

Traffic Management User Guide for NFX Series Devices

Published
2023-06-12

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

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

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

Traffic Management User Guide for NFX Series Devices
Copyright © 2023 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://support.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 This Guide | ix

1

CoS Overview

Basic Concepts | 2

Overview of Junos OS CoS | 2

Configuring CoS | 5

Understanding Junos CoS Components | 12

Assigning CoS Components to Interfaces | 17

Understanding CoS Packet Flow | 19

Understanding Default CoS Settings | 23

CoS Inputs and Outputs Overview | 37

Overview of Policers | 38

2

Classifying and Rewriting Traffic

Using Classifiers, Forwarding Classes, and Rewrite Rules | 49

Understanding CoS Classifiers | 49

Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p) | 59

Example: Configuring Classifiers | 62

Requirements | 63

Overview | 63

Verification | 65

Understanding Default CoS Scheduling and Classification | 66

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces | 77

Understanding CoS Code-Point Aliases | 92

Defining CoS Code-Point Aliases | 95

Understanding CoS Forwarding Classes | 96

Defining CoS Forwarding Classes | 102

Example: Configuring Forwarding Classes | 105

Requirements | 105

Overview | 106

Example 1: Configuring Forwarding Classes for Switches Except QFX10000 | 108

Verification | 109

Example 2: Configuring Forwarding Classes for QFX10000 Switches | 110

Verification | 111

Understanding CoS Rewrite Rules | 112

Defining CoS Rewrite Rules | 114

Troubleshooting an Unexpected Rewrite Value | 117

3

Scheduling Traffic

Using Schedulers | 121

Understanding CoS Scheduling Behavior and Configuration Considerations | 121

Defining CoS Queue Schedulers for Port Scheduling | 128

Defining CoS Queue Scheduling Priority | 132

Example: Configuring Queue Scheduling Priority | 133

Requirements | 134

Overview | 135

Verification | 137

Understanding CoS Traffic Control Profiles | 138

Understanding CoS Priority Group Scheduling | 140

Defining CoS Traffic Control Profiles (Priority Group Scheduling) | 144

Example: Configuring Traffic Control Profiles (Priority Group Scheduling) | 145

Requirements | 146

Overview | 147

Verification | 148

Understanding CoS Priority Group and Queue Guaranteed Minimum Bandwidth | 149

Example: Configuring Minimum Guaranteed Output Bandwidth | 152

Requirements | 154

Overview | 154

Verification | 156

Understanding CoS Priority Group Shaping and Queue Shaping (Maximum Bandwidth) | 159

Example: Configuring Maximum Output Bandwidth | 162

Requirements | 164

Overview | 164

Verification | 165

Understanding CoS Explicit Congestion Notification | 168

Configuration Statements and Operational Commands

Configuration Statements (Basic Concepts) | 179

class-of-service | 179

traceoptions (Class of Service) | 184

Configuration Statements (Classifiers and Rewrite Rules) | 188

class (Forwarding Classes) | 189

class (Forwarding Class Sets) | 191

classifiers | 193

code-point (Rewrite Rules) | 196

code-point-aliases | 197

code-points (CoS) | 199

dscp | 200

dscp-ipv6 | 204

exp | 207

forwarding-class | 209

forwarding-classes | 213

ieee-802.1 | 218

import | 221

interfaces (Class of Service) | 223

loss-priority (Classifiers) | 225

loss-priority (Rewrite Rules) | 227

queue-num | 228

rewrite-rules | 231

unit | 233

Configuration Statements (Scheduling) | 235

buffer-size | 235

excess-rate | 241

fill-level | 243

forwarding-class | 245

guaranteed-rate | 248

interpolate | 251

output-traffic-control-profile | 252

priority (Schedulers) | 253

scheduler | 256

scheduler-map | 257

scheduler-maps | 258

schedulers | 260

shaping-rate | 261

traffic-control-profiles | 264

transmit-rate | 268

Operational Commands (Basic Concepts) | 274

Monitoring Interfaces That Have CoS Components | 274

show class-of-service | 276

show class-of-service interface | 282

show class-of-service shared-buffer | 326

show pfe filter hw summary | 330

show pfe next-hop | 332

show pfe route | 339

show pfe terse | 353

show pfe version | 356

show interfaces voq | 357

Operational Commands (Classifiers and Rewrite Rules) | 376

Monitoring CoS Classifiers | 376

Monitoring CoS Forwarding Classes | 378

Monitoring CoS Rewrite Rules | 382

Monitoring CoS Code-Point Value Aliases | 384

show class-of-service classifier | 385

show class-of-service code-point-aliases | 389

show class-of-service forwarding-class | 391

show class-of-service forwarding-table | 396

show class-of-service forwarding-table classifier | 401

show class-of-service forwarding-table classifier mapping | 403

show class-of-service forwarding-table rewrite-rule | 406

show class-of-service forwarding-table rewrite-rule mapping | 408

show class-of-service interface | 410

show class-of-service rewrite-rule | 453

Operational Commands (Scheduling) | 457

Monitoring CoS Scheduler Maps | 457

show class-of-service drop-profile | 459

show class-of-service forwarding-table | 464

show class-of-service forwarding-table drop-profile | 469

show class-of-service forwarding-table scheduler-map | **471**

show class-of-service interface | **475**

show class-of-service scheduler-map | **518**

show interfaces queue | **523**

show interfaces voq | **583**

About This Guide

Use this guide to understand and configure class of service (CoS) features in Junos OS to define service levels that provide different delay, jitter, and packet loss characteristics to particular applications served by specific traffic flows. Applying CoS features to each device in your network ensures quality of service (QoS) for traffic throughout your entire network.

1

PART

CoS Overview

Basic Concepts | 2

CHAPTER 1

Basic Concepts

IN THIS CHAPTER

- Overview of Junos OS CoS | 2
- Configuring CoS | 5
- Understanding Junos CoS Components | 12
- Assigning CoS Components to Interfaces | 17
- Understanding CoS Packet Flow | 19
- Understanding Default CoS Settings | 23
- CoS Inputs and Outputs Overview | 37
- Overview of Policers | 38

Overview of Junos OS CoS

IN THIS SECTION

- CoS Standards | 3
- How Junos OS CoS Works | 4
- Default CoS Behavior | 5

When a network experiences congestion and delay, some packets must be dropped. Junos OS *class of service* (CoS) enables you to divide traffic into classes and set various levels of throughput and packet loss when congestion occurs. You have greater control over packet loss because you can configure rules tailored to your needs.

You can configure CoS features to provide multiple classes of service for different applications. CoS also allows you to rewrite the Differentiated Services code point (DSCP) or IEEE 802.1p code-point bits of

packets leaving an interface, thus allowing you to tailor packets for the network requirements of the remote peers.

CoS provides multiple classes of service for different applications. You can configure multiple forwarding classes for transmitting packets, define which packets are placed into each output queue, schedule the transmission service level for each queue, and manage congestion using a weighted random early detection (WRED) algorithm.

In designing CoS applications, you must carefully consider your service needs, and you must thoroughly plan and design your CoS configuration to ensure consistency and interoperability across all platforms in a CoS domain.

Because CoS is implemented in hardware rather than in software, you can experiment with and deploy CoS features without affecting packet forwarding and switching performance.

NOTE: CoS policies can be enabled or disabled on each switch interface. Also, each physical and *logical interface* on the switch can have associated custom CoS rules.

When you change or when you deactivate and then reactivate the class-of-service configuration, the system experiences packet drops because the system momentarily blocks traffic to change the mapping of incoming traffic to input queues.

This topic describes:

CoS Standards

The following RFCs define the standards for CoS capabilities:

- RFC 2474, *Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers*
- RFC 2597, *Assured Forwarding PHB Group*
- RFC 2598, *An Expedited Forwarding PHB*
- RFC 2698, *A Two Rate Three Color Marker*
- RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*

The following data center bridging (DCB) standards are also supported to provide the CoS (and other characteristics) that Fibre Channel over Ethernet (FCoE) requires for transmitting storage traffic over an Ethernet network:

- IEEE 802.1Qbb, *priority-based flow control (PFC)*
- IEEE 802.1Qaz, *enhanced transmission selection (ETS)*

- IEEE 802.1AB (LLDP) extension called Data Center Bridging Capability Exchange Protocol (DCBX)

NOTE: OCX Series switches and NFX250 Network Services platforms do not support PFC and DCBX.

Juniper Networks QFX10000 switches support both enhanced transmission selection (ETS) hierarchical port scheduling and direct port scheduling.

How Junos OS CoS Works


Junos OS CoS works by examining traffic entering the edge of your network. The switch classifies traffic into defined service groups to provide the special treatment of traffic across the network. For example, you can send voice traffic across certain links and data traffic across other links. In addition, the data traffic streams can be serviced differently along the network path to ensure that higher-paying customers receive better service. As the traffic leaves the network at the far edge, you can reclassify the traffic to meet the policies of the targeted peer by rewriting the DSCP or IEEE 802.1 code-point bits.

To support CoS, you must configure each switch in the network. Generally, each switch examines the packets that enter it to determine their CoS settings. These settings dictate which packets are transmitted first to the next downstream switch. Switches at the edges of the network might be required to alter the CoS settings of the packets that enter the network to classify the packets into the appropriate service groups.

In [Figure 1 on page 5](#), Switch A is receiving traffic. As each packet enters, Switch A examines the packet's current CoS settings and classifies the traffic into one of the groupings defined on the switch. This definition allows Switch A to prioritize its resources for servicing the traffic streams it receives. Switch A might alter the CoS settings (forwarding class and loss priority) of the packets to better match the defined traffic groups.

When Switch B receives the packets, it examines the CoS settings, determines the appropriate traffic groups, and processes the packet according to those settings. It then transmits the packets to Switch C, which performs the same actions. Switch D also examines the packets and determines the appropriate groups. Because Switch D sits at the far end of the network, it can reclassify (rewrite) the CoS code-point bits of the packets before transmitting them.

Figure 1: Packet Flow Across the Network



Default CoS Behavior

If you do not configure CoS settings, the software performs some CoS functions to ensure that the system forwards traffic and protocol packets with minimum delay when the network is experiencing congestion. Some CoS settings, such as classifiers, are automatically applied to each logical interface that you configure. Other settings, such as *rewrite rules*, are applied only if you explicitly associate them with an interface.

RELATED DOCUMENTATION

Overview of Policers

Understanding Junos CoS Components

Understanding CoS Packet Flow

Understanding CoS Hierarchical Port Scheduling (ETS)

Configuring CoS

The traffic management class-of-service topics describe how to configure the Junos OS class-of-service (CoS) components. Junos CoS provides a flexible set of tools that enable you to fine tune control over the traffic on your network.

- Define classifiers that classify incoming traffic into forwarding classes to place traffic in groups for transmission.
- Map forwarding classes to output queues to define the type of traffic on each output queue.

- Configure schedulers for each output queue to control the service level (priority, bandwidth characteristics) of each type of traffic.
- Provide different service levels for the same forwarding classes on different interfaces.
- On switches that support data center bridging standards, configure lossless transport across the Ethernet network using priority-based flow control (PFC), Data Center Bridging Exchange protocol (DCBX), and enhanced transmission selection (ETS) hierarchical scheduling (OCX Series switches and NFX250 Network Services platform do not support lossless transport, PFC, and DCBX).
- Configure various CoS components individually or in combination to define CoS services.

NOTE: When you change the CoS configuration or when you deactivate and then reactivate the CoS configuration, the system experiences packet drops because the system momentarily blocks traffic to change the mapping of incoming traffic to input queues.

[Table 1 on page 7](#) lists the primary CoS configuration tasks by platform and provides links to those tasks.

NOTE: Links to features that are not supported on the platform for which you are looking up information might not be functional.

Table 1: CoS Configuration Tasks

CoS Configuration Task	Platforms Supported	Links
<p>Basic CoS Configuration:</p> <ul style="list-style-type: none"> • Configure code-point aliases to assign a name to a pattern of code-point bits that you can use instead of the bit pattern when you configure CoS components such as classifiers and rewrite rules • Configure classifiers and multidestination classifiers <ul style="list-style-type: none"> • Set the forwarding class and loss priority of a packet based on the incoming CoS value and assign packets to output queues based on the associated forwarding class • Change the host default output queue and mapping of DSCP bits used in the type of service (ToS) field • Configure forwarding classes • Configure rewrite rules to alter code point bit values in outgoing packets on the outbound interfaces of a switch so that the CoS treatment matches the policies of a targeted peer • Configure Ethernet PAUSE flow control, a congestion relief feature that provides link-level flow control for all traffic on a full-duplex Ethernet link, including those that belong to Ethernet link aggregated (LAG) interfaces. On any particular interface, symmetric and asymmetric flow control are mutually exclusive. • Assign the following CoS components to physical or logical interfaces: 	<ul style="list-style-type: none"> • QFX3500 • QFX3600 • EX4600 • NFX250 • QFX5100 • QFX5200 • QFX5210 • QFX10000 • OCX1100 switches • QFabric systems 	<ul style="list-style-type: none"> • <i>Defining CoS Code-Point Aliases</i> • (QFX10000 only) <i>Example: Configuring Classifiers</i> • (Except QFX10000) <i>Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)</i> • (Except NFX250 and QFX10000) <i>Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers</i> • <i>Changing the Host Outbound Traffic Default Queue Mapping</i> • <i>Example: Configuring Forwarding Classes</i> • <i>Defining CoS Rewrite Rules</i> • (Except NFX250) <i>Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control</i> • (Except NFX250 and OCX1100) <i>Configuring CoS Asymmetric Ethernet PAUSE Flow Control</i> • <i>Assigning CoS Components to Interfaces</i>

Table 1: CoS Configuration Tasks (*Continued*)

CoS Configuration Task	Platforms Supported	Links
<ul style="list-style-type: none"> Classifiers Congestion notification profiles Forwarding classes Forwarding class sets Output traffic control profiles Port schedulers Rewrite rules 		
<p>Configure Weighted random early detection (WRED) drop profiles that define the drop probability of packets of different packet loss probabilities (PLPs) as the output queue fills:</p> <ul style="list-style-type: none"> Configure WRED drop profiles where you associate WRED drop profiles with loss priorities in a scheduler. When you map the scheduler to a forwarding class (queue), you apply the interpolated drop profile to traffic of the specified loss priority on that queue. Configure drop profile maps that map a drop profile to a packet loss priority, and associate the drop profile and packet loss priority with a scheduler Configure explicit congestion notification (ECN) to enable end-to-end congestion notification between two endpoints on TCP/IP based networks. Apply WRED drop profiles to forwarding classes to control how the switch marks ECN-capable packets. 	<ul style="list-style-type: none"> QFX3500 QFX3600 EX4600 QFX5100 QFX5200 QFX5210 QFX10000 OCX1100 switches QFabric systems 	<ul style="list-style-type: none"> <i>Example: Configuring WRED Drop Profiles</i> <i>Example: Configuring Drop Profile Maps</i> <i>Example: Configuring ECN</i>

Table 1: CoS Configuration Tasks (Continued)

CoS Configuration Task	Platforms Supported	Links
<p>Configure queue schedulers and the bandwidth scheduling priority of individual queues. Schedulers define the CoS properties of output queues (output queues are mapped to forwarding classes, and classifiers map traffic into forwarding classes based on IEEE 802.1p or DSCP code points). Queue scheduling works with priority group scheduling to create a two-tier hierarchical scheduler. CoS scheduling properties include the amount of interface bandwidth assigned to the queue, the priority of the queue, whether explicit congestion notification (ECN) is enabled on the queue, and the WRED packet drop profiles associated with the queue.</p>	<ul style="list-style-type: none"> • QFX3500 • QFX3600 • EX4600 • NFX250 • QFX5100 • QFX5200 • QFX5210 • QFX10000 • OCX1100 switches • QFabric systems 	<ul style="list-style-type: none"> • (Except QFX10000) <i>Example: Configuring Queue Schedulers</i> • <i>Example: Configuring Queue Scheduling Priority</i> • (QFX10000 only) <i>Example: Configuring Queue Schedulers for Port Scheduling</i>
<p>Configure traffic control profiles to define the output bandwidth and scheduling characteristics of forwarding class sets (priority groups). The forwarding classes (queues) mapped to a forwarding class set share the bandwidth resources that you configure in the traffic control profile.</p>	<ul style="list-style-type: none"> • QFX3500 • QFX3600 • EX4600 • NFX250 • QFX5100 • QFX5200 • QFX5210 • QFX10000 • OCX1100 switches • QFabric systems 	<ul style="list-style-type: none"> • (Except NFX250) <i>Defining CoS Traffic Control Profiles (Priority Group Scheduling)</i> • (Except NFX250) <i>Example: Configuring Traffic Control Profiles (Priority Group Scheduling)</i> • <i>Example: Configuring Minimum Guaranteed Output Bandwidth</i> • (Except NFX250) <i>Example: Configuring Maximum Output Bandwidth</i>

Table 1: CoS Configuration Tasks (Continued)

CoS Configuration Task	Platforms Supported	Links
<p>Configure enhanced transmission selection (ETS) and forwarding class sets, and disable the ETS recommendation TLV. Hierarchical port scheduling, the Junos OS implementation of ETS, enables you to group priorities that require similar CoS treatment into priority groups. You define the port bandwidth resources for a priority group, and you define the amount of the priority group's resources that each priority in the group can use.</p>	<ul style="list-style-type: none"> • QFX3500 • QFX3600 • EX4600 • QFX5100 • OCX1100 switches • QFX10000 • QFabric systems 	<ul style="list-style-type: none"> • <i>Example: Configuring Forwarding Class Sets</i> • <i>Example: Configuring CoS Hierarchical Port Scheduling (ETS)</i> • <i>(Except OCX1100)Disabling the ETS Recommendation TLV</i>
<p>Configure Data Center Bridging Capability Exchange protocol (DCBX), which discovers the data center bridging (DCB) capabilities of peers by exchanging feature configuration information and is an extension of the Link Layer Discovery Protocol (LLDP)</p> <ul style="list-style-type: none"> • Configure the DCBX mode that an interface uses to communicate with the connected peer • Configure DCBX autonegotiation on a per-interface basis for each supported feature or application • Define each application for which you want DCBX to exchange application protocol information • Map applications to IEEE 802.1p code points • Apply an application map to a DCBX interface 	<ul style="list-style-type: none"> • QFX3500 • QFX3600 • EX4600 • QFX5100 • QFX5200 • QFX5210 • QFX10000 • QFabric systems 	<ul style="list-style-type: none"> • <i>Example: Configuring DCBX Application Protocol TLV Exchange</i> • <i>Configuring the DCBX Mode</i> • <i>Configuring DCBX Autonegotiation</i> • <i>Defining an Application for DCBX Application Protocol TLV Exchange</i> • <i>Configuring an Application Map for DCBX Application Protocol TLV Exchange</i> • <i>Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange</i>

Table 1: CoS Configuration Tasks *(Continued)*

CoS Configuration Task	Platforms Supported	Links
<p>Configure CoS for FCoE:</p> <ul style="list-style-type: none"> Configure priority-based flow control (PFC) to divide traffic on one physical link into eight priorities Configure a congestion notification profile (CNP) that enables priority-based flow control (PFC) on specified IEEE 802.1p priorities Configure Multichassis link aggregation groups (MC-LAGs) to provide redundancy and load balancing between two switches Configure two or more lossless forwarding classes and map them to different priorities Configure lossless FCoE transport if your network uses a different priority than 3 Configure multiple lossless FCoE priorities on a converged Ethernet network If the FCoE network uses a different priority than priority 3 for FCoE traffic, configure a rewrite value to remap incoming traffic from the FC SAN to that priority after the interface encapsulates the FC packets in Ethernet Configure lossless priorities for multiple types of traffic, such as FCoE and iSCSI 	<ul style="list-style-type: none"> QFX3500 QFX3600 EX4600 QFX5100 QFX5200 QFX5210 QFX10000 QFabric systems 	<ul style="list-style-type: none"> <i>Example: Configuring CoS PFC for FCoE Traffic</i> Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG <i>Configuring CoS PFC (Congestion Notification Profiles)</i> (QFX3500 and QFabric only) <i>Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway</i> <i>Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces</i> <i>Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch)</i> <i>Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface</i> (QFX3500, NFX250, and QFabric only) <i>Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces (NP_Ports)</i> <i>Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI)</i>

Understanding Junos CoS Components

IN THIS SECTION

- [Code-Point Aliases | 12](#)
- [Policers | 12](#)
- [Classifiers | 12](#)
- [Forwarding Classes | 13](#)
- [Forwarding Class Sets | 14](#)
- [Flow Control \(Ethernet PAUSE, PFC, and ECN\) | 15](#)
- [WRED Profiles and Tail Drop | 16](#)
- [Schedulers | 16](#)
- [Rewrite Rules | 17](#)

This topic describes the Junos OS class-of-service (CoS) components:

Code-Point Aliases

A *code-point alias* assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers and *rewrite rules*.

Policers

Policers limit traffic of a certain class to a specified bandwidth and burst size. Packets exceeding the policer limits can be discarded, or can be assigned to a different forwarding class, a different loss priority, or both. You define policers with filters that you can associate with input interfaces.

Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. In Junos OS, *classifiers* associate packets with a forwarding class and loss priority and assign packets to output queues based on the associated forwarding class. Junos OS supports two general types of classifiers:

- Behavior aggregate (BA) or CoS value traffic classifiers—Examine the CoS value in the packet header. The value in this single field determines the CoS settings applied to the packet. BA classifiers allow

you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP) value, IEEE 802.1p value, or MPLS EXP value.

NOTE: OCX Series switches and NFX250 Network Services platform do not support MPLS.

- **Multifield traffic classifiers**—Examine multiple fields in the packet, such as source and destination addresses and source and destination port numbers of the packet. With multifield classifiers, you set the forwarding class and loss priority of a packet based on *firewall filter* rules.

On switches that require the separation of unicast and multideestination (multicast, broadcast, and destination lookup fail) traffic, you create separate unicast classifiers and multideestination classifiers. You cannot assign unicast traffic and multideestination traffic to the same classifier. You can apply unicast classifiers to one or more interfaces. Multideestination classifiers apply to all of the switch interfaces and cannot be applied to individual interfaces. Switches that require the separation of unicast and multideestination traffic have 12 output queues to provide 4 output queues reserved for multideestination traffic.

On switches that do not separate unicast and multideestination traffic, unicast and multideestination traffic use the same classifiers, and you do not create a separate special classifier for multideestination traffic. Switches that do not separate unicast and multideestination traffic have eight output queues because no extra queues are required to separate the traffic.

Forwarding Classes

Forwarding classes group packets for transmission and CoS. You assign each packet to an output queue based on the packet's forwarding class. Forwarding classes affect the forwarding, scheduling, and rewrite marking policies applied to packets as they transit the switch.

Switches provide up to five default forwarding classes:

- **best-effort**—Best-effort traffic
- **fcoe**—Fibre Channel over Ethernet traffic
- **no-loss**—Lossless traffic
- **network-control**—Network control traffic
- **mcast**—Multicast traffic

NOTE: The default `mcast` forwarding class applies only to switches that require the separation of unicast and multideestination (multicast, broadcast, and destination lookup fail) traffic. On these

switches, you create separate forwarding classes for the two types of traffic. The default mcast forwarding class transports only multidestination traffic, and the default best-effort, fcoe, no-loss, and network-control forwarding classes transport only unicast traffic. Unicast forwarding classes map to unicast output queues, and multidestination forwarding classes map to multidestination output queues. You cannot assign unicast traffic and multidestination traffic to the same forwarding class or to the same output queue. Switches that require the separation of unicast and multidestination traffic have 12 output queues, 8 for unicast traffic and 4 for multidestination traffic.

On switches that do not separate unicast and multidestination traffic, unicast and multidestination traffic use the same forwarding classes and output queues, so the mcast forwarding class is not valid. You do not create separate forwarding classes for multidestination traffic. Switches that do not separate unicast and multidestination traffic have eight output queues because no extra queues are required to separate the traffic.

NOTE: On OCX Series switches only, do not map traffic to the default fcoe and no-loss forwarding classes. By default, the DSCP default classifier does not map traffic to the fcoe and no-loss forwarding classes, so by default, OCX Series switches do not classify traffic into those forwarding classes. (On other switches, the fcoe and no-loss forwarding classes provide lossless transport for Layer 2 traffic. OCX Series switches do not support lossless Layer 2 transport.)

Switches support a total of either 12 forwarding classes (8 unicast forwarding classes and 4 multicast forwarding classes), or 8 forwarding classes (unicast and multidestination traffic use the same forwarding classes), which provides flexibility in classifying traffic.

NFX250 Network Services platform provide the following forwarding classes:

- best-effort (be)—Provides no service profile. Loss priority is typically not carried in a CoS value.
- expedited-forwarding (ef)—Provides a low loss, low latency, low jitter, assured bandwidth, end-to-end service.
- assured-forwarding (af)—Provides a group of values you can define and includes four subclasses: AF1, AF2, AF3, and AF4, each with two drop probabilities: low and high.
- network-control (nc)—Supports protocol control and thus is typically high priority.

Forwarding Class Sets

You can group forwarding classes (output queues) into *forwarding class sets* to apply CoS to groups of traffic that require similar treatment. Forwarding class sets map traffic into priority groups to support enhanced transmission selection (ETS), which is described in IEEE 802.1Qaz.

You can configure up to three unicast forwarding class sets and one multicast forwarding class set. For example, you can configure different forwarding class sets to apply CoS to unicast groups of local area network (LAN) traffic, storage area network (SAN) traffic, and high-performance computing (HPC) traffic, and configure another group for multicast traffic.

Within each forwarding class set, you can configure special CoS treatment for the traffic mapped to each individual queue. This provides the ability to configure CoS in a two-tier hierarchical manner. At the forwarding class set tier, you configure CoS for groups of traffic using a *traffic control profile*. At the queue tier, you configure CoS for individual output queues within a forwarding class set using a *scheduler* that you map to a queue (forwarding class) using a *scheduler map*.

Flow Control (Ethernet PAUSE, PFC, and ECN)

Ethernet PAUSE (described in IEEE 802.3X) is a link-level flow control mechanism. During periods of network congestion, Ethernet PAUSE stops all traffic on a full-duplex Ethernet link for a period of time specified in the PAUSE message.

NOTE: QFX10000 switches do not support Ethernet PAUSE.

Priority-based flow control (PFC) is described in IEEE 802.1Qbb as part of the IEEE data center bridging (DCB) specifications for creating a lossless Ethernet environment to transport loss-sensitive flows such as Fibre Channel over Ethernet (FCoE) traffic.

NOTE: OCX Series switches do not support PFC.

PFC is a link-level flow control mechanism similar to Ethernet PAUSE. However, Ethernet PAUSE stops all traffic on a link for a period of time. PFC decouples the pause function from the physical link and divides the traffic on the link into eight priorities (3-bit IEEE 802.1p code points). You can think of the eight priorities as eight “lanes” of traffic. You can apply pause selectively to the traffic on any priority without pausing the traffic on other priorities on the same link.

The granularity that PFC provides allows you to configure different levels of CoS for different types of traffic on the link. You can create lossless lanes for traffic such as FCoE, LAN backup, or management, while using standard frame-drop methods of congestion management for IP traffic on the same link.

NOTE: If you transport FCoE traffic, you must enable PFC on the priority assigned to FCoE traffic (usually IEEE 802.1p code point 011 on interfaces that carry FCoE traffic).

Explicit congestion notification (ECN) enables end-to-end congestion notification between two endpoints on TCP/IP based networks. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. Any device in the transmission path that does not support ECN breaks the end-to-end ECN functionality. ECN notifies networks about congestion with the goal of reducing packet loss and delay by making the sending device decrease the transmission rate until the congestion clears, without dropping packets. RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*, defines ECN.

WRED Profiles and Tail Drop

A weighted random early detection (WRED) profile (drop profile) defines parameters that enable the network to drop packets during periods of congestion. A *drop profile* defines the conditions under which packets of different loss priorities drop, by determining the probability of dropping a packet for each loss priority when output queues become congested. Drop profiles essentially set a value for a level of queue fullness—when the queue fills to the level of the queue fullness value, packets drop. The combination of queue fill level, the probability of dropping a packet at that fill level, and loss priority of the packet, determine whether a packet is dropped or forwarded. Each pairing of a fill level with a drop probability creates a point on a drop profile curve.

You can associate different drop profiles with different loss priorities to set the probability of dropping packets. You can apply a drop profile for each loss priority to a forwarding class (output queue) by applying a drop profile to a scheduler, and then mapping the scheduler to a forwarding class using a scheduler map. When the queue mapped to the forwarding class experiences congestion, the drop profile determines the level of packet drop for traffic of each loss priority in that queue.

Loss priority affects the scheduling of a packet without affecting the packet's relative ordering. Typically you mark packets exceeding a particular service level with a high loss priority.

Tail drop is a simple drop mechanism that drops all packets indiscriminately during periods of congestion, without differentiating among the packet loss priorities of traffic flows. Tail drop requires only one curve point that corresponds to the maximum depth of the output queue, and drop probability when traffic exceeds the buffer depth is 100 percent (all packets that cannot be stored in the queue are dropped). WRED is superior to tail-drop because WRED enables you to treat traffic of different priorities in a differentiated manner, so that higher priority traffic receives preference, and because of the ability to set multiple points on the drop curve.

Schedulers

Each switch interface has multiple queues assigned to store packets. The switch determines which queue to service based on a particular method of scheduling. This process often involves determining the sequence in which different types of packets should be transmitted.

You can define the scheduling priority (priority), minimum guaranteed bandwidth (transmit-rate), maximum bandwidth (shaping-rate), and WRED profiles to be applied to a particular queue (forwarding

class) for packet transmission. By default, extra bandwidth is shared among queues in proportion to the minimum guaranteed bandwidth of each queue. On switches that support the `excess-rate` statement, you can configure the percentage of shared extra bandwidth an output queue receives independently from the minimum guaranteed bandwidth transmit rate, or you can use default bandwidth sharing based on the transmit rate.

A scheduler map associates a specified forwarding class with a scheduler configuration. You can associate up to four user-defined scheduler maps with the interfaces.

Rewrite Rules

A *rewrite rule* sets the appropriate CoS bits in the outgoing packet. This allows the next downstream device to classify the packet into the appropriate service group. Rewriting (marking) outbound packets is useful when the switch is at the border of a network and must change the CoS values to meet the policies of the targeted peer.

NOTE: Ingress firewall filters can also rewrite forwarding class and loss priority values.

RELATED DOCUMENTATION

| *Understanding CoS Packet Flow*

Assigning CoS Components to Interfaces

After you define the following CoS components, you assign them to physical or logical interfaces. Components that you assign to physical interfaces are valid for all of the logical interfaces configured on the physical interface. Components that you assign to a logical interface are valid only for that logical interface.

- Classifiers—Assign only to logical interfaces; on some switches, you can apply classifiers to physical Layer 3 interfaces and the classifiers are applied to all logical interfaces on the physical interface.
- Congestion notification profiles—Assign only to physical interfaces.

NOTE: OCX Series switches and NFX250 Network Services platform do not support congestion notification profiles.

- Forwarding classes—Assign to interfaces by mapping to forwarding class sets.
- Forwarding class sets—Assign only to physical interfaces.
- Output traffic control profiles—Assign only to physical interfaces (with a forwarding class set).
- Port schedulers—Assign only to physical interfaces on switches that support port scheduling. Associate the scheduler with a forwarding class in a scheduler map and apply the scheduler map to the physical interface.
- Rewrite rules—Assign only to logical interfaces; on some switches, you can apply classifiers to physical Layer 3 interfaces and the classifiers are applied to all logical interfaces on the physical interface.

You can assign a CoS component to a single interface or to multiple interfaces using wildcards. You can also assign a congestion notification profile or a forwarding class set globally to all interfaces.

To assign CoS components to interfaces:

Assign a CoS component to a physical interface by associating a CoS component (for example, a forwarding class set named `be-priority-group`) with an interface:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/7 forwarding-class-set be-priority-group
```

Assign a CoS component to a logical interface by associating a CoS component (for example, a classifier named `be_classifier`) with a logical interface:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/7 unit 0 classifiers dscp be_classifier
```

Assign a CoS component to multiple interfaces by associating a CoS component (for example, a rewrite rule named `customup-rw`) to all 10-Gigabit Ethernet interfaces on the switch, use wildcard characters for the interface name and logical interface (unit) number:

```
[edit class-of-service interfaces]
user@switch# set xe-* unit * rewrite-rules ieee-802.1 customup-rw xe-* unit * rewrite-rules
ieee-802.1 customup-rw
```

Assign a congestion notification profile or a forwarding class set globally to all interfaces using the `set class-of-service interfaces all` statement. For example, to assign a forwarding class set named `be-priority-group` to all interfaces:

```
[edit class-of-service interfaces]
user@switch# set all forwarding-class-set be-priority-group
```

NOTE: If there is an existing CoS configuration of any type on an interface, the global configuration is not applied to that particular interface. The global configuration is applied to all interfaces that do not have an existing CoS configuration.

For example, if you configure a rewrite rule, assign it to interfaces `xe-0/0/20.0` and `xe-0/0/22.0`, and then configure a forwarding class set and apply it to all interfaces, the forwarding class set is applied to every interface except `xe-0/0/20` and `xe-0/0/22`.

RELATED DOCUMENTATION

[Monitoring Interfaces That Have CoS Components | 274](#)

Understanding Junos CoS Components

Understanding CoS Packet Flow

When a packet traverses a switch, the switch provides the appropriate level of service to the packet using either default *class-of-service* (CoS) settings or CoS settings that you configure. On ingress ports, the switch classifies packets into appropriate forwarding classes and assigns a loss priority to the packets. On egress ports, the switch applies packet scheduling and (if you have configured them) *rewrite rules* to re-mark packets.

You can configure CoS on Layer 2 logical interfaces, and you can configure CoS on Layer 3 physical interfaces if you have defined at least one *logical interface* on the Layer 3 physical interface. You cannot configure CoS on Layer 2 physical interfaces and Layer 3 logical interfaces.

For Layer 2 traffic, either use the default CoS settings or configure CoS on each logical interface. You can apply different CoS settings to different Layer 2 logical interfaces.

NOTE: OCX Series switches do not support Layer 2 interfaces (family ethernet-switching).

For Layer 3 traffic, either use the default CoS settings or configure CoS on the physical interface (not on the logical unit). The switch uses the CoS applied on the physical Layer 3 interface for all logical Layer 3 interfaces configured on the physical Layer 3 interface.

The switch applies CoS to packets as they flow through the system:

- An interface has one or more classifiers of different types applied to it (configure this at the [edit class-of-service interfaces] hierarchy level). The classifier types are based on the portion of the incoming packet that the classifier examines (IEEE 802.1p code point bits or DSCP code point bits).
- When a packet enters an ingress port, the classifier assigns the packet to a forwarding class and a loss priority based on the code point bits of the packet (configure this at the [edit class-of-service classifiers] hierarchy level).
- The switch assigns each forwarding class to an output queue (configure this at the [edit class-of-service forwarding-classes] hierarchy level).
- Input (and output) policers meter traffic and can change the forwarding class and loss priority if a traffic flow exceeds its service level.
- A scheduler map is applied to each interface. When a packet exits an egress port, the scheduler map controls how it is treated (configure this at the [edit class-of-service interfaces] hierarchy level). A scheduler map assigns schedulers to forwarding classes (configure this at the [edit class-of-service scheduler-maps] hierarchy level).
- A scheduler defines how traffic is treated at the egress interface output queue (configure this at the [edit class-of-service schedulers] hierarchy level). You control the transmit rate, shaping rate, priority, and drop profile of each forwarding class by mapping schedulers to forwarding classes in scheduler maps, then applying scheduler maps to interfaces.
- A drop-profile defines how aggressively to drop packets that are mapped to a particular scheduler (configure this at the [edit class-of-service drop-profiles] hierarchy level).
- A rewrite rule takes effect as the packet leaves an interface that has a rewrite rule configured (configure this at the [edit class-of-service rewrite-rules] hierarchy level). The rewrite rule writes information to the packet (for example, a rewrite rule can re-mark the code point bits of outgoing traffic) according to the forwarding class and loss priority of the packet.

Figure 2 on page 21 is a high-level flow diagram of how packets from various sources enter switch interfaces, are classified at the ingress, and then scheduled (provided bandwidth) at the egress queues.

