffic-manager

---

**List of Syntax**    [Syntax \(MX Series\) Configure Queue Monitoring on page 404](#)  
[Syntax \(MX Series, T Series\) on page 404](#)  
[Syntax \(M Series\) on page 404](#)  
[Syntax \(vSRX\) on page 404](#)

**Syntax (MX Series)  
Configure Queue  
Monitoring**

```
traffic-manager {  
    egress-shaping-overhead number;  
    ingress-shaping-overhead number;  
    mode {  
        egress-only;  
        ingress-and-egress;  
        session-shaping;  
    }  
    packet-timestamp {  
        enable;  
    }  
    queue-threshold {  
        fabric-queue {  
            priority high/low {  
                threshold threshold-percentage;  
            }  
        }  
        wan-queue {  
            priority high/medium-high/medium-low/low {  
                threshold threshold-percentage;  
            }  
        }  
    }  
}
```

**Syntax (MX Series, T  
Series)**

```
traffic-manager {  
    egress-shaping-overhead number;  
    ingress-shaping-overhead number;  
    mode {  
        egress-only;  
        ingress-and-egress;  
    }  
}
```

**Syntax (M Series)**

```
traffic-manager {  
    egress-shaping-overhead number;  
    ingress-shaping-overhead number;  
    mode {  
        egress-only;  
        ingress-and-egress;  
        session-shaping;  
    }  
}
```

**Syntax (vSRX)**    traffic-manager {

```

    egress-shaping-overhead number;
}

```

**Hierarchy Level** [edit chassis fpc *slot-number*],  
 [edit chassis fpc *slot-number* pic *pic-number*],  
 [edit chassis lcc *number* fpc *slot-number* pic *pic-number*] (Routing Matrix)

**Release Information** Statement introduced in Junos OS Release 8.3.

**Description** Enable CoS queuing, scheduling, and shaping on an L2TP session.



**NOTE:** Committing changes to `traffic-manager` automatically restarts any necessary components (PICs, DPCs, or FPCs).

**Options** **queue-threshold**—Enable monitoring of Fabric and WAN queues. When the **fabric-queue** statement is configured, an SNMP trap is generated whenever the fabric power utilization exceeds the configured threshold value.

When **wan-queue** is configured, an SNMP trap is generated whenever the WAN queue depth exceeds the configured threshold value.

**egress-shaping-overhead number**—When traffic management (queueing and scheduling) is configured on the egress side, the number of CoS shaping overhead bytes to add to the packets on the egress interface.

Replace **number** with a value from **-63** through **192** bytes.

For vSRX, replace **number** with a value from **-62** through **192** bytes.



**NOTE:** The L2 headers (DA/SA + VLAN tags) are automatically a part of the shaping calculation.

---

**ingress-shaping-overhead number**—When L2TP session shaping is configured, the number of CoS shaping overhead bytes to add to the packets on the ingress side of the L2TP tunnel to determine the shaped session packet length.

When session shaping is not configured and traffic management (queueing and scheduling) is configured on the ingress side, the number of CoS shaping overhead bytes to add to the packets on the ingress interface.

Replace **number** with a value from **-63** through **192** bytes.

**mode**—Configure CoS traffic manager mode of operation. This option has the following suboptions:

- **egress-only**—Enable CoS queueing and scheduling on the egress side for the PIC that houses the interface. This is the default mode for an Enhanced Queueing (EQ) DPC on MX Series routers.



**NOTE:** If ingress packet drops are observed at a high rate for an IQ2 or IQ2E PIC, configure the **traffic-manager** statement to work in the **egress-only** mode.

---

- **ingress-and-egress**—Enable CoS queueing and scheduling on both the egress and ingress sides for the PIC. This is the default mode for IQ2 and IQ2E PICs on M Series and T Series routers.



**NOTE:**

- For EQ DPCs, you must configure the **traffic-manager** statement with **ingress-and-egress** mode to enable ingress CoS on the EQ DPC.

- EQ DPCs have 250 ms of buffering, with only egress queueing (default mode). When **ingress-and-egress** is configured, the buffer is partitioned as 50 ms for the ingress direction and 200 ms for the egress direction.

- 
- **session-shaping**—(M Series routers only) Configure the IQ2 PIC mode for session-aware traffic shaping to enable L2TP session shaping.

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

**Related Documentation**

- *Configuring CoS for L2TP Tunnels on ATM Interfaces*
- *Enabling a Timestamp for Ingress and Egress Queue Packets*
- [show interfaces queue on page 573](#)

## transmit-rate (Schedulers)

<b>Syntax</b>	<code>transmit-rate (rate   percent <i>percentage</i>   remainder) &lt;exact   rate-limit&gt;;</code>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">schedulers</a> <i>scheduler-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>rate-limit</b> option introduced in Junos OS Release 8.3. Applied to the Multiservices PICs in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series routers. Statement introduced in Junos OS Release 12.2 for ACX Series routers.
<b>Description</b>	Specify the transmit rate or percentage for a scheduler.
<b>Default</b>	If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent, respectively.
<b>Options</b>	<b>exact</b> —(Optional) Enforce the exact transmission rate. Under sustained congestion, a rate-controlled queue that goes into negative credit fills up and eventually drops packets. This value should never exceed the rate-controlled amount. For PTX Series routers, this option is allowed only on the non-strict-high (high, medium-high, medium-low, or low) queues.  <b>percent <i>percentage</i></b> —Percentage of transmission capacity. A percentage of zero drops all packets in the queue unless additional bandwidth is available from other queues. <b>Range:</b> 0 through 100 percent for M, MX and T Series routers and EX Series switches; 1 through 100 percent for PTX Series routers; 0 through 200 percent for the SONET/SDH OC48/STM16 IQE PIC



### NOTE:

- On M Series Multiservice Edge Routers, for interfaces configured on 4-port E1 and 4-port T1 PICs only, you can configure a *percentage* value only from 11 through 100. These two PICs do not support transmission rates less than 11 percent.
- The configuration of the `transmit-rate percent 0 exact` statement at the [edit class-of-service `schedulers` *scheduler-name*] hierarchy is ineffective on T4000 routers with Type 5 FPC.
- On MIC and MPC interfaces on MX Series routers, when the transmit rate is configured as a percentage and `exact` or `rate-limit` is enabled on a queue, the shaping rate of the parent node is used to compute the transmit rate. If `exact` or `rate-limit` is not configured, the guaranteed rate of the parent node is used to compute the transmit rate.

- On PTX Series routers, unconfigured interfaces are equivalent to percent 0. This means the system offers no guaranteed rate on the interface, and the queue will always be scheduled in the excess priority.

**rate**—Transmission rate, in bps. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

**Range:** 3200 through 6,400,000,000,000 bps



**NOTE:** For all MX Series interfaces, the rate can be from 65,535 through 6,400,000,000,000 bps.

**rate-limit**—(Optional) Limit the transmission rate to the rate-controlled amount by applying a policing action to the queue. Packets are hard-dropped when traffic exceeds the specified maximum transmission rate.



**NOTE:** For PTX Series routers, this option is allowed only on the strict-high queue. We recommend that you configure rate limit on strict-high queues because the other queues may not meet their guaranteed bandwidths. The **rate-limit** option cannot rate limit the queue if strict-priority scheduling is configured with the *strict-priority-scheduler* statement.



**NOTE:** The configuration of the **rate-limit** statement is supported on T4000 routers only with a Type 5 FPC.

**remainder**—Use the remaining rate available.

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

**Related Documentation**

- *Configuring Schedulers*
- *Configuring Scheduler Transmission Rate*
- *Understanding Scheduling on PTX Series Routers*


## transparent

---

<b>Syntax</b>	transparent;
<b>Hierarchy Level</b>	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> classifiers ieee802.1 vlan-tag]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2
<b>Description</b>	Packet classification based on the transparent VLAN tag.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.



## 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 demux0</a> <i>interface-name</i> unit unit <i>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><i>\$junos-underlying-interface</i></b> when configuring dynamic IP demux interfaces.</li> <li>• <b><i>\$junos-interface-ifd-name</i></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 <a href="#">unit</a> "\$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>	<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>Configuring a PPPoE Dynamic Profile</i></li><li>• <i>Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview</i></li></ul>

## unit

<b>Syntax</b>	<pre> unit logical-unit-number {   classifiers {     type (classifier-name   default) family (mpls   all);   }   forwarding-class class-name;   fragmentation-map map-name;   input-traffic-control-profile profile-name shared-instance instance-name;   output-traffic-control-profile profile-name shared-instance instance-name;   per-session-scheduler;   rewrite-rules {     dscp (rewrite-name   default);     dscp-ipv6 (rewrite-name   default);     exp (rewrite-name   default) protocol protocol-types;     exp-push-push-push default;     exp-swap-push-push default;     ieee-802.1 (rewrite-name   default) vlan-tag (outer   outer-and-inner);     inet-precedence (rewrite-name   default);   }   scheduler-map map-name;   shaping-rate rate; } </pre>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">interfaces</a> interface-name]
<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 16,384</p> <p>The remaining statements are explained separately. See <a href="#">CLI Explorer</a>.</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>Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic</i></li> <li><i>Configuring Rewrite Rules</i></li> </ul>

## 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. See [CLI Explorer](#).

**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 PPPoE)

---

```
Syntax  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 ];
            mru size;
            mtu (size | use-lower-layer);
            chap {
                challenge-length minimum minimum-length maximum maximum-length;
            }
            ignore-magic-number-mismatch;
            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;
            }
            output filter-name {
                precedence precedence;
            }
        }
    }
```

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

```
}
}
```

**Hierarchy Level** [edit [dynamic-profiles](#) *profile-name* [interfaces](#) pp0]

**Release Information** Statement introduced in Junos OS Release 10.1.

**Description** In a dynamic profile, configure a logical unit number for the dynamic PPPoE logical interface. You must configure a logical interface to be able to use the router.

**Options** *logical-unit-number*—Variable used to specify the unit number when the PPPoE logical interface is dynamically created. In the **unit** *logical-unit-number* statement for dynamic PPPoE logical interfaces, you must use the predefined variable **\$junos-interface-unit** in place of *logical-unit-number*. The **\$junos-interface-unit** predefined variable is dynamically replaced with the unit number supplied by the router when the subscriber logs in.

The remaining statements are explained separately. See [CLI Explorer](#).

**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 PPPoE)

---

<b>Syntax</b>	<code>unnumbered-address <i>interface-name</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">dynamic-profiles</a> <i>profile-name</i> <a href="#">interfaces</a> pp0 <a href="#">unit</a> "\$junos-interface-unit" <a href="#">family</a> inet]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	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.
<b>Options</b>	<i>interface-name</i> —Interface from which the local address is derived. The interface name must include a logical unit number and must have a configured address.
<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>



## 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</code> unit <code>logical-unit-number family family</code> ], [edit <code>dynamic-profiles profile-name interfaces demux0</code> unit <code>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> unit <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> unit <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>• <i>address</i></li> <li>• <i>Junos System Basics Configuration Guide</i></li> </ul>

## user-prefix (DHCP Relay Agent)

<b>Syntax</b>	<code>user-prefix <i>user-prefix-string</i>;</code>
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay authentication username-include],          [edit forwarding-options dhcp-relay dhcpv6 authentication username-include],          [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],          [edit forwarding-options dhcp-relay dual-stack-group <i>dual-stack-group-name</i> authentication username-include],          [edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],          [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication username-include],          [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],          [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],          [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],          [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],          [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],          [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],          [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.          Statement introduced in Junos OS Release 12.3R2 for EX Series switches.          Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.          Support at the [edit ... dual-stack-group <i>dual-stack-group-name</i>] hierarchy level introduced in Junos OS Release 15.1.</p>
<b>Description</b>	Specify the user prefix that is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.
<b>Options</b>	<i>user-prefix-string</i> —User prefix string.

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

**Related Documentation** • *Using External AAA Authentication Services with DHCP*

## vlan-id (Dynamic Profiles)

**Syntax** `vlan-id (number | none);`

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

**Release Information** Statement introduced in Junos OS Release 9.5.  
 VLAN demux interface support introduced in Junos OS Release 10.2.

**Description** For VLAN demux, Fast Ethernet, Gigabit Ethernet, and Aggregated Ethernet interfaces only, bind a 802.1Q VLAN tag ID to a logical interface.

**Options** **number**—A valid VLAN identifier. When used in the **dynamic-profiles** hierarchy, specify the **\$junos-vlan-id** predefined variable to dynamically obtain the VLAN identifier.

**none**—Enable the use of untagged pseudo-wire frames on dynamic interfaces.

- For aggregated Ethernet, 4-port, 8-port, and 12-port Fast Ethernet PICs, and for management and internal Ethernet interfaces, 1 through 1023.
- For 48-port Fast Ethernet and Gigabit Ethernet PICs, 1 through 4094.
- VLAN ID 0 is reserved for tagging the priority of frames.

**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 VLAN Demux Interfaces in Dynamic Profiles*

## 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>	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>Enabling VLAN 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>• <i>Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers</i></li> <li>• <i>Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers</i></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-tag

---

<b>Syntax</b>	<code>vlan-tag (outer   outer-and-inner);</code>
<b>Hierarchy Level</b>	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules ieee-802.1 ( <i>rewrite-name</i>   default)]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1.
<b>Description</b>	Apply this IEEE-802.1 rewrite rule to the outer VLAN tag or, if available, both outer and inner VLAN tags.
<b>Default</b>	If you do not include this statement, the rewrite rule applies to the outer VLAN tag only.
<b>Options</b>	<b>outer</b> —Apply the rewrite rule to the outer VLAN tag only. <b>outer-and-inner</b> —Apply the rewrite rule to both the outer and inner VLAN tags.
<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>Applying IEEE 802.1p Rewrite Rules to Dual VLAN Tags</i></li></ul>

## vlan-tags

<b>Syntax</b>	<code>vlan-tags outer [<i>tpid</i>].<i>vlan-id</i> [inner [<i>tpid</i>].<i>vlan-id</i>];</code>
<b>Hierarchy Level</b>	<code>[edit dynamic-profiles <i>profile-name</i> <b>interfaces</b> <i>interface-name</i> unit <i>logical-unit-number</i>]</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 <b>interfaces</b> <i>interface-name</i>]</code> hierarchy level.
<div>  <b>NOTE:</b> The inner-range <i>vid1–vid2</i> option is supported on IQE PICs only. </div>	
<b>Options</b>	<b>inner [<i>tpid</i>].<i>vlan-id</i></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.
<div>  <b>NOTE:</b> On the network-to-network (NNI) or egress interfaces of provider edge (PE) routers, you cannot configure the inner-range <i>tpid. vid1–vid2</i> option with the <b>vlan-tags</b> statement for ISP-facing interfaces. </div>	
<b>Range:</b> For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.	
<b>outer [<i>tpid</i>].<i>vlan-id</i></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-stacked-vlan-id</code> predefined variable.	
<b>Range:</b> 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>Configuring Dual VLAN Tags</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>	Bind TPIDs and 802.1Q VLAN tag IDs to a logical interface. TPID fields are used to identify the frame as an IEEE 802.1Q-tagged frame.
<b>Options</b>	<p><b>inner <i>tpid.vlan-id</i></b>—A TPID and a valid VLAN identifier. TPID is a 16-bit field set to a value of 0x8100 in order to identify the frame as an IEEE 802.1Q-tagged frame.</p> <p><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.</p> <p><b>inner-list <i>value</i></b>— List or a set of VLAN identifiers.</p>



**NOTE:** This is supported on MX Series routers 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. The inner-range must 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 through 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 398</a></li></ul>
------------------------------	--

## CHAPTER 11

# Operational Commands

- `show access-cac interface-set`
- `show class-of-service adjustment-control-profile`
- `show class-of-service classifier`
- `show class-of-service interface`
- `show class-of-service scheduler-map`
- `show class-of-service scheduler-hierarchy interface`
- `show class-of-service scheduler-hierarchy interface-set`
- `show interfaces`
- `show interfaces (PPPoE)`
- `show interfaces demux0 (Demux Interfaces)`
- `show interfaces queue`
- `show subscribers`

## show access-cac interface-set

<b>Syntax</b>	show access-cac interface-set <interface-set-name   detail>
<b>Release Information</b>	Command introduced in Junos OS Release 17.2 for the MX Series.
<b>Description</b>	Display interface-set adjustment information.
<b>Options</b>	<p><b>none</b>—List all interface-set entries with the total aggregate adjustment value for each interface-set.</p> <p><b>interface-set-name</b>—Restrict the output to a specific interface-set.</p> <p><b>detail</b>—List all multicast groups, with adjustment rates, that provide an update on the specified interface-set.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using Hierarchical CoS to Adjust Shaping Rates Based on Multicast Traffic on page 159</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show access-cac interface-set on page 433</a> <a href="#">show access-cac interface-set (detail) on page 433</a>
<b>Output Fields</b>	<a href="#">Table 31 on page 432</a> describes the output fields for the <b>show show access-cac interface-set</b> command. Output fields are listed in the approximate order in which they appear.

Table 31: show class-of-service interface-set Output Fields

Field Name	Field Description
Access cac interface-set	Then name of the interface set.
Total current adjustment on interface-set	The total current adjustment on the interface set, in bps, equal to the sum of the measured rates on all configured multicast channels.
Mcast channels	The list of all configured multicast channels.
Configured rate	The shaping rate configured for each multicast channel.
Measured rate	The current data rate for each multicast channel. This rate for a multicast channel adjusts the shaping rate of the interface set by an equal amount.

## Sample Output

### show access-cac interface-set

```
user@host> show access-cac interface-set
Access cac interface-sets
  Access cac interface-set: ge-5/0/0-1
  Access cac interface-set: ge-5/0/1-1
```

### show access-cac interface-set (detail)

```
user@host> show access-cac interface-set detail
Access cac interface-sets
  Access cac interface-set: ge-5/0/0-1
  Total current adjustment on interface-set: 4080908
  Mcast channels:
    0.0.0.0      , 225.0.0.1    Configured rate: 4000000 Measured rate: 560268
    0.0.0.0      , 225.0.0.2    Configured rate: 4000000 Measured rate: 560280
    0.0.0.0      , 225.0.0.3    Configured rate: 4000000 Measured rate: 560280
    ::           , ff8e::e100:1  Configured rate: 4000000 Measured rate: 2400080

  Access cac interface-set: ge-5/0/1-1
  Total current adjustment on interface-set: 500222379
  Mcast channels:
    0.0.0.0      , 225.0.0.1    Configured rate: 4000000 Measured rate: 166739966
    0.0.0.0      , 225.0.0.2    Configured rate: 4000000 Measured rate: 166741131
    0.0.0.0      , 225.0.0.3    Configured rate: 4000000 Measured rate: 166741282
    ::           , ff8e::e100:1  Configured rate: 4000000 Measured rate: 0
```

## show class-of-service adjustment-control-profile

<b>Syntax</b>	<code>show class-of-service adjustment-control-profile</code> <code>&lt;profile-name&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 13.1 for MX Series Routers.
<b>Description</b>	For MPC/MIC interfaces only, display the adjustment control profiles.
<b>Options</b>	<b>none</b> —Display all profiles.  <b>profile-name</b> —(Optional) Display information about a single profile.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Verifying the CoS Adjustment Control Profile Configuration</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show class-of-service adjustment-control-profile on page 435</a>
<b>Output Fields</b>	<a href="#">Table 32 on page 434</a> describes the output fields for the <b>show class-of-service adjustment-control-profile</b> command. Output fields are listed in the approximate order in which they appear.

**Table 32: show class-of-service adjustment-control-profile Output Fields**

Field Name	Field Description
<b>Name</b>	<p>Name of the adjusting application. Possible values:</p> <ul style="list-style-type: none"> <li>• RADIUS-CoA—RADIUS CoA application.</li> <li>• ANCP—ANCP application.</li> <li>• PPPoE IA tags—PPPoE IA tag application.</li> </ul>
<b>Priority</b>	<p>Priority of the adjusting application. Possible values are 1 through 10; 1 being the highest priority.</p> <p>The lower the priority value, the higher the priority</p>
<b>Algorithm</b>	<p>Algorithm the adjusting application uses to make adjustments.</p> <ul style="list-style-type: none"> <li>• adjust-never—Never perform rate adjustments.</li> <li>• adjust-always—Adjust the shaping rate unconditionally.</li> <li>• adjust-less—Adjust the shaping rate if it is less than the configured value.</li> <li>• adjust-less-or-equal—Adjust the shaping rate if it is less than or equal to the configured value.</li> <li>• adjust-greater—Adjust the shaping rate if it is greater than the configured value.</li> <li>• adjust-greater-or-equal—Adjust the shaping rate if it is greater than or equal to the configured value.</li> </ul>



## Sample Output

### show class-of-service adjustment-control-profile

```
user@host> show class-of-service adjustment-control-profile

name: ANCP, priority: 1, algorithm: less
name: RADIUS CoA, priority: 1, algorithm: always
name: PPPoE IA tags, priority: 2, algorithm: less
```

## show class-of-service classifier

<b>Syntax</b>	<pre>show class-of-service classifier &lt;name <i>name</i>&gt; &lt;type dscp   type dscp-ipv6   type exp   type ieee-802.1   type inet-precedence&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	For each class-of-service (CoS) classifier, display the mapping of code point value to forwarding class and loss priority.
<b>Options</b>	<p><b>none</b>—Display all classifiers.</p> <p><b>name <i>name</i></b>—(Optional) Display named classifier.</p> <p><b>type dscp</b>—(Optional) Display all classifiers of the Differentiated Services code point (DSCP) type.</p> <p><b>type dscp-ipv6</b>—(Optional) Display all classifiers of the DSCP for IPv6 type.</p> <p><b>type exp</b>—(Optional) Display all classifiers of the MPLS experimental (EXP) type.</p> <p><b>type ieee-802.1</b>—(Optional) Display all classifiers of the ieee-802.1 type.</p> <p><b>type inet-precedence</b>—(Optional) Display all classifiers of the inet-precedence type.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show class-of-service classifier type ieee-802.1 on page 437</a></p> <p><a href="#">show class-of-service classifier type ieee-802.1 (QFX Series) on page 437</a></p>
<b>Output Fields</b>	<p><a href="#">Table 33 on page 436</a> describes the output fields for the <b>show class-of-service classifier</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 33: show class-of-service classifier Output Fields**

Field Name	Field Description
Classifier	Name of the classifier.
Code point type	Type of the classifier: <b>exp</b> (not on EX Series switch), <b>dscp</b> , <b>dscp-ipv6</b> (not on EX Series switch), <b>ieee-802.1</b> , or <b>inet-precedence</b> .
Index	Internal index of the classifier.
Code point	Code point value used for classification

Table 33: show class-of-service classifier Output Fields (*continued*)

Field Name	Field Description
<b>Forwarding class</b>	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.
<b>Loss priority</b>	Loss priority value used for classification. For most platforms, the value is <b>high</b> or <b>low</b> . For some platforms, the value is <b>high</b> , <b>medium-high</b> , <b>medium-low</b> , or <b>low</b> .

## Sample Output

### show class-of-service classifier type ieee-802.1

```

user@host> show class-of-service classifier type ieee-802.1
Classifier: ieee802.1-default, Code point type: ieee-802.1, Index: 3
Code Point      Forwarding Class      Loss priority
000             best-effort           low
001             best-effort           high
010             expedited-forwarding  low
011             expedited-forwarding  high
100             assured-forwarding    low
101             assured-forwarding    medium-high
110             network-control       low
111             network-control       high

Classifier: users-ieee802.1, Code point type: ieee-802.1
Code point      Forwarding class      Loss priority
100             expedited-forwarding  low

```

### show class-of-service classifier type ieee-802.1 (QFX Series)

```

user@switch> show class-of-service classifier type ieee-802.1
Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11
Code point      Forwarding class      Loss priority
000             best-effort           low
001             best-effort           low
010             best-effort           low
011             fcoe                  low
100             no-loss               low
101             best-effort           low
110             network-control       low
111             network-control       low

Classifier: ieee8021p-untrust, Code point type: ieee-802.1, Index: 16
Code point      Forwarding class      Loss priority
000             best-effort           low
001             best-effort           low
010             best-effort           low
011             best-effort           low
100             best-effort           low
101             best-effort           low
110             best-effort           low
111             best-effort           low

Classifier: ieee-mcast, Code point type: ieee-802.1, Index: 46
Code point      Forwarding class      Loss priority

```

000	mcast	low
001	mcast	low
010	mcast	low
011	mcast	low
100	mcast	low
101	mcast	low
110	mcast	low
111	mcast	low

## show class-of-service interface

**Syntax** `show class-of-service interface`  
`<comprehensive | detail> <interface-name>`

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
 Forwarding class map information added in Junos OS Release 9.4.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.  
 Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Routers.  
 Command introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.  
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.  
 Options **detail** and **comprehensive** introduced in Junos OS Release 11.4.  
 Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

**Description** Display the logical and physical interface associations for the classifier, rewrite rules, and scheduler map objects.



**NOTE:** On routing platforms with dual Routing Engines, running this command on the backup Routing Engine, with or without any of the available options, is not supported and produces the following error message:

**error: the class-of-service subsystem is not running**

**Options** **none**—Display CoS associations for all physical and logical interfaces.

**comprehensive**—(M Series, MX Series, and T Series routers) (Optional) Display comprehensive quality-of-service (QoS) information about all physical and logical interfaces.

**detail**—(M Series, MX Series, and T Series routers) (Optional) Display QoS and CoS information based on the interface.

If the **interface** *interface-name* is a physical interface, the output includes:

- Brief QoS information about the physical interface
- Brief QoS information about the logical interface
- CoS information about the physical interface
- Brief information about filters or policers of the logical interface
- Brief CoS information about the logical interface

If the **interface** *interface-name* is a logical interface, the output includes:

- Brief QoS information about the logical interface
- Information about filters or policers for the logical interface
- CoS information about the logical interface

**interface-name**—(Optional) Display class-of-service (CoS) associations for the specified interface.

**none**—Display CoS associations for all physical and logical interfaces.

**Required Privilege Level**

view

**Related Documentation**

- *Verifying and Managing Junos OS Enhanced Subscriber Management*

**List of Sample Output**

[show class-of-service interface \(Physical\) on page 451](#)  
[show class-of-service interface \(Logical\) on page 452](#)  
[show class-of-service interface \(Gigabit Ethernet\) on page 452](#)  
[show class-of-service interface \(ANCP\) on page 452](#)  
[show class-of-service interface \(PPPoE Interface\) on page 452](#)  
[show class-of-service interface \(DHCP Interface\) on page 453](#)  
[show class-of-service interface \(T4000 Routers with Type 5 FPCs\) on page 453](#)  
[show class-of-service interface detail on page 453](#)  
[show class-of-service interface comprehensive on page 454](#)  
[show class-of-service interface \(ACX Series Routers\) on page 465](#)  
[show class-of-service interface \(PPPoE Subscriber Interface for Enhanced Subscriber Management\) on page 467](#)

**Output Fields**

[Table 34 on page 440](#) describes the output fields for the **show class-of-service interface** command. Output fields are listed in the approximate order in which they appear.

**Table 34: show class-of-service interface Output Fields**

Field Name	Field Description
Physical interface	Name of a physical interface.
Index	Index of this interface or the internal index of this object.  (Enhanced subscriber management for MX Series routers) Index values for dynamic CoS traffic control profiles and dynamic scheduler maps are larger for enhanced subscriber management than they are for legacy subscriber management.
Dedicated Queues	Status of dedicated queues configured on an interface. Supported only on Trio MPC/MIC interfaces on MX Series routers.  (Enhanced subscriber management for MX-Series routers) This field is not displayed for enhanced subscriber management.

Table 34: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Maximum usable queues	Number of queues you can configure on the interface.
Maximum usable queues	Maximum number of queues you can use.
Total non-default queues created	Number of queues created in addition to the default queues. Supported only on Trio MPC/MIC interfaces on MX Series routers.  (Enhanced subscriber management for MX Series routers) This field is not displayed for enhanced subscriber management.
Rewrite Input IEEE Code-point	(QFX3500 switches only) IEEE 802.1p code point (priority) rewrite value. Incoming traffic from the Fibre Channel (FC) SAN is classified into the forwarding class specified in the native FC interface (NP_Port) fixed classifier and uses the priority specified as the IEEE 802.1p rewrite value.
Shaping rate	Maximum transmission rate on the physical interface. You can configure the shaping rate on the physical interface, or on the logical interface, but not on both. Therefore, the <b>Shaping rate</b> field is displayed for either the physical interface or the logical interface.
Scheduler map	Name of the output scheduler map associated with this interface.  (Enhanced subscriber management for MX Series routers) The name of the dynamic scheduler map object is associated with a generated UID (for example, <b>SMAP-1_UID1002</b> ) instead of with a subscriber interface.
Scheduler map forwarding class sets	(QFX Series only) Name of the output fabric scheduler map associated with a QFabric system Interconnect device interface.
Input shaping rate	For Gigabit Ethernet IQ2 PICs, maximum transmission rate on the input interface.
Input scheduler map	For Gigabit Ethernet IQ2 PICs, name of the input scheduler map associated with this interface.
Chassis scheduler map	Name of the scheduler map associated with the packet forwarding component queues.
Rewrite	Name and type of the rewrite rules associated with this interface.
Traffic-control-profile	Name of the associated traffic control profile.  (Enhanced subscriber management for MX Series routers) The name of the dynamic traffic control profile object is associated with a generated UID (for example, <b>TC_PROF_100_199_SERIES_UID1006</b> ) instead of with a subscriber interface.
Classifier	Name and type of classifiers associated with this interface.
Forwarding-class-map	Name of the forwarding map associated with this interface.
Congestion-notification	(QFX Series and EX4600 switches only) Congestion notification state, <b>enabled</b> or <b>disabled</b> .
Logical interface	Name of a logical interface.

Table 34: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
<b>Object</b>	Category of an object: <b>Classifier</b> , <b>Fragmentation-map</b> (for LSQ interfaces only), <b>Scheduler-map</b> , <b>Rewrite</b> , <b>Translation Table</b> (for IQE PICs only), or <b>traffic-class-map</b> (for T4000 routers with Type 5 FPCs).
<b>Name</b>	Name of an object.
<b>Type</b>	Type of an object: <b>dscp</b> , <b>dscp-ipv6</b> , <b>exp</b> , <b>ieee-802.1</b> , <b>ip</b> , <b>inet-precedence</b> , or <b>ieee-802.1ad</b> (for traffic class map on T4000 routers with Type 5 FPCs)..
<b>Link-level type</b>	Encapsulation on the physical interface.
<b>MTU</b>	MTU size on the physical interface.
<b>Speed</b>	Speed at which the interface is running.
<b>Loopback</b>	Whether loopback is enabled and the type of loopback.
<b>Source filtering</b>	Whether source filtering is enabled or disabled.
<b>Flow control</b>	Whether flow control is enabled or disabled.
<b>Auto-negotiation</b>	(Gigabit Ethernet interfaces) Whether autonegotiation is enabled or disabled.
<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>
<b>Device flags</b>	The <b>Device flags</b> field provides information about the physical device and displays one or more of the following values: <ul style="list-style-type: none"> <li>• <b>Down</b>—Device has been administratively disabled.</li> <li>• <b>Hear-Own-Xmit</b>—Device receives its own transmissions.</li> <li>• <b>Link-Layer-Down</b>—The link-layer protocol has failed to connect with the remote endpoint.</li> <li>• <b>Loopback</b>—Device is in physical loopback.</li> <li>• <b>Loop-Detected</b>—The link layer has received frames that it sent, thereby detecting a physical loopback.</li> <li>• <b>No-Carrier</b>—On media that support carrier recognition, no carrier is currently detected.</li> <li>• <b>No-Multicast</b>—Device does not support multicast traffic.</li> <li>• <b>Present</b>—Device is physically present and recognized.</li> <li>• <b>Promiscuous</b>—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media.</li> <li>• <b>Quench</b>—Transmission on the device is quenched because the output buffer is overflowing.</li> <li>• <b>Recv-All-Multicasts</b>—Device is in multicast promiscuous mode and therefore provides no multicast filtering.</li> <li>• <b>Running</b>—Device is active and enabled.</li> </ul>



Table 34: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
<b>Interface flags</b>	<p>The <b>Interface flags</b> field provides information about the physical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> <li>• <b>Admin-Test</b>—Interface is in test mode and some sanity checking, such as loop detection, is disabled.</li> <li>• <b>Disabled</b>—Interface is administratively disabled.</li> <li>• <b>Down</b>—A hardware failure has occurred.</li> <li>• <b>Hardware-Down</b>—Interface is nonfunctional or incorrectly connected.</li> <li>• <b>Link-Layer-Down</b>—Interface keepalives have indicated that the link is incomplete.</li> <li>• <b>No-Multicast</b>—Interface does not support multicast traffic.</li> <li>• <b>No-receive No-transmit</b>—Passive monitor mode is configured on the interface.</li> <li>• <b>Point-To-Point</b>—Interface is point-to-point.</li> <li>• <b>Pop all MPLS labels from packets of depth</b>—MPLS labels are removed as packets arrive on an interface that has the <b>pop-all-labels</b> statement configured. The depth value can be one of the following: <ul style="list-style-type: none"> <li>• <b>1</b>—Takes effect for incoming packets with one label only.</li> <li>• <b>2</b>—Takes effect for incoming packets with two labels only.</li> <li>• <b>[ 1 2 ]</b>—Takes effect for incoming packets with either one or two labels.</li> </ul> </li> <li>• <b>Promiscuous</b>—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses.</li> <li>• <b>Recv-All-Multicasts</b>—Interface is in multicast promiscuous mode and provides no multicast filtering.</li> <li>• <b>SNMP-Traps</b>—SNMP trap notifications are enabled.</li> <li>• <b>Up</b>—Interface is enabled and operational.</li> </ul>
<b>Flags</b>	<p>The <b>Logical interface flags</b> field provides information about the logical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> <li>• <b>ACFC Encapsulation</b>—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer).</li> <li>• <b>Device-down</b>—Device has been administratively disabled.</li> <li>• <b>Disabled</b>—Interface is administratively disabled.</li> <li>• <b>Down</b>—A hardware failure has occurred.</li> <li>• <b>Clear-DF-Bit</b>—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit.</li> <li>• <b>Hardware-Down</b>—Interface protocol initialization failed to complete successfully.</li> <li>• <b>PFC</b>—Protocol field compression is enabled for the PPP session.</li> <li>• <b>Point-To-Point</b>—Interface is point-to-point.</li> <li>• <b>SNMP-Traps</b>—SNMP trap notifications are enabled.</li> <li>• <b>Up</b>—Interface is enabled and operational.</li> </ul>
<b>Encapsulation</b>	Encapsulation on the logical interface.
<b>Admin</b>	Administrative state of the interface ( <b>Up</b> or <b>Down</b> ).
<b>Link</b>	Status of physical link ( <b>Up</b> or <b>Down</b> ).
<b>Proto</b>	Protocol configured on the interface.

Table 34: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
<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.
<b>Link flags</b>	Provides information about the physical link and displays one or more of the following values: <ul style="list-style-type: none"> <li>• <b>ACFC</b>—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option.</li> <li>• <b>Give-Up</b>—Link protocol does not continue connection attempts after repeated failures.</li> <li>• <b>Loose-LCP</b>—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational.</li> <li>• <b>Loose-LMI</b>—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational.</li> <li>• <b>Loose-NCP</b>—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational.</li> <li>• <b>Keepalives</b>—Link protocol keepalives are enabled.</li> <li>• <b>No-Keepalives</b>—Link protocol keepalives are disabled.</li> <li>• <b>PFC</b>—Protocol field compression is configured. The PPP session negotiates the PFC option.</li> </ul>
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds.
<b>CoS queues</b>	Number of CoS queues configured.
<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>Statistics last cleared</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>IPv6 transit statistics</b>	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.

Table 34: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
<b>Input errors</b>	<p>Input errors on the interface. The labels are explained in the following list:</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>Giants</b>—Number of frames received that are larger than the giant threshold.</li> <li>• <b>Bucket Drops</b>—Drops resulting from the traffic load exceeding the interface transmit or receive leaky bucket configuration.</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. Layer 3 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>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>HS link FIFO overflows</b>—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces.</li> </ul>
<b>Output errors</b>	<p>Output errors on the interface. The labels are explained in the following list:</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 up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning.</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>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>HS link FIFO underflows</b>—Number of FIFO underflows 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 exceeds the MTU of the interface.</li> </ul>
<b>Egress queues</b>	Total number of egress Maximum usable queues on the specified interface.

Table 34: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
<b>Queue counters</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> <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>
<b>SONET alarms</b> <b>SONET defects</b>	<p>(SONET) SONET media-specific alarms and defects that prevent the interface from passing packets. When a defect persists for a certain period, 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 light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: <b>SONET PHY</b>, <b>SONET section</b>, <b>SONET line</b>, and <b>SONET path</b>.</p>
<b>SONET PHY</b>	<p>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. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET PHY</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>PLL Lock</b>—Phase-locked loop</li> <li>• <b>PHY Light</b>—Loss of optical signal</li> </ul>
<b>SONET section</b>	<p>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. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET section</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>BIP-BI</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>SEF</b>—Severely errored framing</li> <li>• <b>LOS</b>—Loss of signal</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>

Table 34: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
<b>SONET line</b>	<p>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. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET line</b> field has the following subfields:</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>SONET path</b>	<p>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. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET path</b> field has the following subfields:</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>ES-PFE</b>—Errored seconds (far-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>

Table 34: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Received SONET overhead	<p>Values of the received and transmitted SONET overhead:</p> <ul style="list-style-type: none"> <li><b>C2</b>—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P.</li> <li><b>F1</b>—Section user channel byte. This byte is set aside for the purposes of users.</li> <li><b>K1 and K2</b>—These bytes are allocated for APS signaling for the protection of the multiplex section.</li> <li><b>J0</b>—Section trace. This byte is defined for STS-1 number 1 of an STS-N signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter.</li> <li><b>S1</b>—Synchronization status. The S1 byte is located in the first STS-1 number of an STS-N signal.</li> <li><b>Z3 and Z4</b>—Allocated for future use.</li> </ul>
Transmitted SONET overhead	
Received path trace	<p>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>
Transmitted path trace	
HDLC configuration	<p>Information about the HDLC configuration.</p> <ul style="list-style-type: none"> <li><b>Policing bucket</b>—Configured state of the receiving policer.</li> <li><b>Shaping bucket</b>—Configured state of the transmitting shaper.</li> <li><b>Giant threshold</b>—Giant threshold programmed into the hardware.</li> <li><b>Runt threshold</b>—Runt threshold programmed into the hardware.</li> </ul>
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> <li><b>PLP byte</b>—Packet Level Protocol byte.</li> </ul>
CoS information	<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>
Forwarding classes	Total number of forwarding classes supported on the specified interface.
Egress queues	Total number of egress Maximum usable queues on the specified interface.

Table 34: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Queue	Queue number.
Forwarding classes	Forwarding class name.
Queued Packets	Number of packets queued to this queue.
Queued Bytes	Number of bytes queued to this queue. The byte counts vary by PIC type.
Transmitted Packets	Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the <b>Packet Forwarding Engine Chassis Queues</b> field) shows the prefragmentation values.
Transmitted Bytes	Number of bytes transmitted by this queue. The byte counts vary by PIC type.
Tail-dropped packets	Number of packets dropped because of tail drop.
RED-dropped packets	<p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> <li>• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li>• <b>Low, non-TCP</b>—Number of low-loss priority non-TCP packets dropped because of RED.</li> <li>• <b>Low, TCP</b>—Number of low-loss priority TCP packets dropped because of RED.</li> <li>• <b>High, non-TCP</b>—Number of high-loss priority non-TCP packets dropped because of RED.</li> <li>• <b>High, TCP</b>—Number of high-loss priority TCP packets dropped because of RED.</li> </ul> </li> <li>• (MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li>• <b>Low</b>—Number of low-loss priority packets dropped because of RED.</li> <li>• <b>Medium-low</b>—Number of medium-low loss priority packets dropped because of RED.</li> <li>• <b>Medium-high</b>—Number of medium-high loss priority packets dropped because of RED.</li> <li>• <b>High</b>—Number of high-loss priority packets dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>

Table 34: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
RED-dropped bytes	<p>Number of bytes dropped because of RED. The byte counts vary by PIC type.</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP bytes dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP bytes dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP bytes dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP bytes dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>
Transmit rate	Configured transmit rate of the scheduler. The rate is a percentage of the total interface bandwidth.
Rate Limit	<p>Rate limiting configuration of the queue. Possible values are :</p> <ul style="list-style-type: none"> <li><b>None</b>—No rate limit.</li> <li><b>exact</b>—Queue transmits at the configured rate.</li> </ul>
Buffer size	Delay buffer size in the queue.
Priority	Scheduling priority configured as <b>low</b> or <b>high</b> .
Excess Priority	Priority of the excess bandwidth traffic on a scheduler: <b>low</b> , <b>medium-low</b> , <b>medium-high</b> , <b>high</b> , or <b>none</b> .
Drop profiles	<p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> <li><b>Loss priority</b>—Packet loss priority for drop profile assignment.</li> <li><b>Protocol</b>—Transport protocol for drop profile assignment.</li> <li><b>Index</b>—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles.</li> <li><b>Name</b>—Name of the drop profile.</li> <li><b>Type</b>—Type of the drop profile: <b>discrete</b> or <b>interpolated</b>.</li> <li><b>Fill Level</b>—Percentage fullness of a queue.</li> <li><b>Drop probability</b>—Drop probability at this fill level.</li> </ul>
Excess Priority	Priority of the excess bandwidth traffic on a scheduler.



Table 34: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
<b>Drop profiles</b>	<p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> <li>• <b>Loss priority</b>—Packet loss priority for drop profile assignment.</li> <li>• <b>Protocol</b>—Transport protocol for drop profile assignment.</li> <li>• <b>Index</b>—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles.</li> <li>• <b>Name</b>—Name of the drop profile.</li> <li>• <b>Type</b>—Type of the drop profile: <b>discrete</b> or <b>interpolated</b>.</li> <li>• <b>Fill Level</b>—Percentage fullness of a queue.</li> <li>• <b>Drop probability</b>—Drop probability at this fill level.</li> </ul>
<b>Adjustment information</b>	<p>Display the assignment of shaping-rate adjustments on a scheduler node or queue.</p> <ul style="list-style-type: none"> <li>• <b>Adjusting application</b>—Application that is performing the shaping-rate adjustment. <ul style="list-style-type: none"> <li>• The adjusting application can appear as <b>ancp LS-0</b>, which is the Junos OS Access Node Control Profile process (<b>ancpd</b>) that performs shaping-rate adjustments on schedule nodes.</li> <li>• The adjusting application can appear as <b>DHCP</b>, which adjusts the shaping-rate and overhead-accounting class-of-service attributes based on DHCP option 82, suboption 9 (Vendor Specific Information). The shaping rate is based on the actual-data-rate-downstream attribute. The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).</li> <li>• The adjusting application can also appear as <b>pppoe</b>, which adjusts the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network based on access line parameters in Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101]. This feature is supported on MPC/MIC interfaces on MX Series routers. The shaping rate is based on the actual-data-rate-downstream attribute. The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).</li> </ul> </li> <li>• <b>Adjustment type</b>—Type of adjustment: <b>absolute</b> or <b>delta</b>.</li> <li>• <b>Configured shaping rate</b>—Shaping rate configured for the scheduler node or queue.</li> <li>• <b>Adjustment value</b>—Value of adjusted shaping rate.</li> <li>• <b>Adjustment target</b>—Level of shaping-rate adjustment performed: <b>node</b> or <b>queue</b>.</li> <li>• <b>Adjustment overhead-accounting mode</b>—Configured shaping mode: <b>frame</b> or <b>cell</b>.</li> <li>• <b>Adjustment overhead bytes</b>—Number of bytes that the ANCP agent adds to or subtracts from the actual downstream frame overhead before reporting the adjusted values to CoS.</li> <li>• <b>Adjustment target</b>—Level of shaping-rate adjustment performed: <b>node</b> or <b>queue</b>.</li> <li>• <b>Adjustment multicast index</b>—</li> </ul>

## Sample Output

### show class-of-service interface (Physical)

```

user@host> show class-of-service interface so-0/2/3
Physical interface: so-0/2/3, Index: 135
Maximum usable queues: 8, Queues in use: 4
Total non-default queues created: 4
  Scheduler map: <default>, Index: 2032638653

Logical interface: fe-0/0/1.0, Index: 68, Dedicated Queues: no

```

```

Shaping rate: 32000
Object      Name      Type      Index
Scheduler-map <default>
Rewrite     exp-default exp        21
Classifier   exp-default exp         5
Classifier   ipprec-compatibility ip          8
Forwarding-class-map exp-default exp         5

```

### show class-of-service interface (Logical)

```

user@host> show class-of-service interface so-0/2/3.0
Logical interface: so-0/2/3.0, Index: 68, Dedicated Queues: no
Shaping rate: 32000
Object      Name      Type      Index
Scheduler-map <default>
Rewrite     exp-default exp        21
Classifier   exp-default exp         5
Classifier   ipprec-compatibility ip          8
Forwarding-class-map exp-default exp         5

```

### show class-of-service interface (Gigabit Ethernet)

```

user@host> show class-of-service interface ge-6/2/0
Physical interface: ge-6/2/0, Index: 175
Maximum usable queues: 4, Queues in use: 4
Scheduler map: <default>, Index: 2
Input scheduler map: <default>, Index: 3
Chassis scheduler map: <default-chassis>, Index: 4

```

### show class-of-service interface (ANCP)

```

user@host> show class-of-service interface pp0.1073741842
Logical interface: pp0.1073741842, Index: 341
Object      Name      Type      Index
Traffic-control-profile TCP-CVLAN Output    12408
Classifier   dscp-ipv6-compatibility dscp-ipv6 9
Classifier   ipprec-compatibility ip        13

Adjusting application: ancp LS-0
Adjustment type: absolute
Configured shaping rate: 4000000
Adjustment value: 11228000
Adjustment overhead-accounting mode: Frame Mode
Adjustment overhead bytes: 50
Adjustment target: node

```

### show class-of-service interface (PPPoE Interface)

```

user@host> show class-of-service interface pp0.1
Logical interface: pp0.1, Index: 85
Object      Name      Type      Index
Traffic-control-profile tcp-pppoe.o.pp0.1 Output    27264446535
Classifier   ipprec-compatibility ip        13

Adjusting application: PPPoE
Adjustment type: absolute
Adjustment value: 5000000
Adjustment overhead-accounting mode: cell
Adjustment target: node

```

**show class-of-service interface (DHCP Interface)**

```

user@host> show class-of-service interface demux0.1
Logical interface: pp0.1, Index: 85
  Object      Name                                     Type      Index
  Traffic-control-profile tcp-dhcp.o.demux0.1      Output    2726446535
  Classifier   ipprec-compatibility ip         13

Adjusting application: DHCP
Adjustment type: absolute
Adjustment value: 5000000
Adjustment overhead-accounting mode: cell
Adjustment target: node

```

**show class-of-service interface (T4000 Routers with Type 5 FPCs)**

```

user@host> show class-of-service interface xe-4/0/0
Physical interface: xe-4/0/0, Index: 153
  Maximum usable queues: 8, Queues in use: 4
  Shaping rate: 5000000000 bps
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

Logical interface: xe-4/0/0.0, Index: 77
  Object      Name                                     Type
Index
  Classifier   ipprec-compatibility ip
13

```

**show class-of-service interface detail**

```

user@host> show class-of-service interface ge-0/3/0 detail

Physical interface: ge-0/3/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1518, 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

Physical interface: ge-0/3/0, Index: 138
  Maximum usable queues: 4, Queues in use: 5
  Shaping rate: 50000 bps
  Scheduler map: interface-scheduler-map, Index: 58414
  Input shaping rate: 10000 bps
  Input scheduler map: scheduler-map, Index: 15103
  Chassis scheduler map: <default-chassis>, Index: 4
  Congestion-notification: Disabled

Logical interface ge-0/3/0.0
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2
  inet
  mpls
Interface      Admin Link Proto Input Filter      Output Filter
ge-0/3/0.0     up    up    inet
               mpls
Interface      Admin Link Proto Input Policer      Output Policer
ge-0/3/0.0     up    up    inet
               mpls

```

```

Logical interface: ge-0/3/0.0, Index: 68
Object      Name      Type      Index
Rewrite     exp-default  exp (mpls-any)  33
Classifier   exp-default  exp           10
Classifier   ipprec-compatibility  ip           13

```

```

Logical interface ge-0/3/0.1
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ] Encapsulation: ENET2
inet
Interface      Admin Link Proto Input Filter      Output Filter
ge-0/3/0.1     up    up    inet
Interface      Admin Link Proto Input Policer      Output Policer
ge-0/3/0.1     up    up    inet

```

```

Logical interface: ge-0/3/0.1, Index: 69
Object      Name      Type      Index
Classifier   ipprec-compatibility  ip           13

```

### show class-of-service interface comprehensive

```

user@host> show class-of-service interface ge-0/3/0 comprehensive
Physical interface: ge-0/3/0, Enabled, Physical link is Up
Interface index: 138, SNMP ifIndex: 601, Generation: 141
Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, BPDU Error: None,
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
CoS queues : 4 supported, 4 maximum usable queues
Schedulers : 256
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:14:f6:f4:b4:5d, Hardware address: 00:14:f6:f4:b4:5d
Last flapped : 2010-09-07 06:35:22 PDT (15:14:42 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
IPv6 total statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
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
Label-switched interface (LSI) traffic statistics:
Input bytes : 0 0 bps
Input 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: 5, 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, 5 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 af3                0                0                0
  1 af2                0                0                0
  2 ef2                0                0                0
  3 ef1                0                0                0

Egress queues: 4 supported, 5 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 af3                0                0                0
  1 af2                0                0                0
  2 ef2                0                0                0
  3 ef1                0                0                0

Active alarms : None
Active defects : None
MAC statistics:
  Receive          Transmit
  Total octets      0                0
  Total packets     0                0
  Unicast packets   0                0
  Broadcast packets 0                0
  Multicast packets 0                0
  CRC/Align errors  0                0
  FIFO errors       0                0
  MAC control frames 0                0
  MAC pause frames   0                0
  Oversized frames   0
  Jabber frames      0
  Fragment frames    0
  VLAN tagged frames 0
  Code violations     0
Filter statistics:
  Input packet count 0
  Input packet rejects 0
  Input DA rejects    0
  Input SA rejects    0
  Output packet count 0
  Output packet pad count 0
  Output packet error count 0
  CAM destination filters: 0, CAM source filters: 0
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: 0
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit

```

	%	bps	%	usec	
2 ef2	39	19500	0	120	high
none					
Direction : Input					
CoS transmit queue		Bandwidth		Buffer	Priority
Limit					
	%	bps	%	usec	
0 af3	30	3000	45	0	low
none					

Physical interface: ge-0/3/0, Enabled, Physical link is Up

Interface index: 138, SNMP ifIndex: 601

Forwarding classes: 16 supported, 5 in use

Ingress queues: 4 supported, 5 in use

Queue: 0, Forwarding classes: af3

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Tail-dropped packets : Not Available

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

Queue: 1, Forwarding classes: af2

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Tail-dropped packets : Not Available

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

Queue: 2, Forwarding classes: ef2

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Tail-dropped packets : Not Available

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

Queue: 3, Forwarding classes: ef1

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Tail-dropped packets : Not Available

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

Forwarding classes: 16 supported, 5 in use

Egress queues: 4 supported, 5 in use

Queue: 0, Forwarding classes: af3

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

```

Packets          :          0          0 pps
Bytes            :          0          0 bps
Tail-dropped packets : Not Available
RL-dropped packets :          0          0 pps
RL-dropped bytes   :          0          0 bps
RED-dropped packets :          0          0 pps
RED-dropped bytes   :          0          0 bps
Queue: 1, Forwarding classes: af2
  Queued:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
  Transmitted:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets :          0          0 pps
    RL-dropped bytes   :          0          0 bps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
Queue: 2, Forwarding classes: ef2
  Queued:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
  Transmitted:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets :          0          0 pps
    RL-dropped bytes   :          0          0 bps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
Queue: 3, Forwarding classes: ef1
  Queued:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
  Transmitted:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets :          0          0 pps
    RL-dropped bytes   :          0          0 bps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
  Queued:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
  Transmitted:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets : Not Available
    RED-dropped bytes   : Not Available
Queue: 1, Forwarding classes: af2
  Queued:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
  Transmitted:

```

```

Packets          : 0 0 pps
Bytes            : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : Not Available
RED-dropped bytes  : Not Available
Queue: 2, Forwarding classes: ef2
  Queued:
    Packets          : 0 0 pps
    Bytes            : 0 0 bps
  Transmitted:
    Packets          : 0 0 pps
    Bytes            : 0 0 bps
    Tail-dropped packets : 0 0 pps
    RED-dropped packets : Not Available
    RED-dropped bytes  : Not Available
Queue: 3, Forwarding classes: ef1
  Queued:
    Packets          : 108546 0 pps
    Bytes            : 12754752 376 bps
  Transmitted:
    Packets          : 108546 0 pps
    Bytes            : 12754752 376 bps
    Tail-dropped packets : 0 0 pps
    RED-dropped packets : Not Available
    RED-dropped bytes  : Not Available

```

```

Physical interface: ge-0/3/0, Index: 138
Maximum usable queues: 4, Queues in use: 5
Shaping rate: 50000 bps

```

```
Scheduler map: interface-scheduler-map, Index: 58414
```

```

Scheduler: ef2, Forwarding class: ef2, Index: 39155
  Transmit rate: 39 percent, Rate Limit: none, Buffer size: 120 us, Buffer
  Limit: none, Priority: high
  Excess Priority: unspecified
  Drop profiles:
    Loss priority  Protocol  Index  Name
    Low           any       1      < default-drop-profile>
    Medium low    any       1      < default-drop-profile>
    Medium high   any       1      < default-drop-profile>
    High          any       1      < default-drop-profile>
  Drop profile: < default-drop-profile>, Type: discrete, Index: 1
    Fill level  Drop probability
    100         100
  Drop profile: < default-drop-profile>, Type: discrete, Index: 1
    Fill level  Drop probability
    100         100
  Drop profile: < default-drop-profile>, Type: discrete, Index: 1
    Fill level  Drop probability
    100         100
  Drop profile: < default-drop-profile>, Type: discrete, Index: 1
    Fill level  Drop probability
    100         100
  Input shaping rate: 10000 bps
  Input scheduler map: scheduler-map

```

```
Scheduler map: scheduler-map, Index: 15103
```

```

Scheduler: af3, Forwarding class: af3, Index: 35058
  Transmit rate: 30 percent, Rate Limit: none, Buffer size: 45 percent, Buffer

```



```

Limit: none, Priority: low
Excess Priority: unspecified
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           any       40582  green
  Medium low    any       1      < default-drop-profile>
  Medium high   any       1      < default-drop-profile>
  High          any      18928  yellow
Drop profile: green, Type: discrete, Index: 40582
  Fill level    Drop probability
  50            0
  100           100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
  100           100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
  100           100
Drop profile: yellow, Type: discrete, Index: 18928
  Fill level    Drop probability
  50            0
  100           100
Chassis scheduler map: < default-drop-profile>
Scheduler map: < default-drop-profile>, Index: 4

Scheduler: < default-drop-profile>, Forwarding class: af3, Index: 25
  Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
Limit: none, Priority: low
Excess Priority: low
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           any       1      < default-drop-profile>
  Medium low    any       1      < default-drop-profile>
  Medium high   any       1      < default-drop-profile>
  High          any       1      < default-drop-profile>
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
  100           100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
  100           100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
  100           100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
  100           100

Scheduler: < default-drop-profile>, Forwarding class: af2, Index: 25
  Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
Limit: none, Priority: low
Excess Priority: low
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           any       1      < default-drop-profile>
  Medium low    any       1      < default-drop-profile>
  Medium high   any       1      < default-drop-profile>
  High          any       1      < default-drop-profile>
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
  100           100

```

```

Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level  Drop probability
    100        100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level  Drop probability
    100        100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level  Drop probability
    100        100

Scheduler: < default-drop-profile>, Forwarding class: ef2, Index: 25
  Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
  Limit: none, Priority: low
  Excess Priority: low
  Drop profiles:
    Loss priority  Protocol  Index  Name
    Low           any       1      < default-drop-profile>
    Medium low    any       1      < default-drop-profile>
    Medium high   any       1      < default-drop-profile>
    High          any       1      < default-drop-profile>
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level  Drop probability
    100        100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level  Drop probability
    100        100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level  Drop probability
    100        100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level  Drop probability
    100        100

Scheduler: < default-drop-profile>, Forwarding class: ef1, Index: 25
  Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
  Limit: none, Priority: low
  Excess Priority: low
  Drop profiles:
    Loss priority  Protocol  Index  Name
    Low           any       1      < default-drop-profile>
    Medium low    any       1      < default-drop-profile>
    Medium high   any       1      < default-drop-profile>
    High          any       1      < default-drop-profile>
Drop profile: , Type: discrete, Index: 1
  Fill level  Drop probability
    100        100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level  Drop probability
    100        100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level  Drop probability
    100        100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level  Drop probability
    100        100
  Congestion-notification: Disabled
Forwarding class
priority Policing priority
af3      normal
af2

ID      Queue  Restricted queue  Fabric
0       0       0                low
1       1       1                low

```

ef2	normal	2	2	2	high
ef1	normal	3	3	3	high
af1	normal	4	4	0	low

Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152) (Generation 159)  
 Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] 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

Protocol inet, MTU: 1500, Generation: 172, Route table: 0

Flags: Sendbcst-pkt-to-re

Input Filters: filter-in-ge-0/3/0.0-i,

Policer: Input: p1-ge-0/3/0.0-inet-i

Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 173, Route table: 0

Flags: Is-Primary

Output Filters: exp-filter,,,,,

Logical interface ge-1/2/0.0 (Index 347) (SNMP ifIndex 638) (Generation 156)