Figure 2: CoS Classifier, Queues, and Scheduler

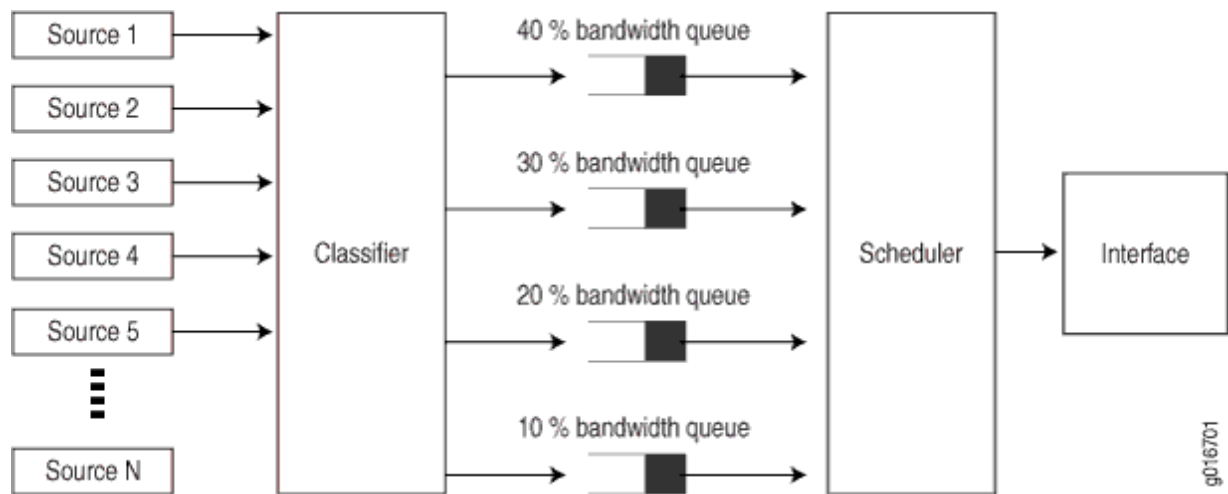



Figure 3 on page 22 shows the packet flow through the CoS components that you can configure.

Figure 3: Packet Flow Through Configurable CoS Components



The middle box (Forwarding Class and Loss Priority) represents two values that you can use on ingress and egress interfaces. The system uses these values for classifying traffic on ingress interfaces and for rewrite rule re-marking on egress interfaces. Each outer box represents a process component. The components in the top row apply to incoming packets. The components in the bottom row apply to outgoing packets.

The solid-line arrows show the direction of packet flow from ingress to egress. The dotted-line arrows that point to the forwarding class and loss priority box indicate processes that configure (set) the forwarding class and loss priority. The dotted-line arrows that point away from the forwarding class and loss priority box indicate processes that use forwarding class and loss priority as input values on which to base actions.

For example, the BA classifier sets the forwarding class and loss priority of incoming packets, so the forwarding class and loss priority are outputs of the classifier and the arrow points away from the classifier. The scheduler receives the forwarding class and loss priority settings, and queues the outgoing packets based on those settings, so the arrow points toward the scheduler.

Understanding Default CoS Settings

IN THIS SECTION

- [Default Forwarding Classes and Queue Mapping | 23](#)
- [Default Forwarding Class Sets \(Priority Groups\) | 24](#)
- [Default Code-Point Aliases | 25](#)
- [Default Classifiers | 27](#)
- [Default Rewrite Rules | 32](#)
- [Default Drop Profile | 32](#)
- [Default Schedulers | 32](#)
- [Default Scheduler Maps | 36](#)
- [Default Shared Buffer Configuration | 36](#)

If you do not configure CoS settings, Junos OS performs some CoS functions to ensure that traffic and protocol packets are forwarded with minimum delay when the network experiences congestion. Some default mappings are automatically applied to each *logical interface* that you configure.

You can display default CoS settings by issuing the `show class-of-service` *operational mode command*.

This topic describes the default configurations for the following CoS components:

Default Forwarding Classes and Queue Mapping

[Table 2 on page 23](#) shows the default mapping of the default forwarding classes to queues and packet drop attribute.

Table 2: Default Forwarding Classes and Queue Mapping

Default Forwarding Class	Description	Default Queue Mapping	Packet Drop Attribute
best-effort (be)	Best-effort traffic class (priority 0, IEEE 802.1p code point 000)	0	drop

Table 2: Default Forwarding Classes and Queue Mapping *(Continued)*

Default Forwarding Class	Description	Default Queue Mapping	Packet Drop Attribute
fcoe	Guaranteed delivery for FCoE traffic (priority 3, IEEE 802.1p code point 011)	3	no-loss
no-loss	Guaranteed delivery for TCP no-loss traffic (priority 4, IEEE 802.1p code point 100)	4	no-loss
network-control (nc)	Network control traffic (priority 7, IEEE 802.1p code point 111)	7	drop
(Excluding QFX10000) mcast	Multidestination traffic	8	drop NOTE: You cannot configure multidestination forwarding classes as no-loss (lossless) traffic classes.

NOTE: On the QFX10000 switch, unicast and multidestination (multicast, broadcast, and destination lookup fail) traffic use the same forwarding classes and output queues 0 through 7.

Default Forwarding Class Sets (Priority Groups)

If you do not explicitly configure forwarding class sets, the system automatically creates a default forwarding class set that contains all of the forwarding classes on the switch. The system assigns 100 percent of the port output bandwidth to the default forwarding class set.

Ingress traffic is classified based on the default classifier settings. The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default scheduler settings. Forwarding classes that are not part of the default scheduler receive no bandwidth.

The default forwarding class set is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange (DCBX) protocol advertisement.

Default Code-Point Aliases

[Table 3 on page 25](#) shows the default mapping of code-point aliases to IEEE code points.

Table 3: Default IEEE 802.1 Code-Point Aliases

CoS Value Types	Mapping
be	000
be1	001
ef	010
ef1	011
af11	100
af12	101
nc1	110
nc2	111

[Table 4 on page 25](#) shows the default mapping of code-point aliases to DSCP and DSCP IPv6 code points.

Table 4: Default DSCP and DCSP IPv6 Code-Point Aliases

CoS Value Types	Mapping
ef	101110
af11	001010

Table 4: Default DSCP and DCSP IPv6 Code-Point Aliases *(Continued)*

CoS Value Types	Mapping
af12	001100
af13	001110
af21	010010
af22	010100
af23	010110
af31	011010
af32	011100
af33	011110
af41	100010
af42	100100
af43	100110
be	000000
cs1	001000
cs2	010000
cs3	011000

Table 4: Default DSCP and DCSP IPv6 Code-Point Aliases (Continued)

CoS Value Types	Mapping
cs4	100000
cs5	101000
nc1	110000
nc2	111000

Default Classifiers

The switch applies default unicast IEEE 802.1, unicast DSCP, and multidestination classifiers to each interface that does not have explicitly configured classifiers. If you explicitly configure one type of classifier but not other types of classifiers, the system uses only the configured classifier and does not use default classifiers for other types of traffic.

NOTE: The QFX10000 switch applies the default MPLS EXP classifier to a logical interface if you enable the MPLS protocol family on that interface.

There are two different default unicast IEEE 802.1 classifiers, a trusted classifier for ports that are in trunk mode or tagged-access mode, and an untrusted classifier for ports that are in access mode. [Table 5 on page 27](#) shows the default mapping of IEEE 802.1 code-point values to forwarding classes and loss priorities for ports in trunk mode or tagged-access mode.

Table 5: Default IEEE 802.1 Classifiers for Ports in Trunk Mode or Tagged Access Mode (Trusted Classifier)

Code Point	Forwarding Class	Loss Priority
be (000)	best-effort	low
be1 (001)	best-effort	low

Table 5: Default IEEE 802.1 Classifiers for Ports in Trunk Mode or Tagged Access Mode (Trusted Classifier) (Continued)

Code Point	Forwarding Class	Loss Priority
ef (010)	best-effort	low
ef1 (011)	fcoe	low
af11 (100)	no-loss	low
af12 (101)	best-effort	low
nc1 (110)	network-control	low
nc2 (111)	network-control	low

[Table 6 on page 28](#) shows the default mapping of IEEE 802.1p code-point values to forwarding classes and loss priorities for ports in access mode (all incoming traffic is mapped to best-effort forwarding classes).

Table 6: Default IEEE 802.1 Classifiers for Ports in Access Mode (Untrusted Classifier)

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

Table 6: Default IEEE 802.1 Classifiers for Ports in Access Mode (Untrusted Classifier) (Continued)

Code Point	Forwarding Class	Loss Priority
110	best-effort	low
111	best-effort	low

[Table 7 on page 29](#) shows the default mapping of IEEE 802.1 code-point values to multidestination (multicast, broadcast, and destination lookup fail traffic) forwarding classes and loss priorities.

Table 7: Default IEEE 802.1 Multidestination Classifiers

Code Point	Forwarding Class	Loss Priority
be (000)	mcast	low
be1 (001)	mcast	low
ef (010)	mcast	low
ef1 (011)	mcast	low
af11 (100)	mcast	low
af12 (101)	mcast	low
nc1 (110)	mcast	low
nc2 (111)	mcast	low

[Table 8 on page 30](#) shows the default mapping of DSCP code-point values to forwarding classes and loss priorities for DSCP IP and DCSP IPv6.

NOTE: There are no default DSCP IP classifiers for multideestination traffic. DSCP IPv6 classifiers are not supported for multideestination traffic.

Table 8: Default DSCP IP and IPv6 Classifiers

Code Point	Forwarding Class	Loss Priority
ef (101110)	best-effort	low
af11 (001010)	best-effort	low
af12 (001100)	best-effort	low
af13 (001110)	best-effort	low
af21 (010010)	best-effort	low
af22 (010100)	best-effort	low
af23 (010110)	best-effort	low
af31 (011010)	best-effort	low
af32 (011100)	best-effort	low
af33 (011110)	best-effort	low
af41 (100010)	best-effort	low
af42 (100100)	best-effort	low
af43 (100110)	best-effort	low

Table 8: Default DSCP IP and IPv6 Classifiers (Continued)

Code Point	Forwarding Class	Loss Priority
be (000000)	best-effort	low
cs1 (001000)	best-effort	low
cs2 (010000)	best-effort	low
cs3 (011000)	best-effort	low
cs4 (100000)	best-effort	low
cs5 (101000)	best-effort	low
nc1 (110000)	network-control	low
nc2 (111000)	network-control	low

On QFX10000 switches, [Table 9 on page 31](#) shows the default mapping of MPLS EXP code-point values to forwarding classes and loss priorities.

Table 9: Default EXP Classifiers on QFX10000 Switches

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

Table 9: Default EXP Classifiers on QFX10000 Switches (Continued)

Code Point	Forwarding Class	Loss Priority
101	assured-forwarding	high
110	network-control	low
111	network-control	high

Default Rewrite Rules

There are no default *rewrite rules*. If you do not explicitly configure rewrite rules, the switch does not reclassify egress traffic.

Default Drop Profile

[Table 10 on page 32](#) shows the default drop profile configuration.

Table 10: Default Drop Profile

Fill Level	Drop Probability
100	100

Default Schedulers

[Table 11 on page 33](#) shows the default scheduler configuration.

Table 11: Default Schedulers

Default Scheduler and Queue Number	Transmit Rate (Guaranteed Minimum Bandwidth)	Shaping Rate (Maximum Bandwidth)	Excess Bandwidth Sharing	Priority	Buffer Size
best-effort forwarding class scheduler (queue 0)	5% (QFX10000 15%)	None	5% (QFX10000 15%)	low	5% (QFX10000 15%)
fcoe forwarding class scheduler (queue 3)	35%	None	35%	low	35%
no-loss forwarding class scheduler (queue 4)	35%	None	35%	low	35%
network-control forwarding class scheduler (queue 7)	5% (QFX10000 15%)	None	5% (QFX10000 15%)	low	5% (QFX10000 15%)
(Excluding QFX10000) mcast forwarding class scheduler (queue 8)	20%	None	20%	low	20%

NOTE: The minimum guaranteed bandwidth (transmit rate) also determines the amount of excess (extra) bandwidth that the queue can share. Extra bandwidth is allocated to queues in proportion to the transmit rate of each queue. On QFX10000 switches, you can use the `excess-rate` statement to override the default transmit rate setting and configure the excess bandwidth percentage independently of the transmit rate.

By default, only the five default schedulers shown in [Table 11 on page 33](#), excluding the mcast scheduler on QFX10000 switches, have traffic mapped to them. Only the queues associated with the default schedulers, and forwarding classes on QFX10000 switches, receive default bandwidth, based on the default scheduler transmit rate. (You can configure schedulers and forwarding classes to allocate bandwidth to other queues or to change the default bandwidth of a default queue.) In addition, other than on QFX5200, QFX5210, and QFX10000 switches, multidestination queue 11 receives enough bandwidth from the default multidestination scheduler to handle CPU-generated multidestination

traffic. If a forwarding class does not transport traffic, the bandwidth allocated to that forwarding class is available to other forwarding classes.

NOTE: On QFX10000 switches, unicast and multdestination (multicast, broadcast, and destination lookup fail) traffic use the same forwarding classes and output queues.

Default hierarchical scheduling, known as enhanced transmission selection (ETS, defined in IEEE 802.1Qaz), divides the total port bandwidth between two groups of traffic: unicast traffic and multdestination traffic. By default, unicast traffic consists of queue 0 (best-effort forwarding class), queue 3 (fcoe forwarding class), queue 4 (no-loss forwarding class), and queue 7 (network-control forwarding class). Unicast traffic receives and shares a total of 80 percent of the port bandwidth. By default, multdestination traffic (mcast queue 8) receives a total of 20 percent of the port bandwidth. So on a 10-Gigabit port, default scheduling provides unicast traffic 8-Gbps of bandwidth and multdestination traffic 2-Gbps of bandwidth.

NOTE: Except on QFX5200, QFX5210, and QFX10000 switches, multdestination queue 11 also receives a small amount of default bandwidth from the multdestination scheduler. CPU-generated multdestination traffic uses queue 11, so you might see a small number of packets egress from queue 11. In addition, in the unlikely case that firewall filter match conditions map multdestination traffic to a unicast forwarding class, that traffic uses queue 11.

On QFX10000 switches, default scheduling is port scheduling. Default hierarchical scheduling, known as ETS, allocates the total port bandwidth to the four default forwarding classes served by the four default schedulers, as defined by the four default schedulers. The result is the same as direct port scheduling. Configuring hierarchical port scheduling, however, enables you to group forwarding classes that carry similar types of traffic into forwarding class sets (also called priority groups), and to assign port bandwidth to each forwarding class set. The port bandwidth assigned to the forwarding class set is then assigned to the forwarding classes within the forwarding class set. This hierarchy enables you to control port bandwidth allocation with greater granularity, and enables hierarchical sharing of extra bandwidth to better utilize link bandwidth.

Default scheduling for all switches uses weighted round-robin (WRR) scheduling. Each queue receives a portion (weight) of the total available interface bandwidth. The scheduling weight is based on the transmit rate of the default scheduler for that queue. For example, queue 7 receives a default scheduling weight of 5 percent, 15 percent on QFX10000 switches, of the available bandwidth, and queue 4 receives a default scheduling weight of 35 percent of the available bandwidth. Queues are mapped to forwarding classes (for example, queue 7 is mapped to the network-control forwarding class and queue 4 is mapped to the no-loss forwarding class), so forwarding classes receive the default bandwidth for the queues to which they are mapped. Unused bandwidth is shared with other default queues.

If you want non-default (unconfigured) queues to forward traffic, you should explicitly map traffic to those queues (configure the forwarding classes and queue mapping) and create schedulers to allocate bandwidth to those queues. For example, except on QFX5200, QFX5210, and QFX10000 switches, by default, queues 1, 2, 5, and 6 are unconfigured, and multidestination queues 9, 10, and 11 are unconfigured. Unconfigured queues have a default scheduling weight of 1 so that they can receive a small amount of bandwidth in case they need to forward traffic. (However, queue 11 can use more of the default multidestination scheduler bandwidth if necessary to handle CPU-generated multidestination traffic.)

NOTE: Except on QFX10000 switches, all four multidestination queues, or two for QFX5200 and QFX5210, switches, have a scheduling weight of 1. Because by default multidestination traffic goes to queue 8, queue 8 receives almost all of the multidestination bandwidth. (There is no default traffic on queue 9 and queue 10, and very little default traffic on queue 11, so there is almost no competition for multidestination bandwidth.)

However, if you explicitly configure queue 9, 10, or 11 (by mapping code points to the unconfigured multidestination forwarding classes using the multidestination classifier), the explicitly configured queues share the multidestination scheduler bandwidth equally with default queue 8, because all of the queues have the same scheduling weight (1). To ensure that multidestination bandwidth is allocated to each queue properly and that the bandwidth allocation to the default queue (8) is not reduced too much, we strongly recommend that you configure a scheduler if you explicitly classify traffic into queue 9, 10, or 11.

If you map traffic to an unconfigured queue, the queue receives only the amount of group bandwidth proportional to its default weight (1). The actual amount of bandwidth an unconfigured queue receives depends on how much bandwidth the other queues in the group are using.

On QFX 10000 switches, if you map traffic to an unconfigured queue and do not schedule port resources for the queue (configure a scheduler, map it to the forwarding class that is mapped to the queue, and apply the scheduler mapping to the port), the queue receives only the amount of excess bandwidth proportional to its default weight (1). The actual amount of bandwidth an unconfigured queue gets depends on how much bandwidth the other queues on the port are using.

If the other queues use less than their allocated amount of bandwidth, the unconfigured queues can share the unused bandwidth. Configured queues have higher priority for bandwidth than unconfigured queues, so if a configured queue needs more bandwidth, then less bandwidth is available for unconfigured queues. Unconfigured queues always receive a minimum amount of bandwidth based on their scheduling weight (1). If you map traffic to an unconfigured queue, to allocate bandwidth to that queue, configure a scheduler for the forwarding class that is mapped to the queue and apply it to the port.

Default Scheduler Maps

Table 12 on page 36 shows the default mapping of forwarding classes to schedulers.

Table 12: Default Scheduler Maps

Forwarding Class	Scheduler
best-effort	Default BE scheduler
fcoe	Default FCoE scheduler
no-loss	No-loss scheduler
network-control	Default network-control scheduler
(Excluding QFX10000) mcast-be	Default multidestination scheduler

Default Shared Buffer Configuration

Table Table 13 on page 36 and Table 14 on page 37 show the default shared buffer allocations:


**NOTE:** Shared buffers do not apply to QFX10000 switches.

Table 13: Default Ingress Shared Buffer Configuration

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
100%	9%	45%	46%

Table 14: Default Egress Shared Buffer Configuration

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
100%	50%	31%	19%

RELATED DOCUMENTATION

<i>Overview of Junos OS CoS</i>
<i>Understanding Junos CoS Components</i>
<i>Understanding Default CoS Scheduling and Classification</i>
<i>Understanding CoS Classifiers</i>
Understanding CoS Classifiers
<i>Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces</i>
<i>Understanding CoS Code-Point Aliases</i>
<i>Understanding CoS Forwarding Classes</i>
<i>Understanding CoS Rewrite Rules</i>
<i>Understanding CoS Output Queue Schedulers</i>
<i>Understanding CoS Port Schedulers on QFX Switches</i>
<i>Understanding CoS WRED Drop Profiles</i>

CoS Inputs and Outputs Overview

Some CoS components map one set of values to another set of values. Each mapping contains one or more inputs and one or more outputs. When you configure a mapping, you set the outputs for a given set of inputs, as shown in [Table 15 on page 38](#).

Table 15: CoS Mappings—Inputs and Outputs

CoS Mappings	Inputs	Outputs	Comments
classifiers	code-points	forwarding-class, loss-priority	The map sets the forwarding class and packet loss priority (PLP) for a specific set of code points.
drop-profile-map	loss-priority, protocol	drop-profile	The map sets the drop profile for a specific PLP and protocol type.
rewrite-rules	loss-priority, forwarding-class	code-points	The map sets the code points for a specific forwarding class and PLP.
rewrite-value (Fibre Channel Interfaces)	<i>forwarding-class</i>	<i>code-point</i>	(Systems that support native Fibre Channel interfaces only) The map sets the code point for the forwarding class specified in the fixed classifier attached to the native Fibre Channel (NP_Port) interface.

RELATED DOCUMENTATION

| *Understanding CoS Packet Flow*

Overview of Policers

IN THIS SECTION

- [Policer Overview | 39](#)
- [Policer Types | 42](#)
- [Policer Actions | 43](#)
- [Policer Colors | 44](#)
- [Filter-Specific Policers | 44](#)
- [Suggested Naming Convention for Policers | 44](#)

- [Policer Counters | 45](#)
- [Policer Algorithms | 45](#)
- [How Many Policers Are Supported? | 45](#)
- [Policers Can Limit Egress Firewall Filters | 46](#)

A switch polices traffic by limiting the input or output transmission rate of a class of traffic according to user-defined criteria. Policing (or rate-limiting) traffic allows you to control the maximum rate of traffic sent or received on an interface and to provide multiple priority levels or classes of service.

Policing is also an important component of firewall filters. You can achieve policing by including policers in *firewall filter* configurations.

Policer Overview

You use policers to apply limits to traffic flow and set consequences for packets that exceed these limits—usually applying a higher loss priority—so that if packets encounter downstream congestion, they can be discarded first. Policers apply only to unicast packets.

Policers provide two functions: metering and marking. A policer meters (measures) each packet against traffic rates and burst sizes that you configure. It then passes the packet and the metering result to the marker, which assigns a packet loss priority that corresponds to the metering result. [Figure 4 on page 41](#) illustrates this process.


NOTE: A policer restricts traffic at the configured transmission rate per PFE. In QFX10016, QFX10002, QFX10002-60C, and QFX10008 switches, when aggregated ethernet (AE) interface bundles span multiple PFEs, the overall transmission rate of the policer for the subscriber could exceed the configured transmission rate of the policer (depending on the number of PFEs involved).

As an example:

- Policer with bandwidth-limit 100 mbps configured on an AE interface that has member links xe-1/0/0 (fpc1-pfe0) and xe-1/0/30 (fpc1-pfe1) . Here, the two member links belong to FPC1, but are on different PFEs. When the policer is applied to the AE interface, this will result in a total bandwidth of 200 Mbps as policer is configured for two PFEs.

- Policer with bandwidth-limit 100 mbps configured on an AE interface that has member links xe-1/0/0 (fpc1-pfe0), et-2/0/1 (fpc2-pfe1) and xe-2/0/18:0 (fpc2-pfe2) . Here, one member link belongs to FPC1 and PFE0 on this FPC. The rest two member links belong to FPC2, but different PFEs. When the policer is applied to the AE interface, this will result in a total bandwidth of 300 Mbps as policer is configured for three PFEs.
- Policer with bandwidth-limit 100 mbps configured on an AE interface that has member links xe-1/0/0 and xe-1/0/1 on a single PFE (fpc1-pfe0) . Here, the member links belong to FPC1 and to the same PFE. When the policer is applied to the AE interface, this will result in a total bandwidth of 100 Mbps as policer is configured on a per PFE basis.

Figure 4: Flow of Tricolor Marking Policer Operation



g017049

After you name and configure a policer, you can use it by specifying it as an action in one or more firewall filters.

Policer Types

A switch supports three types of policers:

- **Single-rate two-color marker**—A two-color policer (or “policer” when used without qualification) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit with a specified PLP or simply discard them.

You can specify this type of policer in an ingress or egress firewall.

NOTE: A two-color policer is most useful for metering traffic at the port (physical interface) level.

- **Single-rate three-color marker**—This type of policer is defined in RFC 2697, *A Single Rate Three Color Marker*, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on one rate—the configured committed information rate (CIR) as well as the committed burst size (CBS) and the excess burst size (EBS). The CIR specifies the average rate at which bits are admitted to the switch. The CBS specifies the usual burst size in bytes and the EBS specifies the maximum burst size in bytes. The EBS must be greater than or equal to the CBS, and neither can be 0.

You can specify this type of policer in an ingress or egress firewall.

NOTE: A single-rate three-color marker (TCM) is most useful when a service is structured according to packet length and not peak arrival rate.

- **Two-rate three-color marker**—This type of policer is defined in RFC 2698, *A Two Rate Three Color Marker*, as part of an assured forwarding per-hop-behavior classification system for a Differentiated Services environment. This type of policer meters traffic based on two rates—the CIR and peak information rate (PIR) along with their associated burst sizes, the CBS and peak burst size (PBS). The PIR specifies the maximum rate at which bits are admitted to the network and must be greater than or equal to the CIR.

You can specify this type of policer in an ingress or egress firewall.

NOTE: A two-rate three-color policer is most useful when a service is structured according to arrival rates and not necessarily packet length.

See [Table 16 on page 43](#) for information about how metering results are applied for each of these policer types.

Policer Actions

Policer actions are implicit or explicit and vary by policer type. *Implicit* means that Junos OS assigns the loss priority automatically. [Table 16 on page 43](#) describes the policer actions.

Table 16: Policer Actions

Policer	Marking	Implicit Action	Configurable Action
Single-rate two-color	Green (conforming)	Assign low loss priority	None
	Red (nonconforming)	None	Discard
Single-rate three-color	Green (conforming)	Assign low loss priority	None
	Yellow (above the CIR and CBS)	Assign medium-high loss priority	None
	Red (above the EBS)	Assign high loss priority	Discard
Two-rate three-color	Green (conforming)	Assign low loss priority	None
	Yellow (above the CIR and CBS)	Assign medium-high loss priority	None
	Red (above the PIR and PBS)	Assign high loss priority	Discard

NOTE: If you specify a policer in an egress *firewall filter*, the only supported action is discard.

Policer Colors

Single-rate and two-rate three-color policers can operate in two modes:

- **Color-blind**—In color-blind mode, the three-color policer assumes that all packets examined have not been previously marked or metered. In other words, the three-color policer is “blind” to any previous coloring a packet might have had.
- **Color-aware**—In color-aware mode, the three-color policer assumes that all packets examined have been previously marked or metered. In other words, the three-color policer is “aware” of the previous coloring a packet might have had. In color-aware mode, the three-color policer can increase the PLP of a packet but cannot decrease it. For example, if a color-aware three-color policer meters a packet with a medium PLP marking, it can raise the PLP level to high but cannot reduce the PLP level to low.

Filter-Specific Policers

You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this on some QFX switches, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in TCAM. If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps. (This behavior does not occur in QFX10000 switches.)

To prevent this unexpected behavior from occurring, use the information about TCAM slices presented in *Planning the Number of Firewall Filters to Create* to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

Suggested Naming Convention for Policers

We recommend that you use the naming convention *policertypeTCM#-color type* when configuring three-color policers and *policer#* when configuring two-color policers. TCM stands for three-color marker. Because policers can be numerous and must be applied correctly to work, a simple naming convention makes it easier to apply the policers properly. For example, the first single-rate, color-aware three-color

policer configured would be named `srTCM1-ca`. The second two-rate, color-blind three-color configured would be named `trTCM2-cb`. The elements of this naming convention are explained below:

- `sr` (single-rate)
- `tr` (two-rate)
- `TCM` (tricolor marking)
- `1` or `2` (number of marker)
- `ca` (color-aware)
- `cb` (color-blind)

Policer Counters

On some QFX switches, each policer that you configure includes an implicit counter that counts the number of packets that exceed the rate limits that are specified for the policer. If you use the same policer in multiple terms—either within the same filter or in different filters—the implicit counter counts all the packets that are policed in all of these terms and provides the total amount. (This does not apply to QFX10000 switches.) If you want to obtain separate packet counts for each term on an affected switch, use these options:

- Configure a unique policer for each term.
- Configure only one policer, but use a unique, explicit counter in each term.

Policer Algorithms

Policing uses the *token-bucket algorithm*, which enforces a limit on average bandwidth while allowing bursts up to a specified maximum value. It offers more flexibility than the *leaky bucket algorithm* in allowing a certain amount of bursty traffic before it starts discarding packets.

NOTE: In an environment of light bursty traffic, QFX5200 might not replicate all multicast packets to two or more downstream interfaces. This occurs only at a line rate burst—if traffic is consistent, the issue does not occur. In addition, the issue occurs only when packet size increases beyond 6k in a one gigabit traffic flow.

How Many Policers Are Supported?

QFX10000 switches support 8K policers (all policer types). QFX5100 and QFX5200 switches support 1535 ingress policers and 1024 egress policers (assuming one policer per firewall filter term). QFX5110

switches support 6144 ingress policers and 1024 egress policers (assuming one policer per firewall filter term).

QFX3500 and QFX3600 standalone switches and QFabric Node devices support the following numbers of policers (assuming one policer per firewall filter term):

- Two-color policers used in ingress firewall filters: 767
- Three-color policers used in ingress firewall filters: 767
- Two-color policers used in egress firewall filters: 1022
- Three-color policers used in egress firewall filters: 512

Policers Can Limit Egress Firewall Filters

On some switches, the number of egress policers you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that take up two entries in a 1024-entry TCAM. These are used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you are unable to commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters get used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
 - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
 - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits

policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

RELATED DOCUMENTATION

Understanding Color-Blind Mode for Single-Rate Tricolor Marking

Understanding Color-Blind Mode for Two-Rate Tricolor Marking

Understanding Color-Aware Mode for Single-Rate Tricolor Marking

Understanding Color-Aware Mode for Two-Rate Tricolor Marking

Configuring Two-Color and Three-Color Policers to Control Traffic Rates

2

PART

Classifying and Rewriting Traffic

Using Classifiers, Forwarding Classes, and Rewrite Rules | 49

CHAPTER 2

Using Classifiers, Forwarding Classes, and Rewrite Rules

IN THIS CHAPTER

- [Understanding CoS Classifiers | 49](#)
- [Defining CoS BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) | 59](#)
- [Example: Configuring Classifiers | 62](#)
- [Understanding Default CoS Scheduling and Classification | 66](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces | 77](#)
- [Understanding CoS Code-Point Aliases | 92](#)
- [Defining CoS Code-Point Aliases | 95](#)
- [Understanding CoS Forwarding Classes | 96](#)
- [Defining CoS Forwarding Classes | 102](#)
- [Example: Configuring Forwarding Classes | 105](#)
- [Understanding CoS Rewrite Rules | 112](#)
- [Defining CoS Rewrite Rules | 114](#)
- [Troubleshooting an Unexpected Rewrite Value | 117](#)

Understanding CoS Classifiers

IN THIS SECTION

- [Interfaces and Output Queues | 50](#)
- [Output Queues for Unicast and Multidestination Traffic | 51](#)
- [Classifier Support by Type | 52](#)
- [Behavior Aggregate Classifiers | 53](#)

- Fixed Classifiers on Ethernet Interfaces | 57
- Fixed Classifiers on Native Fibre Channel Interfaces (NP_Ports) | 58
- Multifield Classifiers | 58
- MPLS EXP Classifiers | 58
- Packet Classification for IRB Interfaces and RVIs | 59

Packet classification maps incoming packets to a particular class-of-service (CoS) servicing level. Classifiers map packets to a forwarding class and a loss priority, and they assign packets to output queues based on the forwarding class. There are three general types of classifiers:

- Behavior aggregate (BA) classifiers—DSCP and DSCP IPv6 classify IP and IPv6 traffic, EXP classifies MPLS traffic, and IEEE 802.1p classifies all other traffic. (Although this topic covers EXP classifiers, for more details, see *Understanding CoS MPLS EXP Classifiers and Rewrite Rules*. EXP classifiers are applied only on family `mpls` interfaces.)
- Fixed classifiers—Fixed classifiers classify all ingress traffic on a physical interface into one forwarding class, regardless of the CoS bits in the packet header.
- Multifield (MF) classifiers—MF classifiers classify traffic based on more than one field in the packet header and take precedence over BA and fixed classifiers.

Classifiers assign incoming unicast and multideestination (multicast, broadcast, and destination lookup fail) traffic to forwarding classes, so that different classes of traffic can receive different treatment. Classification is based on CoS bits, DSCP bits, EXP bits, a forwarding class (fixed classifier), or packet headers (multifield classifiers). Each classifier assigns all incoming traffic that matches the classifier configuration to a particular forwarding class. Except on QFX10000 switches, classifiers and forwarding classes handle either unicast or multideestination traffic. You cannot mix unicast and multideestination traffic in the same classifier or forwarding class. On QFX10000 switches, a classifier can assign both unicast and multideestination traffic to the same forwarding class.

Interfaces and Output Queues

You can apply classifiers to Layer 2 *logical interface* unit 0 (but not to other logical interfaces), and to Layer 3 physical interfaces if the Layer 3 physical interface has at least one defined logical interface. Classifiers applied to Layer 3 physical interfaces are used on all logical interfaces on that physical interface. *Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces* describes the interaction between classifiers and interfaces in greater detail.

NOTE: On QFX10000 switches you can apply different classifiers to different Layer 3 logical interfaces. You cannot apply classifiers to physical interfaces.

You can configure both a BA classifier and an MF classifier on an interface. If you do this, the BA classification is performed first, and then the MF classification is performed. If the two classification results conflict, the MF classification result overrides the BA classification result.

You cannot configure a fixed classifier and a BA classifier on the same interface.

Except on QFX10000 switches, you can configure both a DSCP or DSCP IPv6 classifier and an IEEE 802.1p classifier on the same interface. IP traffic uses the DSCP or DSCP IPv6 classifier. All other traffic uses the IEEE classifier (except when you configure a global EXP classifier; in that case, MPLS traffic uses the EXP classifier providing that the interface is configured as `family mpls`). You can configure only one DSCP classifier on a physical interface (either one DSCP classifier or one DSCP IPv6 classifier, but not both).

On QFX10000 switches, you can configure either a DSCP or a DSCP IPv6 classifier and also an IEEE 802.1p classifier on the same interface. IP traffic uses the DSCP or DSCP IPv6 classifier. If you configure an interface as `family mpls`, then the interface uses the default MPLS EXP classifier. If you configure an MPLS EXP classifier, then all MPLS traffic on the switch uses the global EXP classifier. All other traffic uses the IEEE classifier. You can configure up to 64 EXP classifiers with up to 8 entries per classifier (one entry for each forwarding class) and apply them to logical interfaces.

Except on QFX10000 switches, although you can configure as many EXP classifiers as you want, the switch uses only one MPLS EXP classifier as a global classifier on all interfaces.

After you configure an MPLS EXP classifier, you can configure it as the global EXP classifier by including the EXP classifier at the `[edit class-of-service system-defaults classifiers exp]` hierarchy level. All switch interfaces that are configured as `family mpls` use the EXP classifier, on QFX10000 switches either the default or the global EXP classifier, specified in this configuration statement to classify MPLS traffic.

Output Queues for Unicast and Multidestination Traffic

NOTE: This section applies to switches except QFX10000.

You can create unicast BA classifiers for unicast traffic and multicast BA classifiers for multidestination traffic, which includes multicast, broadcast, and destination lookup fail (DLF) traffic. You cannot assign unicast traffic and multidestination traffic to the same BA classifier.

On each interface, the switch has separate output queues for unicast traffic and for multidestination traffic:

NOTE: QFX5200 switches support 10 output queues, with 8 queues dedicated to unicast traffic and 2 queues dedicated to multdestination traffic.

- The switch supports 12 output queues, with 8 queues dedicated to unicast traffic and 4 queues dedicated to multdestination traffic.
- Queues 0 through 7 are unicast traffic queues. You can apply only unicast BA classifiers to unicast queues. A unicast BA classifier should contain only forwarding classes that are mapped to unicast queues.
- Queues 8 through 11 are multdestination traffic queues. You can apply only multdestination BA classifiers to multdestination queues. A multdestination BA classifier should contain only forwarding classes that are mapped to multdestination queues.

You can apply unicast classifiers to one or more interfaces. Multdestination classifiers and EXP classifiers apply to all of the switch interfaces and cannot be applied to individual interfaces. Use the DSCP multdestination classifier for both IP and IPv6 multdestination traffic. The DSCP IPv6 classifier is not supported for multdestination traffic.

Classifier Support by Type

NOTE: This section applies only to QFX10000 switches.

You can configure enough classifiers to handle most, if not all, network scenarios. [Table 17 on page 52](#) shows how many of each type of classifiers you can configure, and how many entries you can configure per classifier.

Table 17: Classifier Support by Classifier Type

Classifier Type	Default Classifier Name	Maximum Number of Classifiers	Maximum Number of Entries per Classifier
IEEE 802.1p (Layer 2)	ieee8021p-default (for ports in trunk mode) ieee8021p-untrust (for ports in access mode)	64	16

Table 17: Classifier Support by Classifier Type (Continued)

Classifier Type	Default Classifier Name	Maximum Number of Classifiers	Maximum Number of Entries per Classifier
DSCP (Layer 3)	dscp-default	64	64
DSCP IPv6 (Layer 3)	dscp-ipv6-default	64	64
EXP (MPLS)	exp-default	64	8
Fixed	There is no default fixed classifier	8	16

The number of fixed classifiers supported (8) equals the number of supported forwarding classes (fixed classifiers assign all incoming traffic on an interface to one forwarding class).

Behavior Aggregate Classifiers

Behavior aggregate classifiers map a class-of-service (CoS) value to a forwarding class and loss priority. The forwarding class determines the output queue. A scheduler uses the loss priority to control packet discard during periods of congestion by associating different drop profiles with different loss priorities.