Forwarding class ID	Queue	Restricted queue	Fabric priority	Policing priority
SPU priority				
best-effort	0	0	0	low
low				normal

Aggregate Forwarding-class statistics per forwarding-class

Aggregate Forwarding-class statistics:

Forwarding-class statistics:

Forwarding-class best-effort statistics:

Input unicast bytes: 0  
 Output unicast bytes: 0  
 Input unicast packets: 0  
 Output unicast packets: 0

Input multicast bytes: 0  
 Output multicast bytes: 0  
 Input multicast packets: 0  
 Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:

Input unicast bytes: 0  
 Output unicast bytes: 0  
 Input unicast packets: 0  
 Output unicast packets: 0

```
Input multicast bytes:    0
Output multicast bytes:  0
Input multicast packets: 0
Output multicast packets: 0
```

IPv4 protocol forwarding-class statistics:

Forwarding-class statistics:

Forwarding-class best-effort statistics:

```
Input unicast bytes:      0
Output unicast bytes:     0
Input unicast packets:    0
Output unicast packets:   0
```

```
Input multicast bytes:    0
Output multicast bytes:   0
Input multicast packets:  0
Output multicast packets: 0
```

Forwarding-class expedited-forwarding statistics:

```
Input unicast bytes:      0
Output unicast bytes:     0
Input unicast packets:    0
Output unicast packets:   0
```

```
Input multicast bytes:    0
Output multicast bytes:   0
Input multicast packets:  0
Output multicast packets: 0
```

IPv6 protocol forwarding-class statistics:

Forwarding-class statistics:

Forwarding-class best-effort statistics:

```
Input unicast bytes:      0
Output unicast bytes:     0
Input unicast packets:    0
Output unicast packets:   0
```

```
Input multicast bytes:    0
Output multicast bytes:   0
Input multicast packets:  0
Output multicast packets: 0
```

Forwarding-class expedited-forwarding statistics:

```
Input unicast bytes:      0
Output unicast bytes:     0
Input unicast packets:    0
Output unicast packets:   0
```

```
Input multicast bytes:    0
Output multicast bytes:   0
Input multicast packets:  0
Output multicast packets: 0
```

Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152)

Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2

Input packets : 0

Output packets: 0

```

Interface      Admin Link Proto Input Filter      Output Filter
ge-0/3/0.0    up    up    inet  filter-in-ge-0/3/0.0-i
                mpls
Interface      Admin Link Proto Input Policer      Output Policer
ge-0/3/0.0    up    up    inet  p1-ge-0/3/0.0-inet-i
                mpls

```

Filter: filter-in-ge-0/3/0.0-i

Counters:

Name	Bytes	Packets
count-filter-in-ge-0/3/0.0-i	0	0

Filter: exp-filter

Counters:

Name	Bytes	Packets
count-exp-seven-match	0	0
count-exp-zero-match	0	0

Policers:

Name	Packets
p1-ge-0/3/0.0-inet-i	0

Logical interface: ge-0/3/0.0, Index: 68

Object	Name	Type	Index
Rewrite	exp-default	exp (mpls-any)	33

Rewrite rule: exp-default, Code point type: exp, Index: 33

Forwarding class	Loss priority	Code point
af3	low	000
af3	high	001
af2	low	010
af2	high	011
ef2	low	100
ef2	high	101
ef1	low	110
ef1	high	111

Object	Name	Type	Index
Classifier	exp-default	exp	10

Classifier: exp-default, Code point type: exp, Index: 10

Code point	Forwarding class	Loss priority
000	af3	low
001	af3	high
010	af2	low
011	af2	high
100	ef2	low
101	ef2	high
110	ef1	low
111	ef1	high

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13

Code point	Forwarding class	Loss priority
000	af3	low
001	af3	high
010	af3	low

011	af3	high
100	af3	low
101	af3	high
110	ef1	low
111	ef1	high
Forwarding class		
priority	Policing priority	ID Queue Restricted queue Fabric
af3		0 0 0 low
	normal	
af2		1 1 1 low
	normal	
ef2		2 2 2 high
	normal	
ef1		3 3 3 high
	normal	
af1		4 4 0 low
	normal	

Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154) (Generation 160)

Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ] 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

Protocol inet, MTU: 1500, Generation: 174, Route table: 0

Flags: Sendbcst-pkt-to-re

Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154)

Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ] Encapsulation: ENET2

Input packets : 0

Output packets: 0

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/3/0.1	up	up	mpls		
Interface	Admin	Link	Proto	Input Policer	Output Policer
ge-0/3/0.1	up	up	mpls		

Logical interface: ge-0/3/0.1, Index: 69

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13

Code point	Forwarding class	Loss priority
000	af3	low
001	af3	high
010	af3	low
011	af3	high

100	af3	low		
101	af3	high		
110	ef1	low		
111	ef1	high		
Forwarding class	ID	Queue	Restricted queue	Fabric
priority				
af3	0	0	0	low
normal				
af2	1	1	1	low
normal				
ef2	2	2	2	high
normal				
ef1	3	3	3	high
normal				
af1	4	4	0	low
normal				

### show class-of-service interface (ACX Series Routers)

```

user@host-g11# show class-of-service interface
Physical interface: at-0/0/0, Index: 130
Maximum usable queues: 4, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

Logical interface: at-0/0/0.0, Index: 69

Logical interface: at-0/0/0.32767, Index: 70

Physical interface: at-0/0/1, Index: 133
Maximum usable queues: 4, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

Logical interface: at-0/0/1.0, Index: 71

Logical interface: at-0/0/1.32767, Index: 72

Physical interface: ge-0/1/0, Index: 146
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object      Name                Type                Index
Rewrite     dscp-default        dscp                31
Classifier  d1                  dscp                11331
Classifier  ci                   ieee8021p           583

Logical interface: ge-0/1/0.0, Index: 73
Object      Name                Type                Index
Rewrite     custom-exp          exp (mpls-any)      46413

Logical interface: ge-0/1/0.1, Index: 74

Logical interface: ge-0/1/0.32767, Index: 75

Physical interface: ge-0/1/1, Index: 147
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

```

Object Classifier	Name ipprec-compatibility	Type ip	Index 13
Logical interface: ge-0/1/1.0, Index: 76			
Physical interface: ge-0/1/2, Index: 148			
Maximum usable queues: 8, Queues in use: 5			
Scheduler map: <default>, Index: 2			
Congestion-notification: Disabled			
Object Rewrite Classifier	Name ri ci	Type ieee8021p (outer) ieee8021p	Index 35392 583
Physical interface: ge-0/1/3, Index: 149			
Maximum usable queues: 8, Queues in use: 5			
Scheduler map: <default>, Index: 2			
Congestion-notification: Disabled			
Object Classifier	Name ipprec-compatibility	Type ip	Index 13
Logical interface: ge-0/1/3.0, Index: 77			
Object Rewrite	Name custom-exp2	Type exp (mpls-any)	Index 53581
Physical interface: ge-0/1/4, Index: 150			
Maximum usable queues: 8, Queues in use: 5			
Scheduler map: <default>, Index: 2			
Congestion-notification: Disabled			
Object Classifier	Name ipprec-compatibility	Type ip	Index 13
Physical interface: ge-0/1/5, Index: 151			
Maximum usable queues: 8, Queues in use: 5			
Scheduler map: <default>, Index: 2			
Congestion-notification: Disabled			
Object Classifier	Name ipprec-compatibility	Type ip	Index 13
Physical interface: ge-0/1/6, Index: 152			
Maximum usable queues: 8, Queues in use: 5			
Scheduler map: <default>, Index: 2			
Congestion-notification: Disabled			
Object Classifier	Name ipprec-compatibility	Type ip	Index 13
Physical interface: ge-0/1/7, Index: 153			
Maximum usable queues: 8, Queues in use: 5			
Scheduler map: <default>, Index: 2			
Congestion-notification: Disabled			
Object Classifier	Name d1	Type dscp	Index 11331
Physical interface: ge-0/2/0, Index: 154			
Maximum usable queues: 8, Queues in use: 5			
Scheduler map: <default>, Index: 2			
Congestion-notification: Disabled			
Object Classifier	Name ipprec-compatibility	Type ip	Index 13
Physical interface: ge-0/2/1, Index: 155			
Maximum usable queues: 8, Queues in use: 5			



```

    Scheduler map: <default>, Index: 2
    Congestion-notification: Disabled
Object      Name      Type      Index
Classifier  ipprec-compatibility  ip      13

    Logical interface: ge-0/2/1.0, Index: 78

    Logical interface: ge-0/2/1.32767, Index: 79

Physical interface: xe-0/3/0, Index: 156
Maximum usable queues: 8, Queues in use: 5
    Scheduler map: <default>, Index: 2
    Congestion-notification: Disabled
Object      Name      Type      Index
Classifier  ipprec-compatibility  ip      13

    Logical interface: xe-0/3/0.0, Index: 80

Physical interface: xe-0/3/1, Index: 157
Maximum usable queues: 8, Queues in use: 5
    Scheduler map: <default>, Index: 2
    Congestion-notification: Disabled
Object      Name      Type      Index
Classifier  ipprec-compatibility  ip      13

    Logical interface: xe-0/3/1.0, Index: 81

[edit]
user@host-g11#

```

#### show class-of-service interface (PPPoE Subscriber Interface for Enhanced Subscriber Management)

```

user@host> show class-of-service interface pp0.3221225474
    Logical interface: pp0.3221225475, Index: 3221225475
Object      Name      Type      Index
Traffic-control-profile TC_PROF_100_199_SERIES_UID1006 Output  4294967312
Scheduler-map SMAP-1_UID1002 Output  4294967327
Rewrite-Output ieee-rewrite 60432
Rewrite-Output rule1 ip 50463

    Adjusting application: PPPoE IA tags
    Adjustment type: absolute
    Configured shaping rate: 11000000
    Adjustment value: 5000000
    Adjustment target: node

    Adjusting application: ucac
    Adjustment type: delta
    Configured shaping rate: 5000000
    Adjustment value: 100000
    Adjustment target: node

```

## show class-of-service scheduler-map

<b>Syntax</b>	<code>show class-of-service scheduler-map</code> <code>&lt;name&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.
<b>Description</b>	Display the mapping of schedulers to forwarding classes and a summary of scheduler parameters for each entry.
<b>Options</b>	<b>none</b> —Display all scheduler maps.  <b>name</b> —(Optional) Display a summary of scheduler parameters for each forwarding class to which the named scheduler is assigned.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Verifying and Managing Junos OS Enhanced Subscriber Management</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show class-of-service scheduler-map on page 469</a> <a href="#">show class-of-service scheduler-map (QFX Series) on page 470</a>
<b>Output Fields</b>	<a href="#">Table 35 on page 468</a> describes the output fields for the <b>show class-of-service scheduler-map</b> command. Output fields are listed in the approximate order in which they appear.

**Table 35: show class-of-service scheduler-map Output Fields**

Field Name	Field Description
<b>Scheduler map</b>	<p>Name of the scheduler map.</p> <p>(Enhanced subscriber management for MX Series routers) The name of the dynamic scheduler map object is associated with a generated UID (for example, <b>SMAP-1_UID1002</b>) instead of with a subscriber interface.</p>
<b>Index</b>	<p>Index of the indicated object. Objects having indexes in this output include scheduler maps, schedulers, and drop profiles.</p> <p>(Enhanced subscriber management for MX Series routers) Index values for dynamic CoS traffic control profiles are larger for enhanced subscriber management than they are for legacy subscriber management.</p>
<b>Scheduler</b>	Name of the scheduler.

Table 35: show class-of-service scheduler-map Output Fields (*continued*)

Field Name	Field Description
<b>Forwarding class</b>	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.
<b>Transmit rate</b>	Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword <b>remainder</b> , which indicates that the scheduler receives the remaining bandwidth of the interface.
<b>Rate Limit</b>	Rate limiting configuration of the queue. Possible values are <b>none</b> , meaning no rate limiting, and <b>exact</b> , meaning the queue only transmits at the configured rate.
<b>Maximum buffer delay</b>	Amount of transmit delay (in milliseconds) or the buffer size of the queue. The buffer size is shown as a percentage of the total interface buffer allocation, or by the keyword <b>remainder</b> to indicate that the buffer is sized according to what remains after other scheduler buffer allocations.
<b>Priority</b>	Scheduling priority: <b>low</b> or <b>high</b> .
<b>Excess priority</b>	Priority of excess bandwidth: <b>low</b> , <b>medium-low</b> , <b>medium-high</b> , <b>high</b> , or <b>none</b> .
<b>Explicit Congestion Notification</b>	(QFX Series, OCX Series, and EX4600 switches only) Explicit congestion notification (ECN) state: <ul style="list-style-type: none"> <li>• Disable—ECN is disabled on the specified scheduler</li> <li>• Enable—ECN is enabled on the specified scheduler</li> </ul> ECN is disabled by default.
<b>Adjust minimum</b>	Minimum shaping rate for an adjusted queue, in bps.
<b>Adjust percent</b>	Bandwidth adjustment applied to a queue, in percent.
<b>Drop profiles</b>	Table displaying the assignment of drop profiles by name and index to a given loss priority and protocol pair.
<b>Loss priority</b>	Packet loss priority for drop profile assignment.
<b>Protocol</b>	Transport protocol for drop profile assignment.
<b>Name</b>	Name of the drop profile.

## Sample Output

### show class-of-service scheduler-map

```

user@host> show class-of-service scheduler-map
Scheduler map: dd-scheduler-map, Index: 84

Scheduler: aa-scheduler, Index: 8721, Forwarding class: aa-forwarding-class
Transmit rate: 30 percent, Rate Limit: none, Maximum buffer delay: 39 ms,
Priority: high
Drop profiles:
  Loss priority  Protocol  Index  Name

```

Low	non-TCP	8724	aa-drop-profile
Low	TCP	9874	bb-drop-profile
High	non-TCP	8833	cc-drop-profile
High	TCP	8484	dd-drop-profile

Scheduler: bb-scheduler, Forwarding class: aa-forwarding-class  
Transmit rate: 40 percent, Rate limit: none, Maximum buffer delay: 68 ms,  
Priority: high  
Drop profiles:

Loss priority	Protocol	Index	Name
Low	non-TCP	8724	aa-drop-profile
Low	TCP	9874	bb-drop-profile
High	non-TCP	8833	cc-drop-profile
High	TCP	8484	dd-drop-profile

### show class-of-service scheduler-map (QFX Series)

```
user@switch# show class-of-service scheduler-map  
Scheduler map: be-map, Index: 12240
```

Scheduler:be-sched, Forwarding class: best-effort, Index: 115  
Transmit rate: 30 percent, Rate Limit: none, Buffer size: remainder,  
Buffer Limit: none, Priority: low  
Excess Priority: unspecified, Explicit Congestion Notification: disable

Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	3312	lan-dp
Medium-high	any	2714	be-dp1
High	any	3178	be-dp2

## show class-of-service scheduler-hierarchy interface

<b>Syntax</b>	<code>show class-of-service scheduler-hierarchy interface <i>interface-name</i> &lt;detail&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 13.3 for MX Series Routers. Support for up to four hierarchy levels added in Junos OS Release 16.1.
<b>Description</b>	For MPC/MIC interfaces only, display the scheduler hierarchy as well as the shaping rate, guaranteed rate, priorities, and queue weight information for each forwarding class at each hierarchy level.
<b>Options</b>	<b>detail</b> —(Optional) Display scheduler hierarchies based on the interface set.  <b><i>interface-name</i></b> —Display information about a specific interface.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) on page 311</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show class-of-service scheduler-hierarchy interface on page 472</a>
<b>Output Fields</b>	<a href="#">Table 36 on page 471</a> describes the output fields for the <b>show class-of-service scheduler-hierarchy interface</b> command. Output fields are listed in the approximate order in which they appear.

**Table 36: show class-of-service scheduler-hierarchy interface Output Fields**

Field Name	Field Description
<b>interface</b>	Interface name
<b>resource</b>	Traffic resource associated with the logical interface
<b>shaping-rate</b>	Actual shaping rate in bits per second
<b>guaranteed rate</b>	Actual guaranteed rate in bits per second
<b>guaranteed priority</b>	Actual queue priority in the guaranteed region (high, low, or none)
<b>excess priority</b>	Actual queue priority in the excess region (high, low, or none)
<b>queue weight</b>	Actual queue weight for excess CoS weighted round-robin
<b>excess weight</b>	Actual interface unit per priority weights for excess weighted round-robin

## Sample Output

### show class-of-service scheduler-hierarchy interface

```
user@host> show class-of-service scheduler-hierarchy interface xe-1/0/0
```

Interface/ resource name	shaping rate kbits	guaranteed rate kbits	guaranteed/ excess priority		queue weight	excess weight high/low	
xe-1/0/0	12000						
<<< L1							
xe-1/0/0 RTP	12000	0				1	1
best-effort	12000	0	Low	Low	950		
network-control	12000	0	Low	Low	50		
ifset1	12000	0				500	500
<<< L2							
ifset1 RTP	12000	0				1	1
be1	720	0	Low	Low	250		
nc1	12000	0	Low	Low	250		
demux0.96	3000	0				1	1
<<< L3							
demux0.96 RTP	3000	0				500	500
be1	1000	0	Low	Low	250		
nc1	3000	0	Low	Low	250		
pp0.81	2000	0				1	1
<<< L4							
be1	1000	0	Low	Low	250		
nc1	2000	0	Low	Low	250		

## show class-of-service scheduler-hierarchy interface-set

<b>Syntax</b>	show class-of-service scheduler-hierarchy interface-set <i>interface-set-name</i> <detail>
<b>Release Information</b>	Command introduced in Junos OS Release 13.3 for MX Series Routers.
<b>Description</b>	For MPC/MIC interface sets only, display the scheduler hierarchy.
<b>Options</b>	<b>detail</b> —(Optional) Display scheduler hierarchies based on the interface-set. <b>interface-set-name</b> —Display information about a specific interface-set.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show interfaces queue on page 573</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show class-of-service scheduler-hierarchy interface-set on page 474</a>
<b>Output Fields</b>	Table 37 on page 473 describes the output fields for the <b>show class-of-service scheduler-hierarchy interface-set</b> command. Output fields are listed in the approximate order in which they appear.

Table 37: show class-of-service scheduler-hierarchy interface-set Output Fields

Field Name	Field Description
interface	Type of interface
resource	Traffic resource associated with the logical interface
shaping-rate	Actual shaping rate in bits per second
guaranteed rate	Actual guaranteed rate in bits per second
guaranteed priority	Actual queue priority in the guaranteed region (high, low, or none)
excess priority	Actual queue priority in the excess region (high, low, or none)
queue weight	Actual queue weight for excess CoS weighted round-robin
excess weight	Actual interface-set per priority weights for excess weighted round-robin

## Sample Output

### show class-of-service scheduler-hierarchy interface-set

```
user@host> show class-of-service scheduler-hierarchy interface-set ifset
```

Interface/ resource name	shaping rate kbits	guaranteed rate kbits	guaranteed/ excess priority	queue weight	excess weight high/low
ge-1/0/0	100000				
ge-1/0/0 RTP	100000	0			1 1
be	100000	1000	Low Low	1	
da	9000	2000	Medium High	1	
vi	100000	3000	Medium None	626	
vo	100000	4000	High High	373	
gt	100000	0	High High	1	
ge-1/0/0.20	50000	40000			750 750
be	50000	1000	Low Low	1	
da	9000	2000	Medium High	1	
vi	50000	3000	Medium None	626	
vo	50000	4000	High High	373	
gt	50000	Disabled	High High	1	



## show interfaces

<b>List of Syntax</b>	<a href="#">Syntax (Gigabit Ethernet) on page 475</a> <a href="#">Syntax (10 Gigabit Ethernet) on page 475</a> <a href="#">Syntax (SRX Series Devices) on page 475</a>
<b>Syntax (Gigabit Ethernet)</b>	<pre>show interfaces ge-fpc/pic/port &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index snmp-index&gt; &lt;statistics&gt;</pre>
<b>Syntax (10 Gigabit Ethernet)</b>	<pre>show interfaces xe-fpc/pic/port &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index snmp-index&gt; &lt;statistics&gt;</pre>
<b>Syntax (SRX Series Devices)</b>	<pre>show interfaces (   &lt;interface-name&gt;   &lt;brief   detail   extensive   terse&gt;   &lt;controller interface-name&gt;    &lt;descriptions interface-name&gt;    &lt;destination-class (all   destination-class-name logical-interface-name)&gt;    &lt;diagnostics optics interface-name&gt;    &lt;far-end-interval interface-fpc/pic/port&gt;    &lt;filters interface-name&gt;    &lt;flow-statistics interface-name&gt;    &lt;interval interface-name&gt;    &lt;load-balancing (detail   interface-name)&gt;    &lt;mac-database mac-address mac-address&gt;    &lt;mc-ae id identifier unit number revertive-info&gt;    &lt;media interface-name&gt;    &lt;policers interface-name&gt;    &lt;queue both-ingress-egress egress forwarding-class forwarding-class ingress l2-statistics&gt;    &lt;redundancy (detail   interface-name)&gt;    &lt;routing brief detail summary interface-name&gt;    &lt;routing-instance (all   instance-name)&gt;    &lt;snmp-index snmp-index&gt;    &lt;source-class (all   destination-class-name logical-interface-name)&gt;    &lt;statistics interface-name&gt;    &lt;switch-port switch-port number&gt;    &lt;transport pm (all   optics   otn) (all   current   currentday   interval   previousday) (all       interface-name)&gt;    &lt;zone interface-name&gt; )</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4 for Gigabit interfaces.</p> <p>Command introduced in Junos OS Release 8.0 for 10 Gigabit interfaces.</p> <p>Command modified in Junos OS Release 9.5 for SRX Series devices.</p>

Command introduced in Junos OS Release 18.1 for Gigabit interfaces.

**Description** Display status information about the specified Gigabit Ethernet interface.

(M320, M120, MX Series, and T Series routers only) Display status information about the specified 10-Gigabit Ethernet interface.

Display the IPv6 interface traffic statistics about the specified Gigabit Ethernet interface for MX series routers. The input and output bytes (bps) and packets (pps) rates are not displayed for IFD and local traffic.

Display status information and statistics about interfaces on SRX Series appliance running Junos OS.



**NOTE:** On SRX Series appliances, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

**Options** For Gigabit interfaces:

***ge-fpc/pic/port***—Display standard information about the specified Gigabit Ethernet interface.

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

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

**media**—(Optional) Display media-specific information about network interfaces.

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

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

For 10 Gigabit interfaces:

***xe-fpc/pic/port***—Display standard information about the specified 10-Gigabit Ethernet interface.

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

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

**media**—(Optional) Display media-specific information about network interfaces.

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

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

For SRX interfaces:

- **interface-name**—(Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace *pim* with the PIM slot and port with the port number.
  - **at-*pim*/0/*port***—ATM-over-ADSL or ATM-over-SHDSL interface.
  - **ce1-*pim*/0/ *port***—Channelized E1 interface.
  - **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
  - **ct1-*pim*/0/*port***—Channelized T1 interface.
  - **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
  - **e1-*pim*/0/*port***—E1 interface.
  - **e3-*pim*/0/*port***—E3 interface.
  - **fe-*pim*/0/*port***—Fast Ethernet interface.
  - **ge-*pim*/0/*port***—Gigabit Ethernet interface.
  - **se-*pim*/0/*port***—Serial interface.
  - **t1-*pim*/0/*port***—T1 (also called DS1) interface.
  - **t3-*pim*/0/*port***—T3 (also called DS3) interface.
  - **wx-slot/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).
- **interface-name**—(Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace *pim* with the PIM slot and port with the port number.
  - **at-*pim*/0/*port***—ATM-over-ADSL or ATM-over-SHDSL interface.
  - **ce1-*pim*/0/ *port***—Channelized E1 interface.
  - **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
  - **ct1-*pim*/0/*port***—Channelized T1 interface.
  - **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
  - **e1-*pim*/0/*port***—E1 interface.
  - **e3-*pim*/0/*port***—E3 interface.
  - **fe-*pim*/0/*port***—Fast Ethernet interface.
  - **ge-*pim*/0/*port***—Gigabit Ethernet interface.
  - **se-*pim*/0/*port***—Serial interface.
  - **t1-*pim*/0/*port***—T1 (also called DS1) interface.

- **t3-pim/0/port**—T3 (also called DS3) interface.
- **wx-slot/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).

**Additional Information** In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.

**Required Privilege Level** view

**Related Documentation**

- *Understanding Layer 2 Interfaces on Security Devices*
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*
- *Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers*

**List of Sample Output**

- [show interfaces \(Gigabit Ethernet\) on page 515](#)
- [show interfaces \(Gigabit Ethernet on MX Series Routers\) on page 515](#)
- [show interfaces \(link degrade status\) on page 516](#)
- [show interfaces extensive \(Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration\) on page 516](#)
- [show interfaces brief \(Gigabit Ethernet\) on page 517](#)
- [show interfaces detail \(Gigabit Ethernet\) on page 517](#)
- [show interfaces extensive \(Gigabit Ethernet IQ2\) on page 519](#)
- [show interfaces \(Gigabit Ethernet Unnumbered Interface\) on page 522](#)
- [show interfaces \(ACI Interface Set Configured\) on page 522](#)
- [show interfaces \(ALI Interface Set\) on page 522](#)
- [show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, IQ2\) on page 523](#)
- [show interfaces extensive \(10-Gigabit Ethernet, WAN PHY Mode\) on page 525](#)
- [show interfaces extensive \(10-Gigabit Ethernet, DWDM OTN PIC\) on page 527](#)
- [show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode\) on page 529](#)
- [show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only\) on page 530](#)
- [show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only\) on page 531](#)
- [Sample Output SRX Gigabit Ethernet on page 532](#)
- [Sample Output SRX Gigabit Ethernet on page 532](#)
- [show interfaces detail \(Gigabit Ethernet\) on page 533](#)
- [show interfaces statistics st0.0 detail on page 535](#)
- [show interfaces extensive \(Gigabit Ethernet\) on page 536](#)
- [show interfaces terse on page 538](#)
- [show interfaces controller \(Channelized E1 IQ with Logical E1\) on page 539](#)
- [show interfaces controller \(Channelized E1 IQ with Logical DSO\) on page 539](#)
- [show interfaces descriptions on page 539](#)
- [show interfaces destination-class all on page 539](#)

[show interfaces diagnostics optics on page 540](#)  
[show interfaces far-end-interval coc12-5/2/0 on page 540](#)  
[show interfaces far-end-interval coc1-5/2/1:1 on page 541](#)  
[show interfaces filters on page 541](#)  
[show interfaces flow-statistics \(Gigabit Ethernet\) on page 541](#)  
[show interfaces interval \(Channelized OC12\) on page 542](#)  
[show interfaces interval \(E3\) on page 543](#)  
[show interfaces interval \(SONET/SDH\) \(SRX devices\) on page 543](#)  
[show interfaces load-balancing \(SRX devices\) on page 543](#)  
[show interfaces load-balancing detail \(SRX devices\) on page 544](#)  
[show interfaces mac-database \(All MAC Addresses on a Port SRX devices\) on page 544](#)  
[show interfaces mac-database \(All MAC Addresses on a Service SRX devices\) on page 544](#)  
[show interfaces mac-database mac-address on page 545](#)  
[show interfaces mc-ae \(SRX devices\) on page 545](#)  
[show interfaces media \(SONET/SDH\) on page 545](#)  
[show interfaces policers \(SRX devices\) on page 546](#)  
[show interfaces policers interface-name \(SRX devices\) on page 546](#)  
[show interfaces queue \(SRX devices\) on page 546](#)  
[show interfaces redundancy \(SRX devices\) on page 547](#)  
[show interfaces redundancy \(Aggregated Ethernet SRX devices\) on page 548](#)  
[show interfaces redundancy detail \(SRX devices\) on page 548](#)  
[show interfaces routing brief \(SRX devices\) on page 548](#)  
[show interfaces routing detail \(SRX devices\) on page 548](#)  
[show interfaces routing-instance all \(SRX devices\) on page 549](#)  
[show interfaces snmp-index \(SRX devices\) on page 549](#)  
[show interfaces source-class all \(SRX devices\) on page 549](#)  
[show interfaces statistics \(Fast Ethernet SRX devices\) on page 550](#)  
[show interfaces switch-port \(SRX devices\) on page 550](#)  
[show interfaces transport pm \(SRX devices\) on page 551](#)  
[show security zones \(SRX devices\) on page 552](#)

**Output Fields** [Table 38 on page 479](#) describes the output fields for the **show interfaces** (Gigabit Ethernet) command. Output fields are listed in the approximate order in which they appear. For Gigabit Ethernet IQ and IQE PICs, the traffic and MAC statistics vary by interface type. For more information, see [Table 39 on page 507](#).

**Table 38: 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>

Table 38: show interfaces (Gigabit Ethernet) 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>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
<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

Table 38: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<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 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). 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>
<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. 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 <a href="#">show interfaces</a> command.</p>	<b>detail extensive</b>

Table 38: show interfaces (Gigabit 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 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 38: 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 must always be 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 must 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 38: 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 38: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>PCS statistics</b>	<p>(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.</p> <ul style="list-style-type: none"> <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>
<b>Link Degrad</b>	<p>Shows the link degrade status of the physical link and the estimated bit error rates (BERs). This field is available only for the PICs supporting the physical link monitoring feature.</p> <ul style="list-style-type: none"> <li>• <b>Link Monitoring</b>—Indicates if physical link degrade monitoring is enabled on the interface. <ul style="list-style-type: none"> <li>• <b>Enable</b>—Indicates that link degrade monitoring has been enabled (using the <b>link-degrade-monitor</b> statement) on the interface.</li> <li>• <b>Disable</b>—Indicates that link degrade monitoring has not been enabled on the interface. If link degrade monitoring has not been enabled, the output does not show any related information, such as BER values and thresholds.</li> </ul> </li> <li>• <b>Link Degrad Set Threshold</b>—The BER threshold value at which the link is considered degraded and a corrective action is triggered.</li> <li>• <b>Link Degrad Clear Threshold</b>—The BER threshold value at which the degraded link is considered recovered and the corrective action applied to the interface is reverted.</li> <li>• <b>Estimated BER</b>—The estimated bit error rate.</li> <li>• <b>Link-degrade event</b>—Shows link degrade event information. <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Time (in seconds) elapsed after a link degrade event occurred.</li> <li>• <b>Count</b>—The number of link degrade events recorded.</li> <li>• <b>State</b>—Shows the link degrade status (example: <b>Defect Active</b>).</li> </ul> </li> </ul>	<b>detail extensive</b>

Table 38: 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 <a href="#">show interfaces</a> 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 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. <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 <a href="#">show interfaces</a> 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 38: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet may 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 must 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 must 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 must be 0.</li> </ul>	extensive
PMA PHY	<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>	extensive

Table 38: 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 38: 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 38: 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 38: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS information	<p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> <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>	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> <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>	brief detail extensive none

Table 38: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<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</b>	<p>Information displayed for agent circuit identifier (ACI) interface set configured with the <b>agent-circuit-id</b> autoconfiguration stanza.</p> <p><b>Dynamic Profile</b>—Name of the dynamic profile that defines the ACI interface set.</p> <p>If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information.</p> <p><b>NOTE:</b> The ACI VLAN field is replaced with the Line Identity field when an ALI interface set is configured with the <b>line-identity</b> autoconfiguration stanza.</p>	<b>brief detail extensive none</b>
<b>Line Identity</b>	<p>Information displayed for access-line-identifier (ALI) interface sets configured with the <b>line-identity</b> autoconfiguration stanza.</p> <ul style="list-style-type: none"> <li><b>Dynamic Profile</b>—Name of the dynamic profile that defines the ALI interface set.</li> <li>Trusted option used to create the ALI interface set: <b>Circuit-id</b>, <b>Remote-id</b>, or <b>Accept-no-ids</b>. More than one option can be configured.</li> </ul> <p>If configured, the ALI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ALI information.</p> <p><b>NOTE:</b> The Line Identity field is replaced with the ACI VLAN field when an ACI interface set is configured with the <b>agent-circuit-id</b> autoconfiguration stanza.</p>	<b>detail</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 38: 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 38: 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>

The following table describes the output fields for the **show interfaces** (10-Gigabit Ethernet) command.

Field Name	Field Description	Level of Output
<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</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>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
<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.	<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 specified
<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 specified
<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>
<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. 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>	<b>detail extensive</b>

<b>Input errors</b>	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:	<b>extensive</b>
	<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>Output errors</b>	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:	<b>extensive</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 (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>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 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>	<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>
<b>Ingress queues</b>	<p>Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.</p>	<b>extensive</b>
<b>Queue counters (Ingress)</b>	<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>	<b>extensive</b>
<b>Active alarms and Active defects</b>	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value <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>	<b>detail extensive</b> none
<b>OTN alarms</b>	<p>Active OTN alarms identified on the interface.</p>	<b>detail extensive</b>
<b>OTN defects</b>	<p>OTN defects received on the interface.</p>	<b>detail extensive</b>
<b>OTN FEC Mode</b>	<p>The FECmode configured on the interface.</p> <ul style="list-style-type: none"> <li>• <b>efec</b>—Enhanced forward error correction (EFEC) is configured to detect and correct bit errors.</li> <li>• <b>gfec</b>—G.709 Forward error correction (GFEC) mode is configured to detect and correct bit errors.</li> <li>• <b>none</b>—FEC mode is not configured.</li> </ul>	<b>detail extensive</b>



<b>OTN Rate</b>	OTN mode.  <ul style="list-style-type: none"> <li>• <b>fixed-stuff-bytes</b>—Fixed stuff bytes 11.0957 Gbps.</li> <li>• <b>no-fixed-stuff-bytes</b>—No fixed stuff bytes 11.0491 Gbps.</li> <li>• <b>pass-through</b>—Enable OTN passthrough mode.</li> <li>• <b>no-pass-through</b>—Do not enable OTN passthrough mode.</li> </ul>	<b>detail extensive</b>
<b>OTN Line Loopback</b>	Status of the line loopback, if configured for the DWDM OTN PIC. Its value can be: <b>enabled</b> or <b>disabled</b> .	<b>detail extensive</b>
<b>OTN FEC statistics</b>	The forward error correction (FEC) counters for the DWDM OTN PIC.  <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>detail extensive</b>
<b>OTN FEC alarms</b>	OTN FEC excessive or degraded error alarms triggered on the interface.  <ul style="list-style-type: none"> <li>• <b>FEC Degrade</b>—OTU FEC Degrade defect.</li> <li>• <b>FEC Excessive</b>—OTU FEC Excessive Error defect.</li> </ul>	<b>detail extensive</b>
<b>OTN OC</b>	OTN OC defects triggered on the interface.  <ul style="list-style-type: none"> <li>• <b>LOS</b>—OC Loss of Signal defect.</li> <li>• <b>LOF</b>—OC Loss of Frame defect.</li> <li>• <b>LOM</b>—OC Loss of Multiframe defect.</li> <li>• <b>Wavelength Lock</b>—OC Wavelength Lock defect.</li> </ul>	<b>detail extensive</b>
<b>OTN OTU</b>	OTN OTU defects detected on the interface  <ul style="list-style-type: none"> <li>• <b>AIS</b>—OTN AIS alarm.</li> <li>• <b>BDI</b>—OTN OTU BDI alarm.</li> <li>• <b>IAE</b>—OTN OTU IAE alarm.</li> <li>• <b>TTIM</b>—OTN OTU TTIM alarm.</li> <li>• <b>SF</b>—OTN ODU bit error rate fault alarm.</li> <li>• <b>SD</b>—OTN ODU bit error rate defect alarm.</li> <li>• <b>TCA-ES</b>—OTN ODU ES threshold alarm.</li> <li>• <b>TCA-SES</b>—OTN ODU SES threshold alarm.</li> <li>• <b>TCA-UAS</b>—OTN ODU UAS threshold alarm.</li> <li>• <b>TCA-BBE</b>—OTN ODU BBE threshold alarm.</li> <li>• <b>BIP</b>—OTN ODU BIP threshold alarm.</li> <li>• <b>BBE</b>—OTN OTU BBE threshold alarm.</li> <li>• <b>ES</b>—OTN OTU ES threshold alarm.</li> <li>• <b>SES</b>—OTN OTU SES threshold alarm.</li> <li>• <b>UAS</b>—OTN OTU UAS threshold alarm.</li> </ul>	<b>detail extensive</b>
<b>Received DAPI</b>	Destination Access Port Interface (DAPI) from which the packets were received.	<b>detail extensive</b>
<b>Received SAPI</b>	Source Access Port Interface (SAPI) from which the packets were received.	<b>detail extensive</b>
<b>Transmitted DAPI</b>	Destination Access Port Interface (DAPI) to which the packets were transmitted.	<b>detail extensive</b>

<b>Transmitted SAPI</b>	Source Access Port Interface (SAPI) to which the packets were transmitted.	<b>detail extensive</b>
<b>PCS statistics</b>	<p>(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.</p> <ul style="list-style-type: none"> <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>
<b>MAC statistics</b>	<p><b>Receive and Transmit</b> statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> <li>• <b>Total octets and total packets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type.</li> <li>• <b>Unicast packets, Broadcast packets, and 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>	<b>extensive</b>
<b>OTN Received Overhead Bytes</b>	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	<b>extensive</b>
<b>OTN Transmitted Overhead Bytes</b>	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	<b>extensive</b>

Filter statistics	<p><b>Receive and 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>	extensive
PMA PHY	<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>	extensive

<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>

<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 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>
<b>Autonegotiation information</b>	<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>—<b>OK</b> 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>	<b>extensive</b>

<b>Received path trace, Transmitted path trace</b>	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the routing device at the other end of the fiber. The transmitted path trace value is the message that this routing device transmits.	<b>extensive</b>
<b>Packet Forwarding Engine configuration</b>	Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number.</li> </ul>	<b>extensive</b>
<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</b> <b>none</b>
<b>SNMP ifIndex</b>	SNMP interface index number for the logical interface.	<b>detail extensive</b> <b>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>	<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</b> <b>extensive</b> none
<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</b> <b>extensive</b> none
<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</b> <b>extensive</b> none
<b>MTU</b>	Maximum transmission unit size on the logical interface.	<b>detail</b> <b>extensive</b> none
<b>Maximum labels</b>	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	<b>detail</b> <b>extensive</b> none
<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</b> <b>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.  <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.	<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</b> none
<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</b> none
<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</b> none
<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</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><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</b> none
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive</b> none
<b>Local</b>	IP address of the logical interface.	<b>detail extensive</b> none
<b>Broadcast</b>	Broadcast address of the logical interlace.	<b>detail extensive</b> none
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>



For Gigabit Ethernet IQ PICs, traffic and MAC statistics output varies. The following table describes the traffic and MAC statistics for two sample interfaces, each of which is sending traffic in packets of 500 bytes (including 478 bytes for the Layer 3 packet, 18 bytes for the Layer 2 VLAN traffic header, and 4 bytes for cyclic redundancy check [CRC] information). The **ge-0/3/0** interface is the inbound physical interface, and the **ge-0/0/0** interface is the outbound physical interface. On both interfaces, traffic is carried on logical unit **.50** (VLAN 50).

**Table 39: Gigabit and 10 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>	Traffic statistics:  Input bytes: 496 bytes per packet, representing the Layer 2 packet  MAC statistics:  Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes	The additional 4 bytes are for the CRC.
Inbound logical interface	<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	

[Table 40 on page 508](#) lists the output fields for the **show interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 40: 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
<b>Enabled</b>	State of the interface.	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>Link-level type</b>	Encapsulation being used on the physical interface.	All levels
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>MTU</b>	Maximum transmission unit size on the physical interface.	All levels
<b>Link mode</b>	Link mode: Full-duplex or Half-duplex.	
<b>Speed</b>	Speed at which the interface is running.	All levels
<b>BPDU error</b>	Bridge protocol data unit (BPDU) error: Detected or None	
<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>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.	All levels
<b>Interface flags</b>	Information about the interface.	All levels
<b>Link flags</b>	Information about the physical link.	All levels
<b>CoS queues</b>	Number of CoS queues configured.	<b>detail extensive none</b>
<b>Current address</b>	Configured MAC address.	<b>detail extensive none</b>

Table 40: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<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
<b>Output Rate</b>	Output rate in bps and pps.	None
<b>Active alarms and Active defects</b>	<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. 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>	<b>detail extensive none</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>	<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>	<b>detail extensive</b>

Table 40: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Input errors</b>	<p>Input errors on the interface.</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>.</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>
<b>Output errors</b>	<p>Output errors on the interface.</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; therefore, for Gigabit Ethernet PICs, this number must 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 must 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 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>

Table 40: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Ingress queues</b>	Total number of ingress queues supported on the specified interface.	<b>extensive</b>
<b>Queue counters and queue number</b>	CoS queue number and its associated user-configured forwarding class name. <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>
<b>MAC statistics</b>	<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 and total packets</b>—Total number of octets and packets.</li> <li>• <b>Unicast packets, Broadcast packets, and 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 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>	<b>extensive</b>

Table 40: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Filter statistics	<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 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 device (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 must be 0.</li> </ul>	extensive
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> </ul>	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 40: show interfaces 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>Interface transmit statistics</b>	Status of the <b>interface-transmit-statistics</b> configuration: Enabled or Disabled.	<b>detail extensive</b>
<b>Queue counters (Egress)</b>	CoS queue number and its associated user-configured forwarding class name. <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>
<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.	All levels
<b>Encapsulation</b>	Encapsulation on the logical interface.	All levels
<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. 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>

Table 40: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Local statistics</b>	Number and rate of bytes and packets destined to the device.	<b>extensive</b>
<b>Transit statistics</b>	Number and rate of bytes and packets transiting the switch.  <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.	<b>extensive</b>
<b>Security</b>	Security zones that interface belongs to.	<b>extensive</b>
<b>Flow Input statistics</b>	Statistics on packets received by flow module.	<b>extensive</b>
<b>Flow Output statistics</b>	Statistics on packets sent by flow module.	<b>extensive</b>
<b>Flow error statistics (Packets dropped due to)</b>	Statistics on errors in the flow module.	<b>extensive</b>
<b>Protocol</b>	Protocol family.	<b>detail extensive none</b>
<b>MTU</b>	Maximum transmission unit 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 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. .	<b>detail extensive</b>
<b>Addresses, Flags</b>	Information about the address flags..	<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 Gigabit Ethernet

### 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: Sendbcst-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 (link degrade status)

```

user@host> show interfaces et-3/0/0
Physical interface: et-3/0/0, Enabled, Physical link is Down
Interface index: 157, SNMP ifIndex: 537
Link-level type: Ethernet, MTU: 1514, MRU: 0, Speed: 100Gbps, BPDU Error: None,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled
Device flags : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
Link flags : None
CoS queues : 8 supported, 8 maximum usable queues
Current address: 54:e0:32:23:9d:38, Hardware address: 54:e0:32:23:9d:38
Last flapped : 2014-06-18 02:36:38 PDT (02:50:50 ago)
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)
Active alarms : LINK
Active defects : LINK
PCS statistics
  Bit errors : 0
  Errored blocks : 0
Link Degrade* :
Link Monitoring : Enable
Link Degrade Set Threshold: : 1E-7
Link Degrade Clear Threshold: : 1E-12
Estimated BER : 1E-7
Link-degrade event : Seconds Count State
                    782 1 Defect Active

```

### 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

user@host> show interfaces ge-5/2/0.0 statistics detail
Logical interface ge-5/2/0.0 (Index 71) (SNMP ifIndex 573) (Generation 135)
Flags: SNMP-Traps 0x4000 Encapsulation: ENET2

```

```

Egress account overhead: 100
Ingress account overhead: 90
Traffic statistics:
  Input bytes :          271524
  Output bytes :        37769598
  Input packets:          3664
  Output packets:       885790
IPv6 transit statistics:
  Input bytes :           0
  Output bytes :       16681118
  Input packets:           0
  Output packets:      362633
Local statistics:
  Input bytes :          271524
  Output bytes :        308560
  Input packets:          3664
  Output packets:         3659
Transit statistics:
  Input bytes :           0                0 bps
  Output bytes :       37461038            0 bps
  Input packets:           0                0 pps
  Output packets:      882131              0 pps
IPv6 transit statistics:
  Input bytes :           0                0 bps
  Output bytes :       16681118            0 bps
  Input packets:           0                0 pps
  Output packets:      362633              0 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:
    Receive      Transmit
Total octets    38910844056      7174605
Total packets   418398473        78903
Unicast packets 408021893366    1026
Broadcast packets      10         12
Multicast packets  418398217       77865
CRC/Align errors      0           0
FIFO errors            0           0
MAC control frames    0           0
MAC pause frames      0           0
Oversized frames      0
Jabber frames          0
Fragment frames        0
VLAN tagged frames    0
Code violations        0 OTN Received Overhead Bytes:
APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58
Payload Type: 0x08
OTN Transmitted Overhead Bytes:
APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
Payload Type: 0x08
Filter statistics:
Input packet count    418398473
Input packet rejects   479
Input DA rejects      479
Input SA rejects       0
Output packet count           78903
Output packet pad count       0
Output packet error count     0
CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link partner:
Link mode: Full-duplex, Flow control: Symmetric/Asymmetric,
Remote fault: OK
Local resolution:

```

```

Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 7
CoS information:
  Direction : Output
  CoS transmit queue
    %      Bandwidth      Buffer      Priority      Limit
    %      bps            %      usec
  0 best-effort      95      950000000      95      0
low  none
  3 network-control  5      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 38 on page 479](#).

#### 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 (ALI Interface Set)

```
user@host> show interfaces ge-1/0/0.10
Logical interface ge-1/0/0.10 (Index 346) (SNMP ifIndex 554) (Generation 155)
  Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.10 ] Encapsulation: ENET2
  Line Identity:
```



```

Dynamic Profile: ali-set-profile
Circuit-id Remote-id Accept-no-ids
PPPoE:
  Dynamic Profile: ali-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

```

## Sample Output Gigabit Ethernet

### show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2)

```

user@host> show interfaces xe-5/0/0 extensive
Physical interface: xe-5/0/0, Enabled, Physical link is Up
  Interface index: 177, SNMP ifIndex: 99, Generation: 178
  Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Speed: 10Gbps, Loopback:
  None, Source filtering: Enabled,
  Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 1024
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:f6, Hardware address: 00:00:5e:00:53:f6
  Last flapped   : Never
  Statistics last cleared: Never
Traffic statistics:
  Input bytes :          6970332384          0 bps
  Output bytes :              0          0 bps
  Input packets:          81050506          0 pps
  Output packets:              0          0 pps
IPv6 transit statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:              0
Ingress traffic statistics at Packet Forwarding Engine:
  Input bytes :          6970299398          0 bps
  Input packets:          81049992          0 pps
  Drop bytes :              0          0 bps
  Drop packets:              0          0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
  L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0, HS link CRC errors: 0,
  MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort          81049992          81049992          0

  1 expedited-fo              0              0          0

```

```

2 assured-forw          0          0          0
3 network-cont          0          0          0

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

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

Active alarms : None
Active defects : None
PCS statistics
  Bit errors           0
  Errored blocks       0
MAC statistics:
  Receive              Transmit
Total octets          6970332384      0
Total packets         81050506       0
Unicast packets       81050000       0
Broadcast packets     506           0
Multicast packets     0             0
CRC/Align errors      0             0
FIFO errors           0             0
MAC control frames    0             0
MAC pause frames      0             0
Oversized frames      0
Jabber frames         0
Fragment frames       0
VLAN tagged frames    0
Code violations        0
Filter statistics:
Input packet count     81050506
Input packet rejects   506
Input DA rejects       0
Input SA rejects       0
Output packet count    0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 5
CoS information:
Direction : Output
CoS transmit queue    Bandwidth      Buffer Priority  Limit
                        %      bps      %      usec
0 best-effort         95    950000000    95      0      low  none
3 network-control      5    50000000     5      0      low  none

Direction : Input
CoS transmit queue    Bandwidth      Buffer Priority  Limit
                        %      bps      %      usec
0 best-effort         95    950000000    95      0      low  none
3 network-control      5    50000000     5      0      low  none

Logical interface xe-5/0/0.0 (Index 71) (SNMP ifIndex 95) (Generation 195)

```

```

Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ] Encapsulation: ENET2
Egress account overhead: 100
Ingress account overhead: 90
Traffic statistics:
  Input bytes : 0
  Output bytes : 46
  Input packets: 0
  Output packets: 1
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 46
  Input packets: 0
  Output packets: 1
Transit statistics:
  Input bytes : 0 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: 253, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 192.0.2/24, Local: 192.0.2.1, Broadcast: 192.0.2.255,
Generation: 265
Protocol multiservice, MTU: Unlimited, Generation: 254, Route table: 0
  Flags: None
  Policer: Input: __default_arp_policer__

```

### show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode)

```

user@host> show interfaces xe-1/0/0 extensive
Physical interface: xe-1/0/0, Enabled, Physical link is Up
Interface index: 141, SNMP ifIndex: 34, Generation: 47
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Loopback: Disabled
WAN-PHY mode
Source filtering: Disabled, Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps 16384
Link flags : None
CoS queues : 4 supported
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:9d, Hardware address: 00:00:5e:00:53:9d
Last flapped : 2005-07-07 11:22:34 PDT (3d 12:28 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  HS Link CRC errors: 0, HS Link FIFO overflows: 0,

```

```

Resource errors: 0
Output errors:
Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0,
Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
Resource errors: 0
Queue counters:
  Queued packets  Transmitted packets  Dropped packets
0 best-effort      0                0                0
1 expedited-fo     0                0                0
2 assured-forw     0                0                0
3 network-cont     0                0                0
Active alarms : LOL, LOS, LBL
Active defects: LOL, LOS, LBL, SEF, AIS-L, AIS-P
PCS statistics
  Seconds  Count
  Bit errors      0      0
  Errored blocks  0      0
MAC statistics:
  Receive  Transmit
Total octets      0      0
Total packets     0      0
Unicast packets   0      0
Broadcast packets 0      0
Multicast packets 0      0
CRC/Align errors  0      0
FIFO errors       0      0
MAC control frames 0      0
MAC pause frames   0      0
Oversized frames   0
Jabber frames      0
Fragment frames    0
VLAN tagged frames 0
Code violations     0
Filter statistics:
Input packet count      0
Input packet rejects    0
Input DA rejects        0
Input SA rejects        0
Output packet count     0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0
PMA PHY:
  Seconds  Count  State
  PLL lock      0      0 OK
  PHY light    63159      1 Light Missing
WIS section:
  BIP-B1      0      0
  SEF      434430      434438 Defect Active
  LOS      434430      1 Defect Active
  LOF      434430      1 Defect Active
  ES-S      434430
  SES-S      434430
  SEFS-S      434430
WIS line:
  BIP-B2      0      0
  REI-L      0      0
  RDI-L      0      0 OK
  AIS-L      434430      1 Defect Active
  BERR-SF     0      0 OK
  BERR-SD     0      0 OK
  ES-L      434430
  SES-L      434430
  UAS-L      434420
  ES-LFE     0

```

```

SES-LFE                0
UAS-LFE                0
WIS path:
BIP-B3                 0          0
REI-P                  0          0
LOP-P                  0          0 OK
AIS-P                  434430      1 Defect Active
RDI-P                  0          0 OK
UNEQ-P                 0          0 OK
PLM-P                  0          0 OK
ES-P                   434430
SES-P                   434430
UAS-P                   434420
ES-PFE                 0
SES-PFE                 0
UAS-PFE                 0
Received path trace:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted path trace: orissa so-1/0/0
6f 72 69 73 73 61 20 73 6f 2d 31 2f 30 2f 30 00 orissa so-1/0/0.
Packet Forwarding Engine configuration:
  Destination slot: 1
CoS information:
  CoS transmit queue      Bandwidth      Buffer      Priority  Limit
                           %             bps        %      bytes
0 best-effort             95      950000000    95         0        low    none
3 network-control         5       50000000     5         0        low    none

```

### show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC)

```

user@host> show interfaces ge-7/0/0 extensive
Physical interface: ge-7/0/0, Enabled, Physical link is Down
Interface index: 143, SNMP ifIndex: 508, Generation: 208
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled
Device flags : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
Link flags : None
Wavelength : 1550.12 nm, Frequency: 193.40 THz
CoS queues : 8 supported, 8 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:72, Hardware address: 00:00:5e:00:53:72
Last flapped : 2011-04-20 15:48:54 PDT (18:39:49 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 2, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

```

```

FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort              0              0              0

  1 expedited-fo            0              0              0

  2 assured-forw            0              0              0

  3 network-cont
Queue number:      Mapped forwarding classes
  0                best-effort
  1                expedited-forwarding
  2                assured-forwarding
  3                network-control
Active alarms  : LINK
Active defects : LINK
MAC statistics:
Total octets      Receive      Transmit
Total packets    0              0
Unicast packets  0              0
Broadcast packets 0              0
Multicast packets 0              0
CRC/Align errors 0              0
FIFO errors       0              0
MAC control frames 0              0
MAC pause frames  0              0
Oversized frames  0
Jabber frames     0
Fragment frames   0
VLAN tagged frames 0
Code violations    0
Total octets      0              0
Total packets    0              0
Unicast packets  0              0
Broadcast packets 0              0
Multicast packets 0              0
CRC/Align errors 0              0
FIFO errors       0              0
MAC control frames 0              0
MAC pause frames  0              0
Oversized frames  0
Jabber frames     0
Fragment frames   0
VLAN tagged frames 0
Code violations    0
OTN alarms       : None
OTN defects      : None
OTN FEC Mode     : GFEC
OTN Rate         : Fixed Stuff Bytes 11.0957Gbps
OTN Line Loopback : Enabled
OTN FEC statistics :
Corrected Errors      0
Corrected Error Ratio ( 0 sec average) 0e-0
OTN FEC alarms:      Seconds      Count  State
FEC Degrade          0           0  OK
FEC Excessive         0           0  OK
OTN OC:              Seconds      Count  State
LOS                   2           1  OK
LOF                   67164       2  Defect Active

```

```

LOM                                67164          71 Defect Active
Wavelength Lock                    0              0 OK
OTN OTU:
AIS                                0              0 OK
BDI                                65919          4814 Defect Active
IAE                                67158          1 Defect Active
TTIM                               7              1 OK
SF                                 67164          2 Defect Active
SD                                 67164          3 Defect Active
TCA-ES                             0              0 OK
TCA-SES                             0              0 OK
TCA-UAS                             80             40 OK
TCA-BBE                             0              0 OK
BIP                                 0              0 OK
BBE                                 0              0 OK
ES                                  0              0 OK
SES                                 0              0 OK
UAS                                587             0 OK
Received DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Received SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
OTN Received Overhead Bytes:
APS/PCC0: 0x02, APS/PCC1: 0x42, APS/PCC2: 0xa2, APS/PCC3: 0x48
Payload Type: 0x03
OTN Transmitted Overhead Bytes:
APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
Payload Type: 0x03
Filter statistics:
Input packet count                  0
Input packet rejects                0
Input DA rejects                    0
Input SA rejects                    0
Output packet count                  0
Output packet pad count              0
Output packet error count            0
CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
Destination slot: 7
CoS information:
Direction : Output
CoS transmit queue                  Bandwidth          Buffer Priority
Limit
0 best-effort                       95      9500000000    95      0      low
none
3 network-control                    5       500000000       5       0      low
none
...

```

### show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode)

```

user@host> show interfaces xe-7/0/0 extensive
Physical interface: xe-7/0/0, Enabled, Physical link is Up
Interface index: 173, SNMP ifIndex: 212, Generation: 174
Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Enabled,

```

```

Loopback: None, Source filtering: Disabled, Flow control: Enabled
Device flags : Present Running
...