The switch supports three types of BA classifiers:

- Differentiated Services code point (DSCP) for IP DiffServ (IP and IPv6)
- IEEE 802.1p CoS bits
- MPLS EXP (applies only to interfaces configured as `family mpls`)

BA classifiers are based on fixed-length fields, which makes them computationally more efficient than MF classifiers. Therefore, core devices, which handle high traffic volumes, are normally configured to perform BA classification.

Unicast and multicast traffic cannot share the same classifier. You can map unicast traffic and multicast traffic to the same classifier CoS value, but the unicast traffic must belong to a unicast classifier and the multicast traffic must belong to a multidestination classifier.

Default Behavior Aggregate Classification

Juniper Networks Junos OS automatically assigns implicit default classifiers to all logical interfaces based on the type of interface. [Table 18 on page 54](#) lists different types of interfaces and the corresponding implicit default BA classifiers.

Table 18: Default BA Classification

Type of Interface	Default BA Classification
Layer 2 interface in trunk mode or, except on QFX10000, tagged-access mode	ieee8021p-default
(QFX10000 only) Layer 2 interface in access mode	ieee8021p-untrusted
Layer 3 interface	dscp-default dscp-ipv6-default
(Except QFX10000) Layer 2 interface in access mode	ieee8021p-untrusted
(QFX10000 only) MPLS interface	exp-default

NOTE: Default BA classifiers assign traffic only to the best-effort, fcoe, no-loss, network-control, and, except on QFX10000 switches, mcast forwarding classes.

NOTE: Except on QFX10000 switches, there is no default MPLS EXP classifier. You must configure an EXP classifier and apply it globally to all interfaces that are configured as `family mpls` by including it in the `[edit class-of-service system-defaults classifiers exp]` hierarchy. On `family mpls` interfaces, if a fixed classifier is present on the interface, the EXP classifier overrides the fixed classifier.

If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the

interface, MPLS traffic is treated as best-effort traffic. DSCP classifiers are not applied to MPLS traffic.

Because the EXP classifier is global, you cannot configure some ports to use a fixed IEEE 802.1p classifier for MPLS traffic on some interfaces and the global EXP classifier for MPLS traffic on other interfaces. When you configure a global EXP classifier, all MPLS traffic on all interfaces uses the EXP classifier, even interfaces that have a fixed classifier.

When you explicitly associate a classifier with a logical interface, you override the default classifier with the explicit classifier. For other than QFX10000 switches, this applies to unicast classifiers.

NOTE: You can apply only one DSCP and one IEEE 802.1p classifier to a Layer 2 interface. If both types of classifiers are present, DSCP classifiers take precedence over IEEE 802.1p classifiers. If on QFX10000 switches you configure an EXP classifier, or on other switches a global EXP classifier, and apply it on interfaces configured as `family mpls`, then MPLS traffic uses that classifier on those interfaces.

Importing a Classifier

You can use any existing classifier, including the default classifiers, as the basis for defining a new classifier. You accomplish this using the `import` statement.

The imported classifier is used as a template and is not modified. The modifications you make become part of a new classifier (and a new template) identified by the name of the new classifier. Whenever you commit a configuration that assigns a new forwarding class-name and loss-priority value to a code-point alias or set of bits, it replaces the old entry in the new classifier template. As a result, you must explicitly specify every CoS value in every packet classification that requires modification.

Multidestination Classifiers

NOTE: This section applies to switches except QFX10000.

Multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces. You can configure both a DSCP multidestination classifier and an IEEE multidestination classifier. IP and IPv6 traffic use the DSCP classifier, and all other traffic uses the IEEE classifier.

DSCP IPv6 multidestination classifiers are not supported, so IPv6 traffic uses the DSCP multidestination classifier.

The default multideestination classifier is the IEEE 802.1p multideestination classifier.

PFC Priorities

The eight IEEE 802.1p code points correspond to the eight priorities that *priority-based flow control* (PFC) uses to differentiate traffic classes for lossless transport. When you map a forwarding class (which maps to an output queue) to an IEEE 802.1p CoS value, the IEEE 802.1p CoS value identifies the PFC priority.

Although you can map a priority to any output queue (by mapping the IEEE 802.1p code point value to a forwarding class), we recommend that the priority and the forwarding class (unicast except for QFX10000 switches) match in a one-to-one correspondence. For example, priority 0 is assigned to queue 0, priority 1 is assigned to queue 1, and so on, as shown in [Table 19 on page 56](#). A one-to-one correspondence of queue and priority numbers makes it easier to configure and maintain the mapping of forwarding classes to priorities and queues.

Table 19: Default IEEE 802.1p Code Point to PFC Priority, Output Queue, and Forwarding Class Mapping

IEEE 802.1p Code Point	PFC Priority	Output Queue (Unicast except for QFX10000)	Forwarding Class and Packet Drop Attribute
000	0	0	best-effort (drop)
001	1	1	best-effort (drop)
010	2	2	best-effort (drop)
011	3	3	fcoe (no-loss)
100	4	4	no-loss (no-loss)
101	5	5	best-effort (drop)
110	6	6	network-control (drop)
111	7	7	network-control (drop)

NOTE: By convention, deployments with converged server access typically use IEEE 802.1p priority 3 (011) for FCoE traffic. The default mapping of the `fcoe` forwarding class is to queue 3. Apply priority-based flow control (PFC) to the entire FCoE data path to configure the end-to-end lossless behavior that FCoE requires. We recommend that you use priority 3 for FCoE traffic unless your network architecture requires that you use a different priority.

Fixed Classifiers on Ethernet Interfaces

Fixed classifiers map all traffic on a physical interface to a forwarding class and a loss priority, unlike BA classifiers, which map traffic into multiple different forwarding classes based on the IEEE 802.1p CoS bits field value in the VLAN header or the DSCP field value in the type-of-service bits in the packet IP header. Each forwarding class maps to an output queue. However, when you use a fixed classifier, regardless of the CoS or DSCP bits, all Incoming traffic is classified into the forwarding class specified in the fixed classifier. A scheduler uses the loss priority to control packet discard during periods of congestion by associating different drop profiles with different loss priorities.

You cannot configure a fixed classifier and a DSCP or IEEE 802.1p BA classifier on the same interface. If you configure a fixed classifier on an interface, you cannot configure a DSCP or an IEEE classifier on that interface. If you configure a DSCP classifier, an IEEE classifier, or both classifiers on an interface, you cannot configure a fixed classifier on that interface.

NOTE: For MPLS traffic on the same interface, you can configure both a fixed classifier and an EXP classifier on QFX10000, or a global EXP classifier on other switches. When both an EXP classifier or global EXP classifier and a fixed classifier are applied to an interface, MPLS traffic on interfaces configured as `family mpls` uses the EXP classifier, and all other traffic uses the fixed classifier.

To switch from a fixed classifier to a BA classifier, or to switch from a BA classifier to a fixed classifier, deactivate the existing classifier attachment on the interface, and then attach the new classifier to the interface.

NOTE: If you configure a fixed classifier that classifies all incoming traffic into the `fcoe` forwarding class (or any forwarding class designed to handle FCoE traffic), you must ensure that all traffic that enters the interface is FCoE traffic and is tagged with the FCoE IEEE 802.1p code point (priority).

Fixed Classifiers on Native Fibre Channel Interfaces (NP_Ports)

NOTE: This section applies to switches except QFX10000.

Applying a fixed classifier to a native Fibre Channel (FC) interface (NP_Port) is a special case. By default, native FC interfaces classify incoming traffic from the FC SAN into the `fcoe` forwarding class and map the traffic to IEEE 802.1p priority 3 (code point 011). When you apply a fixed classifier to an FC interface, you also configure a priority rewrite value for the interface. The FC interface uses the priority rewrite value as the IEEE 802.1p tag value for all incoming packets instead of the default value of 3.

For example, if you specify a priority rewrite value of 5 (code point 101) for an FC interface, the interface tags all incoming traffic from the FC SAN with priority 5 and classifies the traffic into the forwarding class specified in the fixed classifier.

NOTE: The forwarding class specified in the fixed classifier on FC interfaces must be a lossless forwarding class.

Multifield Classifiers

Multifield classifiers examine multiple fields in a packet such as source and destination addresses and source and destination port numbers of the packet. With MF classifiers, you set the forwarding class and loss priority of a packet based on *firewall filter* rules.

MF classification is normally performed at the network edge because of the general lack of DiffServ code point (DSCP) support in end-user applications. On a switch at the edge of a network, an MF classifier provides the filtering functionality that scans through a variety of packet fields to determine the forwarding class for a packet. Typically, a classifier performs matching operations on the selected fields against a configured value.

MPLS EXP Classifiers

You can configure up to 64 EXP classifiers for MPLS traffic and apply them to `family mpls` interfaces. On QFX10000 switches you can use the default MPLS EXP, but on other switches there is no default MPLS classifier. You can configure an EXP classifier and apply it globally to all interfaces that are configured as `family mpls` by including it in the `[edit class-of-service system-defaults classifiers exp]` hierarchy level. On `family mpls` interfaces, if a fixed classifier is present on the interface, the EXP classifier overrides the fixed classifier for MPLS traffic only.

Except on QFX10000 switches, if an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is

applied to the interface, MPLS traffic is treated as best-effort traffic. DSCP classifiers are not applied to MPLS traffic.

Because the EXP classifier is global, you cannot configure some ports to use a fixed IEEE 802.1p classifier for MPLS traffic on some interfaces and the global EXP classifier for MPLS traffic on other interfaces. When you configure a global EXP classifier, all MPLS traffic on all interfaces uses the EXP classifier, even interfaces that have a fixed classifier.

For details about EXP classifiers, see *Understanding CoS MPLS EXP Classifiers and Rewrite Rules*. EXP classifiers are applied only on family `mpls` interfaces.

Packet Classification for IRB Interfaces and RVIs

On QFX10000 switches, you cannot apply classifiers directly to integrated routing and bridging (*IRB*) interfaces. Similarly, on other switches you cannot apply classifiers directly to routed VLAN interfaces (*RVIs*). This results because the members of IRBs and RVIs are VLANs, not ports. However, you can apply classifiers to the VLAN port members of an IRB interface. You can also apply MF classifiers to IRBs and RVIs.

RELATED DOCUMENTATION

<i>Understanding CoS MPLS EXP Classifiers and Rewrite Rules</i>
<i>Understanding CoS Packet Flow</i>
<i>Understanding Default CoS Settings</i>
<i>Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces</i>
<i>Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)</i>
<i>Example: Configuring Unicast Classifiers</i>
<i>Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)</i>
<i>Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers</i>
<i>Configuring a Global MPLS EXP Classifier</i>
<i>Configuring Rewrite Rules for MPLS EXP Classifiers</i>

Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)

Overview

Packet classification associates incoming packets with a particular CoS servicing level. Behavior aggregate (BA) classifiers examine the Differentiated Services code point (DSCP or DSCP IPv6) value,

the IEEE 802.1p CoS value, or the MPLS EXP value in the packet header to determine the CoS settings applied to the packet. (See *Configuring a Global MPLS EXP Classifier* to learn how to define EXP classifiers for MPLS traffic.) BA classifiers allow you to set the forwarding class and loss priority of a packet based on the incoming CoS value.

On most devices, unicast traffic uses different classifiers than multdestination (multicast, broadcast, and destination lookup fail) traffic. You use the `multi-destination` statement at the `[edit class-of-service]` hierarchy level to configure a multdestination BA classifier.

Multdestination classifiers apply to all of the switch interfaces and handle multicast, broadcast, and destination lookup fail (DLF) traffic. You cannot apply a multdestination classifier to a single interface or to a range of interfaces.

Platform-specific Information

- OCX Series switches do not support MPLS EXP classifiers.
- On QFX10000 switches and NFX Series devices, unicast and multdestination traffic use the same classifiers and forwarding classes.
- QFX5130, QFX5700 & QFX5220 switches do not support DSCP IPv6 classifiers and rewrite rules. However, you can apply DSCP classifiers and rewrite rules for IPV6 traffic as well.

Configuring BA Classifiers

To configure a DSCP, DSCP IPv6, or IEEE 802.1p BA classifier using the CLI:

1. Create a BA classifier:

- To create a DSCP, DSCP IPv6, or IEEE 802.1p BA classifier based on the default classifier, import the default DSCP, DSCP IPv6, or IEEE 802.1p classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | dscp-ipv6 | ieee-802.1) classifier-name import default forwarding-
class forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

- To create a BA classifier that is not based on the default classifier, create a DSCP, DSCP IPv6, or IEEE 802.1p classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | dscp-ipv6 | ieee-802.1) classifier-name forwarding-class
forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

2. For multidestination traffic, except on QFX10000 switches or NFX Series devices, configure the classifier as a multidestination classifier:

```
[edit class-of-service]
user@switch# set multi-destination classifiers (dscp | dscp-ipv6 | ieee-802.1 | inet-
precedence) classifier-name
```

3. Apply the classifier to a specific Ethernet interface or to all Ethernet interfaces, or to all Fibre Channel interfaces on the device.
 - To apply the classifier to a specific interface:

```
[edit class-of-service interfaces]
user@switch# set interface-name unit unit classifiers (dscp | dscp-ipv6 | ieee-802.1)
classifier-name
```

- To apply the classifier to all Ethernet interfaces on the switch, use wildcards for the interface name and the logical interface (unit) number:

```
[edit class-of-service interfaces]
user@switch# set xe-* unit * classifiers (dscp | dscp-ipv6 | ieee-802.1) classifier-name
```

RELATED DOCUMENTATION

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Unicast Classifiers

Configuring a Global MPLS EXP Classifier

Configuring Rewrite Rules for MPLS EXP Classifiers

Monitoring CoS Classifiers

Understanding CoS Classifiers

[Understanding CoS Classifiers](#)[Understanding CoS MPLS EXP Classifiers and Rewrite Rules](#)[Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces](#)[Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces](#)

Example: Configuring Classifiers

IN THIS SECTION

- [Requirements | 63](#)
- [Overview | 63](#)
- [Verification | 65](#)

Packet classification associates incoming packets with a particular CoS servicing level. Classifiers associate packets with a forwarding class and loss priority and assign packets to output queues based on the associated forwarding class. You apply classifiers to ingress interfaces.

Configuring Classifiers

Step-by-Step Procedure

To configure an IEEE 802.1 BA classifier named `ba-classifier` as the default IEEE 802.1 classifier:

1. Associate code point `000` with forwarding class `be` and loss priority `low`:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-classifier import default forwarding-class be loss-priority
low code-points 000
```

2. Associate code point 011 with forwarding class fcoe and loss priority low:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-classifier forwarding-class fcoe loss-priority low code-points
011
```

3. Associate code point 100 with forwarding class no-loss and loss priority low:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-classifier forwarding-class no-loss loss-priority low code-
points 100
```

4. Associate code point 110 with forwarding class nc and loss priority low:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-classifier forwarding-class nc loss-priority low code-points
110
```

5. Apply the classifier to ingress interface xe-0/0/10:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/10 unit 0 classifiers ieee-802.1 ba-classifier
```

Requirements

This example uses the following hardware and software components:

- One switch.
- Junos OS Release 15.1X53-D10 or later for the QFX Series.

Overview

Junos OS supports three general types of classifiers:

- Behavior aggregate or CoS value traffic classifiers—Examine the CoS value in the packet header. The value in this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point

(DSCP or DSCP IPv6) value, IEEE 802.1p value, or MPLS EXP value. (EXP classifiers can be applied only to family mpls interfaces.)

- Fixed classifiers. Fixed classifiers classify all ingress traffic on a physical interface into one forwarding class, regardless of the CoS bits in the VLAN header or the DSCP bits in the IP packet header.
- Multifield traffic classifiers—Examine multiple fields in the packet, such as source and destination addresses and source and destination port numbers of the packet. With multifield classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

This example describes how to configure a BA classifier called `ba-classifier` as the default IEEE 802.1 mapping of incoming traffic to forwarding classes, and apply it to ingress interface `xe-0/0/10`. The BA classifier assigns loss priorities, as shown in [Table 20 on page 64](#), to incoming packets in the four default forwarding classes. You can adapt the example to DSCP traffic by specifying a DSCP classifier instead of an IEEE classifier, and by applying DSCP bits instead of CoS bits.

To set multifield classifiers, use firewall filter rules.

Table 20: ba-classifier Loss Priority Assignments

Forwarding Class	CoS Traffic Type	ba-classifier Loss Priority to IEEE 802.1p Code Point Mapping	Packet Drop Attribute
be	Best-effort traffic	Low loss priority code point: 000	drop
fcoe	Guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic	Low loss priority code point: 011	no-loss
no-loss	Guaranteed delivery for TCP traffic	Low loss priority code point: 100	no-loss
nc	Network-control traffic	Low loss priority code point: 110	drop

Verification

IN THIS SECTION

- [Verifying the Classifier Configuration | 65](#)
- [Verifying the Ingress Interface Configuration | 66](#)

To verify the classifier configuration, perform these tasks:

Verifying the Classifier Configuration

Purpose

Verify that you configured the classifier with the correct forwarding classes, loss priorities, and code points.

Action

List the classifier configuration using the operational mode command `show configuration class-of-service classifiers ieee-802.1 ba-classifier`:

```
user@switch> show configuration class-of-service classifiers ieee-802.1 ba-classifier
  forwarding-class be {
    loss-priority low code-points 000;
  }
  forwarding-class fcoe {
    loss-priority low code-points 011;
  }
  forwarding-class no-loss {
    loss-priority low code-points 100;
  }
  forwarding-class nc
    loss-priority low code-points 110;
  }
```

Verifying the Ingress Interface Configuration

Purpose

Verify that the classifier ba-classifier is attached to ingress interface xe-0/0/10.

Action

List the ingress interface using the operational mode command `show configuration class-of-service interfaces xe-0/0/10`:

```
user@switch> show configuration class-of-service interfaces xe-0/0/10
congestion-notification-profile fcoe-cnp;
unit 0 {
    classifiers {
        ieee-802.1 ba-classifier;
    }
}
```

RELATED DOCUMENTATION

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)

Configuring a Global MPLS EXP Classifier

Configuring Rewrite Rules for MPLS EXP Classifiers

Monitoring CoS Classifiers

Understanding CoS Classifiers

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

Understanding Default CoS Scheduling and Classification

IN THIS SECTION

● [Default Classification](#) | 67

- [Default Scheduling | 73](#)
- [Default DCBX Advertisement | 76](#)
- [Default Scheduling and Classification Summary | 76](#)

If you do not explicitly configure classifiers and apply them to interfaces, the switch uses the default classifier to group ingress traffic into forwarding classes. If you do not configure scheduling on an interface, the switch uses the default schedulers to provide egress port resources for traffic. Default classification maps all traffic into default forwarding classes (best-effort, fcoe, no-loss, network-control, and mcast). Each default forwarding class has a default scheduler, so that the traffic mapped to each default forwarding class receives port bandwidth, prioritization, and packet drop characteristics.

The switch supports direct port scheduling and enhanced transmission selection (ETS), also known as hierarchical port scheduling, except on QFX5200 and QFX5210 switches.

Hierarchical scheduling groups IEEE 802.1p priorities (IEEE 802.1p code points, which classifiers map to forwarding classes, which in turn are mapped to output queues) into priority groups (forwarding class sets). If you use only the default traffic scheduling and classification, the switch automatically creates a default priority group that contains all of the priorities (which are mapped to forwarding classes and output queues), and assigns 100 percent of the port output bandwidth to that priority group. The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default classifier settings. The default priority group is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange (DCBX) protocol advertisement.

NOTE: If you explicitly configure one or more priority groups on an interface, any forwarding class that is not assigned to a priority group on that interface receives *no bandwidth*. This means that if you configure hierarchical scheduling on an interface, every forwarding class (priority) that you want to forward traffic on that interface must belong to a forwarding class set (priority group). ETS is not supported on QFX5200 or QFX5210 switches.

The following sections describe:

Default Classification

On switches except QFX10000 and NFX Series devices, the default classifiers assign unicast and multicast best-effort and network-control ingress traffic to default forwarding classes and loss priorities. The switch applies default unicast IEEE 802.1, unicast DSCP, and multidestination classifiers to each interface that does not have explicitly configured classifiers.

On QFX10000 switches and NFX Series devices, the default classifiers assign ingress traffic to default forwarding classes and loss priorities. The switch applies default IEEE 802.1, DSCP, and DSCP IPv6 classifiers to each interface that does not have explicitly configured classifiers. If you do not configure and apply EXP classifiers for MPLS traffic to logical interfaces, MPLS traffic on interfaces configured as family mpls uses the IEEE classifier.

If you explicitly configure one type of classifier but not other types of classifiers, the system uses only the configured classifier and does not use default classifiers for other types of traffic. There are two default IEEE 802.1 classifiers: a trusted classifier for ports that are in trunk mode or tagged-access mode, and an untrusted classifier for ports that are in access mode.

NOTE: The default classifiers apply to unicast traffic except on QFX10000 switches and NFX Series devices. Tagged-access mode does not apply to QFX10000 switches or NFX Series devices.

Table 21 on page 68 shows the default mapping of IEEE 802.1 code-point values to forwarding classes and loss priorities for ports in trunk mode or tagged-access mode.

Table 21: Default IEEE 802.1 Classifiers for Ports in Trunk Mode or Tagged-Access Mode (Trusted Classifier)

Code Point	Forwarding Class	Loss Priority
be (000)	best-effort	low
be1 (001)	best-effort	low
ef (010)	best-effort	low
ef1 (011)	fcoe	low
af11 (100)	no-loss	low
af12 (101)	best-effort	low
nc1 (110)	network-control	low

Table 21: Default IEEE 802.1 Classifiers for Ports in Trunk Mode or Tagged-Access Mode (Trusted Classifier) (Continued)

Code Point	Forwarding Class	Loss Priority
nc2 (111)	network-control	low

Table 22 on page 69 shows the default mapping of IEEE 802.1p code-point values to forwarding classes and loss priorities for ports in access mode (all incoming traffic is mapped to best-effort forwarding classes).

NOTE: Table 22 on page 69 applies only to unicast traffic except on QFX10000 switches and NFX Series devices.

Table 22: Default IEEE 802.1 Classifiers for Ports in Access Mode (Untrusted Classifier)

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

Table 23 on page 70 shows the default mapping of IEEE 802.1 code-point values to multidestination (multicast, broadcast, and destination lookup fail traffic) forwarding classes and loss priorities.

NOTE: Table 23 on page 70 does not apply to QFX10000 switches or NFX Series devices.

Table 23: Default IEEE 802.1 Multidestination Classifiers

Code Point	Forwarding Class	Loss Priority
be (000)	mcast	low
be1 (001)	mcast	low
ef (010)	mcast	low
ef1 (011)	mcast	low
af11 (100)	mcast	low
af12 (101)	mcast	low
nc1 (110)	mcast	low
nc2 (111)	mcast	low

Table 24 on page 70 shows the default mapping of DSCP code-point values to forwarding classes and loss priorities for DSCP IP and DCSP IPv6.

NOTE: Table 24 on page 70 applies only to unicast traffic except on QFX10000 switches and NFX Series devices.

Table 24: Default DSCP IP and IPv6 Classifiers

Code Point	Forwarding Class	Loss Priority
ef (101110)	best-effort	low

Table 24: Default DSCP IP and IPv6 Classifiers (Continued)

Code Point	Forwarding Class	Loss Priority
af11 (001010)	best-effort	low
af12 (001100)	best-effort	low
af13 (001110)	best-effort	low
af21 (010010)	best-effort	low
af22 (010100)	best-effort	low
af23 (010110)	best-effort	low
af31 (011010)	best-effort	low
af32 (011100)	best-effort	low
af33 (011110)	best-effort	low
af41 (100010)	best-effort	low
af42 (100100)	best-effort	low
af43 (100110)	best-effort	low
be (000000)	best-effort	low
cs1 (001000)	best-effort	low
cs2 (010000)	best-effort	low

Table 24: Default DSCP IP and IPv6 Classifiers (Continued)

Code Point	Forwarding Class	Loss Priority
cs3 (011000)	best-effort	low
cs4 (100000)	best-effort	low
cs5 (101000)	best-effort	low
nc1 (110000)	network-control	low
nc2 (111000)	network-control	low

NOTE: There are no default DSCP IP or IPv6 multdestination classifiers for multdestination traffic. DSCP IPv6 multdestination classifiers are not supported for multdestination traffic.

[Table 25 on page 72](#) shows the default mapping of MPLS EXP code-point values to forwarding classes and loss priorities, which apply only on QFX10000 switches and NFX Series devices.

Table 25: Default EXP Classifiers on QFX10000 Switches and NFX Series Devices

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

Table 25: Default EXP Classifiers on QFX10000 Switches and NFX Series Devices (Continued)

Code Point	Forwarding Class	Loss Priority
101	assured-forwarding	high
110	network-control	low
111	network-control	high

Default Scheduling

The default schedulers allocate egress bandwidth resources to egress traffic as shown in [Table 26 on page 73](#):

Table 26: Default Scheduler Configuration

Default Scheduler and Queue Number	Transmit Rate (Guaranteed Minimum Bandwidth)	Shaping Rate (Maximum Bandwidth)	Excess Bandwidth Sharing	Priority	Buffer Size
best-effort forwarding class scheduler (queue 0)	5% 15% (QFX10000, NFX Series)	None	5% 15% (QFX10000, NFX Series)	low	5% 15% (QFX10000, NFX Series)
fcoe forwarding class scheduler (queue 3)	35%	None	35%	low	35%
no-loss forwarding class scheduler (queue 4)	35%	None	35%	low	35%

Table 26: Default Scheduler Configuration (Continued)

Default Scheduler and Queue Number	Transmit Rate (Guaranteed Minimum Bandwidth)	Shaping Rate (Maximum Bandwidth)	Excess Bandwidth Sharing	Priority	Buffer Size
network-control forwarding class scheduler (queue 7)	5% 15% (QFX10000, NFX Series)	None	5% 15% (QFX10000, NFX Series)	low	5% 15% (QFX10000, NFX Series)
(Excluding QFX10000 and NFX Series) mcast forwarding class scheduler (queue 8)	20%	None	20%	low	20%

NOTE: By default, the minimum guaranteed bandwidth (transmit rate) determines the amount of excess (extra) bandwidth that a queue can share. Extra bandwidth is allocated to queues in proportion to the transmit rate of each queue. On switches that support the `excess-rate` statement, you can override the default setting and configure the excess bandwidth percentage independently of the transmit rate on queues that are not strict-high priority queues.

By default, only the four (QFX10000 switches and NFX Series devices) or five (other switches) default schedulers shown in [Table 26 on page 73](#) have traffic mapped to them. Only the forwarding classes and queues associated with the default schedulers receive default bandwidth, based on the default scheduler transmit rate. (You can configure schedulers and forwarding classes to allocate bandwidth to other queues or to change the bandwidth and other scheduling properties of a default queue.)

On QFX10000 switches and NFX Series devices, if a forwarding class does not transport traffic, the bandwidth allocated to that forwarding class is available to other forwarding classes. Unicast and multidestination (multicast, broadcast, and destination lookup fail) traffic use the same forwarding classes and output queues.

On switches other than QFX10000 and NFX Series devices, multidestination queue 11 receives enough bandwidth from the default multidestination scheduler to handle CPU-generated multidestination traffic.

On QFX10000 and NFX Series devices, default scheduling is port scheduling. Default hierarchical scheduling, known as enhanced transmission selection (ETS, defined in IEEE 802.1Qaz), allocates the

total port bandwidth to the four default forwarding classes served by the four default schedulers, as defined by the four default schedulers. The result is the same as direct port scheduling. Configuring hierarchical port scheduling, however, enables you to group forwarding classes that carry similar types of traffic into forwarding class sets (also called priority groups), and to assign port bandwidth to each forwarding class set. The port bandwidth assigned to the forwarding class set is then assigned to the forwarding classes within the forwarding class set. This hierarchy enables you to control port bandwidth allocation with greater granularity, and enables hierarchical sharing of extra bandwidth to better utilize link bandwidth.

Except on QFX10000 switches and NFX Series devices, default hierarchical scheduling divides the total port bandwidth between two groups of traffic: unicast traffic and multdestination traffic. By default, unicast traffic consists of queue 0 (best-effort forwarding class), queue 3 (fcoe forwarding class), queue 4 (no-loss forwarding class), and queue 7 (network-control forwarding class). Unicast traffic receives and shares a total of 80 percent of the port bandwidth. By default, multdestination traffic (mcast queue 8) receives a total of 20 percent of the port bandwidth. So on a 10-Gigabit port, unicast traffic receives 8-Gbps of bandwidth and multdestination traffic receives 2-Gbps of bandwidth.

NOTE: Except on QFX5200, QFX5210, and QFX10000 switches and NFX Series devices, which do not support queue 11, multdestination queue 11 also receives a small amount of default bandwidth from the multdestination scheduler. CPU-generated multdestination traffic uses queue 11, so you might see a small number of packets egress from queue 11. In addition, in the unlikely case that firewall filter match conditions map multdestination traffic to a unicast forwarding class, that traffic uses queue 11.

Default scheduling uses weighted round-robin (WRR) scheduling. Each queue receives a portion (weight) of the total available interface bandwidth. The scheduling weight is based on the transmit rate of the default scheduler for that queue. For example, queue 7 receives a default scheduling weight of 5 percent, or 15 percent on QFX10000 and NFX Series devices, of the available bandwidth, and queue 4 receives a default scheduling weight of 35 percent of the available bandwidth. Queues are mapped to forwarding classes, so forwarding classes receive the default bandwidth for the queues to which they are mapped.

On QFX10000 switches and NFX Series devices, for example, queue 7 is mapped to the network-control forwarding class and queue 4 is mapped to the no-loss forwarding class. Each forwarding class receives the default bandwidth for the queue to which it is mapped. Unused bandwidth is shared with other default queues.

If you want non-default (unconfigured) queues to forward traffic, you should explicitly map traffic to those queues (configure the forwarding classes and queue mapping) and create schedulers to allocate bandwidth to those queues. By default, queues 1, 2, 5, and 6 are unconfigured.

Except on QFX5200, QFX5210, and QFX10000 switches and NFX Series devices, which do not support them, multdestination queues 9, 10, and 11 are unconfigured. Unconfigured queues have a default

scheduling weight of 1 so that they can receive a small amount of bandwidth in case they need to forward traffic. However, queue 11 can use more of the default multidestination scheduler bandwidth if necessary to handle CPU-generated multidestination traffic.

NOTE: All four (two on QFX5200 and QFX5210 switches) multidestination queues have a scheduling weight of 1. Because by default multidestination traffic goes to queue 8, queue 8 receives almost all of the multidestination bandwidth. (There is no traffic on queue 9 and queue 10, and very little traffic on queue 11, so there is almost no competition for multidestination bandwidth.)

However, if you explicitly configure queue 9, 10, or 11 (by mapping code points to the unconfigured multidestination forwarding classes using the multidestination classifier), the explicitly configured queues share the multidestination scheduler bandwidth equally with default queue 8, because all of the queues have the same scheduling weight (1). To ensure that multidestination bandwidth is allocated to each queue properly and that the bandwidth allocation to the default queue (8) is not reduced too much, we strongly recommend that you configure a scheduler if you explicitly classify traffic into queue 9, 10, or 11.

If you map traffic to an unconfigured queue, the queue receives only the amount of excess bandwidth proportional to its default weight (1). The actual amount of bandwidth an unconfigured queue gets depends on how much bandwidth the other queues are using.

If some queues use less than their allocated amount of bandwidth, the unconfigured queues can share the unused bandwidth. Sharing unused bandwidth is one of the key advantages of hierarchical port scheduling. Configured queues have higher priority for bandwidth than unconfigured queues, so if a configured queue needs more bandwidth, then less bandwidth is available for unconfigured queues. Unconfigured queues always receive a minimum amount of bandwidth based on their scheduling weight (1). If you map traffic to an unconfigured queue, to allocate bandwidth to that queue, configure a scheduler for the forwarding class that is mapped to the queue.

Default DCBX Advertisement

When you configure hierarchical scheduling on an interface, DCBX advertises each priority group, the priorities in each priority group, and the bandwidth properties of each priority and priority group.

If you do not configure hierarchical scheduling on an interface, DCBX advertises the automatically created default priority group and its priorities. DCBX also advertises the default bandwidth allocation of the priority group, which is 100 percent of the port bandwidth.

Default Scheduling and Classification Summary

If you do not configure scheduling on an interface:

- Default classifiers classify ingress traffic.
- Default schedulers schedule egress traffic.
- DCBX advertises a single default priority group with 100 percent of the port bandwidth allocated to that priority group. All priorities (forwarding classes) are assigned to the default priority group and receive bandwidth based on their default schedulers. The default priority group is generated automatically and is not user-configurable.

RELATED DOCUMENTATION

Understanding CoS Packet Flow

Understanding CoS Hierarchical Port Scheduling (ETS)

Understanding Default CoS Settings

Understanding CoS Virtual Output Queues (VOQs) on QFX10000 Switches

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

Understanding DCB Features and Requirements

Understanding Default CoS Scheduling on QFabric System Interconnect Devices (Junos OS Release 13.1 and Later Releases)

Example: Configuring Unicast Classifiers

Example: Configuring Queue Schedulers

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

IN THIS SECTION

- Supported Classifier and Rewrite Rule Types | 78
- Ethernet Interfaces Supported for Classifier and Rewrite Rule Configuration | 80
- Default Classifiers | 84
- Default Rewrite Rules | 84
- Classifier Precedence | 85
- Classifier Behavior and Limitations | 86
- Rewrite Rule Precedence and Behavior | 87
- Classifier and Rewrite Rule Configuration Interaction with Ethernet Interface Configuration | 88

At ingress interfaces, classifiers group incoming traffic into classes based on the IEEE 802.1p, DSCP, or MPLS EXP *class of service* (CoS) code points in the packet header. At egress interfaces, you can use *rewrite rules* to change (re-mark) the code point bits before the interface forwards the packets.

You can apply classifiers and rewrite rules to interfaces to control the level of CoS applied to each packet as it traverses the system and the network. This topic describes:

Supported Classifier and Rewrite Rule Types

Table 27 on page 78 shows the supported types of classifiers and rewrite rules supports:

Table 27: Supported Classifiers and Rewrite Rules

Classifier or Rewrite Rule Type	Description
Fixed classifier	Classifies all ingress traffic on a physical interface into one fixed forwarding class, regardless of the CoS bits in the packet header.
DSCP and DSCP IPv6 unicast classifiers	Classifies IP and IPv6 traffic into forwarding classes and assigns loss priorities to the traffic based on DSCP code point bits.
IEEE 802.1p unicast classifier	Classifies Ethernet traffic into forwarding classes and assigns loss priorities to the traffic based on IEEE 802.1p code point bits.
MPLS EXP classifier	<p>Classifies MPLS traffic into forwarding classes and assigns loss priorities to the traffic on interfaces configured as <code>family mpls</code>.</p> <p>QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and QFabric systems, use one global EXP classifier on all <code>family mpls</code> switch interfaces.</p> <p>QFX10000 switches do not support global EXP classifiers. You can apply the same EXP classifier or different EXP classifiers to different <code>family mpls</code> interfaces.</p>

Table 27: Supported Classifiers and Rewrite Rules (*Continued*)

Classifier or Rewrite Rule Type	Description
DSCP multidestination classifier (also used for IPv6 multidestination traffic) NOTE: This applies only to switches that use different classifiers for unicast and multidestination traffic. It does not apply to switches that use the same classifiers for unicast and multidestination traffic.	Classifies IP and IPv6 multicast, broadcast, and destination lookup fail (DLF) traffic into multidestination forwarding classes. Multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces.
IEEE 802.1p multidestination classifier NOTE: This applies only to switches that use different classifiers for unicast and multidestination traffic. It does not apply to switches that use the same classifiers for unicast and multidestination traffic.	Classifies Ethernet multicast, broadcast, and destination lookup fail (DLF) traffic into multidestination forwarding classes. Multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces.
DSCP and DSCP IPv6 rewrite rules	Re-marks the DSCP code points of IP and IPv6 packets before forwarding the packets.
IEEE 802.1p rewrite rule	Re-marks the IEEE 802.1p code points of Ethernet packets before forwarding the packets.
MPLS EXP rewrite rule	Re-marks the EXP code points of MPLS packets before forwarding the packets on interfaces configured as family <code>mpls</code> .

NOTE: On switches that support native Fibre Channel (FC) interfaces, you can specify a rewrite value on native FC interfaces (NP_Ports) to set the IEEE 802.1p code point of incoming FC traffic when the NP_Port encapsulates the FC packet in Ethernet before forwarding it to the FCoE network (see *Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway*).

DSCP, IEEE 802.1p, and MPLS EXP classifiers are behavior aggregate (BA) classifiers. On QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, unlike DSCP and IEEE 802.1p classifiers, EXP classifiers are global and apply only to all interfaces that are configured as family `mpls`. On QFX10000 switches, you apply EXP classifiers to individual logical interfaces, and different interfaces can use different EXP classifiers.

Unlike DSCP and IEEE 802.1p BA classifiers, there is no default EXP classifier. Also unlike DSCP and IEEE 802.1p classifiers, for MPLS traffic on family mpls interfaces only, EXP classifiers overwrite fixed classifiers. (An interface that has a fixed classifier uses the EXP classifier for MPLS traffic, not the fixed classifier, and the fixed classifier is used for all other traffic.)

On switches that use different classifiers for unicast and multidestination traffic, multidestination classifiers are global and apply to all interfaces; you cannot apply a multidestination classifier to individual interfaces.

Classifying packets into forwarding classes assigns packets to the output queues mapped to those forwarding classes. The traffic classified into a forwarding class receives the CoS scheduling configured for the output queue mapped to that forwarding class.

NOTE: In addition to BA classifiers and fixed classifiers, which classify traffic based on the CoS field in the packet header, you can use firewall filters to configure multifield (MF) classifiers. MF classifiers classify traffic based on more than one field in the packet header and take precedence over BA and fixed classifiers.

Ethernet Interfaces Supported for Classifier and Rewrite Rule Configuration

To apply a classifier to incoming traffic or a rewrite rule to outgoing traffic, you need to apply the classifier or rewrite rule to one or more interfaces. When you apply a classifier or rewrite rule to an interface, the interface uses the classifier to group incoming traffic into forwarding classes and uses the rewrite rule to re-mark the CoS code point value of each packet before it leaves the system.

Not all interfaces types support all types of CoS configuration. This section describes:

Interface Types That Support Classifier and Rewrite Rule Configuration

You can apply classifiers and rewrite rules to Ethernet interfaces. For Layer 3 LAGs, configure BA or fixed classifiers on the LAG (ae) interface. The classifier configured on the LAG is valid on all of the LAG member interfaces.

On switches that support native FC interfaces, you can apply fixed classifiers to native FC interfaces (NP_Ports). You cannot apply other types of classifiers or rewrite rules to native FC interfaces. You can rewrite the value of the IEEE 802.1p code point of incoming FC traffic when the interface encapsulates it in Ethernet before forwarding it to the FCoE network as described in *Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway*.

Classifier and Rewrite Rule Physical and Logical Ethernet Interface Support

The Ethernet ports can function as:

- Layer 2 physical interfaces (family ethernet-switching)
- Layer 2 logical interfaces (family ethernet-switching)
- Layer 3 physical interfaces (family inet/inet6)
- Layer 3 logical interfaces (family inet/inet6)
- MPLS interfaces (family mpls)

You can apply CoS classifiers and rewrite rules only to the following interfaces:

- Layer 2 logical interface

NOTE: On a Layer 2 interface, use **unit *** to apply the rule to all of the logical units on that interface.

- On QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, Layer 3 physical interfaces if at least one logical Layer 3 interface is configured on the physical interface

NOTE: The CoS you configure on a Layer 3 physical interface is applied to all of the Layer 3 logical interfaces on that physical interface. This means that each Layer 3 interface uses the same classifiers and rewrite rules for all of the Layer 3 traffic on that interface.

- On QFX10000 switches, Layer 3 logical interfaces. You can apply different classifiers and rewrite rules to different Layer 3 logical interfaces.

Ethernet Interface Support for Most QFX Series Switches, and QFabric Systems

You cannot apply classifiers or rewrite rules to Layer 2 physical interfaces or to Layer 3 logical interfaces. [Table 28 on page 82](#) shows on which interfaces you can configure and apply classifiers and rewrite rules.

NOTE: The CoS feature support listed in this table is identical on single interfaces and aggregated Ethernet interfaces.

Table 28: Ethernet Interface Support for Classifier and Rewrite Rule Configuration (QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems)

CoS Classifiers and Rewrite Rules	Layer 2 Physical Interfaces	Layer 2 Logical Interface (unit * applies rule to all logical interfaces)	Layer 3 Physical Interfaces (If at Least One Logical Layer 3 Interface Is Defined)	Layer 3 Logical Interfaces
Fixed classifier	No	Yes	Yes	No
DSCP classifier	No	Yes	Yes	No
DSCP IPv6 classifier	No	Yes	Yes	No
IEEE 802.1p classifier	No	Yes	Yes	No
EXP classifier	Global classifier, applies only to all switch interfaces that are configured as family mpls. Cannot be configured on individual interfaces.			
DSCP rewrite rule	No	Yes	Yes	No
DSCP IPv6 rewrite rule	No	Yes	Yes	No
IEEE 802.1p rewrite rule	No	Yes	Yes	No
EXP rewrite rule	No	Yes	Yes	No

NOTE: IEEE 802.1p multidestination and DSCP multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces. No DSCP IPv6 multidestination classifier is supported. IPv6 multidestination traffic uses the DSCP multidestination classifier.

Ethernet Interface Support for QFX10000 Switches

You cannot apply classifiers or rewrite rules to Layer 2 or Layer 3 physical interfaces. You can apply classifiers and rewrite rules only to Layer 2 logical interface unit 0. You can apply different classifiers and rewrite rules to different Layer 3 logical interfaces. [Table 29 on page 83](#) shows on which interfaces you can configure and apply classifiers and rewrite rules.

NOTE: The CoS feature support listed in this table is identical on single interfaces and aggregated Ethernet interfaces.

Table 29: Ethernet Interface Support for Classifier and Rewrite Rule Configuration (QFX10000 Switches)

CoS Classifiers and Rewrite Rules	Layer 2 Physical Interfaces	Layer 2 Logical Interface (Unit 0 Only)	Layer 3 Physical Interfaces	Layer 3 Logical Interfaces
Fixed classifier	No	Yes	No	Yes
DSCP classifier	No	Yes	No	Yes
DSCP IPv6 classifier	No	Yes	No	Yes
IEEE 802.1p classifier	No	Yes	No	Yes
EXP classifier	No	Yes	No	Yes
DSCP rewrite rule	No	Yes	No	Yes
DSCP IPv6 rewrite rule	No	Yes	No	Yes
IEEE 802.1p rewrite rule	No	Yes	No	Yes
EXP rewrite rule	No	Yes	No	Yes

Routed VLAN Interfaces (RVIs) and Integrated Routing and Bridging (IRB) Interfaces

You cannot apply classifiers and rewrite rules directly to routed VLAN interfaces (RVIs) or integrated routing and bridging (IRB) interfaces because the members of RVIs and IRBs are VLANs, not ports. However, you can apply classifiers and rewrite rules to the VLAN port members of an *RVI* or an IRB. You can also apply MF classifiers to RVIs and IRBs.

Default Classifiers

If you do not explicitly configure classifiers on an Ethernet interface, the switch applies default classifiers so that the traffic receives basic CoS treatment. The factors that determine the default classifier applied to the interface include the interface type (Layer 2 or Layer 3), the port mode (trunk, tagged-access, or access), and whether logical interfaces have been configured.

The switch applies default classifiers using the following rules:

- If the physical interface has at least one Layer 3 *logical interface* configured, the logical interfaces use the default DSCP classifier.
- If the physical interface has a Layer 2 logical interface in trunk mode or tagged-access mode, it uses the default IEEE 802.1p trusted classifier.

NOTE: Tagged-access mode is available only on QFX3500 and QFX3600 devices when used as standalone switches or as QFabric system Node devices.

- If the physical interface has a Layer 2 logical interface in access mode, it uses the default IEEE 802.1p untrusted classifier.
- If the physical interface has no logical interface configured, no default classifier is applied.
- On switches that use different classifiers for unicast and multidestination traffic, the default multidestination classifier is the IEEE 802.1p multidestination classifier.
- There is no default MPLS EXP classifier. If you want to classify MPLS traffic using EXP bits on these switches, on QFX10000 switches, configure an EXP classifier and apply it to a logical interface that is configured as `family mpls`. On QFX5100, QFX5200, EX4600, QFX3500 and QFX3600 switches, and on QFabric systems, configure an EXP classifier and configure it as the global system default EXP classifier.

Default Rewrite Rules

No default rewrite rules are applied to interfaces. If you want to re-mark packets at the egress interface, you must explicitly configure a rewrite rule.

Classifier Precedence

You can apply multiple classifiers (MF, fixed, IEEE 802.1p, DSCP, or EXP) to an Ethernet interface to handle different types of traffic. (EXP classifiers are global and apply only to all MPLS traffic on all family mpls interfaces.) When you apply more than one classifier to an interface, the system uses an order of precedence to determine which classifier to use on interfaces:

Classifier Precedence on Physical Ethernet Interfaces (QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems)

QFX10000 switches do not support configuring classifiers on physical interfaces. The precedence of classifiers on physical interfaces, from the highest-priority classifier to the lowest-priority classifier, is:

- MF classifier on a logical interface (no classifier has a higher priority than MF classifiers)
- Fixed classifier on the physical interface
- DSCP or DSCP IPv6 classifier on the physical interface
- IEEE 802.1p classifier on the physical interface

NOTE: If an EXP classifier is configured, MPLS traffic uses the EXP classifier on all family mpls interfaces, even if an MF or fixed classifier is applied to the interface. If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic. DSCP classifiers are not applied to MPLS traffic.

You can apply a DSCP classifier, an IEEE 802.1p classifier, and an EXP classifier on a physical interface. When all three classifiers are on an interface, IP traffic uses the DSCP classifier, MPLS traffic on family mpls interfaces uses the EXP classifier, and all other traffic uses the IEEE classifier.

NOTE: You cannot apply a fixed classifier and a DSCP or IEEE classifier to the same interface. If a DSCP classifier, an IEEE classifier, or both are on an interface, you cannot apply a fixed classifier to that interface unless you first delete the DSCP and IEEE classifiers. If a fixed classifier is on an interface, you cannot apply a DSCP classifier or an IEEE classifier unless you first delete the fixed classifier.

Classifier Precedence on Logical Ethernet Interfaces (All Switches)

The precedence of classifiers on logical interfaces, from the highest priority classifier to the lowest priority classifier, is:

- MF classifier on a logical interface (no classifier has a higher priority than MF classifiers).
- Fixed classifier on the logical interface.
- DSCP or DSCP IPv6 classifier on the physical or logical interface..
- IEEE 802.1p classifier on the physical or logical interface.

NOTE: If a global EXP classifier is configured, MPLS traffic uses the EXP classifier on all family `mpls` interfaces, even if a fixed classifier is applied to the interface. If a global EXP classifier is not configured, then:

- If a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic.

You can apply both a DSCP classifier and an IEEE 802.1p classifier on a logical interface. When both a DSCP and an IEEE classifier are on an interface, IP traffic uses the DSCP classifier, and all other traffic uses the IEEE classifier. Only MPLS traffic on interfaces configured as family `mpls` uses the EXP classifier.

Classifier Behavior and Limitations

Consider the following behaviors and constraints when you apply classifiers to Ethernet interfaces. Behaviors for applying classifiers to physical interfaces do not pertain to QFX10000 switches.

- You can configure only one DSCP classifier (IP or IPv6) on a physical interface. You cannot configure both types of DSCP classifier on one physical interface. Both IP and IPv6 traffic use whichever DSCP classifier is configured on the interface.
- When you configure a DSCP or a DSCP IPv6 classifier on a physical interface and the physical interface has at least one logical Layer 3 interface, all packets (IP, IPv6, and non-IP) use that classifier.
- An interface with both a DSCP classifier (IP or IPv6) and an IEEE 802.1p classifier uses the DSCP classifier for IP and IPv6 packets, and uses the IEEE classifier for all other packets.
- Fixed classifiers and BA classifiers (DSCP and IEEE classifiers) are not permitted simultaneously on an interface. If you configure a fixed classifier on an interface, you cannot configure a DSCP or an IEEE classifier on that interface. If you configure a DSCP classifier, an IEEE classifier, or both classifiers on an interface, you cannot configure a fixed classifier on that interface.

- When you configure an IEEE 802.1p classifier on a physical interface and a DSCP classifier is not explicitly configured on that interface, the interface uses the IEEE classifier for all types of packets. No default DSCP classifier is applied to the interface. (In this case, if you want a DSCP classifier on the interface, you must explicitly configure it and apply it to the interface.)
- The system does not apply a default classifier to a physical interface until you create a logical interface on that physical interface. If you configure a Layer 3 logical interface, the system uses the default DSCP classifier. If you configure a Layer 2 logical interface, the system uses the default IEEE 802.1p trusted classifier if the port is in trunk mode or tagged-access mode, or the default IEEE 802.1p untrusted classifier if the port is in access mode.
- MF classifiers configured on logical interfaces take precedence over BA and fixed classifiers, with the exception of the global EXP classifier, which is always used for MPLS traffic on family `mpls` interfaces. (Use firewall filters to configure MF classifiers.) When BA or fixed classifiers are present on an interface, you can still configure an MF classifier on that interface.
- There is no default EXP classifier for MPLS traffic.
- You can configure up to 64 EXP classifiers. On QFX10000 switches, you can apply different EXP classifiers to different interfaces.

However, on On QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, the switch uses only one MPLS EXP classifier as a global classifier on all family `mpls` interfaces. After you configure an MPLS EXP classifier, you can configure it as the global EXP classifier by including the EXP classifier in the `[edit class-of-service system-defaults classifiers exp]` hierarchy level.

All family `mpls` switch interfaces use the EXP classifier specified using this configuration statement to classify MPLS traffic, even on interfaces that have a fixed classifier. No other traffic uses the EXP classifier.

Rewrite Rule Precedence and Behavior

The following rules apply on Ethernet interfaces for rewrite rules:

- If you configure one DSCP (or DSCP IPv6) rewrite rule and one IEEE 802.1p rewrite rule on an interface, both rewrite rules take effect. Traffic with IP and IPv6 headers use the DSCP rewrite rule, and traffic with a VLAN tag uses the IEEE rewrite rule.
- If you do not explicitly configure a rewrite rule, there is no default rewrite rule, so the system does not apply any rewrite rule to the interface.
- You can apply a DSCP rewrite rule or a DSCP IPv6 rewrite rule to an interface, but you cannot apply both a DSCP and a DSCP IPv6 rewrite rule to the same interface. Both IP and IPv6 packets use the same DSCP rewrite rule, regardless of whether the configured rewrite rule is DSCP or DSCP IPv6.

- MPLS EXP rewrite rules apply only to logical interfaces on family `mpls` interfaces. You cannot apply to an EXP rewrite rule to a physical interface. You can configure up to 64 EXP rewrite rules, but you can only use 16 EXP rewrite rules at any time on the switch.
- A logical interface can use both DSCP (or DSCP IPv6) and EXP rewrite rules.
- DSCP and DSCP IPv6 rewrite rules are not applied to MPLS traffic.
- If the switch is performing penultimate hop popping (PHP), EXP rewrite rules do not take effect. If both an EXP classifier and an EXP rewrite rule are configured on the switch, then the EXP value from the last popped label is copied into the inner label. If either an EXP classifier or an EXP rewrite rule (but not both) is configured on the switch, then the inner label EXP value is sent unchanged.

NOTE: On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.

NOTE: Rewrite rules are applied *before* the egress filter is matched to traffic. Because the code point rewrite occurs before the egress filter is matched to traffic, the egress filter match is based on the rewrite value, not on the original code point value in the packet.

Classifier and Rewrite Rule Configuration Interaction with Ethernet Interface Configuration

On QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 switches used as standalone switches or as QFabric system Node devices, you can apply classifiers and rewrite rules only on Layer 2 logical interface unit 0 and Layer 3 physical interfaces (if the Layer 3 physical interface has at least one defined logical interface). On QFX10000 switches, you can apply classifiers and rewrite rules only to Layer 2 logical interface unit 0 and to Layer 3 logical interfaces. This section focuses on BA classifiers, but the interaction between BA classifiers and interfaces described in this section also applies to fixed classifiers and rewrite rules.

NOTE: On QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 switches used as standalone switches or as QFabric system Node devices, EXP classifiers, are global and apply to all switch interfaces. See *Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)* for how to configure multidestination classifiers and see *Configuring a Global MPLS EXP Classifier* for how to configure EXP classifiers.

On switches that use different classifiers for unicast and multidestination traffic, multidestination classifiers are global and apply to all switch interfaces.

There are two components to applying classifiers or rewrite rules to interfaces:

1. Setting the interface family (inet, inet6, or ethernet-switching; ethernet-switching is the default interface family) in the [edit interfaces] configuration hierarchy.
2. Applying a classifier or rewrite rule to the interface in the [edit class-of-service] hierarchy.

These are separate operations that can be set and committed at different times. Because the type of classifier or rewrite rule you can apply to an interface depends on the interface family configuration, the system performs checks to ensure that the configuration is valid. The method the system uses to notify you of an invalid configuration depends on the set operation that causes the invalid configuration.

NOTE: QFX10000 switches cannot be misconfigured in the following two ways because you can configure classifiers only on logical interfaces. Only switches that allow classifier configuration on physical and logical interfaces can experience the following misconfigurations.

If applying the classifier or rewrite rule to the interface in the [edit class-of-service] hierarchy causes an invalid configuration, the system rejects the configuration and returns a commit check error.

If setting the interface family in the [edit interfaces] configuration hierarchy causes an invalid configuration, the system creates a syslog error message. If you receive the error message, you need to remove the classifier or rewrite rule configuration from the logical interface and apply it to the physical interface, or remove the classifier or rewrite rule configuration from the physical interface and apply it to the logical interface. For classifiers, if you do not take action to correct the error, the system programs the default classifier for the interface family on the interface. (There are no default rewrite rules. If the commit check fails, no rewrite rule is applied to the interface.)

Two scenarios illustrate these situations:

- Applying a classifier to an Ethernet interface causes a commit check error
- Configuring the Ethernet interface family causes a syslog error

These scenarios differ on different switches because some switches support classifiers on physical Layer 3 interfaces but not on logical Layer 3 interfaces, while other switches support classifiers on logical Layer 3 interfaces but not on physical Layer 3 interfaces.

Two scenarios illustrate these situations:

NOTE: Both of these scenarios also apply to fixed classifiers and rewrite rules.

QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 Switch Scenarios

The following scenarios also apply the QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 switches when they are used as QFabric system Node devices.

Scenario 1: Applying a Classifier to an Ethernet Interface Causes a Commit Check Error

In Scenario 1, we set the interface family, and then specify an invalid classifier.

1. Set and commit the interface as a Layer 3 (family `inet`) interface:

```
[edit interfaces]
user@switch# set xe-0/0/20 unit 0 family inet
user@switch# commit
```

This commit operation succeeds.

2. Set and commit a DSCP classifier on the logical interface (this example uses a DSCP classifier named `dscp1`):

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers dscp dscp1
user@switch# commit
```

This configuration is not valid, because it attempts to apply a classifier to a Layer 3 logical interface. Because the failure is caused by the class-of-service configuration and not by the interface configuration, the system rejects the commit operation and issues a commit error, not a syslog message.

Note that the commit operation succeeds if you apply the classifier to the physical Layer 3 interface as follows:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 classifiers dscp dscp1
user@switch# commit
```

Because the logical unit is not specified, the classifier is applied to the physical Layer 3 interface in a valid configuration, and the commit check succeeds.

Scenario 2: Configuring the Ethernet Interface Family Causes a Syslog Error

In Scenario 2, we set the classifier first, and then set an invalid interface type.

1. Set and commit a DSCP classifier on a logical interface that has no existing configuration:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers dscp dscp1
user@switch# commit
```

This commit succeeds. Because no explicit configuration existed on the interface, it is by default a Layer 2 (family ethernet-switching) interface. Layer 2 logical interfaces support BA classifiers, so applying the classifier is a valid configuration.

2. Set and commit the interface as a Layer 3 interface (family inet) interface:

```
[edit interfaces]
user@switch# set xe-0/0/20 unit 0 family inet
user@switch# commit
```

This configuration is not valid because it attempts to change an interface from Layer 2 (family ethernet-switching) to Layer 3 (family inet) when a classifier has already been applied to a logical interface. Layer 3 logical interfaces do not support classifiers. Because the failure is caused by the interface configuration and not by the class-of-service configuration, the system does not issue a commit error, but instead issues a syslog message.

When the system issues the syslog message, it programs the default classifier for the interface type on the interface. In this scenario, the interface has been configured as a Layer 3 interface, so the system applies the default DSCP profile to the physical Layer 3 interface.

In this scenario, to install a configured DSCP classifier, remove the misconfigured classifier from the Layer 3 logical interface and apply it to the Layer 3 physical interface. For example:

```
[edit]
user@switch# delete class-of-service interfaces xe-0/0/20 unit 0 classifiers dscp dscp1
user@switch# commit
user@switch# set class-of-service interfaces xe-0/0/20 classifiers dscp dscp1
user@switch# commit
```

RELATED DOCUMENTATION

Understanding CoS Packet Flow
Configuring CoS

Understanding CoS Code-Point Aliases

A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers and *rewrite rules*.

NOTE: This topic applies to all EX Series switches except the EX4600. Because the EX4600 uses a different chipset than other EX Series switches, the code-point aliases on EX4600 match those on QFX Series switches. For EX4600 code-point aliases, see *Understanding CoS Code-Point Aliases*.

Behavior aggregate classifiers use class-of-service (CoS) values such as Differentiated Services Code Points (DSCPs) or IEEE 802.1 bits to associate incoming packets with a particular forwarding class and the CoS servicing level associated with that forwarding class. You can assign a meaningful name or alias to the CoS values and use that alias instead of bits when configuring CoS components. These aliases are not part of the specifications but are well known through usage. For example, the alias for DSCP 101110 is widely accepted as ef (expedited forwarding).

When you configure forwarding classes and define classifiers, you can refer to the markers by alias names. You can configure code point alias names for user-defined classifiers. If the value of an alias changes, it alters the behavior of any classifier that references it.

You can configure code-point aliases for the following type of CoS markers:

- dscp or dscp-ipv6—Handles incoming IP and IPv6 packets.
- ieee-802.1—Handles Layer 2 frames.

[Table 30 on page 92](#) shows the default mapping of code-point aliases to IEEE code points.

Table 30: Default IEEE 802.1 Code-Point Aliases

CoS Value Types	Mapping
be	000

Table 30: Default IEEE 802.1 Code-Point Aliases *(Continued)*

CoS Value Types	Mapping
be1	001
ef	010
ef1	011
af11	100
af12	101
nc1	110
nc2	111

[Table 31 on page 93](#) shows the default mapping of code-point aliases to DSCP and DSCP IPv6 code points.

Table 31: Default DSCP and DSCP IPv6 Code-Point Aliases

CoS Value Types	Mapping
ef	101110
af11	001010
af12	001100
af13	001110
af21	010010
af22	010100

Table 31: Default DSCP and DSCP IPv6 Code-Point Aliases *(Continued)*

CoS Value Types	Mapping
af23	010110
af31	011010
af32	011100
af33	011110
af41	100010
af42	100100
af43	100110
be	000000
cs1	001000
cs2	010000
cs3	011000
cs4	100000
cs5	101000
nc1	110000
nc2	111000

RELATED DOCUMENTATION

Understanding Junos CoS Components

Defining CoS Code-Point Aliases

Defining CoS Code-Point Aliases

You can use code-point aliases to streamline the process of configuring CoS features on your switch. A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers and rewrite rules.

You can configure code-point aliases for the following CoS marker types:

- DSCP or DSCP IPv6—Handles incoming IPv4 or IPv6 packets.
- IEEE 802.1p—Handles Layer 2 frames.

To configure a code-point alias:

1. Specify a CoS marker type (IEEE 802.1 or DSCP).
2. Assign an alias.
3. Specify the code point that corresponds to the alias.

```
[edit class-of-service code-point-aliases]  
user@switch# set (dscp | dscp-ipv6 | ieee-802.1) alias-name code-point-bits
```

For example, to configure a code-point alias for an IEEE 802.1 CoS marker type that has the alias name be2 and maps to the code-point bits 001:

```
[edit class-of-service code-point-aliases]  
user@switch# set ieee-802.1 be2 001
```

RELATED DOCUMENTATION

Monitoring CoS Code-Point Value Aliases

Understanding CoS Code-Point Aliases

Understanding CoS Forwarding Classes

IN THIS SECTION

- [Default Forwarding Classes | 97](#)
- [Forwarding Class Configuration Rules | 99](#)
- [Lossless Transport Support | 101](#)

Forwarding classes group traffic and assign the traffic to output queues. Each forwarding class is mapped to an output queue. Classification maps incoming traffic to forwarding classes based on the code point bits in the packet or frame header. Forwarding class to queue mapping defines the output queue used for the traffic classified into a forwarding class.

Except on NFX Series devices, a classifier must associate each packet with one of the following four (QFX10000 switches) or five (other switches) default forwarding classes or with a user-configured forwarding class to assign an output queue to the packet:

- **fcoe**—Guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic.
- **no-loss**—Guaranteed delivery for TCP lossless traffic.
- **best-effort**—Provides best-effort delivery without a service profile. Loss priority is typically not carried in a class-of-service (CoS) value.
- **network-control**—Supports protocol control and is typically high priority.
- **mcast**—(Except QFX10000) Delivery of multdestination (multicast, broadcast, and destination lookup fail) packets.

On NFX Series devices, a classifier must associate each packet with one of the following four default forwarding classes or with a user-configured forwarding class to assign an output queue to the packet:

- **best-effort (be)**—Provides no service profile. Loss priority is typically not carried in a CoS value.
- **expedited-forwarding (ef)**—Provides a low loss, low latency, low jitter, assured bandwidth, end-to-end service.
- **assured-forwarding (af)**—Provides a group of values you can define and includes four subclasses: AF1, AF2, AF3, and AF4, each with two drop probabilities: low and high.
- **network-control (nc)**—Supports protocol control and thus is typically high priority.

The switch supports up to eight (QFX10000 and NFX Series devices), 10 (QFX5200 switches), or 12 (other switches) forwarding classes, thus enabling flexible, differentiated, packet classification. For example, you can configure multiple classes of best-effort traffic such as **best-effort**, **best-effort1**, and **best-effort2**.

On QFX10000 and NFX Series devices, unicast and multideestination (multicast, broadcast, and destination lookup fail) traffic use the same forwarding classes and output queues.

Except on QFX10000 and NFX Series devices, a switch supports 8 queues for unicast traffic (queues 0 through 7) and 2 (QFX5200 switches) or 4 (other switches) output queues for multideestination traffic (queues 8 through 11). Forwarding classes mapped to unicast queues are associated with unicast traffic, and forwarding classes mapped to multideestination queues are associated with multideestination traffic. You cannot map unicast and multideestination traffic to the same queue. You cannot map a strict-high priority queue to a multideestination forwarding class because queues 8 through 11 do not support strict-high priority configuration.

Default Forwarding Classes

[Table 32 on page 97](#) shows the four default forwarding classes that apply to all switches but not NFX Series devices. Except on QFX10000, these forwarding classes apply to unicast traffic. You can rename the forwarding classes. Assigning a new forwarding class name does not alter the default classification or scheduling applied to the queue that is mapped to that forwarding class. CoS configurations can be complex, so unless it is required by your scenario, we recommend that you use the default class names and queue number associations.

Table 32: Default Forwarding Classes

Forwarding Class Name	Default Queue Mapping	Comments
best-effort	0	<p>The software does not apply any special CoS handling to best-effort traffic. This is a backward compatibility feature. Best-effort traffic is usually the first traffic to be dropped during periods of network congestion.</p> <p>By default, this is a lossy forwarding class with a packet drop attribute of drop.</p>

Table 32: Default Forwarding Classes *(Continued)*

Forwarding Class Name	Default Queue Mapping	Comments
fcoe	3	<p>By default, the fcoe forwarding class is a lossless forwarding class designed to handle Fibre Channel over Ethernet (FCoE) traffic. The no-loss packet drop attribute is applied by default.</p> <p>NOTE: By convention, deployments with converged server access typically use IEEE 802.1p priority 3 (011) for FCoE traffic. The default mapping of the fcoe forwarding class is to queue 3. Apply <i>priority-based flow control</i> (PFC) to the entire FCoE data path to configure the end-to-end lossless behavior that FCoE requires.</p> <p>We recommend that you use priority 3 for FCoE traffic unless your network architecture requires that you use a different priority.</p>
no-loss	4	<p>By default, this is a lossless forwarding class with a packet drop attribute of no-loss.</p>
network-control	7	<p>The software delivers packets in this service class with a high priority. (These packets are not delay-sensitive.)</p> <p>Typically, these packets represent routing protocol hello or keepalive messages. Because loss of these packets jeopardizes proper network operation, packet delay is preferable to packet discard.</p> <p>By default, this is a lossy forwarding class with a packet drop attribute of drop.</p>

NOTE: Table 33 on page 99 applies only to multideestination traffic except on QFX10000 switches and NFX Series devices.

Table 33: Default Forwarding Classes for Multidestination Packets

Forwarding Class Name	Default Queue Mapping	Comments
mcast	8	<p>The software does not apply any special CoS handling to the multidestination packets. These packets are usually dropped under congested network conditions.</p> <p>By default, this is a lossy forwarding class with a packet drop attribute of drop.</p>

NOTE: Mirrored traffic is always sent to the queue that corresponds to the multidestination forwarding class. The switched copy of the mirrored traffic is forwarded with the priority determined by the behavior aggregate classification process.

Forwarding Class Configuration Rules

Take the following rules into account when you configure forwarding classes:

Queue Assignment Rules

The following rules govern queue assignment:

- CoS configurations that specify more queues than the switch can support are not accepted. The commit operation fails with a detailed message that states the total number of queues available.
- All default CoS configurations are based on queue number. The name of the forwarding class that appears in the default configuration is the forwarding class currently mapped to that queue.
- (Except QFX10000 and NFX Series devices) Only unicast forwarding classes can be mapped to unicast queues (0 through 7), and only multidestination forwarding classes can be mapped to multidestination queues (8 through 11).
- (Except QFX10000 and NFX Series devices) Strict-high priority queues cannot be mapped to multidestination forwarding classes. (Strict-high priority traffic cannot be mapped to queues 8 through 11).
- If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the same queue must have the same packet drop attribute: either all of the forwarding classes must be lossy or all of the forwarding classes must be lossless.

You can limit the amount of traffic that receives strict-high priority treatment on a strict-high priority queue by configuring a transmit rate. The transmit rate sets the amount of traffic on the queue that receives strict-high priority treatment. The switch treats traffic that exceeds the transmit rate as low priority traffic that receives the queue excess rate bandwidth. Limiting the amount of traffic that receives strict-high priority treatment prevents other queues from being starved while also ensuring that the amount of traffic specified in the transmit rate receives strict-high priority treatment.

NOTE: Except on QFX10000 and NFX Series devices, you can use the *shaping-rate* statement to throttle the rate of packet transmission by setting a maximum bandwidth. On QFX10000 and NFX Series devices, you can use the transmit rate to set a limit on the amount of bandwidth that receives strict-high priority treatment on a strict-high priority queue.

On QFX10000 and NFX Series devices, if you configure more than one strict-high priority queue on a port, you must configure a transmit rate on each of the strict-high priority queues. If you configure more than one strict-high priority queue on a port and you do not configure a transmit rate on the strict-high priority queues, the switch treats only the first queue you configure as a strict-high priority queue. The switch treats the other queues as low priority queues. If you configure a transmit rate on some strict-high priority queues but not on other strict-high priority queues on a port, the switch treats the queues that have a transmit rate as strict-high priority queues, and treats the queues that do not have a transmit rate as low priority queues.

Scheduling Rules

When you configure a forwarding class and map traffic to it (that is, you are not using a default classifier and forwarding class), you must also define a scheduling policy for the forwarding class.

Defining a scheduling policy means:

- Mapping a scheduler to the forwarding class in a scheduler map
- Including the forwarding class in a forwarding class set
- Associating the scheduler map with a traffic control profile
- Attaching the traffic control profile to a forwarding class set and applying the traffic control profile to an interface

On QFX10000 switches and NFX Series devices, you can define a scheduling policy using port scheduling as follows:

- Mapping a scheduler to the forwarding class in a scheduler map
- Applying the scheduler map to one or more interfaces

Rewrite Rules

On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.

Lossless Transport Support

The switch supports up to six lossless forwarding classes. For lossless transport, you must enable PFC on the IEEE 802.1p code point of lossless forwarding classes. The following limitations apply to support lossless transport:

- The external cable length from the switch or QFabric system Node device to other devices cannot exceed 300 meters.
- The internal cable length from the QFabric system Node device to the QFabric system Interconnect device cannot exceed 150 meters.
- For FCoE traffic, the interface maximum transmission unit (MTU) must be at least 2180 bytes to accommodate the packet payload, headers, and checks.
- Changing any portion of a PFC configuration on a port blocks the entire port until the change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Changing the PFC configuration means any change to a congestion notification profile that is configured on a port (enabling or disabling PFC on a code point, changing the MRU or cable-length value, or specifying an output flow control queue). Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

NOTE: QFX10002-60C does not support PFC and lossless queues; that is, default lossless queues (fcoe and no-loss) will be lossy queues.

NOTE: Junos OS Release 12.2 introduces changes to the way lossless forwarding classes (the fcoe and no-loss forwarding classes) are handled.

In Junos OS Release 12.1, both explicitly configuring the fcoe and no-loss forwarding classes, and using the default configuration for these forwarding classes, resulted in the same lossless behavior for traffic mapped to those forwarding classes.

However, in Junos OS Release 12.2, if you explicitly configure the fcoe or the no-loss forwarding class, that forwarding class is no longer treated as a lossless forwarding class. Traffic mapped to

these forwarding classes is treated as lossy (best-effort) traffic. This is true even if the explicit configuration is exactly the same as the default configuration.

If your CoS configuration from Junos OS Release 12.1 or earlier includes the explicit configuration of the `fcoe` or the `no-loss` forwarding class, then when you upgrade to Junos OS Release 12.2, those forwarding classes are not lossless. To preserve the lossless treatment of these forwarding classes, delete the explicit `fcoe` and `no-loss` forwarding class configuration before you upgrade to Junos OS Release 12.2.

See *Overview of CoS Changes Introduced in Junos OS Release 12.2* for detailed information about this change and how to delete an existing lossless configuration.

In Junos OS Release 12.3, the default behavior of the `fcoe` and `no-loss` forwarding classes is the same as in Junos OS Release 12.2. However, in Junos OS Release 12.3, you can configure up to six lossless forwarding classes. All explicitly configured lossless forwarding classes must include the new `no-loss` packet drop attribute or the forwarding class is lossy.

RELATED DOCUMENTATION

Overview of CoS Changes Introduced in Junos OS Release 12.2

Understanding Junos CoS Components

Understanding CoS Packet Flow

Understanding CoS Flow Control (Ethernet PAUSE and PFC)

Example: Configuring Forwarding Classes

Defining CoS Forwarding Classes

Defining CoS Forwarding Classes

Forwarding classes allow you to group packets for transmission. The switch supports a total of eight (QFX10000 and NFX Series devices), 10 (QFX5200 switches), or 12 (other switches) forwarding classes. To forward traffic, you map (assign) the forwarding classes to output queues. Starting in Junos OS Release 22.1R1, QFX10000 Series devices support 16 forwarding classes.

The QFX10000 switches and NFX Series devices have eight output queues, queues 0 through 7. These queues support both unicast and multdestination traffic.

Except on QFX10000 and NFX Series devices, the switch has 10 output queues (QFX5200) or 12 output queues (other switches). Queues 0 through 7 are for unicast traffic and queues 8 through 11 are

for multicast traffic. Forwarding classes mapped to unicast queues must carry unicast traffic, and forwarding classes mapped to multideestination queues must carry multideestination traffic. There are four default unicast forwarding classes and one default multideestination forwarding class.

The default forwarding classes, except on NFX Series devices, are:

NOTE: Except on QFX10000, these are the default unicast forwarding classes.

- `best-effort`—Best-effort traffic
- `fcoe`—Guaranteed delivery for Fibre Channel over Ethernet traffic (do not use on OCX Series switches)
- `no-loss`—Guaranteed delivery for TCP no-loss traffic (do not use on OCX Series switches)
- `network-control`—Network control traffic

NOTE: QFX10002-60C does not support PFC and lossless queues; that is, default lossless queues (`fcoe` and `no-loss`) will be lossy queues.

The default multideestination forwarding class, except on QFX10000 switches and NFX Series devices, is:

- `mcast`—Multideestination traffic

The NFX Series devices have the following default forwarding classes:

- `best-effort (be)`—Provides no service profile. Loss priority is typically not carried in a CoS value.
- `expedited-forwarding (ef)`—Provides a low loss, low latency, low jitter, assured bandwidth, end-to-end service.
- `assured-forwarding (af)`—Provides a group of values you can define and includes four subclasses: AF1, AF2, AF3, and AF4, each with two drop probabilities: low and high.
- `network-control (nc)`—Supports protocol control and thus is typically high priority.

You can map forwarding classes to queues using the `class` statement. You can map more than one forwarding class to a single queue. Except on QFX10000 or NFX Series devices, all forwarding classes mapped to a particular queue must be of the same type, either unicast or multicast. You cannot mix unicast and multicast forwarding classes on the same queue.

All of the forwarding classes mapped to the same queue must have the same packet drop attribute: either all of the forwarding classes must be lossy or all of the forwarding classes must be lossless. This is

important because the default fcoe and no-loss forwarding classes have the no-loss drop attribute, which is not supported on OCX Series switches. On OCX Series switches, do not map traffic to the default fcoe and no-loss forwarding classes.

```
[edit class-of-service forwarding-classes]
user@switch# set class class-name queue-num queue-number <no-loss>
```

One example is to create a forwarding class named be2 and map it to queue 1:

```
[edit class-of-service forwarding-classes]
user@switch# set class be2 queue-num 1
```

Another example is to create a lossless forwarding class named fcoe2 and map it to queue 5:

```
[edit class-of-service forwarding-classes]
user@switch# set class fcoe2 queue-num 5 no-loss
```

NOTE: On switches that do not run ELS software, if you are using Junos OS Release 12.2 or later, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best-effort) traffic and does *not* receive lossless treatment unless you include the optional no-loss packet drop attribute introduced in Junos OS Release 12.3 in the forwarding class configuration..

NOTE: On switches that do not run ELS software, Junos OS Release 11.3R1 and earlier supported an alternate method of mapping forwarding classes to queues that allowed you to map only one forwarding class to a queue using the statement:

```
[edit class-of-service forwarding-classes]
user@switch# set queue queue-number class-name
```

The queue statement has been deprecated and is no longer valid in Junos OS Release 11.3R2 and later. If you have a configuration that uses the queue statement to map forwarding classes to queues, edit the configuration to replace the queue statement with the class statement.

Release History Table

Release	Description
22.1R1	Starting in Junos OS Release 22.1R1, QFX10000 Series devices support 16 forwarding classes.

RELATED DOCUMENTATION

<i>Example: Configuring CoS Hierarchical Port Scheduling (ETS)</i>
<i>Example: Configuring Forwarding Classes</i>
<i>Monitoring CoS Forwarding Classes</i>
<i>Understanding CoS Forwarding Classes</i>
<i>Understanding CoS Port Schedulers on QFX Switches</i>

Example: Configuring Forwarding Classes

IN THIS SECTION

- [Requirements | 105](#)
- [Overview | 106](#)
- [Example 1: Configuring Forwarding Classes for Switches Except QFX10000 | 108](#)
- [Example 2: Configuring Forwarding Classes for QFX10000 Switches | 110](#)

Forwarding classes group packets for transmission. Forwarding classes map to output queues, so the packets assigned to a forwarding class use the output queue mapped to that forwarding class. Except on QFX10000, unicast traffic and multideestination (multicast, broadcast, and destination lookup fail) traffic use separate forwarding classes and output queues.

Requirements

This example uses the following hardware and software components for two configuration examples:

Configuring forwarding classes for switches except QFX10000

- One switch except QFX10000 (this example was tested on a Juniper Networks QFX3500 Switch)

- Junos OS Release 11.1 or later for the QFX Series or Junos OS Release 14.1X53-D20 or later for the OCX Series

Configuring forwarding classes for QFX10000 switches

- One QFX10000 switch
- Junos OS Release 15.1X53-D10 or later for the QFX Series

Overview

The QFX10000 switch supports eight forwarding classes. Other switches support up to 12 forwarding classes. To forward traffic, you must map (assign) the forwarding classes to output queues. On the QFX10000 switch, queues 0 through 7 are for both unicast and multdestination traffic. On other switches, queues 0 through 7 are for unicast traffic, and queues 8 through 9 (QFX5200 switch) or 8 through 11 (other switches) are for multdestination traffic. Except for OCX Series switches, switches support up to six lossless forwarding classes. (OCX Series switches do not support lossless Layer 2 transport.)

The switch provides four default forwarding classes, and except on QFX10000 switches, these four forwarding classes are unicast, plus one default multdestination forwarding class. You can define the remaining forwarding classes and configure them as unicast or multdestination forwarding classes by mapping them to unicast or multdestination queues. The type of queue, unicast or multdestination, determines the type of forwarding class.

The four default forwarding classes (unicast except on QFX10000) are:

- `be`—Best-effort traffic
- `fcoe`—Guaranteed delivery for Fibre Channel over Ethernet traffic (do not use on OCX Series switches)
- `no-loss`—Guaranteed delivery for TCP no-loss traffic (do not use on OCX Series switches)
- `nc`—Network control traffic

Except on QFX10000 switches, the default multdestination forwarding class is:

- `mcast`—Multdestination traffic

Map forwarding classes to queues using the `class` statement. You can map more than one forwarding class to a single queue, but all forwarding classes mapped to a particular queue must be of the same type:

- Except on QFX10000 switches, all forwarding classes mapped to a particular queue must be either unicast or multicast. You cannot mix unicast and multicast forwarding classes on the same queue.

- On QFX10000 switches, all forwarding classes mapped to a particular queue must have the same packet drop attribute: all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless.

```
[edit class-of-service forwarding-classes]
user@switch# set class class-name queue-num queue-number;
```

NOTE: On switches that do not run ELS software, if you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless `fcoe` and `no-loss` forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best-effort) traffic and does *not* receive lossless treatment. In Junos OS Release 12.3 and later, you can include the *no-loss* packet drop attribute in explicit forwarding class configurations to configure a lossless forwarding class.

NOTE: On switches that do not run ELS software, Junos OS Release 11.3R1 and earlier supported an alternate method of mapping forwarding classes to queues that allowed you to map only one forwarding class to a queue using the statement:

```
[edit class-of-service forwarding-classes]
user@switch# set queue queue-number class-name
```

The `queue` statement has been deprecated and is no longer valid in Junos OS Release 11.3R2 and later. If you have a configuration that uses the `queue` statement to map forwarding classes to queues, edit the configuration to replace the `queue` statement with the `class` statement.

NOTE: Hierarchical scheduling controls output queue forwarding. When you define a forwarding class and classify traffic into it, you must also define a scheduling policy for the forwarding class. Defining a scheduling policy means:

- Mapping a scheduler to the forwarding class in a scheduler map
- Including the forwarding class in a forwarding class set
- Associating the scheduler map with a traffic control profile
- Attaching the traffic control profile to a forwarding class set and applying the traffic control profile to an interface

On QFX10000 switches, you can define a scheduling policy using port scheduling:

- Mapping a scheduler to the forwarding class in a scheduler map.
- Applying the scheduler map to one or more interfaces.

Example 1: Configuring Forwarding Classes for Switches Except QFX10000

IN THIS SECTION

● [Verification](#) | 109

Configuration

Step-by-Step Procedure

[Table 34 on page 108](#) shows the configuration forwarding-class-to-queue mapping for this example:

Table 34: Forwarding-Class-to-Queue Example Configuration Except on QFX10000

Forwarding Class	Queue
best-effort	0
nc	7
mcast	8

To configure CoS forwarding classes for switches except QFX10000:

1. Map the best-effort forwarding class to queue 0:

```
[edit class-of-service forwarding-classes]  
user@switch# set class best-effort queue-num 0
```

2. Map the `nc` forwarding class to queue 7:

```
[edit class-of-service forwarding-classes]  
user@switch# set class nc queue-num 7
```

3. Map the `mcast-be` forwarding class to queue 8:

```
[edit class-of-service forwarding-classes]  
user@switch# set class mcast-be queue-num 8
```

Verification

IN THIS SECTION

- [Verifying the Forwarding-Class-to-Queue Mapping | 109](#)

Verifying the Forwarding-Class-to-Queue Mapping

Purpose

Verify the forwarding-class-to-queue mapping. (The system shows only the explicitly configured forwarding classes; it does not show default forwarding classes such as `fcoe` and `no-loss`.)

Action

Verify the results of the forwarding class configuration using the operational mode command `show configuration class-of-service forwarding-classes`:

```
user@switch> show configuration class-of-service forwarding-classes  
class best-effort queue-num 0;  
class network-control queue-num 7;  
class mcast queue-num 8;
```

Example 2: Configuring Forwarding Classes for QFX10000 Switches

IN THIS SECTION

- [Verification | 111](#)

Configuration

Step-by-Step Procedure

[Table 35 on page 110](#) shows the configuration forwarding-class-to-queue mapping for this example:

Table 35: Forwarding-Class-to-Queue Example Configuration on QFX10000

Forwarding Class	Queue
best-effort	0
be1	1
nc	7

To configure CoS forwarding classes for QFX10000 switches:

1. Map the best-effort forwarding class to queue 0:

```
[edit class-of-service forwarding-classes]
user@switch# set class best-effort queue-num 0
```

2. Map the be1 forwarding class to queue 1:

```
[edit class-of-service forwarding-classes]
user@switch# set class be1 queue-num 1
```

3. Map the nc forwarding class to queue 7:

```
[edit class-of-service forwarding-classes]
user@switch# set class nc queue-num 7
```

Verification

IN THIS SECTION

- [Verifying the Forwarding-Class-to-Queue Mapping | 111](#)

Verifying the Forwarding-Class-to-Queue Mapping

Purpose

Verify the forwarding-class-to-queue mapping. (The system shows only the explicitly configured forwarding classes; it does not show default forwarding classes such as fcoe and no-loss.)

Action

Verify the results of the forwarding class configuration using the operational mode command `show configuration class-of-service forwarding-classes`:

```
user@switch> show configuration class-of-service forwarding-classes
class best-effort queue-num 0;
class be1 queue-num 1;
class network-control queue-num 7;
```

RELATED DOCUMENTATION

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Defining CoS Forwarding Classes

Monitoring CoS Forwarding Classes

Overview of CoS Changes Introduced in Junos OS Release 11.3

Understanding CoS Rewrite Rules

As packets enter or exit a network, edge switches might be required to alter the class-of-service (CoS) settings of the packets. *Rewrite rules* set the value of the code point bits (Layer 3 DSCP bits, Layer 2 CoS bits, or MPLS EXP bits) within the header of the outgoing packet. Each rewrite rule:

1. Reads the current forwarding class and loss priority associated with the packet.
2. Locates the new (rewrite) code point value from a table.
3. Writes that code point value into the packet header, replacing the old code point value.

Rewrite rules must be assigned to an interface for rewrites to take effect.

You can apply (bind) one DSCP or DSCP IPv6 rewrite rule and one IEEE 802.1p rewrite rule to each interface. You can also bind EXP rewrite rules to family `mpls` logical interfaces to rewrite the CoS bits of MPLS traffic.

NOTE: OCX Series switches do not support MPLS and do not support EXP rewrite rules.

You cannot apply both a DSCP and a DSCP IPv6 rewrite rule to the same physical interface. Each physical interface supports only one DSCP rewrite rule. Both IP and IPv6 packets use the same DSCP rewrite rule, regardless if the configured rewrite rule is DSCP or DSCP IPv6. You can apply an EXP rewrite rule on an interface that has DSCP or IEEE rewrite rules. Only MPLS traffic on family `mpls` interfaces uses the EXP rewrite rule.

You *can* apply both a DSCP rewrite rule and a DSCP IPv6 rewrite rule to a logical interface. IPv6 packets are rewritten with DSCP-IPv6 rewrite-rules and IPv4 packets are remarked with DSCP rewrite-rules.

NOTE: There are no default rewrite rules. If you want to apply a rewrite rule to outgoing packets, you must explicitly configure the rewrite rule.

You can look at behavior aggregate (BA) classifiers and rewrite rules as two sides of the same coin. A BA classifier reads the code point bits of incoming packets and classifies the packets into forwarding classes, then the system applies the CoS configured for the forwarding class to those packets. Rewrite rules

change (rewrite) the code point bits just before the packets leave the system so that the next switch or router can apply the appropriate level of CoS to the packets. When you apply a rewrite rule to an interface, the rewrite rule is the last CoS action performed on the packet before it is forwarded.

Rewrite rules alter CoS values in outgoing packets on the outbound interfaces of an edge switch to accommodate the policies of a targeted peer. This allows the downstream switch in a neighboring network to classify each packet into the appropriate service group.

NOTE: On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.

NOTE: Rewrite rules are applied *before* the egress filter is matched to traffic. Because the code point rewrite occurs before the egress filter is matched to traffic, the egress filter match is based on the rewrite value, not on the original code point value in the packet.

For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rule rewrites only the outer VLAN tag.

MPLS EXP rewrite rules apply only to `family mpls` logical interfaces. You cannot apply to an EXP rewrite rule to a physical interface. You can configure up to 64 EXP rewrite rules, but you can only use 16 EXP rewrite rules at any time on the switch. On a given logical interface, all pushed MPLS labels have the same EXP rewrite rule applied to them. You can apply different EXP rewrite rules to different logical interfaces on the same physical interface.

NOTE: If the switch is performing penultimate hop popping (PHP), EXP rewrite rules do not take effect. If both an EXP classifier and an EXP rewrite rule are configured on the switch, then the EXP value from the last popped label is copied into the inner label. If either an EXP classifier or an EXP rewrite rule (but not both) is configured on the switch, then the inner label EXP value is sent unchanged.

You can configure enough rewrite rules to handle most, if not all, network scenarios. [Table 36 on page 114](#) shows how many of each type of rewrite rules you can configure, and how many entries you can configure per rewrite rule.

Table 36: Configuring Rewrite Rules

Rewrite Rule Type	Maximum Number of Rewrite Rules	Maximum Number of Entries per Rewrite Rule
IEEE 802.1p	64	128
DSCP	32	128
DSCP IPv6	32	128
MPLS EXP	64	128

You cannot apply rewrite rules directly to integrated routing and bridging (IRB), also known as routed VLAN interfaces (RVIs), because the members of IRBs/RVIs are VLANs, not ports. However, you can apply rewrite rules to the VLAN port members of an IRB/RVI.

NOTE: OCX Series switches do not support IRBs/RVIs.

RELATED DOCUMENTATION

Understanding Junos CoS Components

Defining CoS Rewrite Rules

Configuring Rewrite Rules for MPLS EXP Classifiers

Defining CoS Rewrite Rules

Overview

Edge switches might need to change the class-of-service (CoS) settings of the packets. You can configure rewrite rules to alter code point bit values in outgoing packets on the outbound interfaces of a switch so that the CoS treatment matches the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.

To configure a CoS rewrite rule, create the rule by giving it a name and associating it with a forwarding class, loss priority, and code point. This creates a rewrite table. After the rewrite rule is created, enable it on an interface (EXP rewrite rules can only be enabled on family `mpls` logical interfaces, not on physical interfaces). You can also apply an existing rewrite rule on an interface.

NOTE: On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.

NOTE: To replace an existing rewrite rule on the interface with a new rewrite rule of the same type, first explicitly remove the existing rewrite rule and then apply the new rule.

NOTE: For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rule rewrites only the outer VLAN tag.

Platform-specific Information

- OCX Series switches do not support MPLS, so they do not support EXP rewrite rules.
- QFX5130, QFX5700 & QFX5220 switches do not support DSCP IPv6 classifiers and rewrite rules. However, you can apply DSCP classifiers and rewrite rules for IPV6 traffic as well.

Configuring Rewrite Rules

To create rewrite rules and enable them on interfaces:

- To create an 802.1p rewrite rule named `customup-rw` in the rewrite table for all Layer 2 interfaces:

```
[edit class-of-service rewrite-rules]
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority low code-point 000
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority high code-point 001
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority low code-point 010
user@switch# set ieee-802.1 customup-rw forwarding-class fcqe loss-priority low code-point 011
user@switch# set ieee-802.1 customup-rw forwarding-class ef-no-loss loss-priority low code-point 100
user@switch# set ieee-802.1 customup-rw forwarding-class ef-no-loss loss-priority high code-
```

```

point 101
user@switch# set ieee-802.1 customup-rw forwarding-class nc loss-priority low code-point 110
user@switch# set ieee-802.1 customup-rw forwarding-class nc loss-priority high code-point 111

```

- To enable an 802.1p rewrite rule named `customup-rw` on a Layer 2 interface:

```

[edit]
user@switch# set class-of-service interfaces xe-0/0/7 unit 0 rewrite-rules ieee-802.1
customup-rw

```

NOTE: All forwarding classes assigned to port `xe-0/0/7` must have rewrite rules. Do not mix forwarding classes that have rewrite rules with forwarding classes that do not have rewrite rules on the same physical interface.

- To enable an 802.1p rewrite rule named `customup-rw` on all 10-Gigabit Ethernet interfaces on the switch, use wildcards for the interface name and logical interface (unit) number:

```

[edit]
user@switch# set class-of-service interfaces xe-* unit * rewrite-rules customup-rw

```

NOTE: In this case, *all* forwarding classes assigned to *all* 10-Gigabit Ethernet ports must have rewrite rules. Do not mix forwarding classes that have rewrite rules with forwarding classes that do not have rewrite rules on the same physical interface.

RELATED DOCUMENTATION

Monitoring CoS Rewrite Rules

Configuring Rewrite Rules for MPLS EXP Classifiers

Understanding CoS Rewrite Rules

Understanding CoS MPLS EXP Classifiers and Rewrite Rules

Troubleshooting an Unexpected Rewrite Value

IN THIS SECTION

- Problem | 117
- Cause | 117
- Solution | 118

Problem

Description

Traffic from one or more forwarding classes on an egress port is assigned an unexpected rewrite value.

NOTE: For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rules rewrite only the outer VLAN tag.

Cause

If you configure a rewrite rule for a forwarding class on an egress port, but you do not configure a rewrite rule for every forwarding class on that egress port, then the forwarding classes that do not have a configured rewrite rule are assigned random rewrite values.

For example:

1. Configure forwarding classes fc1, fc2, and fc3.
2. Configure rewrite rules for forwarding classes fc1 and fc2, but not for forwarding class fc3.
3. Assign forwarding classes fc1, fc2, and fc3 to a port.

When traffic for these forwarding classes flows through the port, traffic for forwarding classes fc1 and fc2 is rewritten correctly. However, traffic for forwarding class fc3 is assigned a random rewrite value.

Solution

If any forwarding class on an egress port has a configured rewrite rule, then all forwarding classes on that egress port must have a configured rewrite rule. Configuring a rewrite rule for any forwarding class that is assigned a random rewrite value solves the problem.

TIP: If you want the forwarding class to use the same code point value assigned to it by the ingress classifier, specify that value as the rewrite rule value. For example, if a forwarding class has the IEEE 802.1 ingress classifier code point value 011, configure a rewrite rule for that forwarding class that uses the IEEE 802.1p code point value 011.

NOTE: There are no default rewrite rules. You can bind one rewrite rule for DSCP traffic and one rewrite rule for IEEE 802.1p traffic to an interface. A rewrite rule can contain multiple forwarding-class-to-rewrite-value mappings.

1. To assign a rewrite value to a forwarding class, add the new rewrite value to the same rewrite rule as the other forwarding classes on the port:

```
[edit class-of-service rewrite-rules]
user@switch# set (dscp | ieee-802.1) rewrite-name forwarding-class class-name loss-priority
priority code-point (alias | bits)
```

For example, if the other forwarding classes on the port use rewrite values defined in the rewrite rule *custom-rw*, the forwarding class *be2* is being randomly rewritten, and you want to use IEEE 802.1 code point 002 for the *be2* forwarding class:

```
[edit class-of-service rewrite-rules]
user@switch# set ieee-802.1 custom-rw forwarding-class be2 loss-priority low code-point 002
```

2. Enable the rewrite rule on an interface if it is not already enabled on the desired interface:

```
[edit]
user@switch# set class-of-service interfaces interface-name unit unit rewrite-rules (dscp |
ieee-802.1) rewrite-rule-name
```

For example, to enable the rewrite rule `custom-rw` on interface `xe-0/0/24.0`:

```
[edit]
user@switch# set class-of-service interfaces xe-0/0/24 unit 0 rewrite-rules ieee-802.1 custom-
rw
```

RELATED DOCUMENTATION

[interfaces](#)

[rewrite-rules](#)

[Defining CoS Rewrite Rules](#)

[Monitoring CoS Rewrite Rules](#)

3

PART

Scheduling Traffic

[Using Schedulers](#) | 121

Using Schedulers

IN THIS CHAPTER

- Understanding CoS Scheduling Behavior and Configuration Considerations | 121
- Defining CoS Queue Schedulers for Port Scheduling | 128
- Defining CoS Queue Scheduling Priority | 132
- Example: Configuring Queue Scheduling Priority | 133
- Understanding CoS Traffic Control Profiles | 138
- Understanding CoS Priority Group Scheduling | 140
- Defining CoS Traffic Control Profiles (Priority Group Scheduling) | 144
- Example: Configuring Traffic Control Profiles (Priority Group Scheduling) | 145
- Understanding CoS Priority Group and Queue Guaranteed Minimum Bandwidth | 149
- Example: Configuring Minimum Guaranteed Output Bandwidth | 152
- Understanding CoS Priority Group Shaping and Queue Shaping (Maximum Bandwidth) | 159
- Example: Configuring Maximum Output Bandwidth | 162
- Understanding CoS Explicit Congestion Notification | 168

Understanding CoS Scheduling Behavior and Configuration Considerations

Many factors affect scheduling configuration and bandwidth requirements, including:

- When you configure bandwidth for a forwarding class (each forwarding class is mapped to a queue) or a forwarding class set (priority group), the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a forwarding class or for a forwarding class set, consider the preamble and the IFG as well as the data in the calculations.
- When you configure a forwarding class to carry traffic on the switch (instead of using only default forwarding classes), you must also define a scheduling policy for the user-configured forwarding

class. Some switches support enhanced transmission selection (ETS) hierarchical port scheduling, some switches support direct port scheduling, and some switches support both methods of scheduling.

For ETS hierarchical port scheduling, defining a hierarchical scheduling policy using ETS means:

- Mapping a scheduler to the forwarding class in a scheduler map
- Including the forwarding class in a forwarding class set
- Associating the scheduler map with a traffic control profile
- Attaching the traffic control profile to a forwarding class set and an interface

On switches that support port scheduling, defining a scheduling policy means:

- Mapping a scheduler to the forwarding class in a scheduler map.
- Applying the scheduler map to one or more interfaces.
- On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.
- For packets that carry both an inner VLAN tag and an outer VLAN tag, rewrite rules rewrite only the outer VLAN tag.
- For ETS hierarchical port scheduling, configuring the minimum guaranteed bandwidth (`transmit-rate`) for a forwarding class does not work unless you also configure the minimum guaranteed bandwidth (`guaranteed-rate`) for the forwarding class set in the traffic control profile.

Additionally, the sum of the transmit rates of the forwarding classes in a forwarding class set should not exceed the guaranteed rate for the forwarding class set. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.) If you configure transmit rates whose sum exceeds the guaranteed rate of the forwarding class set, the commit check fails and the system rejects the configuration.

- For ETS hierarchical port scheduling, the sum of the forwarding class set guaranteed rates cannot exceed the total port bandwidth. If you configure guaranteed rates whose sum exceeds the port bandwidth, the system sends a syslog message to notify you that the configuration is not valid. However, the system does not perform a commit check. If you commit a configuration in which the sum of the guaranteed rates exceeds the port bandwidth, the hierarchical scheduler behaves unpredictably.
- For ETS hierarchical port scheduling, if you configure the `guaranteed-rate` of a forwarding class set as a percentage, configure all of the transmit rates associated with that forwarding class set as

percentages. In this case, if any of the transmit rates are configured as absolute values instead of percentages, the configuration is not valid and the system sends a syslog message.

- There are several factors to consider if you want to configure a strict-high priority queue (forwarding class):
 - On QFX5200, QFX3500, and QFX3600 switches and on QFabric systems, you can configure only one strict-high priority queue (forwarding class).

On QFX5100 and EX4600 switches, you can configure only one forwarding-class-set (priority group) as strict-high priority. All queues which are part of that strict-high forwarding class set then act as strict-high queues.

On QFX10000 switches, there is no limit to the number of strict-high priority queues you can configure.

- You cannot configure a minimum guaranteed bandwidth (transmit-rate) for a strict-high priority queue on QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems.

On QFX5200 and QFX10000 switches, you can set the transmit-rate on strict-high priority queues to set a limit on the amount of traffic that the queue treats as strict-high priority traffic. Traffic in excess of the transmit-rate is treated as best-effort traffic, and receives an excess bandwidth sharing weight of “1”, which is the proportion of extra bandwidth the strict-high priority queue can share on the port. Queues that are not strict-high priority queues use the transmit rate (default) or the configured excess rate to determine the proportion (weight) of extra port bandwidth the queue can share. However, you cannot configure an excess rate on a strict-high priority queue, and you cannot change the excess bandwidth sharing weight of “1” on a strict-high priority queue.

For ETS hierarchical port scheduling, you cannot configure a minimum guaranteed bandwidth (guaranteed-rate) for a forwarding class set that includes a strict-high priority queue.

- Except on QFX10000 switches, for ETS hierarchical port scheduling only, you must create a separate forwarding class set for a strict-high priority queue. On QFX10000 switches, you can mix strict-high priority and low priority queues in the same forwarding class set.
- Except on QFX10000 switches, for ETS hierarchical port scheduling, only one forwarding class set can contain a strict-high priority queue. On QFX10000 switches, this restriction does not apply.
- Except on QFX10000 switches, for ETS hierarchical port scheduling, a strict-high priority queue cannot belong to the same forwarding class set as queues that are not strict-high priority. (You cannot mix a strict-high priority forwarding class with forwarding classes that are not strict-high priority in one forwarding class set.) On QFX10000 switches, you can mix strict-high priority and low priority queues in the same forwarding class set.

- For ETS hierarchical port scheduling on switches that use different forwarding class sets for unicast and multdestination (multicast, broadcast, and destination lookup fail) traffic, a strict-high priority queue cannot belong to a multdestination forwarding class set.
- On QFX10000 systems, we recommend that you always configure a transmit rate on strict-high priority queues to prevent them from starving other queues. If you do not apply a transmit rate to limit the amount of bandwidth strict-high priority queues can use, then strict-high priority queues can use all of the available port bandwidth and starve other queues on the port.

On QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, we recommend that you always apply a shaping rate to the strict-high priority queue to prevent it from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

- On QFabric systems, if any queue that contains outgoing packets does not transmit packets for 12 consecutive seconds, the port automatically resets. Failure of a queue to transmit packets for 12 consecutive seconds might be due to:
 - A strict-high priority queue consuming all of the port bandwidth
 - Several queues consuming all of the port bandwidth
 - Any queue or port receiving continuous *priority-based flow control* (PFC) or 802.3x Ethernet PAUSE messages (received PFC and PAUSE messages prevent a queue or a port, respectively, from transmitting packets because of network congestion)
 - Other conditions that prevent a queue from obtaining port bandwidth for 12 consecutive seconds

If the cause is a strict-high priority queue consuming all of the port bandwidth, use rate shaping to configure a maximum rate for the strict-high priority queue and prevent it from using all of the port bandwidth. To configure rate shaping, include the `shaping-rate (rate | percent percentage)` statement at the `[edit class-of-service schedulers scheduler-name]` hierarchy level and apply the shaping rate to the strict-high priority scheduler. We recommend that you always apply a shaping rate to strict-high priority traffic to prevent the strict-high priority queue from starving other queues.

If several queues consume all of the port bandwidth, you can use a scheduler to rate shape those queues and prevent them from using all of the port bandwidth.

- For transmit rates below 1 Gbps, we recommend that you configure the transmit rate as a percentage instead of as a fixed rate. This is because the system converts fixed rates into percentages and might round small fixed rates to a lower percentage. For example, a fixed rate of 350 Mbps is rounded down to 3 percent instead of 3.5 percent.

- When you set the maximum bandwidth for a queue or for a priority group (shaping-rate) at 100 Kbps or lower, the traffic shaping behavior is accurate only within +/- 20 percent of the configured shaping-rate.
- On QFX10000 switches, configuring rate shaping ([set class-of-service schedulers *scheduler-name* transmit-rate (*rate / percentage*) exact) on a LAG interface using the [edit class-of-service interfaces *lag-interface-name* scheduler-map *scheduler-map-name*] statement can result in scheduled traffic streams receiving more LAG link bandwidth than expected.

You configure rate shaping in a scheduler to set the maximum bandwidth for traffic assigned to a forwarding class on a particular output queue on a port. For example, you can use a scheduler to configure rate shaping on traffic assigned to the best-effort forwarding class mapped to queue 0, and then apply the scheduler to an interface using a scheduler map, to set the maximum bandwidth for best-effort traffic mapped to queue 0 on that port. Traffic in the best-effort forwarding can use no more than the amount of port bandwidth specified by the transmit rate when you use the exact option.

LAG interfaces are composed of two or more Ethernet links bundled together to function as a single interface. The switch can hash traffic entering a LAG interface onto any member link in the LAG interface. When you configure rate shaping and apply it to a LAG interface, the way that the switch applies the rate shaping to traffic depends on how the switch hashes the traffic onto the LAG links.

To illustrate how link hashing affects the way the switch applies a shaping rate to LAG traffic, let's look at a LAG interface (ae0) that has two member links (xe-0/0/20 and xe-0/0/21). On LAG ae0, we configure rate shaping of 2g for traffic assigned to the best-effort forwarding class, which is mapped to output queue 0. When traffic in the best-effort forwarding class reaches the LAG interface, the switch hashes the traffic onto one of the two member links.

If the switch hashes all of the best-effort traffic onto the same LAG link, the traffic receives a maximum of 2g bandwidth on that link. In this case, the intended cumulative limit of 2g for best-effort traffic on the LAG is enforced.

However, if the switch hashes the best-effort traffic onto both of the LAG links, the traffic receives a maximum of 2g bandwidth on *each* LAG link, not 2g as a cumulative total for the entire LAG, so the best-effort traffic receives a maximum of 4g on the LAG, not the 2g set by the rate shaping configuration. When hashing spreads the traffic assigned to an output queue (which is mapped to a forwarding class) across multiple LAG links, the effective rate shaping (cumulative maximum bandwidth) on the LAG is:

(number of LAG member interfaces) x (rate shaping for the output queue) = cumulative LAG rate shaping

- On switches that do not use virtual output queues (VOQs), ingress port congestion can occur during periods of egress port congestion if an ingress port forwards traffic to more than one egress port, and at least one of those egress ports experiences congestion. If this occurs, the congested egress port can cause the ingress port to exceed its fair allocation of ingress buffer resources. When the ingress

port exceeds its buffer resource allocation, frames are dropped at the ingress. Ingress port frame drop affects not only the congested egress ports, but also all of the egress ports to which the congested ingress port forwards traffic.

If a congested ingress port drops traffic that is destined for one or more uncongested egress ports, configure a weighted random early detection (WRED) drop profile and apply it to the egress queue that is causing the congestion. The drop profile prevents the congested egress queue from affecting egress queues on other ports by dropping frames at the egress instead of causing congestion at the ingress port.

NOTE: On systems that support lossless transport, do not configure drop profiles for lossless forwarding classes such as the default `fcoe` and `no-loss` forwarding classes. FCoE and other lossless traffic queues require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

- On systems that use different classifiers for unicast and multdestination traffic and that support lossless transport, on an ingress port, do not configure classifiers that map the same IEEE 802.1p code point to both a multdestination traffic flow and a lossless unicast traffic flow (such as the default lossless `fcoe` or `no-loss` forwarding classes). Any code point used for multdestination traffic on a port should not be used to classify unicast traffic into a lossless forwarding class on the same port.

If a multdestination traffic flow and a lossless unicast traffic flow use the same code point on a port, the multdestination traffic is treated the same way as the lossless traffic. For example, if priority-based flow control (PFC) is applied to the lossless traffic, the multdestination traffic of the same code point is also paused. During periods of congestion, treating multdestination traffic the same as lossless unicast traffic can create ingress port congestion for the multdestination traffic and affect the multdestination traffic on all of the egress ports the multdestination traffic uses.

For example, the following configuration can cause ingress port congestion for the multdestination flow:

1. For unicast traffic, IEEE 802.1p code point 011 is classified into the `fcoe` forwarding class:

```
user@switch# set class-of-service classifiers ieee-802.1 ucast-cl forwarding-class fcoe
loss-priority low code-points 011
```

2. For multdestination traffic, IEEE 802.1p code point 011 is classified into the `mcast` forwarding class:

```
user@switch# set class-of-service classifiers ieee-802.1 mcast-cl forwarding-class mcast
loss-priority low code-points 011
```

3. The unicast classifier that maps traffic with code point 011 to the fcoe forwarding class is mapped to interface xe-0/0/1:

```
user@switch# set class-of-service interfaces xe-0/0/1 unit 0 classifiers ieee-802.1
ucast_cl
```

4. The multdestination classifier that maps traffic with code point 011 to the mcast forwarding class is mapped to all interfaces (multidestination traffic maps to all interfaces and cannot be mapped to individual interfaces):

```
user@switch# set class-of-service multi-destination classifiers ieee-802.1 mcast-cl
```

Because the same code point (011) maps unicast traffic to a lossless traffic flow and also maps multidestination traffic to a multidestination traffic flow, the multidestination traffic flow might experience ingress port congestion during periods of congestion.

To avoid ingress port congestion, do not map the code point used by the multidestination traffic to lossless unicast traffic. For example:

1. Instead of classifying code point 011 into the fcoe forwarding class, classify code point 011 into the best-effort forwarding class:

```
user@switch# set class-of-service classifiers ieee-802.1 ucast_cl forwarding-class best-
effort loss-priority low code-points 011
```

2.

```
user@switch# set class-of-service classifiers ieee-802.1 mcast-cl forwarding-class mcast
loss-priority low code-points 011
```

3.

```
user@switch# set class-of-service interfaces xe-0/0/1 unit 0 classifiers ieee-802.1
ucast_cl
```

4.

```
user@switch# set class-of-service multi-destination classifiers ieee-802.1 mcast-cl
```

Because the code point 011 does not map unicast traffic to a lossless traffic flow, the multidestination traffic flow does not experience ingress port congestion during periods of congestion.

The best practice is to classify unicast traffic with IEEE 802.1p code points that are also used for multdestination traffic into best-effort forwarding classes.

Defining CoS Queue Schedulers for Port Scheduling

Schedulers define the CoS properties of output queues. You configure CoS properties in a scheduler, then map the scheduler to a forwarding class. Forwarding classes are in turn mapped to output queues. Classifiers map incoming traffic into forwarding classes based on IEEE 802.1p, DSCP, or EXP code points. CoS scheduling properties include the amount of interface bandwidth assigned to the queue, the priority of the queue, whether explicit congestion notification (ECN) is enabled on the queue, and the WRED packet drop profiles associated with the queue.

The parameters you configure in a scheduler define the following characteristics for the queues mapped to the scheduler:

- **priority**—One of three bandwidth priorities that queues associated with a scheduler can receive:
 - **low**—The scheduler has low priority.
 - **high**—The scheduler has high priority. High priority traffic takes precedence over low priority traffic.
 - **strict-high**—The scheduler has strict-high priority. Strict-high priority queues receive preferential treatment over low-priority queues and receive all of their configured bandwidth before low-priority queues are serviced. Low-priority queues do not transmit traffic until strict-high priority queues are empty.

NOTE: We strongly recommend that you configure a transmit rate on all strict-high priority queues to limit the amount of traffic the switch treats as strict-high priority traffic and prevent strict-high priority queues from starving other queues on the port. This is especially important if you configure more than one strict-high priority queue on a port. If you do not configure a transmit rate to limit the amount of bandwidth strict-high priority queues can use, then the strict-high priority queues can use all of the available port bandwidth and starve other queues on the port.

The switch treats traffic in excess of the transmit rate as best-effort traffic that receives bandwidth from the leftover (excess) port bandwidth pool. On strict-high priority queues, all traffic that exceeds the transmit rate shares in the port excess bandwidth pool based on the strict-high priority excess bandwidth sharing weight of “1”, which is not configurable. The actual amount of extra bandwidth that traffic exceeding the transmit rate receives

depends on how many other queues consume excess bandwidth and the excess rates of those queues.

- **transmit-rate**—Minimum guaranteed bandwidth, also known as the *committed information rate (CIR)*, set as a percentage rate or as an absolute value in bits per second. By default, the transmit rate also determines the amount of excess (extra) port bandwidth the queue can share if you do not explicitly configure an excess rate. Extra bandwidth is allocated among the queues on the port in proportion to the transmit rate of each queue. Except on QFX10000 switches, you can configure *shaping-rate* to throttle the rate of packet transmission. On QFX10000 switches, on queues that are not strict-high priority queues, you can configure a transmit rate as *exact*, which shapes the transmission by setting the transmit rate as the maximum bandwidth the queue can consume on the port.

NOTE: On QFX10000 switches, oversubscribing all 8 queues configured with the `transmit rate exact (shaping)` statement at the `[edit class-of-service schedulers scheduler-name]` hierarchy level might result in less than 100 percent utilization of port bandwidth.

On strict-high priority queues, the transmit rate sets the amount of bandwidth used for strict-high priority forwarding; traffic in excess of the transmit rate is treated as best-effort traffic that receives the queue excess rate.

NOTE: Include the preamble bytes and interframe gap (IFG) bytes as well as the data bytes in your bandwidth calculations.