```

### show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only)

```

user@host> show interfaces xe-7/0/0-tx extensive
Physical interface: xe-7/0/0-tx, Enabled, Physical link is Up
  Interface index: 176, SNMP ifIndex: 137, Generation: 177
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Tx-Only
  Device flags : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags : None
  CoS queues : 8 supported, 8 maximum usable queues
  Hold-times : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
  Last flapped : 2007-06-01 09:08:19 PDT (3d 02:31 ago)
  Statistics last cleared: Never
Traffic statistics:
  Input bytes : 0 0 bps
  Output bytes : 322891152287160 9627472888 bps
  Input packets: 0 0 pps
  Output packets: 328809727380 1225492 pps
...

Filter statistics:
  Output packet count 328810554250
  Output packet pad count 0
  Output packet error count 0
...

Logical interface xe-7/0/0-tx.0 (Index 73) (SNMP ifIndex 138) (Generation 139)

Flags: SNMP-Traps Encapsulation: ENET2
Egress account overhead: 100
Ingress account overhead: 90
Traffic statistics:
  Input bytes : 0
  Output bytes : 322891152287160
  Input packets: 0
  Output packets: 328809727380
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 322891152287160 9627472888 bps
  Input packets: 0 0 pps
  Output packets: 328809727380 1225492 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0

```



```

Input packets:          0
Output packets:         0
Protocol inet, MTU: 1500, Generation: 147, Route table: 0
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.11.12/24, Local: 10.11.12.13, Broadcast: 10.11.12.255,
Generation: 141
Protocol multiservice, MTU: Unlimited, Generation: 148, Route table: 0
Flags: None
Policer: Input: __default_arp_policer__

```

### show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only)

```

user@host> show interfaces xe-7/0/0-rx extensive
Physical interface: xe-7/0/0-rx, Enabled, Physical link is Up
Interface index: 174, SNMP ifIndex: 118, Generation: 175
Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Rx-Only
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags     : None
CoS queues    : 8 supported, 8 maximum usable queues
Hold-times    : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
Last flapped  : 2007-06-01 09:08:22 PDT (3d 02:31 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes :      322857456303482      9627496104 bps
Output bytes :              0          0 bps
Input packets:      328775413751      1225495 pps
Output packets:              0          0 pps

...

Filter statistics:
Input packet count      328775015056
Input packet rejects    1
Input DA rejects        0

...

Logical interface xe-7/0/0-rx.0 (Index 72) (SNMP ifIndex 120) (Generation 138)

Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
Input bytes :      322857456303482
Output bytes :              0
Input packets:      328775413751
Output packets:              0
IPv6 transit statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:              0
Local statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:              0
Transit statistics:
Input bytes :      322857456303482      9627496104 bps
Output bytes :              0          0 bps

```

```
Input packets:          328775413751          1225495 pps
Output packets:          0                    0 pps
IPv6 transit statistics:
  Input bytes :          0
  Output bytes :          0
  Input packets:          0
  Output packets:          0
Protocol inet, MTU: 1500, Generation: 145, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 192.0.2/24, Local: 192.0.2.1, Broadcast: 192.0.2.255,
Generation: 139
Protocol multiservice, MTU: Unlimited, Generation: 146, Route table: 0
  Flags: None
  Policer: Input: __default_arp_policer__
```

## Sample Output

### Sample Output SRX Gigabit Ethernet

```
user@host> show interfaces ge-0/0/1
Physical interface: ge-0/0/1, Enabled, Physical link is Down
Interface index: 135, SNMP ifIndex: 510
Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,

BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
Active alarms  : LINK
Active defects : LINK
Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Security: Zone: public
  Protocol inet, MTU: 1500
    Flags: Sendbcst-pkt-to-re
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255
```

### Sample Output SRX Gigabit Ethernet

```
user@host> show interfaces ge-0/0/1
Physical interface: ge-0/0/1, Enabled, Physical link is Down
Interface index: 135, SNMP ifIndex: 510
Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,

BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running Down
```

```

Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Link flags      : None
CoS queues      : 8 supported, 8 maximum usable queues
Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
Last flapped    : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)
Active alarms   : LINK
Active defects  : LINK
Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Security: Zone: public
  Protocol inet, MTU: 1500
    Flags: Sendbcst-pkt-to-re
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255

```

### show interfaces detail (Gigabit Ethernet)

```

user@host> show interfaces ge-0/0/1 detail
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w2d 00:00 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
  Egress queues: 8 supported, 4 in use
  Queue counters:

```

	Queued packets	Transmitted packets	Dropped packets
0 best-effort	0	0	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	0	0	0

```

  Queue number:  Mapped forwarding classes
    0             best-effort
    1             expedited-forwarding
    2             assured-forwarding
    3             network-control
  Active alarms  : LINK
  Active defects : LINK

```

## Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)

Flags: Device-Down SNMP-Traps 0x0 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

Security: Zone: public

## Flow Statistics :

## Flow Input statistics :

Self packets :	0
ICMP packets :	0
VPN packets :	0
Multicast packets :	0
Bytes permitted by policy :	0
Connections established :	0

## Flow Output statistics:

Multicast packets :	0
Bytes permitted by policy :	0

## Flow error statistics (Packets dropped due to):

Address spoofing:	0
Authentication failed:	0
Incoming NAT errors:	0
Invalid zone received packet:	0
Multiple user authentications:	0
Multiple incoming NAT:	0
No parent for a gate:	0
No one interested in self packets:	0
No minor session:	0
No more sessions:	0
No NAT gate:	0
No route present:	0
No SA for incoming SPI:	0
No tunnel found:	0
No session for a gate:	0
No zone or NULL zone binding	0
Policy denied:	0
Security association not active:	0
TCP sequence number out of window:	0
Syn-attack protection:	0
User authentication errors:	0

Protocol inet, MTU: 1500, Generation: 150, Route table: 0

Flags: Sendbroadcast-pkt-to-re

Addresses, Flags: Dest-route-down Is-Preferred Is-Primary

Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255, Generation:

150

## show interfaces statistics st0.0 detail

```

user@host> show interfaces statistics st0.0 detail
Logical interface st0.0 (Index 71) (SNMP ifIndex 609) (Generation 136)
Flags: Up Point-To-Point SNMP-Traps Encapsulation: Secure-Tunnel
Traffic statistics:
  Input bytes :      528152756774
  Output bytes :     575950643520
  Input packets:    11481581669
  Output packets:   12520666095
Local statistics:
  Input bytes :      0
  Output bytes :      0
  Input packets:     0
  Output packets:    0
Transit statistics:
  Input bytes :      0          121859888 bps
  Output bytes :     0          128104112 bps
  Input packets:     0          331141 pps
  Output packets:    0          348108 pps
Security: Zone: untrust
Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
ospf ospf3 pgm pim rip ripng router-discovery rsvp
sap vrrp
Flow Statistics :
Flow Input statistics :
  Self packets :      0
  ICMP packets :      0
  VPN packets :      0
  Multicast packets : 0
  Bytes permitted by policy : 525984295844
  Connections established : 7
Flow Output statistics:
  Multicast packets : 0
  Bytes permitted by policy : 576003290222
Flow error statistics (Packets dropped due to):
  Address spoofing:      0
  Authentication failed: 0
  Incoming NAT errors:   0
  Invalid zone received packet: 0
  Multiple user authentications: 0
  Multiple incoming NAT: 0
  No parent for a gate:  0
  No one interested in self packets: 0
  No minor session:      0
  No more sessions:      0
  No NAT gate:           0
  No route present:      2000280
  No SA for incoming SPI: 0
  No tunnel found:       0
  No session for a gate:  0
  No zone or NULL zone binding 0
  Policy denied:         0
  Security association not active: 0
  TCP sequence number out of window: 0
  Syn-attack protection: 0
  User authentication errors: 0
Protocol inet, MTU: 9192
Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0,
NH drop cnt: 0
Generation: 155, Route table: 0

```

Flags: Sendbcst-pkt-to-re

### show interfaces extensive (Gigabit Ethernet)

```

user@host> show interfaces ge-0/0/1.0 extensive
Physical interface: ge-0/0/1, Enabled, Physical link is Down
Interface index: 135, SNMP ifIndex: 510, Generation: 138
Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,

BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:57 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   :                0                0 bps
Output bytes  :                0                0 bps
Input packets:                0                0 pps
Output packets:              0                0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

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

Queue number:      Mapped forwarding classes
0                  best-effort
1                  expedited-forwarding
2                  assured-forwarding
3                  network-control
Active alarms  : LINK
Active defects : LINK
MAC statistics:
Total octets      Receive      Transmit
Total packets     0            0
Unicast packets   0            0
Broadcast packets 0            0
Multicast packets 0            0
CRC/Align errors  0            0
FIFO errors        0            0
MAC control frames 0            0

```

```

MAC pause frames          0          0
Oversized frames          0
Jabber frames             0
Fragment frames           0
VLAN tagged frames        0
Code violations            0
Filter statistics:
  Input packet count       0
  Input packet rejects     0
  Input DA rejects         0
  Input SA rejects         0
  Output packet count      0
  Output packet pad count  0
  Output packet error count 0
  CAM destination filters: 2, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
      0 best-effort        95      950000000    95      usec      low
none
      3 network-control    5      50000000    5      0      low
none
Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
Flags: Device-Down SNMP-Traps 0x0 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
Security: Zone: public
Flow Statistics :
Flow Input statistics :
  Self packets : 0
  ICMP packets : 0
  VPN packets : 0
  Multicast packets : 0
  Bytes permitted by policy : 0
  Connections established : 0
Flow Output statistics:
  Multicast packets : 0
  Bytes permitted by policy : 0
Flow error statistics (Packets dropped due to):
  Address spoofing: 0

```

```

Authentication failed:          0
Incoming NAT errors:            0
Invalid zone received packet:    0
Multiple user authentications:   0
Multiple incoming NAT:           0
No parent for a gate:           0
No one interested in self packets: 0
No minor session:               0
No more sessions:               0
No NAT gate:                    0
No route present:               0
No SA for incoming SPI:         0
No tunnel found:                0
No session for a gate:          0
No zone or NULL zone binding    0
Policy denied:                  0
Security association not active: 0
TCP sequence number out of window: 0
Syn-attack protection:          0
User authentication errors:      0
Protocol inet, MTU: 1500, Generation: 150, Route table: 0
Flags: Sendbcast-pkt-to-re
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255,
Generation: 150

```

### show interfaces terse

```

user@host> show interfaces terse

```

Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	inet	10.209.4.61/18	
gr-0/0/0	up	up			
ip-0/0/0	up	up			
st0	up	up			
st0.1	up	ready	inet		
ls-0/0/0	up	up			
lt-0/0/0	up	up			
mt-0/0/0	up	up			
pd-0/0/0	up	up			
pe-0/0/0	up	up			
e3-1/0/0	up	up			
t3-2/0/0	up	up			
e1-3/0/0	up	up			
se-4/0/0	up	down			
t1-5/0/0	up	up			
br-6/0/0	up	up			
dc-6/0/0	up	up			
dc-6/0/0.32767	up	up			
bc-6/0/0:1	down	up			
bc-6/0/0:1.0	up	down			
d10	up	up			
d10.0	up	up	inet		
dsc	up	up			
gre	up	up			
ipip	up	up			
lo0	up	up			
lo0.16385	up	up	inet	10.0.0.1 10.0.0.16	--> 0/0 --> 0/0



```

lsi                up    up
mtun               up    up
pimd              up    up
pime              up    up
pp0               up    up

```

### show interfaces controller (Channelized E1 IQ with Logical E1)

```
user@host> show interfaces controller ce1-1/2/6
```

Controller	Admin	Link
ce1-1/2/6	up	up
e1-1/2/6	up	up

### show interfaces controller (Channelized E1 IQ with Logical DSO)

```
user@host> show interfaces controller ce1-1/2/3
```

Controller	Admin	Link
ce1-1/2/3	up	up
ds-1/2/3:1	up	up
ds-1/2/3:2	up	up

### show interfaces descriptions

```
user@host> show interfaces descriptions
```

Interface	Admin	Link	Description
so-1/0/0	up	up	M20-3#1
so-2/0/0	up	up	GSR-12#1
ge-3/0/0	up	up	SMB-OSPF_Area300
so-3/3/0	up	up	GSR-13#1
so-3/3/1	up	up	GSR-13#2
ge-4/0/0	up	up	T320-7#1
ge-5/0/0	up	up	T320-7#2
so-7/1/0	up	up	M160-6#1
ge-8/0/0	up	up	T320-7#3
ge-9/0/0	up	up	T320-7#4
so-10/0/0	up	up	M160-6#2
so-13/0/0	up	up	M20-3#2
so-14/0/0	up	up	GSR-12#2
ge-15/0/0	up	up	SMB-OSPF_Area100
ge-15/0/1	up	up	GSR-13#3

### show interfaces destination-class all

```
user@host> show interfaces destination-class all
```

```
Logical interface so-4/0/0.0
```

Destination class	Packets (packet-per-second)	Bytes (bits-per-second)
gold	0	0
(silver	0)	0)
(	0)	0)

```
Logical interface so-0/1/3.0
```

Destination class	Packets (packet-per-second)	Bytes (bits-per-second)
gold	0	0
(	0)	0)

```

silver                                0                                0
(                                     0) (                               0)

```

### show interfaces diagnostics optics

```

user@host> show interfaces diagnostics optics ge-2/0/0
Physical interface: ge-2/0/0
Laser bias current                : 7.408 mA
Laser output power                : 0.3500 mW / -4.56 dBm
Module temperature                : 23 degrees C / 73 degrees F
Module voltage                    : 3.3450 V
Receiver signal average optical power : 0.0002 mW / -36.99 dBm
Laser bias current high alarm     : Off
Laser bias current low alarm      : Off
Laser bias current high warning   : Off
Laser bias current low warning    : Off
Laser output power high alarm     : Off
Laser output power low alarm      : Off
Laser output power high warning   : Off
Laser output power low warning    : Off
Module temperature high alarm     : Off
Module temperature low alarm      : Off
Module temperature high warning   : Off
Module temperature low warning    : Off
Module voltage high alarm         : Off
Module voltage low alarm          : Off
Module voltage high warning       : Off
Module voltage low warning        : Off
Laser rx power high alarm         : Off
Laser rx power low alarm          : On
Laser rx power high warning       : Off
Laser rx power low warning        : On
Laser bias current high alarm threshold : 17.000 mA
Laser bias current low alarm threshold : 1.000 mA
Laser bias current high warning threshold : 14.000 mA
Laser bias current low warning threshold : 2.000 mA
Laser output power high alarm threshold : 0.6310 mW / -2.00 dBm
Laser output power low alarm threshold : 0.0670 mW / -11.74 dBm
Laser output power high warning threshold : 0.6310 mW / -2.00 dBm
Laser output power low warning threshold : 0.0790 mW / -11.02 dBm
Module temperature high alarm threshold : 95 degrees C / 203 degrees F
Module temperature low alarm threshold : -25 degrees C / -13 degrees F
Module temperature high warning threshold : 90 degrees C / 194 degrees F
Module temperature low warning threshold : -20 degrees C / -4 degrees F
Module voltage high alarm threshold : 3.900 V
Module voltage low alarm threshold : 2.700 V
Module voltage high warning threshold : 3.700 V
Module voltage low warning threshold : 2.900 V
Laser rx power high alarm threshold : 1.2590 mW / 1.00 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 0.7940 mW / -1.00 dBm
Laser rx power low warning threshold : 0.0158 mW / -18.01 dBm

```

### show interfaces far-end-interval coc12-5/2/0

```

user@host> show interfaces far-end-interval coc12-5/2/0
Physical interface: coc12-5/2/0, SNMP ifIndex: 121
05:30-current:
ES-L: 1, SES-L: 1, UAS-L: 0

```

```

05:15-05:30:
  ES-L: 0, SES-L: 0, UAS-L: 0
05:00-05:15:
  ES-L: 0, SES-L: 0, UAS-L: 0
04:45-05:00:
  ES-L: 0, SES-L: 0, UAS-L: 0
04:30-04:45:
  ES-L: 0, SES-L: 0, UAS-L: 0
04:15-04:30:
  ES-L: 0, SES-L: 0, UAS-L: 0
04:00-04:15:
...

```

#### show interfaces far-end-interval coc1-5/2/1:1

```

user@host> run show interfaces far-end-interval coc1-5/2/1:1
Physical interface: coc1-5/2/1:1, SNMP ifIndex: 342
05:30-current:
  ES-L: 1, SES-L: 1, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
05:15-05:30:
  ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
05:00-05:15:
  ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:45-05:00:
  ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:30-04:45:
  ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:15-04:30:
  ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:00-04:15:

```

#### show interfaces filters

```

user@host> show interfaces filters

```

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	inet		
			iso		
ge-5/0/0	up	up			
ge-5/0/0.0	up	up	any		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 flow-statistics (Gigabit Ethernet)

```

user@host> show interfaces flow-statistics ge-0/0/1.0

```

```

Logical interface ge-0/0/1.0 (Index 70) (SNMP ifIndex 49)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 5161
  Output packets: 83
  Security: Zone: zone2
  Allowed host-inbound traffic : bootp bfd bgp dns dvmrp ldp msdp nhrp ospf
pgm
  pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http
https ike
  netconf ping rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
xnm-ssl
  lsping
  Flow Statistics :
  Flow Input statistics :
    Self packets : 0
    ICMP packets : 0
    VPN packets : 2564
    Bytes permitted by policy : 3478
    Connections established : 1
  Flow Output statistics:
    Multicast packets : 0
    Bytes permitted by policy : 16994
  Flow error statistics (Packets dropped due to):
    Address spoofing: 0
    Authentication failed: 0
    Incoming NAT errors: 0
    Invalid zone received packet: 0
    Multiple user authentications: 0
    Multiple incoming NAT: 0
    No parent for a gate: 0
    No one interested in self packets: 0
    No minor session: 0
    No more sessions: 0
    No NAT gate: 0
    No route present: 0
    No SA for incoming SPI: 0
    No tunnel found: 0
    No session for a gate: 0
    No zone or NULL zone binding 0
    Policy denied: 0
    Security association not active: 0
    TCP sequence number out of window: 0
    Syn-attack protection: 0
    User authentication errors: 0
  Protocol inet, MTU: 1500
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255

```

### show interfaces interval (Channelized OC12)

```

user@host> show interfaces interval t3-0/3/0:0
Physical interface: t3-0/3/0:0, SNMP ifIndex: 23
17:43-current:
  LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
  SEFS: 0, UAS: 0
17:28-17:43:
  LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
  SEFS: 0, UAS: 0
17:13-17:28:

```

```

LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
16:58-17:13:
LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
16:43-16:58:
LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
...
Interval Total:
LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
CES: 230, CSES: 230, SEFS: 230, UAS: 238

```

### show interfaces interval (E3)

```

user@host> show interfaces interval e3-0/3/0
Physical interface: e3-0/3/0, SNMP ifIndex: 23
17:43-current:
LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
17:28-17:43:
LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
17:13-17:28:
LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
16:58-17:13:
LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
16:43-16:58:
LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
....
Interval Total:
LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
CES: 230, CSES: 230, SEFS: 230, UAS: 238

```

### show interfaces interval (SONET/SDH) (SRX devices)

```

user@host> show interfaces interval so-0/1/0
Physical interface: so-0/1/0, SNMP ifIndex: 19
20:02-current:
ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
SES-P: 0, UAS-P: 0
19:47-20:02:
ES-S: 267, SES-S: 267, SEFS-S: 267, ES-L: 267, SES-L: 267, UAS-L: 267,
ES-P: 267, SES-P: 267, UAS-P: 267
19:32-19:47:
ES-S: 56, SES-S: 56, SEFS-S: 56, ES-L: 56, SES-L: 56, UAS-L: 46, ES-P: 56,
SES-P: 56, UAS-P: 46
19:17-19:32:
ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
SES-P: 0, UAS-P: 0
19:02-19:17:
.....

```

### show interfaces load-balancing (SRX devices)

```

user@host> show interfaces load-balancing
Interface State      Last change  Member count
ams0      Up              1d 00:50      2
ams1      Up              00:00:59      2

```

**show interfaces load-balancing detail (SRX devices)**

```

user@host>show interfaces load-balancing detail
Load-balancing interfaces detail
Interface      : ams0
State          : Up
Last change    : 1d 00:51
Member count   : 2
Members        :
  Interface    Weight  State
  mams-2/0/0   10     Active
  mams-2/1/0   10     Active

```

**show interfaces mac-database (All MAC Addresses on a Port SRX devices)**

```

user@host> show interfaces mac-database xe-0/3/3
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2

```

MAC address	Input frames	Input bytes	Output frames	Output bytes
00:00:00:00:00:00	1	56	0	0
00:00:c0:01:01:02	7023810	323095260	0	0
00:00:c0:01:01:03	7023810	323095260	0	0
00:00:c0:01:01:04	7023810	323095260	0	0
00:00:c0:01:01:05	7023810	323095260	0	0
00:00:c0:01:01:06	7023810	323095260	0	0
00:00:c0:01:01:07	7023810	323095260	0	0
00:00:c0:01:01:08	7023809	323095214	0	0
00:00:c0:01:01:09	7023809	323095214	0	0
00:00:c0:01:01:0a	7023809	323095214	0	0
00:00:c0:01:01:0b	7023809	323095214	0	0
00:00:c8:01:01:02	30424784	1399540064	37448598	1722635508
00:00:c8:01:01:03	30424784	1399540064	37448598	1722635508
00:00:c8:01:01:04	30424716	1399536936	37448523	1722632058
00:00:c8:01:01:05	30424789	1399540294	37448598	1722635508
00:00:c8:01:01:06	30424788	1399540248	37448597	1722635462
00:00:c8:01:01:07	30424783	1399540018	37448597	1722635462
00:00:c8:01:01:08	30424783	1399540018	37448596	1722635416
00:00:c8:01:01:09	8836796	406492616	8836795	406492570
00:00:c8:01:01:0a	30424712	1399536752	37448521	1722631966
00:00:c8:01:01:0b	30424715	1399536890	37448523	1722632058

```

Number of MAC addresses : 21

```

**show interfaces mac-database (All MAC Addresses on a Service SRX devices)**

```

user@host> show interfaces mac-database xe-0/3/3
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2

```

MAC address	Input frames	Input bytes	Output frames	Output bytes
00:00:00:00:00:00	1	56	0	0

00:00:c0:01:01:02	7023810	323095260	0	0
00:00:c0:01:01:03	7023810	323095260	0	0
00:00:c0:01:01:04	7023810	323095260	0	0
00:00:c0:01:01:05	7023810	323095260	0	0
00:00:c0:01:01:06	7023810	323095260	0	0
00:00:c0:01:01:07	7023810	323095260	0	0
00:00:c0:01:01:08	7023809	323095214	0	0
00:00:c0:01:01:09	7023809	323095214	0	0
00:00:c0:01:01:0a	7023809	323095214	0	0
00:00:c0:01:01:0b	7023809	323095214	0	0
00:00:c8:01:01:02	31016568	1426762128	38040381	1749857526
00:00:c8:01:01:03	31016568	1426762128	38040382	1749857572
00:00:c8:01:01:04	31016499	1426758954	38040306	1749854076
00:00:c8:01:01:05	31016573	1426762358	38040381	1749857526
00:00:c8:01:01:06	31016573	1426762358	38040381	1749857526
00:00:c8:01:01:07	31016567	1426762082	38040380	1749857480
00:00:c8:01:01:08	31016567	1426762082	38040379	1749857434
00:00:c8:01:01:09	9428580	433714680	9428580	433714680
00:00:c8:01:01:0a	31016496	1426758816	38040304	1749853984
00:00:c8:01:01:0b	31016498	1426758908	38040307	1749854122

### show interfaces mac-database mac-address

```

user@host> show interfaces mac-database xe-0/3/3 mac-address (SRX devices)
00:00:c8:01:01:09
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
MAC address: 00:00:c8:01:01:09, Type: Configured,
  Input bytes   : 202324652
  Output bytes  : 202324560
  Input frames  : 4398362
  Output frames : 4398360
Policer statistics:
Policer type    Discarded frames  Discarded bytes
Output aggregate      3992386        183649756

```

### show interfaces mc-ae (SRX devices)

```

user@host> show interfaces mc-ae ae0 unit 512
Member Links   : ae0
Local Status   : active
Peer Status    : active
Logical Interface      : ae0.512
Core Facing Interface : Label Ethernet Interface
ICL-PL          : Label Ethernet Interface

```

### show interfaces media (SONET/SDH)

The following example displays the output fields unique to the **show interfaces media** command for a SONET interface (with no level of output specified):

```

user@host> show interfaces media so-4/1/2
Physical interface: so-4/1/2, Enabled, Physical link is Up
  Interface index: 168, SNMP ifIndex: 495
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1783 (00:00:00 ago), Output: 1786 (00:00:08 ago)
  LCP state: Opened
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Not-configured
  CoS queues      : 8 supported
  Last flapped   : 2005-06-15 12:14:59 PDT (04:31:29 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)
  SONET alarms    : None
  SONET defects   : None
  SONET errors:
    BIP-B1: 121, BIP-B2: 916, REI-L: 0, BIP-B3: 137, REI-P: 16747, BIP-BIP2: 0
  Received path trace: routerb so-1/1/2
  Transmitted path trace: routera so-4/1/2

```

#### show interfaces policers (SRX devices)

```

user@host> show interfaces policers
Interface      Admin Link Proto Input Policer      Output Policer
ge-0/0/0       up    up
ge-0/0/0.0     up    up    inet
               up    up    iso
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
...
so-2/0/0       up    up
so-2/0/0.0     up    up    inet so-2/0/0.0-in-policer so-2/0/0.0-out-policer
               up    up    iso
so-2/1/0       up    down
...

```

#### show interfaces policers interface-name (SRX devices)

```

user@host> show interfaces policers so-2/1/0
Interface      Admin Link Proto Input Policer      Output Policer
so-2/1/0       up    down
so-2/1/0.0     up    down inet so-2/1/0.0-in-policer so-2/1/0.0-out-policer
               up    down iso
               up    down inet6

```

#### show interfaces queue (SRX devices)

The following truncated example shows the CoS queue sizes for queues 0, 1, and 3. Queue 1 has a queue buffer size (guaranteed allocated memory) of 9192 bytes.



```

user@host> show interfaces queue
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 509
  Forwarding classes: 8 supported, 8 in use
  Egress queues: 8 supported, 8 in use
  Queue: 0, Forwarding classes: class0
    Queued:
      Packets      :                0          0 pps
      Bytes       :                0          0 bps
    Transmitted:
      Packets      :                0          0 pps
      Bytes       :                0          0 bps
      Tail-dropped packets :          0          0 pps
      RL-dropped packets  :          0          0 pps
      RL-dropped bytes   :          0          0 bps
      RED-dropped packets :          0          0 pps
        Low             :          0          0 pps
        Medium-low      :          0          0 pps
        Medium-high     :          0          0 pps
        High            :          0          0 pps
      RED-dropped bytes  :          0          0 bps
        Low             :          0          0 bps
        Medium-low      :          0          0 bps
        Medium-high     :          0          0 bps
        High            :          0          0 bps
    Queue Buffer Usage:
      Reserved buffer   :          118750000 bytes
      Queue-depth bytes :
      Current           :                0
  ..
  ..
  Queue: 1, Forwarding classes: class1
  ..
  ..
  Queue Buffer Usage:
    Reserved buffer    :                9192 bytes
    Queue-depth bytes  :
    Current            :                0
  ..
  ..
  Queue: 3, Forwarding classes: class3
    Queued:
    ..
    ..
  Queue Buffer Usage:
    Reserved buffer    :                6250000 bytes
    Queue-depth bytes  :
    Current            :                0
  ..
  ..