- **excess-rate**—Percentage of extra bandwidth (bandwidth that is not used by other queues) a low-priority queue can receive. If not set, the switch uses the transmit rate to determine extra bandwidth sharing. You cannot set an excess rate on a strict-high priority queue.
- **drop-profile-map**—Drop profile mapping to a packet loss priority to apply WRED to the scheduler and control packet drop for different packet loss priorities during periods of congestion.
- **buffer-size**—Size of the queue buffer as a percentage of the dedicated buffer space on the port, or as a proportional share of the dedicated buffer space on the port that remains after the explicitly configured queues are served.
- **explicit-congestion-notification**—ECN enable on a best-effort queue. ECN enables end-to-end congestion notification between two ECN-enabled endpoints on TCP/IP based networks. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. ECN is disabled by default.

NOTE: Do not configure drop profiles for the fcoe and no-loss forwarding classes. FCoE and other lossless traffic queues require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

To apply scheduling properties to traffic, map schedulers to forwarding classes using a scheduler map, and then apply the scheduler map to interfaces. Using different scheduler maps, you can map different schedulers to the same forwarding class on different interfaces, to apply different scheduling to that traffic on different interfaces.

To configure a scheduler using the CLI:

1. Name the scheduler and set the minimum guaranteed bandwidth for the queue; optionally, set a maximum bandwidth limit (shaping rate) on a low priority queue by configuring either *shaping-rate* (except on QFX10000 switches) or the *exact* option (only on QFX10000 switches):

```
[edit class-of-service]
user@switch# set schedulers scheduler-name transmit-rate (rate | percent percentage)
<exact>
```

2. Set the amount of excess bandwidth a low-priority queue can share:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name excess-rate percent percentage
```

3. Set the queue priority:

```
[edit class-of-service schedulers scheduler-name]
user@switch# set priority level
```

4. Specify drop profiles for packet loss priorities using a drop profile map:

```
[edit class-of-service schedulers scheduler-name]
user@switch# set drop-profile-map loss-priority (low | medium-high | high) drop-profile drop-profile-name
```

5. Configure the size of the buffer space for the queue:

```
[edit class-of-service schedulers scheduler-name]
user@switch# set buffer-size (percent percent | remainder)
```

6. Enable ECN, if desired (on best-effort traffic only):

```
[edit class-of-service schedulers scheduler-name]
user@switch# set explicit-congestion-notification
```

7. Configure a scheduler map to map the scheduler to a forwarding class, which applies the scheduler's properties to the traffic in that forwarding class:

```
[edit class-of-service]
user@switch# set scheduler-maps scheduler-map-name forwarding-class forwarding-class-name
scheduler scheduler-name
```

8. Assign the scheduler map and its associated schedulers to one or more interfaces.

```
[edit class-of-service]
user@switch# set interfaces interface-name scheduler-map scheduler-map-name
```

RELATED DOCUMENTATION

Example: Configuring Queue Schedulers for Port Scheduling

Example: Configuring ECN

Defining CoS Queue Scheduling Priority

Configuring CoS WRED Drop Profiles

Monitoring CoS Scheduler Maps

Understanding CoS Port Schedulers on QFX Switches

Understanding CoS Explicit Congestion Notification

Defining CoS Queue Scheduling Priority

You can configure the scheduling priority of individual queues by specifying the priority in a scheduler, and then associating the scheduler with a queue by using a scheduler map. On QFX5100, QFX5200, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, queues can have one of two bandwidth scheduling priorities, strict-high priority or low priority. On QFX10000 Series switches, queues can also be configured as high priority.

NOTE: By default, all queues are low priority queues.

The switch services low priority queues after servicing any queue that has strict-high priority traffic or high priority traffic. Strict-high priority queues receive preferential treatment over all other queues and receive all of their configured bandwidth before other queues are serviced. Low-priority queues do not transmit traffic until strict-high priority queues are empty, and receive the bandwidth that remains after the strict-high queues have been serviced. High priority queues receive preference over low priority queues.

Different switches handle traffic configured as strict-high priority traffic in different ways:

- QFX5100, QFX5200, QFX3500, QFX3600, and EX4600 switches, and QFabric systems—You can configure only one queue as a strict-high priority queue.

On these switches, we recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

- QFX10000 switches—You can configure as many queues as you want as strict-high priority. However, keep in mind that too much strict-high priority traffic can starve low priority queues on the port.

NOTE: We strongly recommend that you configure a transmit rate on all strict-high priority queues to limit the amount of traffic the switch treats as strict-high priority traffic and prevent strict-high priority queues from starving other queues on the port. This is especially important if you configure more than one strict-high priority queue on a port. If you do not configure a transmit rate to limit the amount of bandwidth strict-high priority queues can use, then the strict-high priority queues can use all of the available port bandwidth and starve other queues on the port.

The switch treats traffic in excess of the transmit rate as best-effort traffic that receives bandwidth from the leftover (excess) port bandwidth pool. On strict-high priority queues, all traffic that exceeds the transmit rate shares in the port excess bandwidth pool based on the

strict-high priority excess bandwidth sharing weight of “1”, which is not configurable. The actual amount of extra bandwidth that traffic exceeding the transmit rate receives depends on how many other queues consume excess bandwidth and the excess rates of those queues.

- To configure queue priority using the CLI:

```
[edit class-of-service]
```

```
user@switch# set schedulers scheduler-name priority level
```

RELATED DOCUMENTATION

Example: Configuring Queue Scheduling Priority

Monitoring CoS Scheduler Maps

Example: Configuring Queue Scheduling Priority

IN THIS SECTION

- [Requirements | 134](#)
- [Overview | 135](#)
- [Verification | 137](#)

You can configure the bandwidth scheduling priority of individual queues by specifying the priority in a scheduler, and then using a scheduler map to associate the scheduler with a queue.

Configuring Queue Scheduling Priority

CLI Quick Configuration

To quickly configure queue scheduling priority, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

```
[edit class-of-service]
set schedulers fcoe-sched priority low
set schedulers nl-sched priority low
set scheduler-maps schedmap1 forwarding-class fcoe scheduler fcoe-sched
set scheduler-maps schedmap1 forwarding-class no-loss scheduler nl-sched
```

Step-by-Step Procedure

To configure queue priority using the CLI:

1. Create the FCoE scheduler with low priority:

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low
```

2. Create the no-loss scheduler with low priority:

```
[edit class-of-service]
user@switch# set schedulers nl-sched priority low
```

3. Associate the schedulers with the desired queues in the scheduler map:

```
[edit class-of-service]
user@switch# set scheduler-maps schedmap1 forwarding-class fcoe scheduler fcoe-sched
user@switch# set scheduler-maps schedmap1 forwarding-class no-loss scheduler nl-sched
```

Requirements

This example uses the following hardware and software components:

- One switch.
- Junos OS Release 11.1 or later for the QFX Series or Junos OS Release 14.1X53-D20 or later for the OCX Series.

Overview

Queues can have one of several bandwidth priorities:

- **strict-high**—Strict-high priority allocates bandwidth to the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced. On QFX10000 switches, you can configure as many queues as you want as strict-high priority queues. On QFX5200, QFX3500, and QFX3600 switches and on QFabric systems, you can configure only one queue as a strict-high queue. On QFX5100 and EX4600 switches, you can configure only one forwarding-class-set (priority group) as strict-high priority. All queues which are part of that strict-high forwarding class set then act as strict-high queues.

NOTE: On QFX5200 switches, it is not possible to support multiple queues with strict-high priority because QFX5200 doesn't support flexible hierarchical scheduling. When multiple strict-high priority queues are configured, all of those queues are treated as strict-high priority but the higher number queue among them is given highest priority.

On QFX10000 switches, if you configure strict-high priority queues on a port, we strongly recommend that you configure a transmit rate on those queues. The transmit rate sets the amount of traffic that the switch forwards as strict-high priority; traffic in excess of the transmit rate is treated as best-effort traffic that receives the queue excess rate. Even if you configure only one strict-high priority queue, we strongly recommend that you configure a transmit rate the queue to prevent it from starving other queues. If you do not configure a transmit rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

On QFX5200, QFX5100, QFX3500, QFX3600, and EX4600 switches and on QFabric systems, we recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

NOTE: On switches that support enhanced transmission selection (ETS) hierarchical scheduling, if you use ETS and you configure a strict-high priority queue, you must create a forwarding class set that is dedicated only to strict-high priority traffic. Only one forwarding

class set can contain a strict-high priority queue. Queues that are not strict-high priority cannot belong to the same forwarding class set as strict-high priority queues.

On switches that use different output queues for unicast and multdestination traffic, the multdestination forwarding class set cannot contain strict-high priority queues.

- `high` (QFX10000 Series switches only)—High priority. Traffic with high priority is serviced after any queue that has a `strict-high` priority, and before queues with low priority.
- `low`—Low priority. Traffic with low priority is serviced after any queue that has a `strict-high` priority.

NOTE: By default, all queues are low priority queues.

Table 37 on page 136 shows the configuration components for this example.

This example describes how to set the queue priority for two forwarding classes (queues) named `fcoe` and `no-loss`. Both queues have a priority of `low`. The scheduler for the `fcoe` queue is named `fcoe-sched` and the scheduler for the `no-loss` queue is named `n1-sched`. One scheduler map, `schedmap1`, associates the schedulers to the queues.

Table 37: Components of the Queue Scheduler Priority Configuration Example

Component	Settings
Hardware	One switch
Schedulers	<code>fcoe-sched</code> for FCoE traffic <code>n1-sched</code> for no-loss traffic
Priority	<code>low</code> for FCoE traffic <code>low</code> for no-loss traffic
Scheduler map	<code>schedmap1</code> : FCoE mapping: scheduler <code>fcoe-sched</code> to forwarding class <code>fcoe</code> No-loss mapping: scheduler <code>n1-sched</code> to forwarding class <code>no-loss</code>

NOTE: OCX Series switches do not support lossless transport. On OCX Series switches, the default DSCP classifier does not map traffic to the default fcoe and no-loss forwarding classes. On an OCX Series switch, you could use this example by substituting other forwarding classes (for example, best-effort or network-control) for the fcoe and no-loss forwarding classes, and naming the schedulers appropriately. The active forwarding classes (best-effort, network-control, and mcast) share the unused bandwidth assigned to the fcoe and no-loss forwarding classes.

Verification

IN THIS SECTION

- [Verifying the Queue Scheduling Priority | 137](#)
- [Verifying the Scheduler-to-Forwarding-Class Mapping | 138](#)

To verify that you configured the queue scheduling priority for bandwidth and mapped the schedulers to the correct forwarding classes, perform these tasks:

Verifying the Queue Scheduling Priority

Purpose

Verify that you configured the queue schedulers fcoe-sched and nl-sched with low queue scheduling priority.

Action

Display the fcoe-sched scheduler priority configuration using the operational mode command `show configuration class-of-service schedulers fcoe-sched priority`:

```
user@switch> show configuration class-of-service schedulers fcoe-sched priority
priority low;
```

Display the nl-sched scheduler priority configuration using the operational mode command `show configuration class-of-service schedulers nl-sched priority`:

```
user@switch> show configuration class-of-service schedulers nl-sched priority
priority low;
```

Verifying the Scheduler-to-Forwarding-Class Mapping

Purpose

Verify that you configured the scheduler map `schedmap1` to map scheduler `fcoe-sched` to forwarding class `fcoe` and schedule `nl-sched` to forwarding class `no-loss`.

Action

Display the scheduler map `schedmap1` using the operational mode command `show configuration class-of-service scheduler-maps schedmap1`:

```
user@switch> show configuration class-of-service scheduler-maps schedmap1
forwarding-class fcoe scheduler fcoe-sched;
forwarding-class no-loss scheduler nl-sched;
```

RELATED DOCUMENTATION

Defining CoS Queue Scheduling Priority

Monitoring CoS Scheduler Maps

Understanding CoS Traffic Control Profiles

A traffic control profile defines the output bandwidth and scheduling characteristics of forwarding class sets (priority groups). The forwarding classes (which are mapped to output queues) that belong to a forwarding class set (fc-set) share the bandwidth that you assign to the fc-set in the traffic control profile.

This two-tier hierarchical scheduling architecture provides flexibility in allocating resources among forwarding classes, and also:

- Assigns a portion of port bandwidth to an fc-set. You define the port resources for the fc-set in a traffic control profile.
- Allocates fc-set bandwidth among the forwarding classes (queues) that belong to the fc-set. A scheduler map attached to the traffic control profile defines the amount of the fc-set's resources that each forwarding class can use.

Attaching an fc-set and a traffic control profile to a port defines the hierarchical scheduling properties of the group and the forwarding classes that belong to the group.

The ability to create fc-sets supports enhanced transmission selection (ETS), which is described in IEEE 802.1Qaz. When an fc-set does not use its allocated port bandwidth, ETS shares the excess port bandwidth among other fc-sets on the port in proportion to their guaranteed minimum bandwidth (guaranteed rate). This utilizes the port bandwidth better than scheduling schemes that reserve bandwidth for groups even if that bandwidth is not used. ETS shares unused port bandwidth, so traffic groups that need extra bandwidth can use it if the bandwidth is available, while preserving the ability to specify the minimum guaranteed bandwidth for traffic groups.

Traffic control profiles define the following CoS properties for fc-sets:

- Minimum guaranteed bandwidth—Also known as the *committed information rate (CIR)*. This is the minimum amount of port bandwidth the priority group receives. Priorities in the priority group receive their minimum guaranteed bandwidth as a portion of the priority group's minimum guaranteed bandwidth. The guaranteed-rate statement defines the minimum guaranteed bandwidth.

NOTE: You cannot apply a traffic control profile with a minimum guaranteed bandwidth to a priority group that includes strict-high priority queues.

- Shared excess (extra) bandwidth—When the priority groups on a port do not consume the full amount of bandwidth allocated to them or there is unallocated link bandwidth available, priority groups can contend for that extra bandwidth if they need it. Priorities in the priority group contend for extra bandwidth as a portion of the priority group's extra bandwidth. The amount of extra bandwidth for which a priority group can contend is proportional to the priority group's guaranteed minimum bandwidth (guaranteed rate).
- Maximum bandwidth—Also known as *peak information rate (PIR)*. This is the maximum amount of port bandwidth the priority group receives. Priorities in the priority group receive their maximum bandwidth as a portion of the priority group's maximum bandwidth. The shaping-rate statement defines the maximum bandwidth.
- Queue scheduling—Each traffic control profile includes a scheduler map. The scheduler map maps forwarding classes (priorities) to schedulers to define the scheduling characteristics of the individual forwarding classes in the fc-set. The resources scheduled for each forwarding class represent portions of the resources that the traffic control profile schedules for the entire fc-set, not portions

of the total link bandwidth. The `scheduler-maps` statement defines the mapping of forwarding classes to schedulers.

RELATED DOCUMENTATION

Understanding CoS Hierarchical Port Scheduling (ETS)

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

Defining CoS Traffic Control Profiles (Priority Group Scheduling)

Understanding CoS Priority Group Scheduling

IN THIS SECTION

- [Priority Group Scheduling Components | 141](#)
- [Default Traffic Control Profile | 142](#)
- [Guaranteed Rate \(Minimum Guaranteed Bandwidth\) | 142](#)
- [Sharing Extra Bandwidth | 142](#)
- [Shaping Rate \(Maximum Bandwidth\) | 143](#)
- [Scheduler Maps | 143](#)

Priority group scheduling defines the class-of-service (CoS) properties of a group of output queues (priorities). Priority group scheduling works with output queue scheduling to create a two-tier hierarchical scheduler. The hierarchical scheduler allocates bandwidth to a group of queues (a priority group, called a forwarding class set in Junos OS configuration). Queue scheduling determines the portion of the priority group bandwidth that the particular queue can use.

You configure priority group scheduling in a traffic control profile and then associate the traffic control profile with a forwarding class set and an interface. You attach a scheduler map to the traffic control profile to specify the queue scheduling characteristics.

NOTE: When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.

Priority Group Scheduling Components

[Table 38 on page 141](#) provides a quick reference to the traffic control profile components you can configure to determine the bandwidth properties of priority groups, and [Table 39 on page 141](#) provides a quick reference to some related scheduling configuration components.

Table 38: Priority Group Scheduler Components

Traffic Control Profile Component	Description
Guaranteed rate	Sets the minimum guaranteed port bandwidth for the priority group. Extra port bandwidth is shared among priority groups in proportion to the guaranteed rate of each priority group on the port.
Shaping rate	Sets the maximum port bandwidth the priority group can consume.
Scheduler map	Maps schedulers to queues (forwarding classes, also called priorities). This determines the portion of the priority group bandwidth that a queue receives.

Table 39: Other Scheduling Components

Other Scheduling Components	Description
Forwarding class	Maps traffic to a queue (priority).
Forwarding class set	Name of a priority group. You map forwarding classes to priority groups. A forwarding class set consists of one or more forwarding classes.

Table 39: Other Scheduling Components (Continued)

Other Scheduling Components	Description
Scheduler	Sets the bandwidth and scheduling priority of individual queues (forwarding classes).

Default Traffic Control Profile

There is no default traffic control profile.

Guaranteed Rate (Minimum Guaranteed Bandwidth)

The guaranteed rate determines the minimum guaranteed bandwidth for each priority group. It also determines how much excess (extra) port bandwidth the priority group can share; each priority group shares extra port bandwidth in proportion to its guaranteed rate. You specify the rate in bits per second as a fixed value such as 3 Mbps or as a percentage of the total port bandwidth.

The minimum transmission bandwidth can exceed the configured rate if additional bandwidth is available from other priority groups on the port. In case of congestion, the configured guaranteed rate is guaranteed for the priority group. This property enables you to ensure that each priority group receives the amount of bandwidth appropriate to its level of service.

NOTE: Configuring the minimum guaranteed bandwidth (transmit rate) for a forwarding class does not work unless you also configure the minimum guaranteed bandwidth (guaranteed rate) for the forwarding class set in the traffic control profile.

Additionally, the sum of the transmit rates of the queues in a forwarding class set should not exceed the guaranteed rate for the forwarding class set. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)

You cannot configure a guaranteed rate for forwarding class sets that include strict-high priority queues.

Sharing Extra Bandwidth

Extra bandwidth is available to priority groups when the priority groups do not use the full amount of available port bandwidth. This extra port bandwidth is shared among the priority groups based on the minimum guaranteed bandwidth of each priority group.

For example, Port A has three priority groups: fc-set-1, fc-set-2, and fc-set-3. Fc-set-1 has a guaranteed rate of 2 Gbps, fc-set-2 has a guaranteed rate of 2 Gbps, and fc-set-3 has a guaranteed rate of 4 Gbps. After servicing the minimum guaranteed bandwidth of these priority groups, the port has an extra 2 Gbps of available bandwidth, and all three priority groups have still have packets to forward. The priority groups receive the extra bandwidth in proportion to their guaranteed rates, so fc-set-1 receives an extra 500 Mbps, fc-set-2 receives an extra 500 Mbps, and fc-set-3 receives an extra 1 Gbps.

Shaping Rate (Maximum Bandwidth)

The shaping rate determines the maximum bandwidth the priority group can consume. You specify the rate in bits per second as a fixed value such as 5 Mbps or as a percentage of the total port bandwidth.

The maximum bandwidth for a priority group depends on the total bandwidth available on the port and how much bandwidth the other priority groups on the port consume.

Scheduler Maps

A scheduler map maps schedulers to queues. When you associate a scheduler map with a traffic control profile, then associate the traffic control profile with an interface and a forwarding class set, the scheduling defined by the scheduler map determines the portion of the priority group resources that each individual queue can use.

You can associate up to four user-defined scheduler maps with traffic control profiles.

RELATED DOCUMENTATION

Understanding Junos CoS Components

Understanding CoS Output Queue Schedulers

Understanding CoS Hierarchical Port Scheduling (ETS)

Understanding CoS Scheduling Behavior and Configuration Considerations

Understanding CoS Scheduling on QFabric System Node Device Fabric (fte) Ports

Understanding Default CoS Scheduling on QFabric System Interconnect Devices (Junos OS Release 13.1 and Later Releases)

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Minimum Guaranteed Output Bandwidth

Example: Configuring Maximum Output Bandwidth

Example: Configuring Queue Schedulers

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

Example: Configuring WRED Drop Profiles

Defining CoS Traffic Control Profiles (Priority Group Scheduling)

A traffic control profile defines the output bandwidth and scheduling characteristics of forwarding class sets (priority groups). The forwarding classes (which are mapped to output queues) contained in a forwarding class set (fc-set) share the bandwidth resources that you configure in the traffic control profile. A scheduler map associates forwarding classes with schedulers to define how the individual forwarding classes that belong to an fc-set share the bandwidth allocated to that fc-set.

The parameters you configure in a traffic control profile define the following characteristics for the fc-set:

- **guaranteed-rate**—Minimum bandwidth, also known as the *committed information rate (CIR)*. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the fc-set can share. Extra port bandwidth is allocated among the fc-sets on a port in proportion to the guaranteed rate of each fc-set.

NOTE: You cannot configure a guaranteed rate for a fc-set that includes strict-high priority queues. If the traffic control profile is for an fc-set that contains strict-high priority queues, do not configure a guaranteed rate.

- **shaping-rate**—Maximum bandwidth, also known as the *peak information rate (PIR)*.
- **scheduler-map**—Bandwidth and scheduling characteristics for the queues, defined by mapping forwarding classes to schedulers. (The queue scheduling characteristics represent amounts or percentages of the fc-set bandwidth, not the amounts or percentages of total link bandwidth.)

NOTE: Because a port can have more than one fc-set, when you assign resources to an fc-set, keep in mind that the total port bandwidth must serve all of the queues associated with that port.

To configure a traffic control profile using the CLI:

1. Name the traffic control profile and define the minimum guaranteed bandwidth for the fc-set:

```
[edit class-of-service ]
user@switch# set traffic-control-profiles traffic-control-profile-name guaranteed-rate (rate
| percent percentage)
```

2. Define the maximum bandwidth for the fc-set:

```
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
user@switch# set shaping-rate (rate | percent percentage)
```

3. Attach a scheduler map to the traffic control profile:

```
[edit class-of-service traffic-control-profiles ]
user@switch# set scheduler-map scheduler-map-name
```

RELATED DOCUMENTATION

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

Example: Configuring Minimum Guaranteed Output Bandwidth

Example: Configuring Maximum Output Bandwidth

Defining CoS Queue Schedulers

Understanding CoS Traffic Control Profiles

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

IN THIS SECTION

● [Requirements | 146](#)

● [Overview | 147](#)

A traffic control profile defines the output bandwidth and scheduling characteristics of forwarding class sets (priority groups). The forwarding classes (queues) mapped to a forwarding class set share the bandwidth resources that you configure in the traffic control profile. A scheduler map associates forwarding classes with schedulers to define how the individual queues in a forwarding class set share the bandwidth allocated to that forwarding class set.

Configuring a Traffic Control Profile

Step-by-Step Procedure

This example describes how to configure a traffic control profile named `san-tcp` with a scheduler map named `san-map1` and allocate to it a minimum bandwidth of 4 Gbps and a maximum bandwidth of 8 Gbps:

1. Create the traffic control profile and set the `guaranteed-rate` (minimum guaranteed bandwidth) to 4g:

```
[edit class-of-service]
user@switch# set traffic-control-profiles san-tcp guaranteed-rate 4g
```

2. Set the `shaping-rate` (maximum guaranteed bandwidth) to 8g:

```
[edit class-of-service]
user@switch# set traffic-control-profiles san-tcp shaping-rate 8g
```

3. Associate the scheduler map `san-map1` with the traffic control profile:

```
[edit class-of-service]
user@switch# set traffic-control-profiles san-tcp scheduler-map san-map1
```

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

The parameters you configure in a traffic control profile define the following characteristics for the priority group:

- **guaranteed-rate**—Minimum bandwidth, also known as the *committed information rate (CIR)*. Each fc-set receives a minimum of either the configured amount of absolute bandwidth or the configured percentage of bandwidth. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the fc-set can share. Extra port bandwidth is allocated among the fc-sets on a port in proportion to the guaranteed rate of each fc-set.

NOTE: In order for the *transmit-rate* option (minimum bandwidth for a queue that you set using scheduler configuration) to work properly, you must configure the **guaranteed-rate** for the fc-set. If an fc-set does not have a guaranteed minimum bandwidth, the forwarding classes that belong to the fc-set cannot have a guaranteed minimum bandwidth.

NOTE: Include the preamble bytes and interframe gap bytes as well as the data bytes in your bandwidth calculations.

- **shaping-rate**—Maximum bandwidth, also known as the *peak information rate (PIR)*. Each fc-set receives a maximum of the configured amount of absolute bandwidth or the configured percentage of bandwidth, even if more bandwidth is available.

NOTE: Include the preamble bytes and interframe gap bytes as well as the data bytes in your bandwidth calculations.

- **scheduler-map**—Bandwidth and scheduling characteristics for the queues, defined by mapping forwarding classes to schedulers. (The queue scheduling characteristics represent amounts or percentages of the fc-set bandwidth, not the amounts or percentages of total link bandwidth.)

NOTE: Because a port can have more than one fc-set, when you assign resources to an fc-set, keep in mind that the total port bandwidth must serve all of the queues associated with that port.

For example, if you map three fc-sets to a 10-Gigabit Ethernet port, the queues associated with all three of the fc-sets share the 10-Gbps bandwidth as defined by the traffic control profiles. Therefore, the total combined **guaranteed-rate** value of the three fc-sets should not exceed 10

Gbps. If you configure guaranteed rates whose sum exceeds the port bandwidth, the system sends a syslog message to notify you that the configuration is not valid. However, the system does not perform a commit check. If you commit a configuration in which the sum of the guaranteed rates exceeds the port bandwidth, the hierarchical scheduler behaves unpredictably.

The sum of the forwarding class (queue) transmit rates cannot exceed the total guaranteed-rate of the fc-set to which the forwarding classes belong. If you configure transmit rates whose sum exceeds the fc-set guaranteed rate, the commit check fails and the system rejects the configuration.

If you configure the `guaranteed-rate` of an fc-set as a percentage, configure all of the transmit rates associated with that fc-set as percentages. In this case, if any of the transmit rates are configured as absolute values instead of percentages, the configuration is not valid and the system sends a syslog message.

Verification

IN THIS SECTION

- [Verifying the Traffic Control Profile Configuration | 148](#)

Verifying the Traffic Control Profile Configuration

Purpose

Verify that you created the traffic control profile `san-tcp` with a minimum guaranteed bandwidth of 4 Gbps, a maximum bandwidth of 8 Gbps, and the scheduler map `san-map1`.

Action

List the traffic control profile using the operational mode command `show configuration class-of-service traffic-control-profiles san-tcp`:

```
user@switch> show configuration class-of-service traffic-control-profiles san-tcp
scheduler-map san-map1;
shaping-rate percent 8g;
guaranteed-rate 4g;
```

RELATED DOCUMENTATION

<i>Example: Configuring CoS Hierarchical Port Scheduling (ETS)</i>
<i>Example: Configuring Minimum Guaranteed Output Bandwidth</i>
<i>Example: Configuring Maximum Output Bandwidth</i>
<i>Example: Configuring Queue Schedulers</i>
<i>Defining CoS Traffic Control Profiles (Priority Group Scheduling)</i>
<i>Understanding CoS Traffic Control Profiles</i>
<i>Understanding CoS Hierarchical Port Scheduling (ETS)</i>

Understanding CoS Priority Group and Queue Guaranteed Minimum Bandwidth

IN THIS SECTION

- [Guaranteeing Bandwidth Using Hierarchical Scheduling | 149](#)
- [Priority Group Guaranteed Rate \(Guaranteed Minimum Bandwidth\) | 151](#)
- [Queue Transmit Rate \(Guaranteed Minimum Bandwidth\) | 151](#)

You can set a guaranteed minimum bandwidth for individual forwarding classes (queues) and for groups of forwarding classes called *forwarding class sets* (priority groups). Setting a minimum guaranteed bandwidth ensures that priority groups and queues receive the bandwidth required to support the expected traffic.

Guaranteeing Bandwidth Using Hierarchical Scheduling

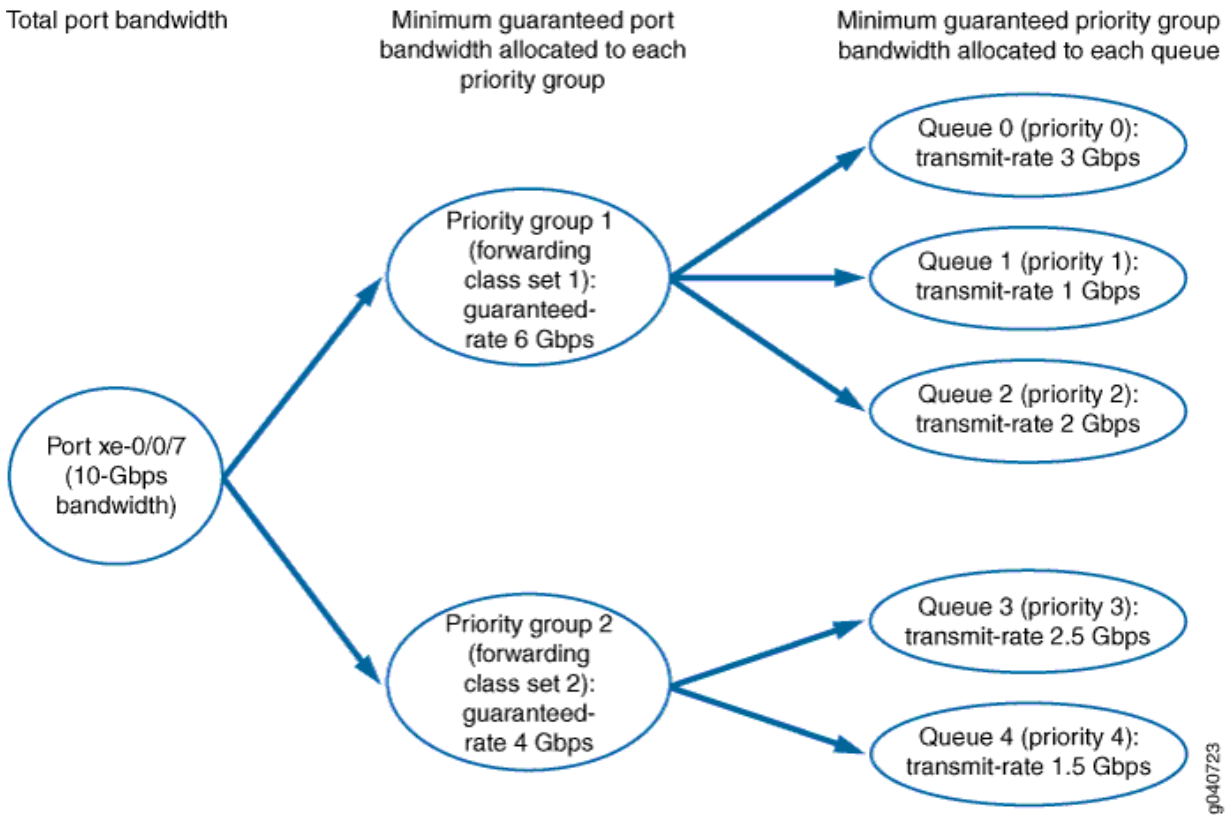
The *guaranteed-rate* value for the priority group (configured in a traffic control profile) defines the minimum amount of bandwidth allocated to a forwarding class set on a port, whereas the *transmit-rate* value of the queue (configured in a scheduler) defines the minimum amount of bandwidth allocated to a particular queue in a priority group. The queue bandwidth is a portion of the priority group bandwidth.

NOTE: You cannot configure a minimum guaranteed bandwidth (transmit rate) for a forwarding class that is mapped to a strict-high priority queue, and you cannot configure a minimum

guaranteed bandwidth (guaranteed rate) for a priority group that includes strict-high priority queues.

Figure 5 on page 150 shows how the total port bandwidth is allocated to priority groups (forwarding class sets) based on the guaranteed rate of each priority group. It also shows how the guaranteed bandwidth of each priority group is allocated to the queues in the priority group based on the transmit rate of each queue.

Figure 5: Allocating Guaranteed Bandwidth Using Hierarchical Scheduling



The sum of the priority group guaranteed rates cannot exceed the total port bandwidth. If you configure guaranteed rates whose sum exceeds the port bandwidth, the system sends a syslog message to notify you that the configuration is not valid. However, the system does not perform a commit check. If you commit a configuration in which the sum of the guaranteed rates exceeds the port bandwidth, the hierarchical scheduler behaves unpredictably.

The sum of the queue transmit rates cannot exceed the total guaranteed rate of the priority group to which the queues belong. If you configure transmit rates whose sum exceeds the priority group guaranteed rate, the commit check fails and the system rejects the configuration.

NOTE: You must set both the priority group `guaranteed-rate` value and the queue `transmit-rate` value in order to configure the minimum bandwidth for individual queues. If you set the `transmit-rate` value but do not set the `guaranteed-rate` value, the configuration fails.

You can set the `guaranteed-rate` value for a priority group without setting the `transmit-rate` value for individual queues in the priority group. However, queues that do not have a configured `transmit-rate` value can become starved for bandwidth if other higher-priority queues need the priority group's bandwidth. To avoid starving a queue, it is a good practice to configure a `transmit-rate` value for most queues.

If you configure the guaranteed rate of a priority group as a percentage, configure all of the transmit rates associated with that priority group as percentages. In this case, if any of the transmit rates are configured as absolute values instead of percentages, the configuration is not valid and the system sends a syslog message.

Priority Group Guaranteed Rate (Guaranteed Minimum Bandwidth)

Setting a priority group (forwarding class set) `guaranteed-rate` enables you to reserve a portion of the port bandwidth for the forwarding classes (queues) in that forwarding class set. The minimum bandwidth (`guaranteed-rate`) that you configure for a priority group sets the minimum bandwidth available to all of the forwarding classes in the forwarding class set.

The combined `guaranteed-rate` value of all of the forwarding class sets associated with an interface cannot exceed the amount of bandwidth available on that interface.

You configure the priority group `guaranteed-rate` in the traffic control profile. You cannot apply a traffic control profile that has a guaranteed rate to a priority group that includes a strict-high priority queue.

Queue Transmit Rate (Guaranteed Minimum Bandwidth)

Setting a queue (forwarding class) `transmit-rate` enables you to reserve a portion of the priority group bandwidth for the individual queue. For example, a queue that handles Fibre Channel over Ethernet (FCoE) traffic might require a minimum rate of 4 Gbps to ensure the *class of service* that storage area network (SAN) traffic requires.

The priority group `guaranteed-rate` sets the aggregate minimum amount of bandwidth available to the queues that belong to the priority group. The cumulative total minimum bandwidth the queues consume cannot exceed the minimum bandwidth allocated to the priority group to which they belong. (The combined transmit rates of the queues in a priority group cannot exceed the priority group's guaranteed rate.)

You must configure the `guaranteed-rate` value of the priority group in order to set a `transmit-rate` value for individual queues that belong to the priority group. The reason is that if there is no guaranteed bandwidth for a priority group, there is no way to guarantee bandwidth for queues in that priority group.

You configure the queue `transmit-rate` in the scheduler configuration. You cannot configure a `transmit-rate` for a strict-high priority queue.

RELATED DOCUMENTATION

Understanding CoS Output Queue Schedulers

Understanding CoS Traffic Control Profiles

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Queue Schedulers

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

Defining CoS Queue Schedulers

Defining CoS Traffic Control Profiles (Priority Group Scheduling)

Example: Configuring Minimum Guaranteed Output Bandwidth

IN THIS SECTION

- [Requirements | 154](#)
- [Overview | 154](#)
- [Verification | 156](#)

Scheduling the minimum guaranteed output bandwidth for a queue (forwarding class) requires configuring both tiers of the two-tier hierarchical scheduler. One tier is scheduling the resources for the individual queue. The other tier is scheduling the resources for the priority group (forwarding class set) to which the queue belongs. You set a minimum guaranteed bandwidth to ensure that priority groups and queues receive the bandwidth required to support the expected traffic.