```

#### show interfaces redundancy (SRX devices)

```

user@host> show interfaces redundancy
Interface  State      Last change  Primary  Secondary  Current status
rsp0       Not present
rsp1       On secondary  1d 23:56    sp-1/0/0 sp-0/2/0    both down
rsp2       On primary   10:10:27    sp-1/3/0 sp-0/2/0    secondary down
rlsq0      On primary   00:06:24    lsq-0/3/0 lsq-1/0/0    both up

```

**show interfaces redundancy (Aggregated Ethernet SRX devices)**

```
user@host> show interfaces redundancy
Interface State      Last change Primary      Secondary    Current status
r1sq0      On secondary  00:56:12    1sq-4/0/0    1sq-3/0/0    both up

ae0
ae1
ae2
ae3
ae4
```

**show interfaces redundancy detail (SRX devices)**

```
user@host> show interfaces redundancy detail
Interface      : r1sq0
State          : On primary
Last change    : 00:45:47
Primary        : 1sq-0/2/0
Secondary      : 1sq-1/2/0
Current status : both up
Mode           : hot-standby

Interface      : r1sq0:0
State          : On primary
Last change    : 00:45:46
Primary        : 1sq-0/2/0:0
Secondary      : 1sq-1/2/0:0
Current status : both up
Mode           : warm-standby
```

**show interfaces routing brief (SRX devices)**

```
user@host> show interfaces routing brief
Interface      State Addresses
so-5/0/3.0     Down  ISO    enabled
so-5/0/2.0     Up    MPLS   enabled
               ISO    enabled
               INET   192.168.2.120
               INET   enabled
so-5/0/1.0     Up    MPLS   enabled
               ISO    enabled
               INET   192.168.2.130
               INET   enabled
at-1/0/0.3     Up    CCC    enabled
at-1/0/0.2     Up    CCC    enabled
at-1/0/0.0     Up    ISO    enabled
               INET   192.168.90.10
               INET   enabled
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 detail (SRX devices)**

```
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-instance all (SRX devices)

```

user@host> show interfaces terse routing-instance all
Interface  Admin  Link  Proto  Local          Remote Instance
at-0/0/1   up     up    inet   10.0.0.1/24
ge-0/0/0.0 up     up    inet   192.168.4.28/24      sample-a
at-0/1/0.0 up     up    inet6  fe80::a:0:0:4/64     sample-b
so-0/0/0.0 up     up    inet   10.0.0.1/32

```

#### show interfaces snmp-index (SRX devices)

```

user@host> show interfaces snmp-index 33
Physical interface: so-2/1/1, Enabled, Physical link is Down
Interface index: 149, SNMP ifIndex: 33
Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: 0C48,
Loopback: None, FCS: 16, Payload scrambler: Enabled
Device flags   : Present Running Down
Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
Link flags     : Keepalives
CoS queues     : 8 supported
Last flapped   : 2005-06-15 11:45:57 PDT (05:38:43 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
SONET alarms   : LOL, PLL, LOS
SONET defects  : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P

```

#### show interfaces source-class all (SRX devices)

```

user@host> show interfaces source-class all
Logical interface so-0/1/0.0

Source class          Packets          Bytes
                    (packet-per-second) (bits-per-second)
gold                  1928095          161959980

```

```

( 889) ( 597762)
bronze 0 0
( 0) ( 0)
silver 0 0
( 0) ( 0)
Logical interface so-0/1/3.0
Source class Packets Bytes
(packet-per-second) (bits-per-second)
gold 0 0
( 0) ( 0)
bronze 0 0
( 0) ( 0)
silver 116113 9753492
( 939) ( 631616)

```

### show interfaces statistics (Fast Ethernet SRX devices)

```

user@host> show interfaces fe-1/3/1 statistics
Physical interface: fe-1/3/1, Enabled, Physical link is Up
Interface index: 144, SNMP ifIndex: 1042
Description: ford fe-1/3/1
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: 00:90:69:93:04:dc, Hardware address: 00:90:69:93:04:dc
Last flapped : 2006-04-18 03:08:59 PDT (00:01:24 ago)
Statistics last cleared: Never
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)
Input errors: 0, Output errors: 0
Active alarms : None
Active defects : None
Logical interface fe-1/3/1.0 (Index 69) (SNMP ifIndex 50)
Flags: SNMP-Traps Encapsulation: ENET2
Protocol inet, MTU: 1500
Flags: Is-Primary, DCU, SCU-in
Destination class Packets Bytes
(packet-per-second) (bits-per-second)
silver1 0 0
( 0) ( 0)
silver2 0 0
( 0) ( 0)
silver3 0 0
( 0) ( 0)
Addresses, Flags: Is-Default Is-Preferred Is-Primary
Destination: 10.27.245/24, Local: 10.27.245.2,
Broadcast: 10.27.245.255
Protocol iso, MTU: 1497
Flags: Is-Primary

```

### show interfaces switch-port (SRX devices)

```

user@host# show interfaces ge-slot/0/0 switch-port port-number
Port 0, Physical link is Up
Speed: 100mbps, Auto-negotiation: Enabled
Statistics:
Total bytes 28437086 21792250
Total packets 409145 88008

```

```

Unicast packets          9987          83817
Multicast packets        145002         0
Broadcast packets        254156        4191
Multiple collisions       23           10
FIFO/CRC/Align errors    0           0
MAC pause frames         0           0
Oversized frames         0
Runt frames              0
Jabber frames            0
Fragment frames          0
Discarded frames         0
Autonegotiation information:
Negotiation status: Complete
Link partner:
Link mode: Full-duplex, Flow control: None, Remote fault: OK, Link
partner Speed: 100 Mbps
Local resolution:
Flow control: None, Remote fault: Link OK

```

### show interfaces transport pm (SRX devices)

```

user@host> show interfaces transport pm all current et-0/1/0
Physical interface: et-0/1/0, SNMP ifIndex 515
14:45-current Elapse time:900 Seconds
Near End      Suspect Flag:False      Reason:None
PM            COUNT            THRESHOLD      TCA-ENABLED    TCA-RAISED

OTU-BBE       0            800            No             No
OTU-ES        0            135            No             No
OTU-SES       0            90             No             No
OTU-UAS       427          90             No             No
Far End      Suspect Flag:True      Reason:Unknown
PM            COUNT            THRESHOLD      TCA-ENABLED    TCA-RAISED

OTU-BBE       0            800            No             No
OTU-ES        0            135            No             No
OTU-SES       0            90             No             No
OTU-UAS       0            90             No             No
Near End      Suspect Flag:False      Reason:None
PM            COUNT            THRESHOLD      TCA-ENABLED    TCA-RAISED

ODU-BBE       0            800            No             No
ODU-ES        0            135            No             No
ODU-SES       0            90             No             No
ODU-UAS       427          90             No             No
Far End      Suspect Flag:True      Reason:Unknown
PM            COUNT            THRESHOLD      TCA-ENABLED    TCA-RAISED

ODU-BBE       0            800            No             No
ODU-ES        0            135            No             No
ODU-SES       0            90             No             No
ODU-UAS       0            90             No             No
FEC           Suspect Flag:False      Reason:None
PM            COUNT            THRESHOLD      TCA-ENABLED    TCA-RAISED

FEC-CorrectedErr 2008544300    0             NA             NA
FEC-UncorrectedWords 0             0             NA             NA
BER           Suspect Flag:False      Reason:None
PM            MIN      MAX      AVG      THRESHOLD    TCA-ENABLED
TCA-RAISED
BER           3.6e-5   5.8e-5   3.6e-5   10.0e-3      No

```

```

Yes
Physical interface: et-0/1/0, SNMP ifIndex 515
14:45-current
Suspect Flag:True          Reason:Object Disabled
PM          CURRENT  MIN      MAX      AVG      THRESHOLD
TCA-ENABLED  TCA-RAISED
(MIN)
(MAX) (MIN) (MAX) (MIN) (MAX)
Lane chromatic dispersion 0 0 0 0 0
0 NA NA NA NA
Lane differential group delay 0 0 0 0 0
0 NA NA NA NA
q Value 120 120 120 120 0
0 NA NA NA NA
SNR 28 28 29 28 0
0 NA NA NA NA
Tx output power(0.01dBm) -5000 -5000 -5000 -5000 -300
-100 No No No No
Rx input power(0.01dBm) -3642 -3665 -3626 -3637 -1800
-500 No No No No
Module temperature(Celsius) 46 46 46 46 -5
75 No No No No
Tx laser bias current(0.1mA) 0 0 0 0 0
0 NA NA NA NA
Rx laser bias current(0.1mA) 1270 1270 1270 1270 0
0 NA NA NA NA
Carrier frequency offset(MHz) -186 -186 -186 -186 -5000
5000 No No No No

```

### show security zones (SRX devices)

```

user@host> show security zones
Functional zone: management
  Description: This is the management zone.
  Policy configurable: No
  Interfaces bound: 1
  Interfaces:
    ge-0/0/0.0
Security zone: Host
  Description: This is the host zone.
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
    fxp0.0
Security zone: abc
  Description: This is the abc zone.
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
    ge-0/0/1.0
Security zone: def
  Description: This is the def zone.
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
    ge-0/0/2.0

```

## show interfaces (PPPoE)

**Syntax** `show interfaces pp0.logical`  
`<brief | detail | extensive | terse>`  
`<descriptions>`  
`<media>`  
`<snmp-index snmp-index>`  
`<statistics>`

**Release Information** Command introduced before Junos OS Release 7.4.

**Description** (M120 routers, M320 routers, and MX Series routers only). Display status information about the PPPoE interface.

**Options** **pp0.logical**—Display standard status information about the PPPoE interface.

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

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

**media**—(Optional) Display media-specific information about PPPoE interfaces.

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

**statistics**—(Optional) Display PPPoE interface statistics.

**Required Privilege Level** view

**List of Sample Output** [show interfaces \(PPPoE\) on page 559](#)  
[show interfaces \(PPPoE over Aggregated Ethernet\) on page 559](#)  
[show interfaces brief \(PPPoE\) on page 560](#)  
[show interfaces detail \(PPPoE\) on page 560](#)  
[show interfaces extensive \(PPPoE on M120 and M320 Routers\) on page 561](#)

**Output Fields** [Table 41 on page 553](#) lists the output fields for the **show interfaces (PPPoE)** command. Output fields are listed in the approximate order in which they appear.

**Table 41: 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>

Table 41: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<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>	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>



Table 41: show interfaces (PPPoE) 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 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 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>	<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 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>	<b>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 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>	<b>extensive</b>

---

#### Logical Interface

---

Table 41: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<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 none</b>
<b>SNMP ifIndex</b>	Logical interface SNMP interface index number.	<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>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>

Table 41: 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 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>	detail extensive
Local statistics	<p>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.</p>	detail extensive
Transit statistics	<p>Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize.</p> <p><b>NOTE:</b> The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p>	detail extensive
Keepalive settings	<p>(PPP and HDLC) Configured settings for keepalives.</p> <ul style="list-style-type: none"> <li>• <b>interval seconds</b>—The time in seconds between successive keepalive requests. The range is 10 seconds through 32,767 seconds, with a default of 10 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 1 through 255, with a default of 3.</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 1 through 255, with a default of 1.</li> </ul>	detail extensive
Keepalive statistics	<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>	detail extensive
Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified

Table 41: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>LCP state</b>	(PPP) Link 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>—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>
<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</b> none
<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>	none <b>detail extensive</b>
<b>Protocol</b>	Protocol family configured on the logical interface.	<b>detail extensive</b> none
<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</b> none
<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</b> none
<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</b> none

Table 41: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<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,
  mp1s: 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, mp1s:
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>	<pre>show interfaces demux0.logical-interface-number &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 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 snmp-index</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>	<p><a href="#">show interfaces demux0 (Demux) on page 569</a></p> <p><a href="#">show interfaces demux0 (PPPoE over Aggregated Ethernet) on page 570</a></p> <p><a href="#">show interfaces demux0 extensive (Targeted Distribution for Aggregated Ethernet Links) on page 571</a></p> <p><a href="#">show interfaces demux0 (ACI Interface Set Configured) on page 571</a></p>
<b>Output Fields</b>	Table 42 on page 563 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 42: show interfaces demux0 (Demux Interfaces) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		

Table 42: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Physical interface	Name of the physical interface.	brief detail extensive none
Interface index	Index number of the physical interface, which reflects its initialization sequence.	brief detail extensive none
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	brief detail extensive none
Physical link	Status of the physical link (Up or Down).	detail extensive none
Admin	Administrative state of the interface (Up or Down).	terse
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
Link	Status of the physical link (Up or Down).	terse
Targeting summary	Status of aggregated Ethernet links that are configured with targeted distribution (primary or backup)	extensive
Bandwidth	Bandwidth allocated to the aggregated Ethernet links that are configured with targeted distribution.	extensive
Proto	Protocol family configured on the interface.	terse
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Type	Type of interface. <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

Table 42: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<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, in milliseconds.	<b>detail extensive</b>
<b>Current address</b>	Configured MAC address.	<b>detail extensive</b>
<b>Hardware address</b>	Hardware MAC address.	<b>detail extensive</b>
<b>Alternate link address</b>	Backup address of the link.	<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 none</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>	<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>	<b>detail extensive</b>

Table 42: 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 42: 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 42: 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 42: 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: Sendbcst-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 queue

**Syntax** show interfaces queue  
 <aggregate | remaining-traffic>  
 <both-ingress-egress>  
 <egress>  
 <forwarding-class *forwarding-class*>  
 <ingress>  
 <interface-name *interface-name*>  
 <l2-statistics>

**Release Information** Command introduced before Junos OS Release 7.4.  
**both-ingress-egress**, **egress**, and **ingress** options introduced in Junos OS Release 7.6.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.  
**l2-statistics** option introduced in Junos OS Release 12.1.  
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Display class-of-service (CoS) queue information for physical interfaces.

**Options** **none**—Show detailed CoS queue statistics for all physical interfaces.

**aggregate**—(Optional) Display the aggregated queuing statistics of all logical interfaces that have traffic-control profiles configured. (Not on the QFX Series.)

**both-ingress-egress**—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics. (Not on the QFX Series.)

**egress**—(Optional) Display egress queue statistics.

**forwarding-class *forwarding-class***—(Optional) Forwarding class name for this queue. Shows detailed CoS statistics for the queue associated with the specified forwarding class.

**ingress**—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics. (Not on the QFX Series.)

**interface-name *interface-name***—(Optional) Show detailed CoS queue statistics for the specified interface.

**l2-statistics**—(Optional) Display Layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles

**remaining-traffic**—(Optional) Display the remaining-traffic queue statistics of all logical interfaces that have traffic-control profiles configured.

### Overhead for Layer 2 Statistics

Transmitted packets and transmitted byte counts are displayed for the Layer 2 level with the addition of encapsulation overheads applied for fragmentation, as shown in [Table 43 on page 574](#). Others counters, such as packets and bytes queued (input) and drop counters, are displayed at the Layer 3 level. In the case of link fragmentation

and interleaving (LFI) for which fragmentation is not applied, corresponding Layer 2 overheads are added, as shown in [Table 43 on page 574](#).

**Table 43: Layer 2 Overhead and Transmitted Packets or Byte Counts**

Protocol	Fragmentation		LFI
	First fragmentation	Second to <i>n</i> fragmentations	
	Bytes	Bytes	
MLPPP (Long)	13	12	8
MLPPP (short)	11	10	8
MLFR (FRF15)	12	10	8
MFR (FRF16)	10	8	-
MCMLPPP(Long)	13	12	-
MCMLPPP(Short)	11	10	-

#### Layer 2 Statistics—Fragmentation Overhead Calculation

MLPPP/MC-MLPPP Overhead details:

=====

Fragment 1:

```

Outer PPP header           : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
Inner PPP header           : 1 byte
HDLC flag and FCS bytes    : 4 bytes

```

Fragments 2 .. n :

```

Outer PPP header           : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
HDLC flag and FCS bytes    : 4 bytes

```

MLFR (FRF15) Overhead details:

=====

Fragment 1:

```

Framereley header         : 2 bytes
Control,NLPID              : 2 bytes
Fragmentaion header       : 2 bytes
Inner proto                : 2 bytes
HDLC flag and FCS         : 4 bytes

```

Fragments 2 ...n :

```

Framereley header         : 2 bytes
Control,NLPID              : 2 bytes
Fragmentaion header       : 2 bytes
HDLC flag and FCS         : 4 bytes

```

```

MFR (FRF16) Overhead details:
=====
Fragment 1:
  Fragmentaion header   : 2 bytes
  Framereelay header    : 2 bytes
  Inner proto          : 2 bytes
  HDLC flag and FCS     : 4 bytes

Fragments 2 ...n :
  Fragmentaion header   : 2 bytes
  Framereelay header    : 2 bytes
  HDLC flag and FCS     : 4 bytes

```

### Overhead with LFI

```

MLPPP(Long & short sequence):
=====
  Outer PPP header      : 4 bytes
  HDLC flag and FCS     : 4 bytes

MLFR (FRF15):
=====
  Framereelay header    : 2 bytes
  Control,NLPID         : 2 bytes
  HDLC flag and FCS     : 4 bytes

```

The following examples show overhead for different cases:

- A 1000-byte packet is sent to a mlppp bundle without any fragmentation. At the Layer 2 level, bytes transmitted is 1013 in 1 packet. This overhead is for MLPPP long sequence encap.
- A 1000-byte packet is sent to a mlppp bundle with a fragment threshold of 250byte. At the Layer 2 level, bytes transmitted is 1061 bytes in 5 packets.
- A 1000-byte LFI packet is sent to an mlppp bundle. At the Layer 2 level, bytes transmitted is 1008 in 1 packet.

**remaining-traffic**—(Optional) Display the queuing statistics of all logical interfaces that do not have traffic-control profiles configured. (Not on the QFX Series.)

**Additional Information** For rate-limited interfaces hosted on Modular Interface Cards (MICs), Modular Port Concentrators (MPCs), or Enhanced Queuing DPCs, rate-limit packet-drop operations occur *before* packets are queued for transmission scheduling. For such interfaces, the statistics for queued traffic do not include the packets that have already been dropped due to rate limiting, and consequently the displayed statistics for queued traffic are the same as the displayed statistics for transmitted traffic.



**NOTE:** For rate-limited interfaces hosted on other types of hardware, rate-limit packet-drop operations occur *after* packets are queued for transmission scheduling. For these other interface types, the statistics for queued traffic include the packets that are later dropped due to rate limiting, and consequently the displayed statistics for queued traffic equals the sum of the statistics for transmitted and rate-limited traffic.

On M Series routers (except for the M320 and M120 routers), this command is valid only for a PIC installed on an enhanced Flexible PIC Concentrator (FPC).

Queue statistics for aggregated interfaces are supported on the M Series and T Series routers only. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can view the statistics for a child interface by using the **show interfaces statistics** command for that child interface.

When you configure tricolor marking on a 10-port 1-Gigabit Ethernet PIC, for queues 6 and 7 only, the output does not display the number of queued bytes and packets, or the number of bytes and packets dropped because of RED. If you do not configure tricolor marking on the interface, these statistics are available for all queues.

For the 4-port Channelized OC12 IQE PIC and 1-port Channelized OC48 IQE PIC, the **Packet Forwarding Engine Chassis Queues** field represents traffic bound for a particular physical interface on the PIC. For all other PICs, the **Packet Forwarding Engine Chassis Queues** field represents the total traffic bound for the PIC.

For Gigabit Ethernet IQ2 PICs, the **show interfaces queue** command output does not display the number of tail-dropped packets. This limitation does not apply to Packet Forwarding Engine chassis queues.

When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (under the **Packet Forwarding Engine Chassis Queues** field) shows the prefragmentation values.

The behavior of the **egress** queues for the **Routing Engine-Generated Traffic** is not same as the configured queue for MLPPP and MFR configurations.

For information about how to configure CoS, see the *Junos OS Network Interfaces Library for Routing Devices*. For related CoS operational mode commands, see the [CLI Explorer](#).

**Required Privilege Level**

view

**List of Sample Output**

[show interfaces queue \(Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC\) on page 582](#)  
[show interfaces queue \(Aggregated Ethernet on a T320 Router\) on page 583](#)  
[show interfaces queue \(Gigabit Ethernet on a T640 Router\) on page 585](#)  
[show interfaces queue aggregate \(Gigabit Ethernet Enhanced DPC\) on page 585](#)  
[show interfaces queue \(Gigabit Ethernet IQ2 PIC\) on page 589](#)

[show interfaces queue both-ingress-egress \(Gigabit Ethernet IQ2 PIC\) on page 592](#)  
[show interfaces queue ingress \(Gigabit Ethernet IQ2 PIC\) on page 594](#)  
[show interfaces queue egress \(Gigabit Ethernet IQ2 PIC\) on page 595](#)  
[show interfaces queue remaining-traffic \(Gigabit Ethernet Enhanced DPC\) on page 597](#)  
[show interfaces queue \(Channelized OC12 IQE Type 3 PIC in SONET Mode\) on page 599](#)  
[show interfaces queue \(QFX Series\) on page 609](#)  
[show interfaces queue l2-statistics \(lsq interface\) on page 610](#)  
[show interfaces queue lsq \(lsq-ifd\) on page 611](#)  
[show interfaces queue \(Aggregated Ethernet on a MX series Router\) on page 612](#)

**Output Fields** Table 44 on page 577 lists the output fields for the **show interfaces queue** command. Output fields are listed in the approximate order in which they appear.

**Table 44: show interfaces queue Output Fields**

Field Name	Field Description
<b>Physical interface</b>	Name of the physical interface.
<b>Enabled</b>	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .
<b>Interface index</b>	Physical interface's index number, which reflects its initialization sequence.
<b>SNMP ifIndex</b>	SNMP index number for the interface.
<b>Forwarding classes supported</b>	Total number of forwarding classes supported on the specified interface.
<b>Forwarding classes in use</b>	Total number of forwarding classes in use on the specified interface.
<b>Ingress queues supported</b>	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues supported on the specified interface.
<b>Ingress queues in use</b>	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues in use on the specified interface.
<b>Output queues supported</b>	Total number of output queues supported on the specified interface.
<b>Output queues in use</b>	Total number of output queues in use on the specified interface.
<b>Egress queues supported</b>	Total number of egress queues supported on the specified interface.
<b>Egress queues in use</b>	Total number of egress queues in use on the specified interface.

Table 44: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
<b>Queue counters (Ingress)</b>	<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> </ul> <p><b>NOTE:</b> This field is not supported on QFX5100, QFX5110, QFX5200, and QFX5210 switches due to hardware limitations.</p> <ul style="list-style-type: none"> <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>Burst size</b>	(Logical interfaces on IQ PICs only) Maximum number of bytes up to which the logical interface can burst. The burst size is based on the shaping rate applied to the interface.
The following output fields are applicable to both interface component and Packet Forwarding component in the <b>show interfaces queue</b> command:	
<b>Queue</b>	Queue number.
<b>Forwarding classes</b>	Forwarding class name.
<b>Queued Packets</b>	<p>Number of packets queued to this queue.</p> <p><b>NOTE:</b> For Gigabit Ethernet IQ2 interfaces, the Queued Packets count is calculated by the Junos OS interpreting one frame buffer as one packet. If the queued packets are very large or very small, the calculation might not be completely accurate for transit traffic. The count is completely accurate for traffic terminated on the router.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see <a href="#">“Additional Information” on page 575</a>.</p> <p><b>NOTE:</b> This field is not supported on QFX5100, QFX5110, QFX5200, and QFX5210 switches due to hardware limitations.</p>
<b>Queued Bytes</b>	<p>Number of bytes queued to this queue. The byte counts vary by interface hardware. For more information, see <a href="#">Table 45 on page 581</a>.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see <a href="#">“Additional Information” on page 575</a>.</p> <p><b>NOTE:</b> This field is not supported on QFX5100, QFX5110, QFX5200, and QFX5210 switches due to hardware limitations.</p>
<b>Transmitted Packets</b>	<p>Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the <b>Packet Forwarding Engine Chassis Queues</b> field) shows the prefragmentation values.</p> <p><b>NOTE:</b> For Layer 2 statistics, see <a href="#">“Overhead for Layer 2 Statistics” on page 573</a></p>



Table 44: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
<b>Transmitted Bytes</b>	<p>Number of bytes transmitted by this queue. The byte counts vary by interface hardware. For more information, see <a href="#">Table 45 on page 581</a>.</p> <p><b>NOTE:</b> On MX Series routers, this number can be inaccurate when you issue the command for a physical interface repeatedly and in quick succession, because the statistics for the child nodes are collected infrequently. Wait ten seconds between successive iterations to avoid this situation.</p> <p><b>NOTE:</b> For Layer 2 statistics, see <a href="#">“Overhead for Layer 2 Statistics” on page 573</a></p>
<b>Tail-dropped packets</b>	<p>Number of packets dropped because of tail drop.</p> <p><b>NOTE:</b> The <b>Tail-dropped packets</b> counter is not supported on the PTX Series Packet Transport Routers.</p>
<b>RL-dropped packets</b>	<p>Number of packets dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see <a href="#">“Additional Information” on page 575</a>.</p> <p><b>NOTE:</b> The <b>RL-dropped packets</b> counter is not supported on the PTX Series Packet Transport Routers, and is omitted from the output.</p>
<b>RL-dropped bytes</b>	<p>Number of bytes dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see <a href="#">“Additional Information” on page 575</a>.</p>
<b>RED-dropped packets</b>	<p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> <li>• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li>• <b>Low, non-TCP</b>—Number of low-loss priority non-TCP packets dropped because of RED.</li> <li>• <b>Low, TCP</b>—Number of low-loss priority TCP packets dropped because of RED.</li> <li>• <b>High, non-TCP</b>—Number of high-loss priority non-TCP packets dropped because of RED.</li> <li>• <b>High, TCP</b>—Number of high-loss priority TCP packets dropped because of RED.</li> </ul> </li> <li>• (MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li>• <b>Low</b>—Number of low-loss priority packets dropped because of RED.</li> <li>• <b>Medium-low</b>—Number of medium-low loss priority packets dropped because of RED.</li> <li>• <b>Medium-high</b>—Number of medium-high loss priority packets dropped because of RED.</li> <li>• <b>High</b>—Number of high-loss priority packets dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>

Table 44: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
RED-dropped bytes	<p>Number of bytes dropped because of RED. The byte counts vary by interface hardware. For more information, see <a href="#">Table 45 on page 581</a>.</p> <ul style="list-style-type: none"> <li>• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li>• <b>Low, non-TCP</b>—Number of low-loss priority non-TCP bytes dropped because of RED.</li> <li>• <b>Low, TCP</b>—Number of low-loss priority TCP bytes dropped because of RED.</li> <li>• <b>High, non-TCP</b>—Number of high-loss priority non-TCP bytes dropped because of RED.</li> <li>• <b>High, TCP</b>—Number of high-loss priority TCP bytes dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>
Queue-depth bytes	Displays queue-depth average, current, peak, and maximum values for RTP queues. Because queue-depth values cannot be aggregated, displays the values for RTP queues regardless of whether <b>aggregate</b> , <b>remaining-traffic</b> , or neither option is selected.
Queue-depth bytes	Displays queue-depth average, current, peak, and maximum values for RTP queues. Because queue-depth values cannot be aggregated, displays the values for RTP queues regardless of whether <b>aggregate</b> , <b>remaining-traffic</b> , or neither option is selected.
Last-packet enqueued	Starting with Junos OS Release 16.1, <b>Last-packet enqueued</b> output field is introduced. If <b>packet-timestamp</b> is enabled for an FPC, shows the day, date, time, and year in the format <i>day-of-the-week month day-date hh:mm:ss yyyy</i> when a packet was enqueued in the CoS queue. When the timestamp is aggregated across all active Packet Forwarding Engines, the latest timestamp for each CoS queue is reported.

Byte counts vary by interface hardware. [Table 45 on page 581](#) shows how the byte counts on the outbound interfaces vary depending on the interface hardware. [Table 45 on page 581](#) is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.

Table 45: Byte Count by Interface Hardware

Interface Hardware	Output Level	Byte Count Includes	Comments
Gigabit Ethernet IQ and IQE PICs	Interface	<p>Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes</p>	<p>The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type.</p> <p>For RED dropped, 6 bytes are added for the source MAC address.</p>
	Packet forwarding component	<p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p>	—
Non-IQ PIC	Interface	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet.</li> <li>• Transmitted: 478 bytes of Layer 3 packet.</li> </ul> <p>T4000 routers with Type 5 FPCs :</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Inter frame Gap.</li> <li>• Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Interframe Gap.</li> </ul> <p>M Series routers:</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet.</li> <li>• Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead.</li> </ul> <p>PTX Series Packet Transport Routers:</p> <ul style="list-style-type: none"> <li>• Queued: The sum of the transmitted bytes and the RED dropped bytes.</li> <li>• Transmitted: Full Layer 2 overhead (including all L2 encapsulation and CRC) + 12 inter-packet gap + 8 for the preamble.</li> <li>• RED dropped: Full Layer 2 overhead (including all L2 encapsulation and CRC) + 12 inter-packet gap + 8 for the preamble (does not include the VLAN header or MPLS pushed bytes).</li> </ul>	<p>The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.</p>

Table 45: Byte Count by Interface Hardware (*continued*)

Interface Hardware	Output Level	Byte Count Includes	Comments
IQ and IQE PICs with a SONET/SDH interface	Interface	<p>Queued: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>Transmitted: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>RED dropped: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p>	The additional 4 bytes are for the Layer 2 Point-to-Point Protocol (PPP) header.
	Packet forwarding component	<p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 486 bytes per packet, representing 478 bytes of Layer 3 packet + 8 bytes</p>	For transmitted packets, the additional 8 bytes includes 4 bytes for the PPP header and 4 bytes for a cookie.
Non-IQ PIC with a SONET/SDH interface	Interface	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 478 bytes of Layer 3 packet.</li> </ul> <p>M Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 483 bytes per packet, representing 478 bytes of Layer 3 packet + 5 bytes</li> <li>RED dropped: 478 bytes per packet, representing 478 bytes of Layer 3 packet</li> </ul>	For transmitted packets, the additional 5 bytes includes 4 bytes for the PPP header and 1 byte for the packet loss priority (PLP).
Interfaces configured with Frame Relay Encapsulation	Interface	The default Frame Relay overhead is 7 bytes. If you configure the Frame Check Sequence (FCS) to 4 bytes, then the overhead increases to 10 bytes.	
1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs	Interface	<p>Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p> <p>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p>	The Layer 2 overhead is 18 bytes for non-VLAN traffic and 22 bytes for VLAN traffic.
4-port 1G IQ2 and IQ2-E PICs	Packet forwarding component	Queued: 478 bytes of Layer 3 packet.	—
8-port 1G IQ2 and IQ2-E PICs		Transmitted: 478 bytes of Layer 3 packet.	

## Sample Output

### show interfaces queue (Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC)

The following example shows queue information for the rate-limited interface ge-4/2/0 on a Gigabit Ethernet MIC in an MPC. For rate-limited queues for interfaces hosted on MICs or MPCs, rate-limit packet drops occur prior to packet output queuing. In the

command output, the nonzero statistics displayed in the **RL-dropped packets** and **RL-dropped bytes** fields quantify the traffic dropped to rate-limit queue 0 output to 10 percent of 1 gigabyte (100 megabits) per second. Because the RL-dropped traffic is not included in the **Queued** statistics, the statistics displayed for queued traffic are the same as the statistics for transmitted traffic.

```
user@host> show interfaces queue ge-4/2/0
Physical interface: ge-4/2/0, Enabled, Physical link is Up
  Interface index: 203, SNMP ifIndex: 1054
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets          :          131300649          141751 pps
    Bytes            :          11287964840        99793248 bps
  Transmitted:
    Packets          :          131300649          141751 pps
    Bytes            :          11287964840        99793248 bps
    Tail-dropped packets :          0          0 pps
    RL-dropped packets :          205050862        602295 pps
    RL-dropped bytes   :          13595326612      327648832 bps
    RED-dropped packets :          0          0 pps
      Low              :          0          0 pps
      Medium-low       :          0          0 pps
      Medium-high      :          0          0 pps
      High             :          0          0 pps
    RED-dropped bytes   :          0          0 bps
      Low              :          0          0 bps
      Medium-low       :          0          0 bps
      Medium-high      :          0          0 bps
      High             :          0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
```

### show interfaces queue (Aggregated Ethernet on a T320 Router)

The following example shows that the aggregated Ethernet interface, **ae1**, has traffic on queues **af1** and **af12**:

```
user@host> show interfaces queue ae1
Physical interface: ae1, Enabled, Physical link is Up
  Interface index: 158, SNMP ifIndex: 33 Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets          :          5          0 pps
    Bytes            :          242          0 bps
  Transmitted:
    Packets          :          5          0 pps
    Bytes            :          242          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
Queue: 1, Forwarding classes: af1
  Queued:
    Packets          :          42603765          595484 pps
```

```

Bytes                :          5453281920          609776496 bps
Transmitted:
Packets              :          42603765           595484 pps
Bytes                :          5453281920          609776496 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 2, Forwarding classes: ef1
Queued:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Transmitted:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 3, Forwarding classes: nc
Queued:
Packets              :              45              0 pps
Bytes                :             3930              0 bps
Transmitted:
Packets              :              45              0 pps
Bytes                :             3930              0 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 4, Forwarding classes: af11
Queued:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Transmitted:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 5, Forwarding classes: ef11
Queued:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Transmitted:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 6, Forwarding classes: af12
Queued:
Packets              :          31296413          437436 pps
Bytes                :          4005940864          447935200 bps
Transmitted:
Packets              :          31296413          437436 pps
Bytes                :          4005940864          447935200 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 7, Forwarding classes: nc2
Queued:
Packets              :              0              0 pps
Bytes                :              0              0 bps

```

```

Transmitted:
Packets      :          0          0 pps
Bytes        :          0          0 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
RED-dropped bytes  :          0          0 bps

```

### show interfaces queue (Gigabit Ethernet on a T640 Router)

```

user@host> show interfaces queue
Physical interface: ge-7/0/1, Enabled, Physical link is Up
Interface index: 150, SNMP ifIndex: 42
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :          13          0 pps
    Bytes        :         622          0 bps
  Transmitted:
    Packets      :          13          0 pps
    Bytes        :         622          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 1, Forwarding classes: af1
  Queued:
    Packets      :      1725947945      372178 pps
    Bytes        :    220921336960    381110432 bps
  Transmitted:
    Packets      :      1725947945      372178 pps
    Bytes        :    220921336960    381110432 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 2, Forwarding classes: ef1
  Queued:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
  Transmitted:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 3, Forwarding classes: nc
  Queued:
    Packets      :          571          0 pps
    Bytes        :         49318        336 bps
  Transmitted:
    Packets      :          571          0 pps
    Bytes        :         49318        336 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps

```

### show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 aggregate

```

```

Physical interface: ge-2/2/9, Enabled, Physical link is Up
Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :          148450735          947295 pps
    Bytes        :          8016344944        409228848 bps
  Transmitted:
    Packets      :          76397439          487512 pps
    Bytes        :         4125461868        210602376 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :          72053285          459783 pps
      Low        :          72053285          459783 pps
      Medium-low :              0              0 pps
      Medium-high :              0              0 pps
      High       :              0              0 pps
    RED-dropped bytes :         3890877444        198626472 bps
      Low        :         3890877444        198626472 bps
      Medium-low :              0              0 bps
      Medium-high :              0              0 bps
      High       :              0              0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps
  Transmitted:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :              0              0 pps
      Low        :              0              0 pps
      Medium-low :              0              0 pps
      Medium-high :              0              0 pps
      High       :              0              0 pps
    RED-dropped bytes :              0              0 bps
      Low        :              0              0 bps
      Medium-low :              0              0 bps
      Medium-high :              0              0 bps
      High       :              0              0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      :          410278257          473940 pps
    Bytes        :         22156199518        204742296 bps
  Transmitted:
    Packets      :          4850003           4033 pps
    Bytes        :         261900162        1742256 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :          405425693          469907 pps
      Low        :          405425693          469907 pps
      Medium-low :              0              0 pps
      Medium-high :              0              0 pps
      High       :              0              0 pps
    RED-dropped bytes :         21892988124        203000040 bps
      Low        :         21892988124        203000040 bps
      Medium-low :              0              0 bps
      Medium-high :              0              0 bps
      High       :              0              0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      :              0              0 pps

```



```

Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : 76605230 485376 pps
Bytes : 5209211400 264044560 bps
Transmitted:
Packets : 76444631 484336 pps
Bytes : 5198235612 263478800 bps
Tail-dropped packets : Not Available
RED-dropped packets : 160475 1040 pps
Low : 160475 1040 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 10912300 565760 bps
Low : 10912300 565760 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 4836136 3912 pps
Bytes : 333402032 2139056 bps
Transmitted:
Packets : 3600866 1459 pps
Bytes : 244858888 793696 bps
Tail-dropped packets : Not Available

```

```

RED-dropped packets :          1225034          2450 pps
  Low                :          1225034          2450 pps
  Medium-low         :              0              0 pps
  Medium-high        :              0              0 pps
  High               :              0              0 pps
RED-dropped bytes   :          83302312        1333072 bps
  Low                :          83302312        1333072 bps
  Medium-low         :              0              0 bps
  Medium-high        :              0              0 bps
  High               :              0              0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets            :              0              0 pps
  Bytes              :              0              0 bps
Transmitted:
  Packets            :              0              0 pps
  Bytes              :              0              0 bps
Tail-dropped packets : Not Available
RED-dropped packets :              0              0 pps
  Low                :              0              0 pps
  Medium-low         :              0              0 pps
  Medium-high        :              0              0 pps
  High               :              0              0 pps
RED-dropped bytes   :              0              0 bps
  Low                :              0              0 bps
  Medium-low         :              0              0 bps
  Medium-high        :              0              0 bps
  High               :              0              0 bps

```

#### Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

```

Queued:
  Packets            :          77059796        486384 pps
  Bytes              :          3544750624      178989576 bps
Transmitted:
  Packets            :          77059797        486381 pps
  Bytes              :          3544750670      178988248 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets :              0              0 pps
  Low                :              0              0 pps
  Medium-low         :              0              0 pps
  Medium-high        :              0              0 pps
  High               :              0              0 pps
RED-dropped bytes   :              0              0 bps
  Low                :              0              0 bps
  Medium-low         :              0              0 bps
  Medium-high        :              0              0 bps
  High               :              0              0 bps

```

Queue: 1, Forwarding classes: expedited-forwarding

```

Queued:
  Packets            :              0              0 pps
  Bytes              :              0              0 bps
Transmitted:
  Packets            :              0              0 pps
  Bytes              :              0              0 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets :              0              0 pps
  Low                :              0              0 pps
  Medium-low         :              0              0 pps
  Medium-high        :              0              0 pps

```

```

      High : 0 0 pps
    RED-dropped bytes : 0 0 bps
      Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
      High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets : 4846580 3934 pps
  Bytes : 222942680 1447768 bps
Transmitted:
  Packets : 4846580 3934 pps
  Bytes : 222942680 1447768 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
  Low : 0 0 pps
  Medium-low : 0 0 pps
  Medium-high : 0 0 pps
  High : 0 0 pps
RED-dropped bytes : 0 0 bps
  Low : 0 0 bps
  Medium-low : 0 0 bps
  Medium-high : 0 0 bps
  High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets : 0 0 pps
  Bytes : 0 0 bps
Transmitted:
  Packets : 0 0 pps
  Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
  Low : 0 0 pps
  Medium-low : 0 0 pps
  Medium-high : 0 0 pps
  High : 0 0 pps
RED-dropped bytes : 0 0 bps
  Low : 0 0 bps
  Medium-low : 0 0 bps
  Medium-high : 0 0 bps
  High : 0 0 bps

```

### show interfaces queue (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-7/1/3
Physical interface: ge-7/1/3, Enabled, Physical link is Up
Interface index: 170, SNMP ifIndex: 70 Forwarding classes: 16 supported, 4 in
use Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets : 418390039 10 pps
  Bytes : 38910269752 7440 bps
Transmitted:
  Packets : 418390039 10 pps
  Bytes : 38910269752 7440 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding

```

```

Queued:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
Transmitted:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :          0          0 pps
  RED-dropped bytes  :          0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
Transmitted:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :          0          0 pps
  RED-dropped bytes  :          0          0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets      :         7055          1 pps
  Bytes       :        451552        512 bps
Transmitted:
  Packets      :         7055          1 pps
  Bytes       :        451552        512 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :          0          0 pps
  RED-dropped bytes  :          0          0 bps
Forwarding classes: 16 supported, 4 in use Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets      :         1031          0 pps
  Bytes       :        143292          0 bps
Transmitted:
  Packets      :         1031          0 pps
  Bytes       :        143292          0 bps
  Tail-dropped packets : Not Available
  RL-dropped packets  :          0          0 pps
  RL-dropped bytes   :          0          0 bps
  RED-dropped packets :          0          0 pps
  RED-dropped bytes  :          0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
Transmitted:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
  Tail-dropped packets : Not Available
  RL-dropped packets  :          0          0 pps
  RL-dropped bytes   :          0          0 bps
  RED-dropped packets :          0          0 pps
  RED-dropped bytes  :          0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps
Transmitted:
  Packets      :          0          0 pps
  Bytes       :          0          0 bps

```

```

Tail-dropped packets : Not Available
RL-dropped packets   :                0                0 pps
RL-dropped bytes     :                0                0 bps
RED-dropped packets   :                0                0 pps
RED-dropped bytes     :                0                0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets              :                77009              11 pps
Bytes                :               6894286             7888 bps
Transmitted:
Packets              :                77009              11 pps
Bytes                :               6894286             7888 bps
Tail-dropped packets : Not Available
RL-dropped packets   :                0                0 pps
RL-dropped bytes     :                0                0 bps
RED-dropped packets   :                0                0 pps
RED-dropped bytes     :                0                0 bps

```

#### Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

```

Queued:
Packets              :                1031                0 pps
Bytes                :               147328                0 bps
Transmitted:
Packets              :                1031                0 pps
Bytes                :               147328                0 bps
Tail-dropped packets :                0                0 pps
RED-dropped packets   :                0                0 pps
Low, non-TCP          :                0                0 pps
Low, TCP              :                0                0 pps
High, non-TCP         :                0                0 pps
High, TCP             :                0                0 pps
RED-dropped bytes     :                0                0 bps
Low, non-TCP          :                0                0 bps
Low, TCP              :                0                0 bps
High, non-TCP         :                0                0 bps
High, TCP             :                0                0 bps

```

Queue: 1, Forwarding classes: expedited-forwarding

```

Queued:
Packets              :                0                0 pps
Bytes                :                0                0 bps
Transmitted:
Packets              :                0                0 pps
Bytes                :                0                0 bps
Tail-dropped packets :                0                0 pps
RED-dropped packets   :                0                0 pps
Low, non-TCP          :                0                0 pps
Low, TCP              :                0                0 pps
High, non-TCP         :                0                0 pps
High, TCP             :                0                0 pps
RED-dropped bytes     :                0                0 bps
Low, non-TCP          :                0                0 bps
Low, TCP              :                0                0 bps
High, non-TCP         :                0                0 bps
High, TCP             :                0                0 bps

```

Queue: 2, Forwarding classes: assured-forwarding

```

Queued:
Packets              :                0                0 pps
Bytes                :                0                0 bps
Transmitted:

```

```

Packets          : 0 0 pps
Bytes            : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
  Low, non-TCP    : 0 0 pps
  Low, TCP        : 0 0 pps
  High, non-TCP   : 0 0 pps
  High, TCP       : 0 0 pps
RED-dropped bytes : 0 0 bps
  Low, non-TCP    : 0 0 bps
  Low, TCP        : 0 0 bps
  High, non-TCP   : 0 0 bps
  High, TCP       : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets          : 94386 12 pps
  Bytes            : 13756799 9568 bps
Transmitted:
  Packets          : 94386 12 pps
  Bytes            : 13756799 9568 bps
  Tail-dropped packets : 0 0 pps
  RED-dropped packets : 0 0 pps
    Low, non-TCP    : 0 0 pps
    Low, TCP        : 0 0 pps
    High, non-TCP   : 0 0 pps
    High, TCP       : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low, non-TCP    : 0 0 bps
    Low, TCP        : 0 0 bps
    High, non-TCP   : 0 0 bps
    High, TCP       : 0 0 bps

```

### show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 both-ingress-egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets          : Not Available
  Bytes            : 0 0 bps
Transmitted:
  Packets          : 254 0 pps
  Bytes            : 16274 0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
  RED-dropped bytes   : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets          : Not Available
  Bytes            : 0 0 bps
Transmitted:
  Packets          : 0 0 pps
  Bytes            : 0 0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
  RED-dropped bytes   : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding

```

```

Queued:
  Packets          : Not Available
  Bytes           :                0                0 bps
Transmitted:
  Packets          :                0                0 pps
  Bytes           :                0                0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                0                0 pps
  RED-dropped bytes  :                0                0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets          : Not Available
  Bytes           :                0                0 bps
Transmitted:
  Packets          :                0                0 pps
  Bytes           :                0                0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                0                0 pps
  RED-dropped bytes  :                0                0 bps
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets          : Not Available
  Bytes           :                0                0 bps
Transmitted:
  Packets          :                3                0 pps
  Bytes           :               126                0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                0                0 pps
  RED-dropped bytes  :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets          : Not Available
  Bytes           :                0                0 bps
Transmitted:
  Packets          :                0                0 pps
  Bytes           :                0                0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                0                0 pps
  RED-dropped bytes  :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets          : Not Available
  Bytes           :                0                0 bps
Transmitted:
  Packets          :                0                0 pps
  Bytes           :                0                0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                0                0 pps
  RED-dropped bytes  :                0                0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets          : Not Available
  Bytes           :                0                0 bps
Transmitted:
  Packets          :                0                0 pps
  Bytes           :                0                0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                0                0 pps
  RED-dropped bytes  :                0                0 bps

```

```

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :      80564692      0 pps
    Bytes        :      3383717100    0 bps
  Transmitted:
    Packets      :      80564692      0 pps
    Bytes        :      3383717100    0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :      80564685      0 pps
    Bytes        :      3383716770    0 bps
  Transmitted:
    Packets      :      80564685      0 pps
    Bytes        :      3383716770    0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      :      0      0 pps
    Bytes        :      0      0 bps
  Transmitted:
    Packets      :      0      0 pps
    Bytes        :      0      0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      :      9397      0 pps
    Bytes        :      3809052      232 bps
  Transmitted:
    Packets      :      9397      0 pps
    Bytes        :      3809052      232 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps

```

### show interfaces queue ingress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 ingress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      : Not Available
    Bytes        :      0      0 bps
  Transmitted:
    Packets      :      288      0 pps
    Bytes        :      18450      0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps

```



```

Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      : Not Available
    Bytes        :                      0          0 bps
  Transmitted:
    Packets      :                      0          0 pps
    Bytes        :                      0          0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                      0          0 pps
    RED-dropped bytes  :                      0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      : Not Available
    Bytes        :                      0          0 bps
  Transmitted:
    Packets      :                      0          0 pps
    Bytes        :                      0          0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                      0          0 pps
    RED-dropped bytes  :                      0          0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      : Not Available
    Bytes        :                      0          0 bps
  Transmitted:
    Packets      :                      0          0 pps
    Bytes        :                      0          0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                      0          0 pps
    RED-dropped bytes  :                      0          0 bps

```

### show interfaces queue egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      : Not Available
    Bytes        :                      0          0 bps
  Transmitted:
    Packets      :                      3          0 pps
    Bytes        :                     126          0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                      0          0 pps
    RED-dropped bytes  :                      0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      : Not Available
    Bytes        :                      0          0 bps
  Transmitted:
    Packets      :                      0          0 pps
    Bytes        :                      0          0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                      0          0 pps
    RED-dropped bytes  :                      0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:

```

```

Packets          : Not Available
Bytes            :                      0          0 bps
Transmitted:
Packets          :                      0          0 pps
Bytes            :                      0          0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets          : Not Available
Bytes            :                      0          0 bps
Transmitted:
Packets          :                      0          0 pps
Bytes            :                      0          0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets          :                      80564692      0 pps
Bytes            :                      3383717100     0 bps
Transmitted:
Packets          :                      80564692      0 pps
Bytes            :                      3383717100     0 bps
Tail-dropped packets :                      0          0 pps
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets          :                      80564685      0 pps
Bytes            :                      3383716770     0 bps
Transmitted:
Packets          :                      80564685      0 pps
Bytes            :                      3383716770     0 bps
Tail-dropped packets :                      0          0 pps
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets          :                      0          0 pps
Bytes            :                      0          0 bps
Transmitted:
Packets          :                      0          0 pps
Bytes            :                      0          0 bps
Tail-dropped packets :                      0          0 pps
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets          :                      9538          0 pps
Bytes            :                      3819840        0 bps
Transmitted:
Packets          :                      9538          0 pps
Bytes            :                      3819840        0 bps
Tail-dropped packets :                      0          0 pps
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps

```

### show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 remaining-traffic
Physical interface: ge-2/2/9, Enabled, Physical link is Up
  Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :          110208969          472875 pps
    Bytes        :          5951284434        204282000 bps
  Transmitted:
    Packets      :          110208969          472875 pps
    Bytes        :          5951284434        204282000 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :
    Low          :              0              0 pps
    Medium-low   :              0              0 pps
    Medium-high  :              0              0 pps
    High         :              0              0 pps
  RED-dropped bytes :
    Low          :              0              0 bps
    Medium-low   :              0              0 bps
    Medium-high  :              0              0 bps
    High         :              0              0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps
  Transmitted:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :
    Low          :              0              0 pps
    Medium-low   :              0              0 pps
    Medium-high  :              0              0 pps
    High         :              0              0 pps
  RED-dropped bytes :
    Low          :              0              0 bps
    Medium-low   :              0              0 bps
    Medium-high  :              0              0 bps
    High         :              0              0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps
  Transmitted:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :
    Low          :              0              0 pps
    Medium-low   :              0              0 pps
    Medium-high  :              0              0 pps
    High         :              0              0 pps
  RED-dropped bytes :
    Low          :              0              0 bps
    Medium-low   :              0              0 bps

```

```

    Medium-high      :          0          0 bps
    High             :          0          0 bps
Queue: 3, Forwarding classes: network-control
Queued:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
Transmitted:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
Tail-dropped packets : Not Available
RED-dropped packets :          0          0 pps
    Low              :          0          0 pps
    Medium-low       :          0          0 pps
    Medium-high      :          0          0 pps
    High             :          0          0 pps
RED-dropped bytes   :          0          0 bps
    Low              :          0          0 bps
    Medium-low       :          0          0 bps
    Medium-high      :          0          0 bps
    High             :          0          0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
    Packets          :      109355853      471736 pps
    Bytes            :      7436199152     256627968 bps
Transmitted:
    Packets          :      109355852      471736 pps
    Bytes            :      7436198640     256627968 bps
Tail-dropped packets : Not Available
RED-dropped packets :          0          0 pps
    Low              :          0          0 pps
    Medium-low       :          0          0 pps
    Medium-high      :          0          0 pps
    High             :          0          0 pps
RED-dropped bytes   :          0          0 bps
    Low              :          0          0 bps
    Medium-low       :          0          0 bps
    Medium-high      :          0          0 bps
    High             :          0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
Transmitted:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
Tail-dropped packets : Not Available
RED-dropped packets :          0          0 pps
    Low              :          0          0 pps
    Medium-low       :          0          0 pps
    Medium-high      :          0          0 pps
    High             :          0          0 pps
RED-dropped bytes   :          0          0 bps
    Low              :          0          0 bps
    Medium-low       :          0          0 bps
    Medium-high      :          0          0 bps
    High             :          0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
    Packets          :          0          0 pps

```

```

      Bytes                :                0                0 bps
Transmitted:
  Packets                :                0                0 pps
  Bytes                  :                0                0 bps
Tail-dropped packets : Not Available
RED-dropped packets :
  Low                    :                0                0 pps
  Medium-low             :                0                0 pps
  Medium-high            :                0                0 pps
  High                   :                0                0 pps
RED-dropped bytes :
  Low                    :                0                0 bps
  Medium-low             :                0                0 bps
  Medium-high            :                0                0 bps
  High                   :                0                0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets                :                0                0 pps
  Bytes                  :                0                0 bps
Transmitted:
  Packets                :                0                0 pps
  Bytes                  :                0                0 bps
Tail-dropped packets : Not Available
RED-dropped packets :
  Low                    :                0                0 pps
  Medium-low             :                0                0 pps
  Medium-high            :                0                0 pps
  High                   :                0                0 pps
RED-dropped bytes :
  Low                    :                0                0 bps
  Medium-low             :                0                0 bps
  Medium-high            :                0                0 bps
  High                   :                0                0 bps

```

#### show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode)

```

user@host> show interfaces queue t3-1/1/0:7
Physical interface: t3-1/1/0:7, Enabled, Physical link is Up

Interface index: 192, SNMP ifIndex: 1948

Description: full T3 interface connect to 6ce13 t3-3/1/0:7 for FR testing -
Lam

Forwarding classes: 16 supported, 9 in use

Egress queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

Queued:

  Packets                :                214886            13449 pps

  Bytes                  :                9884756           5164536 bps

Transmitted:

  Packets                :                214886            13449 pps

```

Bytes	:	9884756	5164536 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 1, Forwarding classes: REALTIME

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 2, Forwarding classes: PRIVATE

## Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

## Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

## Queue: 3, Forwarding classes: CONTROL

## Queued:

Packets	:	60	0 pps
Bytes	:	4560	0 bps

## Transmitted:

Packets	:	60	0 pps
Bytes	:	4560	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps

RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps



RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 6, Forwarding classes: CLASS\_V\_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 7, Forwarding classes: CLASS\_S\_OUTPUT, GETS

Queued:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
-------	---	---	-------

## Transmitted:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
-------	---	---	-------

Tail-dropped packets	:	0	0 pps
----------------------	---	---	-------

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

Low	:	0	0 pps
-----	---	---	-------

Medium-low	:	0	0 pps
------------	---	---	-------

Medium-high	:	0	0 pps
-------------	---	---	-------

High	:	0	0 pps
------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

Low	:	0	0 bps
-----	---	---	-------

Medium-low	:	0	0 bps
------------	---	---	-------

Medium-high	:	0	0 bps
-------------	---	---	-------

High	:	0	0 bps
------	---	---	-------

## Packet Forwarding Engine Chassis Queues:

Queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

## Queued:

Packets	:	371365	23620 pps
---------	---	--------	-----------

Bytes	:	15597330	7936368 bps
-------	---	----------	-------------

## Transmitted:

Packets	:	371365	23620 pps
---------	---	--------	-----------

Bytes	:	15597330	7936368 bps
-------	---	----------	-------------

Tail-dropped packets	:	0	0 pps
----------------------	---	---	-------

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

Low	:	0	0 pps
-----	---	---	-------

Medium-low	:	0	0 pps
------------	---	---	-------

Medium-high	:	0	0 pps
-------------	---	---	-------

High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 1, Forwarding classes: REALTIME

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 2, Forwarding classes: PRIVATE

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 3, Forwarding classes: CONTROL

Queued:

Packets	:	32843	0 pps
Bytes	:	2641754	56 bps

Transmitted:

Packets	:	32843	0 pps
Bytes	:	2641754	56 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

## Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

## Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

## Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

## Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

## Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps

RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 6, Forwarding classes: CLASS\_V\_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 7, Forwarding classes: CLASS\_S\_OUTPUT, GETS

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps

RED-dropped packets :	0	0 pps
Low :	0	0 pps
Medium-low :	0	0 pps
Medium-high :	0	0 pps
High :	0	0 pps
RED-dropped bytes :	0	0 bps
Low :	0	0 bps
Medium-low :	0	0 bps
Medium-high :	0	0 bps
High :	0	0 bps

#### show interfaces queue (QFX Series)

```

user@switch> show interfaces queue xe-0/0/15
Physical interface: xe-0/0/15, Enabled, Physical link is Up
Interface index: 49165, SNMP ifIndex: 539
Forwarding classes: 12 supported, 8 in use
Egress queues: 12 supported, 8 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Tail-dropped packets : Not Available
  Total-dropped packets: 0 0 pps
  Total-dropped bytes  : 0 0 bps
Queue: 3, Forwarding classes: fcoe
  Queued:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Tail-dropped packets : Not Available
  Total-dropped packets: 0 0 pps
  Total-dropped bytes  : 0 0 bps
0 bps
Queue: 4, Forwarding classes: no-loss
  Queued:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Tail-dropped packets : Not Available
  Total-dropped packets: 0 0 pps
  Total-dropped bytes  : 0 0 bps

```

```

Queue: 7, Forwarding classes: network-control
  Queued:
    Packets      :           0          0 pps
    Bytes        :           0          0 bps
  Transmitted:
    Packets      :           0          0 pps
    Bytes        :           0          0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets:           0          0 pps
    Total-dropped bytes  :           0          0 bps
Queue: 8, Forwarding classes: mcast
  Queued:
    Packets      :           0          0 pps
    Bytes        :           0          0 bps
  Transmitted:
    Packets      :           0          0 pps
    Bytes        :           0          0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets:           0          0 pps
    Total-dropped bytes  :           0          0 bps

```

#### show interfaces queue l2-statistics (lsq interface)

```

user@switch> show interfaces queue lsq-2/2/0.2 l2-statistics
Logical interface lsq-2/2/0.2 (Index 69) (SNMP ifIndex 1598)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :           1          0 pps
    Bytes        :        1001          0 bps
  Transmitted:
    Packets      :           5          0 pps
    Bytes        :        1062          0 bps
    Tail-dropped packets :           0          0 pps
    RED-dropped packets :           0          0 pps
    RED-dropped bytes  :           0          0 bps
Queue: 1, Forwarding classes: ef
  Queued:
    Packets      :           1          0 pps
    Bytes        :        1500          0 bps
  Transmitted:
    Packets      :           6          0 pps
    Bytes        :        1573          0 bps
    Tail-dropped packets :           0          0 pps
    RED-dropped packets :           0          0 pps
    RED-dropped bytes  :           0          0 bps
Queue: 2, Forwarding classes: af
  Queued:
    Packets      :           1          0 pps
    Bytes        :         512          0 bps
  Transmitted:
    Packets      :           3          0 pps
    Bytes        :         549          0 bps
    Tail-dropped packets :           0          0 pps
    RED-dropped packets :           0          0 pps
    RED-dropped bytes  :           0          0 bps
Queue: 3, Forwarding classes: nc
  Queued:

```



```

Packets          :          0          0 pps
Bytes            :          0          0 bps
Transmitted:
Packets          :          0          0 pps
Bytes            :          0          0 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
RED-dropped bytes  :          0          0 bps
=====

```

### show interfaces queue lsq (lsq-ifd)

```

user@switch> show interfaces queue lsq-1/0/0
Logical interface lsq-1/0/0 (Index 348) (SNMP ifIndex 660)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
  Queued:
    Packets          :          55576          1206 pps
    Bytes            :        29622008        5145472 bps
  Transmitted:
    Packets          :          55576          1206 pps
    Bytes            :        29622008        5145472 bps
    Tail-dropped packets :          0          0 pps
    RL-dropped packets :          0          0 pps
    RL-dropped bytes  :          0          0 bps
    RED-dropped packets :          0          0 pps
    Low              :          0          0 pps
    Medium-low       :          0          0 pps
    Medium-high      :          0          0 pps
    High             :          0          0 pps
    RED-dropped bytes :          0          0 bps
    Low              :          0          0 bps
    Medium-low       :          0          0 bps
    Medium-high      :          0          0 bps
    High             :          0          0 bps
Queue: 1, Forwarding classes: ef
  Queued:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
  Transmitted:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
    Tail-dropped packets :          0          0 pps
    RL-dropped packets :          0          0 pps
    RL-dropped bytes  :          0          0 bps
    RED-dropped packets :          0          0 pps
    Low              :          0          0 pps
    Medium-low       :          0          0 pps
    Medium-high      :          0          0 pps
    High             :          0          0 pps
    RED-dropped bytes :          0          0 bps
    Low              :          0          0 bps
    Medium-low       :          0          0 bps
    Medium-high      :          0          0 bps
    High             :          0          0 bps
Queue: 2, Forwarding classes: af
  Queued:
    Packets          :          0          0 pps

```

```

      Bytes                :                0                0 bps
Transmitted:
  Packets                :                0                0 pps
  Bytes                  :                0                0 bps
  Tail-dropped packets :                0                0 pps
  RL-dropped packets   :                0                0 pps
  RL-dropped bytes     :                0                0 bps
  RED-dropped packets  :                0                0 pps
  Low                   :                0                0 pps
  Medium-low           :                0                0 pps
  Medium-high          :                0                0 pps
  High                  :                0                0 pps
  RED-dropped bytes    :                0                0 bps
  Low                   :                0                0 bps
  Medium-low           :                0                0 bps
  Medium-high          :                0                0 bps
  High                  :                0                0 bps
Queue: 3, Forwarding classes: nc
Queued:
  Packets                :            22231            482 pps
  Bytes                  :        11849123        2057600 bps
Transmitted:
  Packets                :            22231            482 pps
  Bytes                  :        11849123        2057600 bps
  Tail-dropped packets :                0                0 pps
  RL-dropped packets   :                0                0 pps
  RL-dropped bytes     :                0                0 bps
  RED-dropped packets  :                0                0 pps
  Low                   :                0                0 pps
  Medium-low           :                0                0 pps
  Medium-high          :                0                0 pps
  High                  :                0                0 pps
  RED-dropped bytes    :                0                0 bps
  Low                   :                0                0 bps
  Medium-low           :                0                0 bps
  Medium-high          :                0                0 bps
  High                  :                0                0 bps

```

## Sample Output

### show interfaces queue (Aggregated Ethernet on a MX series Router)

```
user@host> show interfaces queue ae0 remaining-traffic
```

```
Physical interface: ae0      , Enabled, Physical link is Up
```

```
Interface index: 128, SNMP ifIndex: 543
```

```
Forwarding classes: 16 supported, 4 in use
```

```
Egress queues: 8 supported, 4 in use
```

```
Queue: 0, Forwarding classes: best-effort
```

```
Queued:
```

```

  Packets                :                16                0 pps
  Bytes                  :                1896                0 bps

```

```
Transmitted:
```

```

  Packets                :                16                0 pps
  Bytes                  :                1896                0 bps
  Tail-dropped packets :                0                0 pps
  RL-dropped packets   :                0                0 pps
  RL-dropped bytes     :                0                0 bps
  RED-dropped packets  :                0                0 pps
  Low                   :                0                0 pps
  Medium-low           :                0                0 pps

```

```

    Medium-high      :          0          0 pps
    High             :          0          0 pps
    RED-dropped bytes :          0          0 bps
    Low              :          0          0 bps
    Medium-low       :          0          0 bps
    Medium-high      :          0          0 bps
    High             :          0          0 bps
    Queue-depth bytes :
    Average          :          0
    Current          :          0
    Peak             :          0
    Maximum          :      119013376
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
Transmitted:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
    Tail-dropped packets :          0          0 pps
    RL-dropped packets :          0          0 pps
    RL-dropped bytes  :          0          0 bps
    RED-dropped packets :          0          0 pps
    Low              :          0          0 pps
    Medium-low       :          0          0 pps
    Medium-high      :          0          0 pps
    High             :          0          0 pps
    RED-dropped bytes :          0          0 bps
    Low              :          0          0 bps
    Medium-low       :          0          0 bps
    Medium-high      :          0          0 bps
    High             :          0          0 bps
    Queue-depth bytes :
    Average          :          0
    Current          :          0
    Peak             :          0
    Maximum          :      32768
Queue: 2, Forwarding classes: assured-forwarding
Queued:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
Transmitted:
    Packets          :          0          0 pps
    Bytes            :          0          0 bps
    Tail-dropped packets :          0          0 pps
    RL-dropped packets :          0          0 pps
    RL-dropped bytes  :          0          0 bps
    RED-dropped packets :          0          0 pps
    Low              :          0          0 pps
    Medium-low       :          0          0 pps
    Medium-high      :          0          0 pps
    High             :          0          0 pps
    RED-dropped bytes :          0          0 bps
    Low              :          0          0 bps
    Medium-low       :          0          0 bps
    Medium-high      :          0          0 bps
    High             :          0          0 bps
    Queue-depth bytes :
    Average          :          0
    Current          :          0
    Peak             :          0

```

```
Maximum : 32768
Queue: 3, Forwarding classes: network-control
Queued:
  Packets : 0 0 pps
  Bytes : 0 0 bps
Transmitted:
  Packets : 0 0 pps
  Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
  Low : 0 0 pps
  Medium-low : 0 0 pps
  Medium-high : 0 0 pps
  High : 0 0 pps
RED-dropped bytes : 0 0 bps
  Low : 0 0 bps
  Medium-low : 0 0 bps
  Medium-high : 0 0 bps
  High : 0 0 bps
Queue-depth bytes :
  Average : 0
  Current : 0
  Peak : 0
  Maximum : 6258688
```

## 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 *session-id* <accounting-statistics>>  
 <interface *interface* <accounting-statistics>>  
 <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.  
 Enhanced subscriber management supported in Junos OS Release 15.1R3 for the MX Series.  
**accounting-statistics** option added in Junos OS Release 15.1R3 and 17.4R1 for MX Series.