Configuring Guaranteed Minimum Bandwidth

CLI Quick Configuration

To quickly configure the minimum guaranteed bandwidth for a priority group and a queue, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

```
[edit class-of-service]
set schedulers be-sched transmit-rate 2g
set traffic-control-profiles be-tcp guaranteed-rate 4g
set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
set traffic-control-profiles be-tcp scheduler-map be-map
set forwarding-class-sets be-pg class best-effort
set interfaces xe-0/0/7 forwarding-class-set be-pg output-traffic-control-profile be-tcp
```

Step-by-Step Procedure

To configure the minimum guaranteed bandwidth hierarchical scheduling for a queue and a priority group:

1. Configure the minimum guaranteed queue bandwidth of 2 Gbps for scheduler be-sched:

```
[edit class-of-service schedulers]
user@switch# set be-sched transmit-rate 2g
```

2. Configure the minimum guaranteed priority group bandwidth of 4 Gbps for traffic control profile be-tcp:

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp guaranteed-rate 4g
```

3. Associate the scheduler be-sched with the best-effort queue in the scheduler map be-map:

```
[edit class-of-service scheduler-maps]
user@switch# set be-map forwarding-class best-effort scheduler be-sched
```

4. Associate the scheduler map with the traffic control profile:

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp scheduler-map be-map
```

5. Assign the best-effort queue to the priority group be-pg:

```
[edit class-of-service forwarding-class-sets]
user@switch# set be-pg class best-effort
```

6. Apply the configuration to interface xe-0/0/7:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/7 forwarding-class-set be-pg output-traffic-control-profile be-tcp
```

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series or Junos OS Release 14.1X53-D20 or later for the OCX Series

Overview

The priority group minimum guaranteed bandwidth defines the minimum total amount of bandwidth available for all of the queues in the priority group to meet their minimum bandwidth requirements.

The `transmit-rate` setting in the scheduler configuration determines the minimum guaranteed bandwidth for an individual queue. The transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share. Extra priority group bandwidth is allocated among the queues in the priority group in proportion to the transmit rate of each queue.

The `guaranteed-rate` setting in the traffic control profile configuration determines the minimum guaranteed bandwidth for a priority group. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the priority group can share. Extra port bandwidth is allocated among the priority groups on a port in proportion to the guaranteed rate of each priority group.

NOTE: You must configure both the `transmit-rate` value for the queue and the `guaranteed-rate` value for the priority group to set a valid minimum bandwidth guarantee for a queue. (If the priority group does not have a guaranteed minimum bandwidth, there is no guaranteed bandwidth pool from which the queue can take its guaranteed minimum bandwidth.)

The sum of the queue transmit rates in a priority group should not exceed the guaranteed rate for the priority group. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)

NOTE: When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.

NOTE: You cannot configure minimum guaranteed bandwidth on strict-high priority queues or on a priority group that contains strict-high priority queues.

This example describes how to:

- Configure a transmit rate (minimum guaranteed queue bandwidth) of 2 Gbps for queues in a scheduler named `be-sched`.
- Configure a guaranteed rate (minimum guaranteed priority group bandwidth) of 4 Gbps for a priority group in a traffic control profile named `be-tcp`.
- Assign the scheduler to a queue named `best-effort` by using a scheduler map named `be-map`.
- Associate the scheduler map `be-map` with the traffic control profile `be-tcp`.
- Assign the queue `best-effort` to a priority group named `be-pg`.
- Assign the priority group and the minimum guaranteed bandwidth scheduling to the egress interface `xe-0/0/7`.

Table 40 on page 156 shows the configuration components for this example:

Table 40: Components of the Minimum Guaranteed Output Bandwidth Configuration Example

Component	Settings
Hardware	QFX3500 switch
Minimum guaranteed queue bandwidth	Transmit rate: 2g
Minimum guaranteed priority group bandwidth	Guaranteed rate: 4g
Scheduler	be-sched
Scheduler map	be-map
Traffic control profile	be-tcp
Forwarding class set (priority group)	be-pg
Queue (forwarding class)	best-effort
Egress interface	xe-0/0/7

Verification

IN THIS SECTION

- [Verifying the Minimum Guaranteed Queue Bandwidth | 157](#)
- [Verifying the Priority Group Minimum Guaranteed Bandwidth and Scheduler Map Association | 157](#)
- [Verifying the Scheduler Map Configuration | 158](#)
- [Verifying Queue \(Forwarding Class\) Membership in the Priority Group | 158](#)
- [Verifying the Egress Interface Configuration | 158](#)

To verify the minimum guaranteed output bandwidth configuration, perform these tasks:

Verifying the Minimum Guaranteed Queue Bandwidth

Purpose

Verify that you configured the minimum guaranteed queue bandwidth as 2g in the scheduler be-sched.

Action

Display the minimum guaranteed bandwidth in the be-sched scheduler configuration using the operational mode command `show configuration class-of-service schedulers be-sched transmit-rate`:

```
user@switch> show configuration class-of-service schedulers be-sched transmit-rate
2g;
```

Verifying the Priority Group Minimum Guaranteed Bandwidth and Scheduler Map Association

Purpose

Verify that the minimum guaranteed priority group bandwidth is 4g and the attached scheduler map is be-map in the traffic control profile be-tcp.

Action

Display the minimum guaranteed bandwidth in the be-tcp traffic control profile configuration using the operational mode command `show configuration class-of-service traffic-control-profiles be-tcp guaranteed-rate`:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp guaranteed-rate
4g;
```

Display the scheduler map in the be-tcp traffic control profile configuration using the operational mode command `show configuration class-of-service traffic-control-profiles be-tcp scheduler-map`:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp scheduler-map
scheduler-map be-map;
```

Verifying the Scheduler Map Configuration

Purpose

Verify that the scheduler map `be-map` maps the forwarding class `best-effort` to the scheduler `be-sched`.

Action

Display the `be-map` scheduler map configuration using the operational mode command `show configuration class-of-service schedulers maps be-map`:

```
user@switch> show configuration class-of-service scheduler-maps be-map
forwarding-class best-effort scheduler be-sched;
```

Verifying Queue (Forwarding Class) Membership in the Priority Group

Purpose

Verify that the forwarding class set `be-pg` includes the forwarding class `best-effort`.

Action

Display the `be-pg` forwarding class set configuration using the operational mode command `show configuration class-of-service forwarding-class-sets be-pg`:

```
user@switch> show configuration class-of-service forwarding-class-sets be-pg
class best-effort;
```

Verifying the Egress Interface Configuration

Purpose

Verify that the forwarding class set `be-pg` and the traffic control profile `be-tcp` are attached to egress interface `xe-0/0/7`.

Action

Display the egress interface using the operational mode command `show configuration class-of-service interfaces xe-0/0/7`:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7
forwarding-class-set {
    be-pg {
        output-traffic-control-profile be-tcp;
    }
}
```

RELATED DOCUMENTATION

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Queue Schedulers

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

Example: Configuring Queue Scheduling Priority

Example: Configuring Forwarding Class Sets

Understanding CoS Traffic Control Profiles

Understanding CoS Hierarchical Port Scheduling (ETS)

Understanding CoS Priority Group Shaping and Queue Shaping (Maximum Bandwidth)

IN THIS SECTION

- [Priority Group Shaping | 160](#)
- [Queue Shaping | 160](#)
- [Shaping Maximum Bandwidth Using Hierarchical Scheduling | 161](#)

If the amount of traffic on an interface exceeds the maximum bandwidth available on the interface, it leads to congestion. You can use priority group (forwarding class set) shaping and queue (forwarding class) shaping to manage traffic and avoid congestion.

Configuring a maximum bandwidth sets the most bandwidth a priority group or a queue can use after all of the priority group and queue minimum bandwidth requirements are met, even if more bandwidth is available.

Priority Group Shaping

Priority group shaping enables you to shape the aggregate traffic of a forwarding class set on a port to a maximum rate that is less than the line or port rate. The maximum bandwidth (*shaping-rate*) that you configure for a priority group sets the maximum bandwidth available to all of the forwarding classes (queues) in the forwarding class set.

If a port has more than one priority group and the combined *shaping-rate* value of the priority groups is greater than the amount of port bandwidth available, the bandwidth is shared proportionally among the priority groups.

You configure the priority group *shaping-rate* in the traffic control profile.

Queue Shaping

Queue shaping throttles the rate at which queues transmit packets. For example, using queue shaping, you can rate-limit a strict-high priority queue so that the strict-priority queue does not lock out (or starve) low-priority queues.

NOTE: We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

Similarly, for any queue, you can configure queue shaping (*shaping-rate*) to set the maximum bandwidth for a particular queue.

The *shaping-rate* value of the priority group sets the aggregate maximum amount of bandwidth available to the queues that belong to the priority group. On a port, the cumulative total bandwidth the queues consume cannot exceed the maximum bandwidth of the priority group to which they belong.

If a priority group has more than one queue, and the combined *shaping-rate* of the queues is greater than the amount of bandwidth available to the priority group, the bandwidth is shared proportionally among the queues.


You configure the queue shaping-rate in the scheduler configuration, and you set the shaping-rate for priority groups in the traffic control profile configuration.

Shaping Maximum Bandwidth Using Hierarchical Scheduling

Priority group shaping defines the maximum bandwidth allocated to a forwarding class set on a port, whereas queue shaping defines a limit on maximum bandwidth usage per queue. The queue bandwidth is a portion of the priority group bandwidth.

Figure 6 on page 161 shows how the port bandwidth is allocated to priority groups (forwarding class sets) based on the shaping rate of each priority group, and how the bandwidth of each priority group is allocated to the queues in the priority group based on the shaping rate of each queue.

Figure 6: Setting Maximum Bandwidth Using Hierarchical Scheduling



RELATED DOCUMENTATION

Understanding CoS Output Queue Schedulers

Understanding CoS Traffic Control Profiles

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Queue Schedulers

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

Defining CoS Queue Schedulers

Defining CoS Traffic Control Profiles (Priority Group Scheduling)

Example: Configuring Maximum Output Bandwidth

IN THIS SECTION

- [Requirements | 164](#)
- [Overview | 164](#)
- [Verification | 165](#)

Scheduling the maximum output bandwidth for a queue (forwarding class) requires configuring both tiers of the hierarchical scheduler. One tier is scheduling the resources for the individual queue. The other tier is scheduling the resources for the priority group (forwarding class set) to which the queue belongs. You can use priority group and queue shaping to prevent traffic from using more bandwidth than you want the traffic to receive.

Configuring Maximum Bandwidth

CLI Quick Configuration

To quickly configure the maximum bandwidth for a priority group and a queue, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

```
[edit class-of-service]
set schedulers be-sched shaping-rate percent 4g
set traffic-control-profiles be-tcp shaping-rate 6g
set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
set traffic-control-profiles be-tcp scheduler-map be-map
```

```
set forwarding-class-sets be-pg class best-effort
set interfaces xe-0/0/7 forwarding-class-set be-pg output-traffic-control-profile be-tcp
```

Step-by-Step Procedure

To configure the maximum bandwidth hierarchical scheduling for a queue and a priority group:

1. Configure the maximum queue bandwidth of 4 Gbps for scheduler be-sched:

```
[edit class-of-service schedulers]
user@switch# set be-sched shaping-rate 4g
```

2. Configure the maximum priority group bandwidth of 6 Gbps for traffic control profile be-tcp:

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp shaping-rate 6g
```

3. Associate the scheduler be-sched with the best-effort queue in the scheduler map be-map:

```
[edit class-of-service scheduler-maps]
user@switch# set be-map forwarding-class best-effort scheduler be-sched
```

4. Associate the scheduler map with the traffic control profile:

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp scheduler-map be-map
```

5. Assign the best-effort queue to the priority group be-pg:

```
[edit class-of-service forwarding-class-sets]
user@switch# set be-pg class best-effort
```

6. Apply the configuration to interface xe-0/0/7:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/7 forwarding-class-set be-pg output-traffic-control-profile be-tcp
```

Requirements

This example uses the following hardware and software components:

- One switch (this example was tested on a Juniper Networks QFX3500 Switch)
- Junos OS Release 11.1 or later for the QFX Series or Junos OS Release 14.1X53-D20 or later for the OCX Series

Overview

The priority group maximum bandwidth defines the maximum total amount of bandwidth available for all of the queues in the priority group.

The `shaping-rate` setting in the scheduler configuration determines the maximum bandwidth for an individual queue.

The `shaping-rate` setting in the traffic control profile configuration determines the maximum bandwidth for a priority group.

NOTE: When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.

NOTE: When you set the maximum bandwidth (`shaping-rate`) for a queue or for a priority group at 100 Kbps or less, the traffic shaping behavior is accurate only within +/- 20 percent of the configured `shaping-rate` value.

This example describes how to:

- Configure a maximum rate of 4 Gbps for queues in a scheduler named `be-sched`.
- Configure a maximum rate of 6 Gbps for a priority group in a traffic control profile named `be-tcp`.

- Assign the scheduler to a queue named best-effort by using a scheduler map named be-map.
- Associate the scheduler map be-map with the traffic control profile be-tcp.
- Assign the queue best-effort to a priority group named be-pg.
- Assign the priority group and the bandwidth scheduling to the interface xe-0/0/7.

Table 41 on page 165 shows the configuration components for this example:

Table 41: Components of the Maximum Output Bandwidth Configuration Example

Component	Settings
Hardware	QFX3500 switch
Maximum queue bandwidth	Shaping rate: 4g
Maximum priority group bandwidth	Shaping rate: 6g
Scheduler	be-sched
Scheduler map	be-map
Traffic control profile	be-tcp
Forwarding class set (priority group)	be-pg
Queue (forwarding class)	best-effort
Egress interface	xe-0/0/7

Verification

IN THIS SECTION

Verifying the Maximum Queue Bandwidth | 166

- [Verifying the Priority Group Maximum Bandwidth and Scheduler Map Association | 166](#)
- [Verifying the Scheduler Map Configuration | 167](#)
- [Verifying Queue \(Forwarding Class\) Membership in the Priority Group | 167](#)
- [Verifying the Egress Interface Configuration | 168](#)

To verify the maximum output bandwidth configuration, perform these tasks:

Verifying the Maximum Queue Bandwidth

Purpose

Verify that you configured the maximum queue bandwidth as 4g in the scheduler be-sched.

Action

List the maximum bandwidth in the be-sched scheduler configuration using the operational mode command `show configuration class-of-service schedulers be-sched shaping-rate`:

```
user@switch> show configuration class-of-service schedulers be-sched shaping-rate
4g;
```

Verifying the Priority Group Maximum Bandwidth and Scheduler Map Association

Purpose

Verify that the maximum priority group bandwidth is 6g and the attached scheduler map is be-map in the traffic control profile be-tcp.

Action

List the maximum bandwidth in the be-tcp traffic control profile configuration using the operational mode command `show configuration class-of-service traffic-control-profiles be-tcp shaping-rate`:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp shaping-rate
6g;
```

List the scheduler map in the be-tcp traffic control profile configuration using the operational mode command `show configuration class-of-service traffic-control-profiles be-tcp scheduler-map`:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp scheduler-map
scheduler-map be-map;
```

Verifying the Scheduler Map Configuration

Purpose

Verify that the scheduler map `be-map` maps the forwarding class `best-effort` to the scheduler `be-sched`.

Action

List the `be-map` scheduler map configuration using the operational mode command `show configuration class-of-service schedulers maps be-map`:

```
user@switch> show configuration class-of-service scheduler-maps be-map
forwarding-class best-effort scheduler be-sched;
```

Verifying Queue (Forwarding Class) Membership in the Priority Group

Purpose

Verify that the forwarding class set `be-pg` includes the forwarding class `best-effort`.

Action

List the `be-pg` forwarding class set configuration using the operational mode command `show configuration class-of-service forwarding-class-sets be-pg`:

```
user@switch> show configuration class-of-service forwarding-class-sets be-pg
class best-effort;
```


Verifying the Egress Interface Configuration

Purpose

Verify that the forwarding class set be-pg and the traffic control profile be-tcp are attached to egress interface xe-0/0/7.

Action

List the egress interface using the operational mode command `show configuration class-of-service interfaces xe-0/0/7`:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7
forwarding-class-set {
    be-pg {
        output-traffic-control-profile be-tcp;
    }
}
```

RELATED DOCUMENTATION

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Queue Schedulers

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

Example: Configuring Forwarding Class Sets

Understanding CoS Traffic Control Profiles

Understanding CoS Hierarchical Port Scheduling (ETS)

Understanding CoS Explicit Congestion Notification

IN THIS SECTION

- [How ECN Works | 169](#)
- [WRED Drop Profile Control of ECN Thresholds | 175](#)

Explicit congestion notification (ECN) enables end-to-end congestion notification between two endpoints on TCP/IP based networks. The two endpoints are an ECN-enabled sender and an ECN-enabled receiver. ECN must be enabled on both endpoints. However, in the case of an unsupported peer, a NFX device that supports ECN bootstraps the incoming packets from the unsupported peer and marks the packets to signal network congestion when it occurs.

ECN notifies networks about congestion with the goal of reducing packet loss and delay by making the sending device decrease the transmission rate until the congestion clears, without dropping packets. RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*, defines ECN.

ECN is disabled by default. Normally, you enable ECN only on queues that handle best-effort traffic because other traffic types use different methods of congestion notification—lossless traffic uses priority-based flow control (PFC) and strict-high priority traffic receives all of the port bandwidth it requires up to the point of a configured maximum rate.

You enable ECN on individual output queues (as represented by forwarding classes) by enabling ECN in the queue scheduler configuration, mapping the scheduler to forwarding classes (queues), and then applying the scheduler to interfaces.

NOTE: For ECN to work on a queue, you must also apply a weighted random early detection (WRED) packet drop profile to the queue.

How ECN Works

Without ECN, devices respond to network congestion by dropping TCP/IP packets. Dropped packets signal the network that congestion is occurring. Devices on the IP network respond to TCP packet drops by reducing the packet transmission rate to allow the congestion to clear. However, the packet drop method of congestion notification and management has some disadvantages. For example, packets are dropped and must be retransmitted. Also, bursty traffic can cause the network to reduce the transmission rate too much, resulting in inefficient bandwidth utilization.

Instead of dropping packets to signal network congestion, ECN marks packets to signal network congestion, without dropping the packets. For ECN to work, all of the devices in the path between two ECN-enabled endpoints must have ECN enabled. ECN is negotiated during the establishment of the TCP connection between the endpoints.

ECN-enabled devices determine the queue congestion state based on the WRED packet drop profile configuration applied to the queue, so each ECN-enabled queue must also have a WRED drop profile. If a queue fills to the level at which the WRED drop profile has a packet drop probability greater than zero (0), the device marks the packet as experiencing congestion. Whether or not a device marks a packet as experiencing congestion is the same probability as the drop probability of the queue at that fill level.

ECN communicates whether or not congestion is experienced by marking the two least-significant bits in the differentiated services (DiffServ) field in the IP header. The most significant six bits in the DiffServ field contain the Differentiated Services Code Point (DSCP) bits. The state of the two ECN bits signals whether or not the packet is an ECN-capable packet and whether or not congestion has been experienced.

ECN-capable senders mark packets as ECN-capable. If a sender is not ECN-capable, it marks packets as not ECN-capable. If an ECN-capable packet experiences congestion at the egress queue of a device, then the device marks the packet as experiencing congestion. When the packet reaches the ECN-capable receiver (destination endpoint), the receiver echoes the congestion indicator to the sender (source endpoint) by sending a packet marked to indicate congestion.

After receiving the congestion indicator from the receiver, the source endpoint reduces the transmission rate to relieve the congestion. This is similar to the result of TCP congestion notification and management, but instead of dropping the packet to signal network congestion, ECN marks the packet and the receiver echoes the congestion notification to the sender. Because the packet is not dropped, the packet does not need to be retransmitted.

ECN Bits in the DiffServ Field

The two ECN bits in the DiffServ field provide four codes that determine if a packet is marked as an ECN-capable transport (ECT) packet, meaning that both endpoints of the transport protocol are ECN-capable, and if there is congestion experienced (CE), as shown in [Table 42 on page 170](#):

Table 42: ECN Bit Codes

ECN Bits (Code)	Meaning
00	Non-ECT—Packet is marked as not ECN-capable
01	ECT(1)—Endpoints of the transport protocol are ECN-capable
10	ECT(0)—Endpoints of the transport protocol are ECN-capable
11	CE—Congestion experienced

Codes 01 and 10 have the same meaning: the sending and receiving endpoints of the transport protocol are ECN-capable. There is no difference between these codes.

End-to-End ECN Behavior

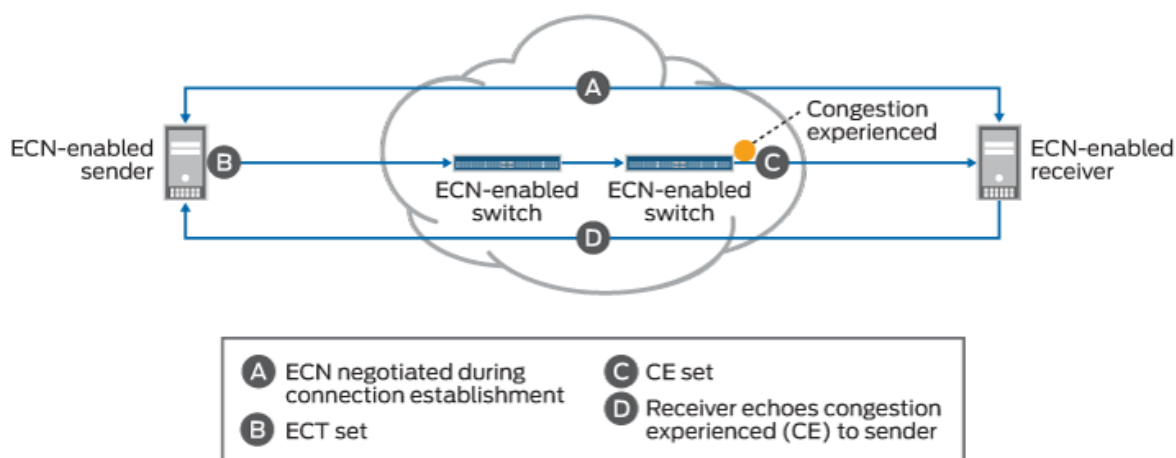
After the sending and receiving endpoints negotiate ECN, the sending endpoint marks packets as ECN-capable by setting the DiffServ ECN field to ECT(1) (01) or ECT(0) (10).

When a packet traverses a device and experiences congestion at an output queue that uses the WRED packet drop mechanism, the device marks the packet as experiencing congestion by setting the DiffServ ECN field to CE (11). Instead of dropping the packet (as with TCP congestion notification), the device forwards the packet.

NOTE: At the egress queue, the WRED algorithm determines whether or not a packet is drop eligible based on the queue fill level (how full the queue is). If a packet is drop eligible and marked as ECN-capable, the packet can be marked CE and forwarded. If a packet is drop eligible and is not marked as ECN-capable, it is dropped. See ["WRED Drop Profile Control of ECN Thresholds"](#) on page 175 for more information about the WRED algorithm.

When the packet reaches the receiver endpoint, the CE mark tells the receiver that there is network congestion. The receiver then sends (echoes) a message to the sender that indicates there is congestion on the network. The sender acknowledges the congestion notification message and reduces its transmission rate. [Figure 7 on page 171](#) summarizes how ECN works to mitigate network congestion:

Figure 7: Explicit Congestion Notification



End-to-end ECN behavior includes:

1. The ECN-capable sender and receiver negotiate ECN capability during the establishment of their connection.

NOTE: If the client is not ECN capable, then the NFX device negotiates ECN on behalf of client during the connection establishment. The NFX device sets the ECE and CWR bits in the TCP header of the SYN packet.

2. After successful negotiation of ECN capability, the ECN-capable sender sends IP packets with the ECT field set to the receiver.
3. If the WRED algorithm on a device egress queue determines that the queue is experiencing congestion and the packet is drop eligible, the device can mark the packet as “congestion experienced” (CE) to indicate to the receiver that there is congestion on the network. If the packet has already been marked CE (congestion has already been experienced at the egress of another device), then the device forwards the packet with CE marked.

If there is no congestion at the device egress queue, then the device forwards the packet and does not change the ECT-enabled marking of the ECN bits, so the packet is still marked as ECN-capable but not as experiencing congestion.

4. The receiver receives a packet marked CE to indicate that congestion was experienced along the congestion path.
5. The receiver echoes (sends) a packet back to the sender with the ECE bit (bit 9) marked in the flag field of the TCP header. The ECE bit is the ECN echo flag bit, which notifies the sender that there is congestion on the network.
6. The sender reduces the data transmission rate and sends a packet to the receiver with the CWR bit (bit 8) marked in the flag field of the TCP header. The CWR bit is the congestion window reduced flag bit, which acknowledges to the receiver that the congestion experienced notification was received.
7. When the receiver receives the CWR flag, the receiver stops setting the ECE bit in replies to the sender.

[Table 43 on page 173](#) summarizes the behavior of traffic on ECN-enabled queues.

Table 43: Traffic Behavior on ECN-Enabled Queues

Incoming IP Packet Marking of ECN Bits	ECN Configuration on the Output Queue	Action if WRED Algorithm Determines Packet is Drop Eligible	Outgoing Packet Marking of ECN Bits	Log Format
Non-ECT (00) SYN	WRED enabled—both scenarios where threshold is crossed and within the threshold limit	Bootstrap to provide ECN support	Set ECE and CWR in TCP header and ECT in IP header	ECT-BIT: 00 WRED-MET: true
Non-ECT (00) Data	WRED enabled	Do not drop. Mark ECN bit to 01/10.	Packet marked ECT 01/10	Not applicable
Non-ECT (00) Data	WRED enabled—threshold met	Do not drop. Mark ECN bit 11.	Packet marked ECT (CE)	ECT-BIT: 00 WRED-MET: true
Non-ECT (00)	WRED disabled	No change	No change	Not applicable
ECT (10 or 01)	WRED enabled	No change	No change	Not applicable
ECT (10 or 01)	WRED enabled—threshold met	Do no drop. Mark ECN bit to 11 and drop according to drop profile.	Packet marked ECT (CE)	ECT-BIT: 10 WRED-MET: true
ECT(10 or 01)	WRED disabled	No change	No change	Not applicable
ECT(11)	WRED enabled	Do not drop. As packet is already marked with CE, send the packet without any change	Packet marked ECT (11) to indicate congestion	ECT-BIT: 11 WRED-MET: false
ECT (11)	WRED disabled	Drop packet	Drop packet	Not applicable

Table 43: Traffic Behavior on ECN-Enabled Queues (Continued)

Incoming IP Packet Marking of ECN Bits	ECN Configuration on the Output Queue	Action if WRED Algorithm Determines Packet is Drop Eligible	Outgoing Packet Marking of ECN Bits	Log Format
ECT (11)	WRED enabled—threshold met	Do not drop. Packet is already marked as experiencing congestion, forward the packet without changing the ECN marking.	Packet marked ECT (11) to indicate congestion	ECT-BIT: 11 WRED-MET: true

When an output queue is not experiencing congestion as defined by the WRED drop profile mapped to the queue, all packets are forwarded, and no packets are dropped.

ECN Compared to PFC and Ethernet PAUSE

ECN is an end-to-end network congestion notification mechanism for IP traffic. Priority-based flow control (PFC) (IEEE 802.1Qbb) and Ethernet PAUSE (IEEE 802.3X) are different types of congestion management mechanisms.

ECN requires that an output queue must also have an associated WRED packet drop profile. Output queues used for traffic on which PFC is enabled should not have an associated WRED drop profile. Interfaces on which Ethernet PAUSE is enabled should not have an associated WRED drop profile.

PFC is a peer-to-peer flow control mechanism to support lossless traffic. PFC enables connected peer devices to pause flow transmission during periods of congestion. PFC enables you to pause traffic on a specified type of flow on a link instead of on all traffic on a link. For example, you can (and should) enable PFC on lossless traffic classes such as the `fcoe` forwarding class. Ethernet PAUSE is also a peer-to-peer flow control mechanism, but instead of pausing only specified traffic flows, Ethernet PAUSE pauses all traffic on a physical link.

With PFC and Ethernet PAUSE, the sending and receiving endpoints of a flow do not communicate congestion information to each other across the intermediate devices. Instead, PFC controls flows between two PFC-enabled peer devices that support data center bridging (DCB) standards. PFC works by sending a pause message to the connected peer when the flow output queue becomes congested. Ethernet PAUSE simply pauses all traffic on a link during periods of congestion and does not require DCB.

WRED Drop Profile Control of ECN Thresholds

You apply WRED drop profiles to forwarding classes (which are mapped to output queues) to control how the device marks ECN-capable packets. A scheduler map associates a drop profile with a scheduler and a forwarding class, and then you apply the scheduler map to interfaces to implement the scheduling properties for the forwarding class on those interfaces.

Drop profiles define queue fill level (the percentage of queue fullness) and drop probability (the percentage probability that a packet is dropped) pairs. When a queue fills to a specified level, traffic that matches the drop profile has the drop probability paired with that fill level. When you configure a drop profile, you configure pairs of fill levels and drop probabilities to control how packets drop at different levels of queue fullness.

The first fill level and drop probability pair is the drop start point. Until the queue reaches the first fill level, packets are not dropped. When the queue reaches the first fill level, packets that exceed the fill level have a probability of being dropped that equals the drop probability paired with the fill level.

The last fill level and drop probability pair is the drop end point. When the queue reaches the last fill level, all packets are dropped unless they are configured for ECN.

NOTE: Lossless queues (forwarding class configured with the `no-loss` packet drop attribute) and strict-high priority queues do not use drop profiles. Lossless queues use PFC to control the flow of traffic.

The drop profile configuration affects ECN packets as follows:

- Drop start point—ECN-capable packets might be marked as congestion experienced (CE).
- Drop end point—ECN-capable packets are always marked CE.

As a queue fills from the drop start point to the drop end point, the probability that an ECN packet is marked CE is the same as the probability that a non-ECN packet is dropped if you apply the drop profile to best-effort traffic. As the queue fills, the probability of an ECN packet being marked CE increases, just as the probability of a non-ECN packet being dropped increases when you apply the drop profile to best-effort traffic.

At the drop end point, all ECN packets are marked CE, but the ECN packets are not dropped. When the queue fill level exceeds the drop end point, all ECN packets are marked CE. ECN packets (and all other packets) are tail-dropped if the queue fills completely.

To configure a WRED packet drop profile and apply it to an output queue (using hierarchical scheduling on devices that support ETS):

1. Configure a drop profile using the statement `set class-of-service drop-profiles profile-name interpolate fill-level drop-start-point fill-level drop-end-point drop-probability 0 drop-probability percentage`.

2. Map the drop profile to a queue scheduler using the statement `set class-of-service schedulers scheduler-name drop-profile-map loss-priority (low | medium-high | high) protocol any drop-profile profile-name`. The name of the drop-profile is the name of the WRED profile configured in Step 1.
3. Map the scheduler, which Step 2 associates with the drop profile, to the output queue using the statement `set class-of-service scheduler-maps map-name forwarding-class forwarding-class-name scheduler scheduler-name`. The forwarding class identifies the output queue. Forwarding classes are mapped to output queues by default, and can be remapped to different queues by explicit user configuration. The scheduler name is the scheduler configured in Step 2.
4. Associate the scheduler map with a traffic control profile using the statement `set class-of-service traffic-control-profiles tcp-name scheduler-map map-name`. The scheduler map name is the name configured in Step 3.
5. Associate the traffic control profile with an interface using the statement `set class-of-service interface interface-name forwarding-class-set forwarding-class-set-name output-traffic-control-profile tcp-name`. The output traffic control profile name is the name of the traffic control profile configured in Step 4.

The interface uses the scheduler map in the traffic control profile to apply the drop profile (and other attributes, including the enable ECN attribute) to the output queue (forwarding class) on that interface. Because you can use different traffic control profiles to map different schedulers to different interfaces, the same queue number on different interfaces can handle traffic in different ways.

You can configure a WRED packet drop profile and apply it to an output queue on devices that support port scheduling (ETS hierarchical scheduling is either not supported or not used). To configure a WRED packet drop profile and apply it to an output queue on devices that support port scheduling (ETS hierarchical scheduling is either not supported or not used):

1. Configure a drop profile using the statement `set class-of-service drop-profiles profile-name interpolate fill-level level1 level2 ... level32 drop-probability probability1 probability2 ... probability32`. You can specify as few as two fill level/drop probability pairs or as many as 32 pairs.
2. Map the drop profile to a queue scheduler using the statement `set class-of-service schedulers scheduler-name drop-profile-map loss-priority (low | medium-high | high) drop-profile profile-name`. The name of the drop-profile is the name of the WRED profile configured in Step 1.
3. Map the scheduler, which Step 2 associates with the drop profile, to the output queue using the statement `set class-of-service scheduler-maps map-name forwarding-class forwarding-class-name scheduler scheduler-name`. The forwarding class identifies the output queue. Forwarding classes are mapped to output queues by default, and can be remapped to different queues by explicit user configuration. The scheduler name is the scheduler configured in Step 2.
4. Associate the scheduler map with an interface using the statement `set class-of-service interfaces interface-name scheduler-map scheduler-map-name`.

The interface uses the scheduler map to apply the drop profile (and other attributes) to the output queue mapped to the forwarding class on that interface. Because you can use different scheduler maps on different interfaces, the same queue number on different interfaces can handle traffic in different ways.

Support, Limitations, and Notes

If the WRED algorithm that is mapped to a queue does not find a packet drop eligible, then the ECN configuration and ECN bits marking does not matter. The packet transport behavior is the same as when ECN is not enabled.

ECN is disabled by default. Normally, you enable ECN only on queues that handle best-effort traffic, and you do not enable ECN on queues that handle lossless traffic or strict-high priority traffic.

ECN supports the following:

- IPv4 and IPv6 packets
- Untagged, single-tagged, and double-tagged packets
- The outer IP header of IP tunneled packets (but not the inner IP header)

ECN does not support the following:

- IP packets with MPLS encapsulation
- The inner IP header of IP tunneled packets (however, ECN works on the outer IP header)
- Multicast, broadcast, and destination lookup fail (DLF) traffic
- Non-IP traffic

4

PART

Configuration Statements and Operational Commands

Configuration Statements (Basic Concepts) | 179

Configuration Statements (Classifiers and Rewrite Rules) | 188

Configuration Statements (Scheduling) | 235

Operational Commands (Basic Concepts) | 274

Operational Commands (Classifiers and Rewrite Rules) | 376

Operational Commands (Scheduling) | 457

CHAPTER 4

Configuration Statements (Basic Concepts)

IN THIS CHAPTER

- [class-of-service | 179](#)
- [traceoptions \(Class of Service\) | 184](#)

class-of-service

IN THIS SECTION

- [Syntax | 179](#)
- [Hierarchy Level | 183](#)
- [Description | 183](#)
- [Default | 183](#)
- [Required Privilege Level | 183](#)
- [Release Information | 184](#)

Syntax

```
class-of-service {
  classifiers {
    (dscp | dscp-ipv6 | ieee-802.1 | exp) classifier-name {
      import (classifier-name | default);
      forwarding-class
    }
  }
  class-name {
    loss-
    level {
      code-points [ aliases ] [ bit-patterns ];
    }
  }
}
```



```

set-name {
    class class-name;
}
forwarding-classes {
    class class-name {
        pfc-priority pfc-priority;
        no-loss;
        queue-num queue-number <no-loss>;
    }
}
host-outbound-traffic{
    forwarding-class class-name;
    dscp-code-point code-point;
}
interfaces interface-name {
    classifiers {
        (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
    }
    congestion-notification-profile profile-name;
    forwarding-class lossless-forwarding-class-name;
    forwarding-class-set forwarding-class-set-name {
        output-traffic-control-profile profile-name;
    }
    rewrite-value {
        input {
            ieee-802.1{
                code-point code-point-bits;
            }
        }
    }
    scheduler-map scheduler-map-name;
    unit logical-unit-number {
        classifiers {
            (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
        }
        forwarding-class
class-name;
        rewrite-rules {
            (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
        }
    }
}
multi-destination {

```

```

        classifiers {
            (dscp | ieee-802.1) classifier-name;
        }
    }
    rewrite-rules {
        (dscp | dscp-ipv6 | ieee-802.1 | exp) classifier-name {
            import (rewrite-name | default);
            forwarding-class
        }
    }
    class-name {
        loss-priority
    }
    priority code-point (alias | bits);
}

scheduler-map-forwarding-class-sets {
    fabric-scheduler-map-name {
        forwarding-class-set fabric-forwarding-class-set-name scheduler
scheduler-name;
    }
}

scheduler-maps {
    map-name {
        forwarding-class class-name scheduler
scheduler-name;
    }
}

schedulers {
    scheduler-name {
        buffer-size (percent percentage | remainder);
        drop-profile-map loss-
priority (low | medium-high | high)
protocol protocol drop-
profile drop-profile-name;
        excess-rate percent percentage;
        explicit-congestion-notification;
        priority priority;
        shaping-rate (rate | percent percentage);
        transmit-rate (percent percentage) <exact>;
    }
}

shared-buffer {
    egress {
        percent percent;
    }
}

```

```

        buffer-partition (lossless | lossy | multicast) {
            percent percent
        }
    }
    ingress {
        percent percent;
        buffer-partition (lossless | lossless-headroom | lossy) {
            percent percent
        }
    }
}
system-defaults {
    classifiers exp classifier-name;
}
traffic-control-profiles profile-name {
    guaranteed-rate(rate| percent percentage);
    scheduler-map map-name;
                                shaping-rate (rate| percent percentage);
}
}

```

Hierarchy Level

[edit]

Description

Configure class-of-service parameters on the switch.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Default

If you do not configure any CoS features, the default CoS settings are used.