**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 session-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.

**id session-id accounting-statistics**—(Optional) Display accurate subscriber accounting statistics for a subscriber session with the specified ID. Requires the **actual-transmit-statistics** statement to be configured in the dynamic profile for the dynamic logical interface.

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

**interface accounting-statistics**—(Optional) Display subscriber accounting statistics for the specified interface. Requires the **actual-transmit-statistics** statement to be configured in the dynamic profile for the dynamic logical 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 65,535.

**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** **stacked-vlan-id** option to match the outer VLAN tag.



**NOTE:** Because of display limitations, logical system and routing instance output values are truncated when necessary.

**Required Privilege Level**    view

- Related Documentation**
- [\*show subscribers summary\*](#)
  - [\*Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration\*](#)
  - [\*Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers\*](#)
  - [\*Verifying and Managing Junos OS Enhanced Subscriber Management\*](#)

- List of Sample Output**
- [show subscribers \(IPv4\) on page 625](#)
  - [show subscribers \(IPv6\) on page 625](#)
  - [show subscribers \(IPv4 and IPv6 Dual Stack\) on page 625](#)
  - [show subscribers \(Single Session DHCP Dual Stack\) on page 625](#)
  - [show subscribers \(Single Session DHCP Dual Stack detail\) on page 626](#)
  - [show subscribers \(LNS on MX Series Routers\) on page 626](#)
  - [show subscribers \(L2TP Switched Tunnels\) on page 626](#)
  - [show subscribers client-type dhcp detail on page 626](#)
  - [show subscribers client-type dhcp extensive on page 627](#)
  - [show subscribers client-type vlan-oob detail on page 627](#)
  - [show subscribers count on page 628](#)
  - [show subscribers address detail \(IPv6\) on page 628](#)
  - [show subscribers detail \(IPv4\) on page 628](#)
  - [show subscribers detail \(IPv6\) on page 629](#)
  - [show subscribers detail \(pseudowire Interface for GRE Tunnel\) on page 629](#)
  - [show subscribers detail \(IPv6 Static Demux Interface\) on page 629](#)
  - [show subscribers detail \(L2TP LNS Subscribers on MX Series Routers\) on page 630](#)
  - [show subscribers detail \(L2TP Switched Tunnels\) on page 630](#)
  - [show subscribers detail \(Tunneled Subscriber\) on page 630](#)
  - [show subscribers detail \(IPv4 and IPv6 Dual Stack\) on page 631](#)
  - [show subscribers detail \(ACI Interface Set Session\) on page 631](#)
  - [show subscribers detail \(PPPoE Subscriber Session with ACI Interface Set\) on page 632](#)
  - [show subscribers extensive on page 632](#)
  - [show subscribers extensive \(Passive Optical Network Circuit Interface Set\) on page 633](#)
  - [show subscribers extensive \(DNS Addresses from Access Profile or Global Configuration\) on page 633](#)
  - [show subscribers extensive \(DNS Addresses from RADIUS\) on page 633](#)
  - [show subscribers extensive \(IPv4 DNS Addresses from RADIUS, IPv6 from Access Profile or Global Configuration\) on page 634](#)
  - [show subscribers extensive \(RPF Check Fail Filter\) on page 634](#)
  - [show subscribers extensive \(L2TP LNS Subscribers on MX Series Routers\) on page 635](#)
  - [show subscribers extensive \(IPv4 and IPv6 Dual Stack\) on page 635](#)
  - [show subscribers extensive \(ADF Rules \) on page 636](#)
  - [show subscribers extensive \(Effective Shaping-Rate\) on page 636](#)
  - [show subscribers extensive \(PPPoE Subscriber Access Line Rates on page 637](#)
  - [show subscribers extensive \(Subscriber Session Using PCEF Profile\) on page 638](#)
  - [show subscribers aci-interface-set-name detail \(Subscriber Sessions Using Specified ACI Interface Set\) on page 639](#)
  - [show subscribers agent-circuit-identifier detail \(Subscriber Sessions Using Specified ACI Substring\) on page 639](#)



[show subscribers id accounting-statistics on page 640](#)  
[show subscribers interface accounting-statistics on page 640](#)  
[show subscribers interface extensive on page 641](#)  
[show subscribers logical-system terse on page 642](#)  
[show subscribers physical-interface count on page 642](#)  
[show subscribers routing-instance inst1 count on page 642](#)  
[show subscribers stacked-vlan-id detail on page 642](#)  
[show subscribers stacked-vlan-id vlan-id detail \(Combined Output\) on page 642](#)  
[show subscribers stacked-vlan-id vlan-id interface detail \(Combined Output for a Specific Interface\) on page 642](#)  
[show subscribers user-name detail on page 642](#)  
[show subscribers vlan-id on page 643](#)  
[show subscribers vlan-id detail on page 643](#)  
[show subscribers vpi vci extensive \(PPPoE-over-ATM Subscriber Session\) on page 643](#)  
[show subscribers address detail \(Enhanced Subscriber Management\) on page 644](#)

**Output Fields** Table 46 on page 619 lists the output fields for the **show subscribers** command. Output fields are listed in the approximate order in which they appear.

**Table 46: show subscribers Output Fields**

Field Name	Field Description
<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.
<b>IP Address/VLAN ID</b>	Subscriber IP address or VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i>  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> .
<b>User Name</b>	Name of subscriber.
<b>LS:RI</b>	Logical system and routing instance associated with the subscriber.
<b>Type</b>	Subscriber client type (DHCP, GRE, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN).
<b>IP Address</b>	Subscriber IPv4 address.
<b>IP Netmask</b>	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.
<b>Primary DNS Address</b>	IP address of primary DNS server.  This field is displayed with the <b>extensive</b> option only when the address is provided by RADIUS.

Table 46: show subscribers Output Fields (*continued*)

Field Name	Field Description
Secondary DNS Address	IP address of secondary DNS server.  This field is displayed with the <b>extensive</b> option only when the address is provided by RADIUS.
IPv6 Primary DNS Address	IPv6 address of primary DNS server.  This field is displayed with the <b>extensive</b> option only when the address is provided by RADIUS.
IPv6 Secondary DNS Address	IPv6 address of secondary DNS server.  This field is displayed with the <b>extensive</b> option only when the address is provided by RADIUS.
Domain name server inet	IP addresses for the DNS server, displayed in order of configuration.  This field is displayed with the <b>extensive</b> option only when the addresses are derived from the access profile or the global access configuration.
Domain name server inet6	IPv6 addresses for the DNS server, displayed in order of configuration.  This field is displayed with the <b>extensive</b> option only when the addresses are derived from the access profile or the global access configuration.
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> .

Table 46: show subscribers Output Fields (*continued*)

Field Name	Field Description
<b>Interface Set</b>	<p>Internally generated name of the dynamic ACI or ALI interface set used by the subscriber session. The prefix of the name indicates the string received in DHCP or PPPoE control packets on which the interface set is based. For ALI interface sets, the prefix indicates that the value is configured as a trusted option to identify the subscriber line.</p> <p>The name of the interface set uses one of the following prefixes:</p> <ul style="list-style-type: none"> <li>• <b>aci</b>—ACI; for example, aci-1033-demux0.3221225524. This is the only prefix allowed for ACI interface sets.</li> <li>• <b>ari</b>—ARI; for example, ari-1033-demux0.3221225524.</li> <li>• <b>aci+ari</b>—Both the ACI and ARI; for example, aci+ari-1033-demux0.3221225524.</li> <li>• <b>noids</b>—Neither the ACI nor the ARI were received; for example, noids-1033-demux0.3221225524.</li> </ul> <p><b>NOTE:</b> ACI interface sets are configured with the <b>agent-circuit-identifier</b> autoconfiguration stanza. ALI interface sets are configured with the <b>line-identity</b> autoconfiguration stanza.</p> <p>Besides dynamic ACI and ALI interface sets, this field can be an interface set based on a substring of the ARI string. This occurs when the dynamic profile includes the predefined variable \$junos-pon-id-interface-set-name, and the profile is applied for a passive optical network (PON). The ARI string is inserted by the optical line terminal (OLT). The final substring in the string, unique for the PON, identifies individual subscriber circuits, and is used as the name of the interface set.</p>
<b>Interface Set Type</b>	Interface type of the ACI interface set: <b>Dynamic</b> . This is the only ACI interface set type currently supported.
<b>Interface Set Session ID</b>	Identifier of the dynamic ACI interface set entry in the session database.
<b>Underlying Interface</b>	Name of the underlying interface for the subscriber session.
<b>Dynamic Profile Name</b>	Dynamic profile used for the subscriber.
<b>Dynamic Profile Version</b>	Version number of the dynamic profile used for the subscriber.
<b>MAC Address</b>	MAC address associated with the subscriber.
<b>State</b>	Current state of the subscriber session ( <b>Init</b> , <b>Configured</b> , <b>Active</b> , <b>Terminating</b> , <b>Tunneled</b> ).
<b>L2TP State</b>	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.
<b>Tunnel switch Profile Name</b>	Name of the L2TP tunnel switch profile that initiates tunnel switching.
<b>Local IP Address</b>	IP address of the local gateway (LAC).
<b>Remote IP Address</b>	IP address of the remote peer (LNS).
<b>VLAN Id</b>	VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> .

Table 46: show subscribers Output Fields (*continued*)

Field Name	Field Description
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>
Accounting Statistics	Actual transmitted subscriber accounting statistics by session ID or interface. Service accounting statistics are not included. These statistics do not include overhead bytes or dropped packets; they are the accurate statistics used by RADIUS. The statistics are counted when the <b>actual-transmit-statistics</b> statement is included in the dynamic profile.
DHCP Relay IP Address	IP address used by the DHCP relay agent.
ATM VPI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual path identifier (VPI) on the subscriber's physical interface.
ATM VCI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual circuit identifier (VCI) for each VPI configured on the subscriber interface.
Login Time	Date and time at which the subscriber logged in.
DHCPV6 Options	<b>len</b> = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCPv6 options.
Server DHCP Options	<b>len</b> = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options.
Server DHCPV6 Options	<b>len</b> = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCPv6 options.
DHCPV6 Header	<b>len</b> = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCPv6 options.
Effective shaping-rate	Actual downstream traffic shaping rate for the subscriber, in kilobits per second.
IPv4 Input Service Set	Input service set in access dynamic profile.

Table 46: show subscribers Output Fields (*continued*)

Field Name	Field Description
IPv4 Output Service Set	Output service set in access dynamic profile.
PCEF Profile	PCEF profile in access dynamic profile.
PCEF Rule/Rulebase	PCC rule or rulebase used in dynamic profile.
Dynamic configuration	Values for variables that are passed into the dynamic profile from RADIUS.
Service activation time	Time at which the first family in this service became active.
IPv4 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv4 packets that fail the RPF check.
IPv6 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv6 packets that fail the RPF check.
DHCP Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options, as defined in RFC 2132.
Session ID	ID number for a subscriber session.
Underlying Session ID	For DHCPv6 subscribers on a PPPoE network, displays the session ID of the underlying PPPoE interface.
Service Sessions	Number of service sessions (that is, a service activated using RADIUS CoA) associated with the subscribers.
Service Session ID	ID number for a subscriber service session.
Service Session Name	Service session profile name.
Session Timeout (seconds)	Number of seconds of access provided to the subscriber before the session is automatically terminated.
Idle Timeout (seconds)	Number of seconds subscriber can be idle before the session is automatically terminated.
IPv6 Delegated Address Pool	Name of the pool used for DHCPv6 prefix delegation.
IPv6 Delegated Network Prefix Length	Length of the prefix configured for the IPv6 delegated address pool.
IPv6 Interface Address	Address assigned by the Framed-Ipv6-Prefix AAA attribute. This field is displayed only when the predefined variable \$junos-ipv6-address is used in the dynamic profile.
IPv6 Framed Interface Id	Interface ID assigned by the Framed-Interface-Id AAA attribute.

Table 46: show subscribers Output Fields (*continued*)

Field Name	Field Description
<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.
<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).
<b>DSL type</b>	PPPoE subscriber's access line type reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute DSL-Type (0x0091). The DSL type is one of the following types: <b>ADSL</b> , <b>ADSL2</b> , <b>ADSL2+</b> , <b>OTHER</b> , <b>SDSL</b> , <b>VDSL</b> , or <b>VDSL2</b> .
<b>Frame/Cell Mode</b>	<p>Mode type of the PPPoE subscriber's access line determined by the PPPoE daemon based on the received subattribute DSL-Type (0x0091):</p> <ul style="list-style-type: none"> <li>• <b>Cell</b>—When the DSL line type is one of the following: <b>ADSL</b>, <b>ADSL2</b>, or <b>ADSL2+</b>.</li> <li>• <b>Frame</b>—When the DSL line type is one of the following: <b>OTHER</b>, <b>SDSL</b>, <b>VDSL</b>, or <b>VDSL2</b>.</li> </ul> <p>The value is stored in the subscriber session database.</p>
<b>Overhead accounting bytes</b>	Number of bytes added to or subtracted from the actual downstream cell or frame overhead to account for the technology overhead of the DSL line type. The value is determined by the PPPoE daemon based on the received subattribute DSL-Type (0x0091). The value is stored in the subscriber session database.
<b>Actual upstream data rate</b>	Unadjusted upstream data rate for the PPPoE subscriber's access line reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute Actual-Net-Data-Rate-Upstream (0x0081).

Table 46: show subscribers Output Fields (*continued*)

Field Name	Field Description
<b>Actual downstream data rate</b>	Unadjusted downstream data rate for the PPPoE subscriber's access line reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute Actual-Net-Data-Rate-Downstream (0x0082).
<b>Adjusted downstream data rate</b>	Adjusted downstream data rate for the PPPoE subscriber's access line, calculated by the PPPoE daemon and stored in the subscriber session database.
<b>Adjusted upstream data rate</b>	Adjusted upstream data rate for the PPPoE subscriber's access line, calculated by the PPPoE daemon and stored in the subscriber session database.

## 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 test1:retailer1
demux0.1073741825    203.0.113.3         RETAILER2-CLIENT test1:retailer1
demux0.1073741826    203.0.113.3         RETAILER2-CLIENT test1:retailer1

```

### 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 (Single Session DHCP Dual Stack)

```
user@host> show subscribers
```

```

Interface      IP Address/VLAN ID  User Name      LS:RI
demux0.1073741364 192.168.10.10        dual-stack-retail35 default:default

```

```

2001:db8::100:0:0:0/74
2001:db8:3ffe:0:4::/64
                                default:default

```

### show subscribers (Single Session DHCP Dual Stack detail)

```

user@host> show subscribers id 27 detail
Type: DHCP
User Name: dual-stack-retail33
IP Address: 10.10.0.53
IPv6 Address: 2001:db8:3000:0:0:8003::2
IPv6 Prefix: 2001:db8:3ffe:0:4::/64
Logical System: default
Routing Instance: default
Interface: ae0.3221225472
Interface type: Static
Underlying Interface: ae0.3221225472
Dynamic Profile Name: dhcp-retail-18
MAC Address: 00:00:5E:00:53:02
State: Active
DHCP Relay IP Address: 10.10.0.1
Radius Accounting ID: 27
Session ID: 27
PFE Flow ID: 2
Stacked VLAN Id: 2000
VLAN Id: 1
Login Time: 2014-05-15 10:12:10 PDT
DHCP Options: len 60
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 00 64 01 01 02
00 06 00 04 00 03 00 19 00 03 00 0c 00 00 00 00 00 00 00 00
00 00 00 00 00 19 00 0c 00 00 00 00 00 00 00 00 00 00 00 00

```

### 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 dhcp extensive

```

user@host> show subscribers client-type dhcp extensive
Type: DHCP
User Name: user
IP Address: 30.11.1.4
IP Netmask: 255.0.0.0
IPv6 Address: 7001:2:3::103
IPv6 Prefix: 7001::/68
Domain name server inet6: abcd::1 abcd::2
Logical System: default
Routing Instance: default
Interface: ge-0/0/0.0
Interface type: Static
Underlying Interface: ge-0/0/0.0
MAC Address: 00:10:94:00:00:01
State: Configured
Radius Accounting ID: 10
Session ID: 10
PFE Flow ID: 2
VLAN Id: 100
Agent Circuit ID: ge-0/0/0:100
Agent Remote ID: ge-0/0/0:100
Login Time: 2017-05-23 12:52:22 IST
DHCPV6 Options: len 69
00 01 00 0e 00 01 00 01 59 23 e3 31 00 10 94 00 00 01 00 08
00 02 00 00 00 19 00 29 00 00 00 00 00 04 9d 40 00 07 62 00
00 1a 00 19 00 09 3a 80 00 27 8d 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00
Server DHCP Options: len 13
3a 04 00 00 00 ff 00 3b 04 00 00 0f 00
Server DHCPV6 Options: len 8
00 0a 00 04 ab cd ef ab
DHCPV6 Header: len 4
01 00 00 04
IP Address Pool: al_pool30
IPv6 Address Pool: ia_na_pool
IPv6 Delegated Address Pool: prefix_delegate_pool

```

### 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 (pseudowire Interface for GRE Tunnel)

```

user@host> show subscribers detail

```

Interface	IP Address/VLAN ID	User Name	LS:RI
ps0.3221225484	30.1.0.2		
ps0.3221225485	30.1.0.3		
demux0.3221225486	1		default:default
demux0.3221225487	1		default:default
demux0.3221225488	100.16.0.1		default:default
demux0.3221225489	100.16.0.2		default:default

### 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
```

### 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 (Passive Optical Network Circuit Interface Set)**

```

user@host> show subscribers client-type dhcp extensive
Type: DHCP
IP Address: 192.0.2.136
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073741842
Interface type: Dynamic
Interface Set: ot101.xyz101-202
Underlying Interface: demux0.1073741841
Dynamic Profile Name: dhcp-profile
MAC Address: 00:10:94:00:00:02
State: Active
Radius Accounting ID: jnpr :19
Session ID: 19
VLAN Id: 1100
Agent Remote ID: ABCD01234|100M|AAAA01234|ot101.xyz101-202

Login Time: 2017-03-29 10:30:46 PDT
DHCP Options: len 97
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 02 33 04 00 00
17 70 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 32 2f
32 2d 31 2d 31 37 05 01 06 0f 21 2c 52 2b 02 29 41 42 43 44
30 31 32 33 34 7c 31 30 30 4d 7c 41 41 41 41 30 31 32 33 34
7c 6f 74 6c 30 31 2e 78 79 7a 31 30 31 2d 32 30 32
IP Address Pool: P00L-V4

```

**show subscribers extensive (DNS Addresses from Access Profile or Global Configuration)**

```

user@host> show subscribers extensive
Type: DHCP
User Name: test-user@example-com
IP Address: 192.0.2.119
IP Netmask: 255.255.255.255
Domain name server inet: 198.51.100.1 198.51.100.2
IPv6 Address: 2001:db8::1:11
Domain name server inet6: 2001:db8:5001::12 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool

```

**show subscribers extensive (DNS Addresses from RADIUS)**

```

user@host> show subscribers extensive
Type: DHCP
User Name: test-user@example-com

```

```
IP Address: 192.0.2.119
IP Netmask: 255.255.255.255
Primary DNS Address: 198.51.100.1
Secondary DNS Address: 198.51.100.2
IPv6 Address: 2001:db8::1:11
IPv6 Primary DNS Address: 2001:db8:5001::12
IPv6 Secondary DNS Address: 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool
```

#### show subscribers extensive (IPv4 DNS Addresses from RADIUS, IPv6 from Access Profile or Global Configuration)

```
user@host> show subscribers extensive
Type: DHCP
User Name: test-user@example-com
IP Address: 192.0.2.119
IP Netmask: 255.255.255.255
Primary DNS Address: 198.51.100.1
Secondary DNS Address: 198.51.100.2
IPv6 Address: 2001:db8::1:11
Domain name server inet6: 2001:db8:5001::12 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool
```

#### 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 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 extensive (PPPoE Subscriber Access Line Rates)

```
user@host> show subscribers extensive
Type: PPPoE
IP Address: 192.85.128.1
IP Netmask: 255.255.255.255
Logical System: default
Routing Instance: default
Interface: pp0.3221225475
Interface type: Dynamic
Underlying Interface: demux0.3221225474
Dynamic Profile Name: pppoe-client-profile-with-cos
MAC Address: 00:00:6c:07:01:02
State: Active
Radius Accounting ID: 4
Session ID: 4
PFE Flow ID: 14
Stacked VLAN Id: 40
VLAN Id: 1
Agent Circuit ID: circuit0
Agent Remote ID: remote0
Login Time: 2017-04-06 15:52:32 PDT

User Name: DAVE-L2BSA-SERVICE
Logical System: default
Routing Instance: isp-1-subscriber
Interface: ge-1/2/4.3221225472
Interface type: Dynamic
Interface Set: ge-1/2/4
Underlying Interface: ge-1/2/4
Core IFL Name: ge-1/3/4.0
Dynamic Profile Name: L2BSA-88a8-400LL1300V0
State: Active
Radius Accounting ID: 1
Session ID: 1
PFE Flow ID: 14
VLAN Id: 13
VLAN Map Id: 102
Inner VLAN Map Id: 1
Agent Circuit ID: circuit-aci-3
Agent Remote ID: remote49-3
Login Time: 2017-04-05 16:59:29 EDT
Service Sessions: 4
IFL Input Filter Name: L2BSA-CP-400LL1300V0-ge-1/2/4.3221225472-in
IFL Output Filter Name: L2BSA-CP-400LL1300V0-ge-1/2/4.3221225472-out
Accounting interval: 900
DSL type: VDSL
Frame/Cell Mode: Frame
Overhead accounting bytes: -10
Actual upstream data rate: 1024 kbps
Actual downstream data rate: 4096 kbps
Adjusted downstream data rate: 3686 kbps
Adjusted upstream data rate: 922 kbps
Dynamic configuration:
  junos-vlan-map-id: 102
  Service Session ID: 5
  Service Session Name: SRL-L1
  State: Active
```

```
Family: inet, inet6
IFL Input Filter Name: L2BSA-FWF-in-10048-ge-1/2/4.3221225472-in
IFL Output Filter Name: L2BSA-FWF-out-25088-ge-1/2/4.3221225472-out
Service Activation time: 2017-04-05 16:59:30 EDT
Dynamic configuration:
  l2bsa-fwf-in: L2BSA-FWF-in-10048
  l2bsa-fwf-out: L2BSA-FWF-out-25088
  rldown: 25088
  rlup: 10048
```

#### show subscribers extensive (Subscriber Session Using PCEF Profile)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.3221225517
Interface type: Dynamic
Underlying Interface: ge-1/0/3
Dynamic Profile Name: svlan-dhcp
State: Active
Session ID: 59
PFE Flow ID: 71
Stacked VLAN Id: 0x8100.1
VLAN Id: 0x8100.2
Login Time: 2017-03-28 08:23:08 PDT

Type: DHCP
User Name: pcefuser
IP Address: 5.0.0.26
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: demux0.3221225518
Interface type: Dynamic
Underlying Interface: demux0.3221225517
Dynamic Profile Name: dhcp-client-prof
MAC Address: 00:11:01:00:00:01
State: Active
Radius Accounting ID: 60
Session ID: 60
PFE Flow ID: 73
Stacked VLAN Id: 1
VLAN Id: 2
Login Time: 2017-03-28 08:23:08 PDT
Service Sessions: 1
DHCP Options: len 9
35 01 01 37 04 01 03 3a 3b
IP Address Pool: pool-ipv4
IPv4 Input Service Set: tdf-service-set
IPv4 Output Service Set: tdf-service-set
PCEF Profile: pcef-prof-1
PCEF Rule/Rulebase: default
Dynamic configuration:
  junos-input-service-filter: svc-filt-1
  junos-input-service-set: tdf-service-set
  junos-output-service-filter: svc-filt-1
  junos-output-service-set: tdf-service-set
  junos-pcef-profile: pcef-prof-1
  junos-pcef-rule: default
```

```

Service Session ID: 61
Service Session Name: pcef-serv-prof
State: Active
Family: inet
IPv4 Input Service Set: tdf-service-set
IPv4 Output Service Set: tdf-service-set
PCEF Profile: pcef-prof-1
PCEF Rule/Rulebase: limit-fb
Service Activation time: 2017-03-28 08:31:19 PDT
Dynamic configuration:
  pcef-prof: pcef-prof-1
  pcef-rule1: limit-fb
  svc-filt: svc-filt-1
  svc-set: tdf-service-set

```

### 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 id accounting-statistics

```
user@host> show subscribers id 601 accounting-statistics
Session ID: 601
Accounting Statistics:
Input bytes : 199994
Output bytes : 121034
Input packets: 5263
Output packets: 5263
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
```

#### show subscribers interface accounting-statistics

```
user@host> show subscribers interface pp0.3221226949 accounting-statistics
Session ID: 501
Accounting Statistics:
Input bytes : 199994
Output bytes : 121034
Input packets: 5263
Output packets: 5263
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Session ID: 502
Accounting Statistics:
Input bytes : 87654
Output bytes : 72108
Input packets: 3322
```

```

Output packets: 3322
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Session ID: 503
Accounting Statistics:
Input bytes : 156528
Output bytes : 123865
Input packets: 7448
Output packets: 7448
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

#### 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
```