Required Privilege Level

interface—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

NOTE: Not all switches support all portions of the class of service hierarchy. For example, some switches use the same classifiers for unicast and multidestination traffic, and those switches do not support the multi-destination classifier hierarchy, and some switches do not support shared buffer configuration, and those switches do not support the shared-buffer hierarchy.

NOTE: OCX Series switches do not support MPLS exp classifiers and rewrite rules (including MPLS system defaults), and they do not support congestion notification profiles.

RELATED DOCUMENTATION

Assigning CoS Components to Interfaces

Overview of Junos OS CoS

traceoptions (Class of Service)

IN THIS SECTION

- [Syntax | 185](#)
- [Hierarchy Level | 185](#)
- [Description | 185](#)
- [Default | 185](#)
- [Options | 185](#)
- [Required Privilege Level | 187](#)
- [Release Information | 187](#)

Syntax

```
traceoptions {
  file filename <size size> <files number>
  <world-readable | no-world-readable>;
  flag flag <flag-modifier>;
  no-remote-trace
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Set class-of-service (CoS) tracing options.

NOTE: The traceoptions statement is not supported on the QFabric system.

Default

Traceoptions is disabled.

Options

- | | |
|--------------------------------|--|
| file
<i>filename</i> | Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Traceoption output files are located in the <code>/var/log/</code> directory. |
| files
<i>number</i> | <p>(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed <i>trace-file.0</i>. The traceoption output continues in a second trace file named <i>trace-file.1</i>. When <i>trace-file.1</i> reaches its maximum size, output continues in a third file named <i>trace-file.2</i>, and so on. When the maximum number of trace files is reached, the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also specify a maximum file size with the size option.</p> |

- **Range:** 2 through 1000 files
- **Default:** 1 trace file

flag Tracing operation to perform. To specify more than one tracing operation, include multiple *flag* statements:

- *all*—Trace all operations.
- *asynch*—Trace asynchronous configuration processing.
- *chassis-scheduler*—Trace chassis stream scheduler processing.
- *cos-adjustment*—Trace CoS rate adjustments.
- *dynamic*—Trace dynamic CoS functions.
- *hardware-database*—Trace the chassis hardware database related processing.
- *init*—Trace initialization events.
- *performance-monitor*—Trace performance monitor countors.
- *process*—Trace configuration processing.
- *restart*—Trace restart processing.
- *route-socket*—Trace route-socket events.
- *show*—Trace show command servicing.
- *snmp*—Trace SNMP-related processing.
- *util*—Trace utilities.

The following are the global tracing options:

- *all*—Perform all tracing operations
- *parse*—Trace parser processing.

<i>no-remote-trace</i>	(Optional) Disable remote tracing.
<i>no-world-readable</i>	(Optional) Prevent any user from reading the log file.
<i>size size</i>	(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed <i>trace-file.0</i> . Incoming tracefile data is logged in the now empty <i>trace-file</i> . When <i>trace-file</i> again

reaches its maximum size, *trace-file.0* is renamed *trace-file.1* and *trace-file* is renamed *trace-file.0*. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the `files` option.

- **Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB
- **Range:** 10 KB through the maximum file size of 4 GB (maximum is lower if 4 GB is not supported on your system)
- **Default:** 1 MB

`world-readable` (Optional) Allow any user to read the log file.

Required Privilege Level

`routing`—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

Configuration Statements (Classifiers and Rewrite Rules)

IN THIS CHAPTER

- [class \(Forwarding Classes\) | 189](#)
- [class \(Forwarding Class Sets\) | 191](#)
- [classifiers | 193](#)
- [code-point \(Rewrite Rules\) | 196](#)
- [code-point-aliases | 197](#)
- [code-points \(CoS\) | 199](#)
- [dscp | 200](#)
- [dscp-ipv6 | 204](#)
- [exp | 207](#)
- [forwarding-class | 209](#)
- [forwarding-classes | 213](#)
- [ieee-802.1 | 218](#)
- [import | 221](#)
- [interfaces \(Class of Service\) | 223](#)
- [loss-priority \(Classifiers\) | 225](#)
- [loss-priority \(Rewrite Rules\) | 227](#)
- [queue-num | 228](#)
- [rewrite-rules | 231](#)
- [unit | 233](#)

class (Forwarding Classes)

IN THIS SECTION

- [Syntax for QFX Series and OCX Series Routers | 189](#)
- [Syntax for M120, M320, MX Series routers, T Series routers, and EX Series switches | 189](#)
- [Hierarchy Level | 190](#)
- [Description | 190](#)
- [Options | 191](#)
- [Required Privilege Level | 191](#)
- [Release Information | 191](#)

Syntax for QFX Series and OCX Series Routers

```
class {  
    class-name {  
        pfc-priority pfc-priority;  
        queue-num queue-number <no-loss>;  
    }  
}
```

Syntax for M120, M320, MX Series routers, T Series routers, and EX Series switches

```
class {  
    class-name {  
        queue-num queue-number ;  
        priority (high | low) ;  
    }  
}
```

Hierarchy Level

```
[edit class-of-service forwarding-classes]
```

Description

On M120 , M320, MX Series routers, T Series routers and EX Series switches only, specify the output transmission queue to which to map all input from an associated forwarding class.

This statement enables you to configure up to 16 forwarding classes with multiple forwarding classes mapped to single queues. If you want to configure up to eight forwarding classes with one-to-one mapping to output queues, use the *queue* statement instead of the class statement at the [edit class-of-service forwarding-classes] hierarchy level.

Map one or more forwarding classes to a single queue. Also, when configuring DSCP-based PFC, map a forwarding class to a PFC priority value to use in pause frames when traffic on a DSCP value becomes congested (see *Configuring DSCP-based PFC for Layer 3 Untagged Traffic* for details).

You can map unicast forwarding classes to a unicast queue (0 through 7) and multidestination forwarding classes to a multicast queue (8 through 11). The queue to which you map a forwarding class determines if the forwarding class is a unicast or multicast forwarding class.

NOTE: On systems that do not use the ELS CLI, if you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless *fcoe* and *no-loss* forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

NOTE: On systems that do not use the ELS CLI, if you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the *fcoe* and *no-loss* forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the *no-loss* option. If you do not specify the *no-loss* option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if you explicitly configure the *fcoe* forwarding class and you do not include the *no-loss* option, the *fcoe* forwarding class is lossy, not lossless.

Options

class-name—Name of the forwarding class.

queue-number—Output queue number.

- **Range:** 0 through 7. Some T Series router PICs are restricted to 0 through 3.

The remaining statements are explained separately. See [CLI Explorer](#) for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 8.1.

No-loss option introduced in Junos OS Release 12.3 for the QFX Series.

PFC-priority statement introduced in Junos OS Release 17.4R1 for the QFX Series.

RELATED DOCUMENTATION

Example: Configuring Forwarding Classes

Understanding CoS Forwarding Classes

[Understanding CoS Forwarding Classes](#)

Configuring a Custom Forwarding Class for Each Queue

queue (Global Queues)

class (Forwarding Class Sets)

IN THIS SECTION

● [Syntax](#) | 192

● [Hierarchy Level](#) | 192

- [Description | 192](#)
- [Options | 192](#)
- [Required Privilege Level | 192](#)
- [Release Information | 192](#)

Syntax

```
class class-name;
```

Hierarchy Level

```
[edit class-of-service forwarding-class-sets forwarding-class-set-name]
```

Description

Group forwarding classes into sets of forwarding classes (priority groups). You can group some or all of the configured forwarding classes into up to three unicast forwarding class sets and one multidestination forwarding class set.

Options

class-name—Name of the forwarding class.

Required Privilege Level

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

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Forwarding Class Sets

Understanding CoS Forwarding Class Sets (Priority Groups)

classifiers

IN THIS SECTION

- [Syntax \(BA Classifiers\) | 193](#)
- [Syntax \(Multidestination BA Classifiers\) | 194](#)
- [Syntax \(Interface Classifier Association: DSCP, DSCP IPv6, IEEE\) | 194](#)
- [Syntax \(Global EXP Interface Classifier Association with Interfaces\) | 194](#)
- [Hierarchy Level \(BA Classifiers\) | 194](#)
- [Hierarchy Level \(Multidestination BA Classifiers\) | 194](#)
- [Hierarchy Level \(Interface Classifier Association: DSCP, DSCP IPv6, IEEE\) | 195](#)
- [Hierarchy Level \(Global EXP Classifier\) | 195](#)
- [Description | 195](#)
- [Options | 195](#)
- [Required Privilege Level | 195](#)
- [Release Information | 195](#)

Syntax (BA Classifiers)

```
classifiers {
  (dscp | dscp-ipv6 | ieee-802.1) classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
      loss-priority level {
        code-points [ aliases ] [ bit-patterns ];
      }
    }
  }
}
```

```
    }
}
```

Syntax (Multidestination BA Classifiers)

Multidestination BA Classifiers

```
classifiers {
    (dscp | ieee-802.1) classifier-name;
}
```

Syntax (Interface Classifier Association: DSCP, DSCP IPv6, IEEE)

Interface Classifier Association

```
classifiers {
    (dscp | dscp-ipv6 | ieee-802.1) (default | classifier-name);
}
```

Syntax (Global EXP Interface Classifier Association with Interfaces)

```
classifiers {
    exp classifier-name;
}
```

Hierarchy Level (BA Classifiers)

```
[edit class-of-service],
```

Hierarchy Level (Multidestination BA Classifiers)

```
[edit class-of-service multi-destination],
```

Hierarchy Level (Interface Classifier Association: DSCP, DSCP IPv6, IEEE)

```
[edit class-of-service          interfaces          interface-name
  unit          logical-unit-number]
```

Hierarchy Level (Global EXP Classifier)

```
[edit class-of-service system-defaults]
```

Description

Define a unicast or multidestination CoS behavior aggregate (BA) classifier.

NOTE: OCX Series switches do not support MPLS, so they do not support EXP classifier configuration.

Options

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.

Release Information

Statement introduced in Junos OS Release 11.1.

EXP statement introduced in Junos OS Release 12.3 for the QFX Series.

RELATED DOCUMENTATION

Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)

Configuring a Global MPLS EXP Classifier

Example: Configuring Unicast Classifiers

Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers

Understanding CoS Classifiers

[Understanding CoS Classifiers](#)

Understanding CoS MPLS EXP Classifiers and Rewrite Rules

code-point (Rewrite Rules)

IN THIS SECTION

- [Syntax | 196](#)
- [Hierarchy Level | 196](#)
- [Description | 196](#)
- [Options | 197](#)
- [Required Privilege Level | 197](#)
- [Release Information | 197](#)

Syntax

```
code-point [ alias ] [ bit-pattern ];
```

Hierarchy Level

```
[edit class-of-service rewrite-rules (dscp | ieee-802.1) forwarding-class class-name loss-priority  
level]
```

Description

Configure a code-point alias or bit set to apply to a forwarding class for a rewrite rule.

NOTE: OCX Series switches do not support MPLS, so they do not support EXP rewrite rules.

Options

alias—Name of the alias.

bit-pattern—Value of the code-point bits, in decimal form.

Required Privilege Level

interfaces—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Defining CoS Rewrite Rules

Understanding CoS Classifiers

[Understanding CoS Classifiers](#)

code-point-aliases

IN THIS SECTION

- [Syntax | 198](#)
- [Hierarchy Level | 198](#)
- [Description | 198](#)
- [Options | 198](#)
- [Required Privilege Level | 198](#)

Syntax

```
code-point-aliases {  
    (dscp| dscp-ipv6 | ieee-802.1 | exp) {  
        alias-name bits;  
    }  
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Define an alias for a CoS marker. You can use the alias instead of the bit pattern when you specify the code point during configuration.

NOTE: OCX Series switches do not support MPLS, so they do not support EXP code-point aliases.

Options

(dscp | dscp-ipv6 | ieee-802.1 | exp)—Set the type of classifier for which you are creating an alias.

alias-name—Name of the code-point alias.

bits—Value of the code-point bits, in decimal form.

Required Privilege Level

interface—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Defining CoS Code-Point Aliases

Understanding CoS Code-Point Aliases

code-points (CoS)

IN THIS SECTION

- [Syntax | 199](#)
- [Hierarchy Level | 199](#)
- [Description | 200](#)
- [Options | 200](#)
- [Required Privilege Level | 200](#)
- [Release Information | 200](#)

Syntax

```
code-points ([ aliases ] | [ bit-patterns ]);
```

Hierarchy Level

```
[edit class-of-service classifiers type classifier-name forwarding-class class-name loss-  
priority level]
```


Description

Specify one or more DSCP code-point aliases or bit sets to apply to a forwarding class..

NOTE: OCX Series switches do not support MPLS, and therefore, do not support EXP code points or code point aliases.

Options

aliases—Name of the DSCP alias.

bit-patterns—Value of the code-point bits, in six-bit binary form.

Required Privilege Level

interface—To view this statement in the configuration.

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

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Understanding Interfaces](#)

Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic

Example: Configuring Behavior Aggregate Classifiers

[Example: Configuring BA Classifiers on Transparent Mode Security Devices](#)

dscp

IN THIS SECTION

● [Syntax \(Classifier\)](#) | 201

- Syntax (Code-Point Alias) | 201
- Syntax (Multidestination Classifier) | 202
- Syntax (Interface Classifier Association) | 202
- Syntax (Rewrite Rule) | 202
- Hierarchy Level (Classifier) | 202
- Hierarchy Level (Code-Point Aliases) | 202
- Hierarchy Level (Multidestination Classifier) | 203
- Hierarchy Level (Interface Classifier Association) | 203
- Hierarchy Level (Rewrite Rule) | 203
- Description | 203
- Options | 203
- Required Privilege Level | 203
- Release Information | 203

Syntax (Classifier)

```
dscp classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-points [ aliases ] [ bit-patterns ];
        }
    }
}
```

Syntax (Code-Point Alias)

Code-Point Alias Configuration

```
dscp alias-name bit-pattern;
```

Syntax (Multidestination Classifier)

Multidestination Classifier Configuration

```
dscp classifier-name;
```

Syntax (Interface Classifier Association)

Interface Classifier Association

```
dscp (classifier-name | default);
```

Syntax (Rewrite Rule)

Rewrite Rule Configuration

```
dscp rewrite-name {
    import (rewrite-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-point [ aliases ] [ bit-patterns ];
        }
    }
}
```

Hierarchy Level (Classifier)

```
[edit class-of-service classifiers],
```

Hierarchy Level (Code-Point Aliases)

```
[edit class-of-service code-point-aliases],
```

Hierarchy Level (Multidestination Classifier)

```
[edit class-of-service multi-destination classifiers],
```

Hierarchy Level (Interface Classifier Association)

```
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers],  
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules],
```

Hierarchy Level (Rewrite Rule)

```
[edit class-of-service                               rewrite-rules]
```

Description

Define the Differentiated Services code point (DSCP) mapping that is applied to the packets.

Options

classifier-name—Name of the classifier.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interfaces—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring Unicast Classifiers

Example: Configuring Classifiers

Defining CoS Code-Point Aliases

Defining CoS Rewrite Rules

Assigning CoS Components to Interfaces

Understanding CoS Classifiers

[Understanding CoS Classifiers](#)

Understanding CoS Rewrite Rules

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

[Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces](#)

dscp-ipv6

IN THIS SECTION

- [Syntax \(Classifier\) | 204](#)
- [Syntax \(Code-Point Alias\) | 205](#)
- [Syntax \(Interface Classifier Association\) | 205](#)
- [Syntax \(Rewrite Rule\) | 205](#)
- [Hierarchy \(Classifier\) | 205](#)
- [Hierarchy \(Code-Point Alias\) | 205](#)
- [Hierarchy \(Interface Classifier Association\) | 206](#)
- [Hierarchy \(Rewrite Rule\) | 206](#)
- [Description | 206](#)
- [Options | 206](#)
- [Required Privilege Level | 206](#)
- [Release Information | 206](#)

Syntax (Classifier)

```
dscp-ipv6 classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
```

```

        loss-priority level {
            code-points [ aliases ] [ bit-patterns ];
        }
    }
}

```

Syntax (Code-Point Alias)

```
dscp-ipv6 alias-name bit-pattern;
```

Syntax (Interface Classifier Association)

```
dscp-ipv6 (classifier-name | default);
```

Syntax (Rewrite Rule)

```

dscp-ipv6 rewrite-name {
    import (rewrite-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-point [ aliases ] [ bit-patterns ];
        }
    }
}

```

Hierarchy (Classifier)

```
[edit class-of-service classifiers],
```

Hierarchy (Code-Point Alias)

```
[edit class-of-service code-point-aliases],
```

Hierarchy (Interface Classifier Association)

```
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers],
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules],
```

Hierarchy (Rewrite Rule)

```
[edit class-of-service rewrite-rules]
```

Description

Define the Differentiated Services code point (DSCP) IPv6 mapping that is applied to the packets.

NOTE: On switches that use different classifiers for unicast and multidestination (multicast, broadcast, and destination lookup fail) traffic, there is no DSCP IPv6 classifier for multidestination (multicast, broadcast, and destination lookup fail) traffic. Multidestination IPv6 traffic uses the multidestination DSCP classifier.

Options

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.

Release Information

Statement introduced in Junos OS Release 12.2.

RELATED DOCUMENTATION

Defining CoS Code-Point Aliases

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

exp

IN THIS SECTION

- [Syntax | 207](#)
- [Hierarchy Level | 208](#)
- [Description | 208](#)
- [Options | 208](#)
- [Required Privilege Level | 208](#)
- [Release Information | 208](#)

Syntax

```
exp classifier-name {  
    import (classifier-name | default);  
    forwarding-class class-name {  
        loss-priority level {  
            code-points [ aliases ] [ bit-patterns ];  
        }  
    }  
}
```

Rewrite Rule Configuration

```
exp rewrite-name {  
    import (rewrite-name | default);  
    forwarding-class class-name {  
        loss-priority level {  
            code-point [ aliases ] [ bit-patterns ];  
        }  
    }  
}
```


Global Classifier Association with Interfaces

```
exp classifier-name;
```

Hierarchy Level

```
[edit class-of-service classifiers]
[edit class-of-service rewrite-rules]
[edit class-of-service system-defaults classifiers]
```

Description

Define the EXP code point mapping that is applied to MPLS packets. EXP classifiers are not applied to any traffic except MPLS traffic. EXP classifiers are applied only to interfaces that are configured as family mpls (for example, set interfaces xe-0/0/35 unit 0 family mpls.)

There are no default EXP classifiers. You can configure up to 64 EXP classifiers.

On QFX10000 switches, you can configure and apply EXP classifiers to interfaces in the same way that you configure and apply DSCP, DSCP IPv6, and IEEE classifiers to interfaces. Different interfaces can have different EXP classifiers. QFX10000 switches do not support global EXP classifiers.

However, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, the switch uses only one EXP classifier as a global MPLS classifier on all interfaces. You specify the global EXP classifier in the [edit class-of-service system-defaults] hierarchy.

Options

classifier-name—Name of the EXP classifier.

Required Privilege Level

interfaces—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 12.3X50.

RELATED DOCUMENTATION

Configuring Rewrite Rules for MPLS EXP Classifiers

Understanding CoS MPLS EXP Classifiers and Rewrite Rules

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

forwarding-class

IN THIS SECTION

- Classifier | [209](#)
- Classifier (EX Series switches) | [210](#)
- Rewrite Rule | [210](#)
- Scheduler Map | [210](#)
- Interface | [210](#)
- Classifier Hierarchy Level | [210](#)
- Classifier Hierarchy Level (EX Series Switches) | [211](#)
- Rewrite Rule Hierarchy Level | [211](#)
- Rewrite Rule Hierarchy Level (EX Series Switches) | [211](#)
- Scheduler Map Hierarchy Level | [211](#)
- Interface Hierarchy Level | [211](#)
- Host-Outbound Traffic Hierarchy Level (EX Series Switches) | [211](#)
- Description | [211](#)
- Options | [212](#)
- Required Privilege Level | [212](#)
- Release Information | [212](#)

Classifier

```
forwarding-class class-name {
    loss-priority level {
        code-points [ aliases ] [ bit-patterns ];
```

```
    }
}
```

Classifier (EX Series switches)

```
forwarding-class class-name {
    loss-priority level {
        code-points [aliases] [6-bit-patterns];
    }
}
```

Rewrite Rule

```
forwarding-class class-name {
    loss-priority level {
        code-point [ aliases ] [ bit-patterns ];
    }
}
```

Scheduler Map

```
forwarding-class class-name {
    scheduler scheduler-name;
}
```

Interface

```
forwarding-class class-name;
```

Classifier Hierarchy Level

```
[edit class-of-service classifiers (dscp | dscp-ipv6 | ieee-802.1 | exp) classifier-name]
```

Classifier Hierarchy Level (EX Series Switches)

```
[edit class-of-service classifiers (dscp | ieee-802.1 | inet-precedence) classifier-name]
```

Rewrite Rule Hierarchy Level

```
[edit class-of-service rewrite-rules] (dscp | dscp-ipv6 | ieee-802.1) rewrite-name | exp]
```

Rewrite Rule Hierarchy Level (EX Series Switches)

```
[edit class-of-service rewrite-rules] (dscp | ieee-802.1 | inet-precedence) rewrite-rule-name]
```

Scheduler Map Hierarchy Level

```
[edit class-of-service scheduler-maps map-name]
```

Interface Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Host-Outbound Traffic Hierarchy Level (EX Series Switches)

```
[edit class-of-service host-outbound-traffic]
```

Description

Define forwarding class name and option values.

- **Classifiers**—Assign incoming traffic to the specified forwarding class based on the specified code point values and assign that traffic the specified loss priority.

- Rewrite rules—At the egress interface, change (rewrite) the value of the code point bits and the loss priority to specified new values for traffic assigned to the specified forwarding class, before forwarding the traffic to the next hop.
- Scheduler maps—Apply the specified scheduler to the specified forwarding class.
- Interfaces—Assign the specified forwarding class to the interface to use as a fixed classifier (all incoming traffic on the interface is classified into that forwarding class).

NOTE: OCX Series switches do not support MPLS, so they do not support EXP classifiers or rewrite rules.

Options

class-name—Name of the forwarding class.

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.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches

Defining CoS Forwarding Classes (CLI Procedure)

Defining CoS Forwarding Classes (J-Web Procedure)

Understanding CoS Forwarding Classes

forwarding-classes

IN THIS SECTION

- [SRX Series | 213](#)
- [QFX Series and OCX Series | 213](#)
- [EX Series \(Except EX4300\) | 214](#)
- [EX4300 | 214](#)
- [M320, MX Series, T Series, and PTX Series | 214](#)
- [Hierarchy Level | 215](#)
- [Description | 215](#)
- [Options | 216](#)
- [Required Privilege Level | 217](#)
- [Release Information | 217](#)

SRX Series

```
forwarding-classes {
  class class-name {
    priority (high | low);
    queue-num number;
    spu-priority (high | low | medium);
  }
  queue queue-number {
    class class-name {
      priority (high | low);
    }
  }
}
```

QFX Series and OCX Series

```
forwarding-classes {
  class class-name {
```

```

    pfc-priority pfc-priority;
    no-loss;
    queue-num queue-number <no-loss>;
  }
}

```

EX Series (Except EX4300)

```

forwarding-classes {
  class class-name {
    queue-num queue-number;
    priority (high | low);
  }
}

```

EX4300

```

forwarding-classes {
  class class-name ;
    queue-num queue-number;
  }
}

```

M320, MX Series, T Series, and PTX Series

```

forwarding-classes {
  class class-name {
    queue queue-number;
    priority (high | low);
  }
  queue queue-number {
    class class-name {
      priority (high | low) [policing-priority (premium | normal)];
    }
  }
}

```

Hierarchy Level

[edit class-of-service]

Description

Command used to associate forwarding classes with class names and queues with queue numbers.

SRX Series Firewalls

All traffic traversing the SRX Series Firewall is passed to an SPC to have service processing applied. Junos OS provides a configuration option to enable packets with specific Differentiated Services (DiffServ) code points (DSCP) precedence bits to enter a high-priority queue, a medium-priority queue, or a low-priority queue on the SPC. The Services Processing Unit (SPU) draws packets from the highest priority queue first, then from the medium priority queue, and last from the low priority queue. The processing of the queue is weighted-based not strict-priority-based. This feature can reduce overall latency for real-time traffic, such as voice traffic.

Initially, the spu-priority queue options were "high" and "low". Then, these options (depending on the devices) were expanded to "high", "medium-high", "medium-low", and "low". The two middle options ("medium-high" and "medium-low") have now been deprecated (again, depending on the devices) and replaced with "medium". So, the available options for spu-priority queue are "high", "medium", and "low".

We recommend that the high-priority queue be selected for real-time and high-value traffic. The other options would be selected based on user judgement on the value or sensitivity of the traffic.

M320, MX Series, and T Series Routers and EX Series Switches

For M320, MX Series, and T Series routers, and EX Series switches only, you can configure fabric priority queuing by including the priority statement. For Enhanced IQ PICs, you can include the policing-priority option.

NOTE: The priority and policing-priority options are not supported on PTX Series routers.

EX Series Switches

For the EX Series switches, this statement associates the forwarding class with a class name and queue number. It can define the fabric queuing priority as high, medium-high, medium-low, or low.

Map one or more forwarding classes to a single output queue. Also, when configuring DSCP-based priority-based flow control (PFC), map a forwarding class to a PFC priority value to use in pause frames when traffic on a DSCP value becomes congested (see *Configuring DSCP-based PFC for Layer 3 Untagged Traffic* for details).

Switches that use different forwarding classes for unicast and multdestination (multicast, broadcast, and destination lookup fail) traffic support 12 forwarding classes and 12 output queues (0 through 11). You map unicast forwarding classes to a unicast queue (0 through 7) and multdestination forwarding classes to a multdestination queue (8 through 11). The queue to which you map a forwarding class determines if the forwarding class is a unicast or multdestination forwarding class.

Switches that use the same forwarding classes for unicast and multdestination traffic support eight forwarding classes and eight output queues (0 through 7). You map forwarding classes to output queues. All traffic classified into one forwarding class (unicast and multdestination) uses the same output queue.

You cannot configure weighted random early detection (WRED) packet drop on forwarding classes configured with the `no-loss` packet drop attribute. Do not associate a drop profile with lossless forwarding classes.

NOTE: If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the same queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).

NOTE: On switches that do not use the Enhanced Layer 2 Software (ELS) CLI, if you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless `fcoe` and `no-loss` forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

NOTE: On switches that do not use the ELS CLI, if you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the `fcoe` and `no-loss` forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the `no-loss` option. If you do not specify the `no-loss` option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if you explicitly configure the `fcoe` forwarding class and you do not include the `no-loss` option, the `fcoe` forwarding class is lossy, not lossless.

Options

`class class-name` Define the forwarding class name.

queue-num *queue-number* Output queue number to associate with forwarding class.

- **Range:** 0 through 7.

priority Fabric priority value:

high Forwarding class fabric queuing has high priority.

low Forwarding class fabric queuing has low priority.

The default priority is low.

spu-priority SPU priority queue, high, medium, or low. The default spu-priority is low.

NOTE: The spu-priority option is supported only on the SRX5000 line of firewalls.

The remaining statements are explained separately. See [CLI Explorer](#) for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

Release Information

The **policing-priority** option was introduced in Junos OS Release 9.5.

Statement updated in Junos OS Release 11.4.

The **spu-priority** option was introduced in Junos OS Release 11.4R2.

The **no-loss** option was introduced in Junos OS Release 12.3 on QFX Series switches.

Change from two to four queues made in Junos OS Release 12.3X48-D40 and in Junos OS Release 15.1X49-D70.

The **pfc-priority** statement was introduced in Junos OS Release 17.4R1 on QFX Series switches.

The **medium-high** and **medium-low** priorities for **spu-priority** were deprecated and **medium** priority was added in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Configuring a Custom Forwarding Class for Each Queue

Forwarding Classes and Fabric Priority Queues

Configuring Hierarchical Layer 2 Policers on IQE PICs

Classifying Packets by Egress Interface

ieee-802.1

IN THIS SECTION

- [Syntax \(Classifier\) | 218](#)
- [Syntax \(Code-Point Alias\) | 219](#)
- [Syntax \(Multidestination Classifier\) | 219](#)
- [Syntax \(Interface Classifier Association\) | 219](#)
- [Syntax \(Rewrite Rule\) | 219](#)
- [Hierarchy Level \(Classifier\) | 220](#)
- [Hierarchy Level \(Code-Point Alias\) | 220](#)
- [Hierarchy Level \(Multidestination Classifier\) | 220](#)
- [Hierarchy Level \(Interface Classifier Association\) | 220](#)
- [Hierarchy Level \(Rewrite Rule\) | 220](#)
- [Description | 220](#)
- [Options | 220](#)
- [Required Privilege Level | 221](#)
- [Release Information | 221](#)

Syntax (Classifier)

```
ieee-802.1 classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-points [ aliases ] [ bit-patterns ];
```

```

    }
  }
}

```

Syntax (Code-Point Alias)

Code-Point Alias Configuration

```
ieee-802.1 alias-name bit-pattern;
```

Syntax (Multidestination Classifier)

Multidestination Classifier Configuration

```
ieee-802.1 classifier-name;
```

Syntax (Interface Classifier Association)

Interface Classifier Association

```
ieee-802.1 (classifier-name | default);
```

Syntax (Rewrite Rule)

Rewrite Rule Configuration

```

ieee-802.1 rewrite-name {
  import (rewrite-name | default);
  forwarding-class class-name {
    loss-priority level {
      code-point [ aliases ] [ bit-patterns ];
    }
  }
}

```

Hierarchy Level (Classifier)

```
[edit class-of-service classifiers],
```

Hierarchy Level (Code-Point Alias)

```
[edit class-of-service code-point-aliases],
```

Hierarchy Level (Multidestination Classifier)

```
[edit class-of-service multi-destination classifiers ],
```

Hierarchy Level (Interface Classifier Association)

```
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers],  
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules],
```

Hierarchy Level (Rewrite Rule)

```
[edit class-of-service rewrite-rules]
```

Description

Configure an IEEE 802.1 classifier, configure an IEEE 802.1 code-point alias, apply a fixed IEEE 802.1 classifier to an interface, or apply an IEEE-802.1 rewrite rule.

Options

classifier-name—Name of the classifier.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interfaces—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring Unicast Classifiers

Defining CoS Code-Point Aliases

Defining CoS Rewrite Rules

Assigning CoS Components to Interfaces

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

[Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces](#)

Understanding CoS Classifiers

[Understanding CoS Classifiers](#)

Understanding CoS Rewrite Rules

import

IN THIS SECTION

- [Syntax | 222](#)
- [Hierarchy Level | 222](#)
- [Description | 222](#)
- [Options | 222](#)
- [Required Privilege Level | 222](#)
- [Release Information | 222](#)

Syntax

```
import (import | default);
```

Hierarchy Level

```
[edit class-of-service classifiers (dscp | dscp-ipv6 | ieee-802.1 | exp) classifier-name],  
[edit class-of-service rewrite-rules (dscp | dscp-ipv6 | ieee-802.1 | exp) classifier-name]
```

Description

Specify a default or previously defined classifier.

Options

import—Name of the classifier mapping configured at the [edit class-of-service classifiers] hierarchy level.

default—Default classifier mapping.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interfaces—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring Unicast Classifiers

Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)

Defining CoS Rewrite Rules

Understanding CoS Classifiers

Understanding CoS Classifiers

Understanding CoS Classifiers

Understanding CoS Rewrite Rules

interfaces (Class of Service)

IN THIS SECTION

- Syntax | 223
- Hierarchy Level | 224
- Description | 224
- Options | 224
- Required Privilege Level | 225
- Release Information | 225

Syntax

```
interfaces interface-name {
    classifiers {
        (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
    }
    congestion-notification-profile profile-name;
    dedicated-buffer-profile profile-name;
    forwarding-class forwarding-class-name;
    forwarding-class-set forwarding-class-set-name {
        output-traffic-control-profile profile-name;
    }
    rewrite-value {
        input {
            ieee-802.1{
                code-point code-point-bits;
            }
        }
    }
}
```



```

scheduler-map scheduler-map-name
unit logical-unit-number {
    classifiers {
        (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
    }
    forwarding-class class-name;
    rewrite-rules {
        (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
    }
}
}

```

Hierarchy Level

[edit class-of-service]

Description

Configure interface-specific CoS properties for incoming packets.

NOTE: Only switches that support direct port scheduling also support applying a scheduler map directly to an interface. When using enhanced transmission selection (ETS) hierarchical port scheduling, you cannot apply a scheduler map directly to an interface; instead, you associate the scheduler map with a traffic control profile and apply the traffic control profile to the interface.

NOTE: Only switches that support native Fibre Channel interfaces support the `rewrite-value` statement, which enables you to rewrite the IEEE 802.1p code points on native Fibre Channel interfaces.

Options

interface-name—Name of the interface.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interfaces—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Assigning CoS Components to Interfaces

Interfaces Overview for Switches

loss-priority (Classifiers)

IN THIS SECTION

- [Syntax | 225](#)
- [Hierarchy Level | 226](#)
- [Description | 226](#)
- [Options | 226](#)
- [Required Privilege Level | 226](#)
- [Release Information | 226](#)

Syntax

```
loss-priority level {  
    code-points [ aliases ] [ bit-patterns ];  
}
```

Hierarchy Level

```
[edit class-of-service classifiers (dscp | dscp-ipv6 | ieee-802.1) classifier-name forwarding-class
class-name]
```

Description

Configure packet loss priority value for a specific set of code-point aliases and bit patterns.

NOTE: OCX Series switches do not support MPLS, so they do not support EXP classifiers.

Options

level—Can be one of the following:

- *low*—Packet has low loss priority.
- *medium-low*—Packet has medium-low loss priority.
- *medium-high*—Packet has medium-high loss priority.
- *high*—Packet has high loss priority.

The remaining statement is explained separately. See [CLI Explorer](#).

Required Privilege Level

interfaces—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring Unicast Classifiers

[Defining CoS BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\)](#) | 59

loss-priority (Rewrite Rules)

IN THIS SECTION

- Syntax | 227
- Hierarchy Level | 227
- Description | 227
- Options | 228
- Required Privilege Level | 228
- Release Information | 228

Syntax

```
loss-priority level {  
    code-point (alias | bit-pattern);  
}
```

Hierarchy Level

```
[edit class-of-service rewrite-rules (dscp | ieee-802.1) rewrite-name forwarding-class class-name]
```

Description

Specify a loss priority to which to apply a rewrite rule. The rewrite rule sets the code-point aliases and bit patterns for a specific forwarding class and loss priority. Packets that match the forwarding class and loss priority are rewritten with the rewrite code-point alias or bit pattern.

NOTE: OCX Series switches do not support MPLS, so they do not support EXP rewrite rules.

Options

level—Can be one of the following:

- *low*—Packet has low loss priority.
- *medium-low*—Packet has medium-low loss priority.
- *medium-high*—Packet has medium-high loss priority.
- *high*—Packet has high loss priority.

The remaining statement is explained separately. See [CLI Explorer](#).

Required Privilege Level

interfaces—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

[Understanding CoS Rewrite Rules | 112](#)

[Defining CoS Rewrite Rules | 114](#)

queue-num

IN THIS SECTION

- [Syntax | 229](#)
- [Hierarchy Level | 229](#)
- [Description | 229](#)
- [Options | 230](#)
- [Required Privilege Level | 230](#)

Syntax

```
queue-num queue-number <no-loss>;
```

Hierarchy Level

```
[edit class-of-service forwarding-classes class class-name]
```

Description

Map a forwarding class to an output queue number. Optionally, configure the forwarding class as a lossless forwarding class. Each switch provides enough output queues so that you can map forwarding classes to queues on a one-to-one basis, so each forwarding class can have a dedicated output queue.

On switches that use different forwarding classes and output queues for unicast and multdestination (multicast, broadcast, destination lookup fail) traffic, the switch supports 12 forwarding classes and 12 output queues, eight of each for unicast traffic and four of each for multdestination traffic. You can map some or all of the eight unicast forwarding classes to a unicast queue (0 through 7) and some or all of the four multdestination forwarding classes to the a multdestination queue (8 through 11). You cannot map a forwarding class to more than one queue (each forwarding class maps to one and only one queue), but you can map multiple forwarding classes to one queue. The queue to which you map a forwarding class determines if the forwarding class is a unicast or multdestination forwarding class.

On switches that use the same forwarding classes and output queues for unicast and multdestination traffic, the switch supports eight forwarding classes and eight output queues. You can map some or all of the eight of the forwarding classes to queues (0 through 7). You cannot map a forwarding class to more than one queue (each forwarding class maps to one and only one queue), but you can map multiple forwarding classes to one queue.

You cannot configure weighted random early detection (WRED) packet drop on forwarding classes configured with the no-loss packet drop attribute. Do not associate a drop profile with lossless forwarding classes. Instead, use priority-based flow control (PFC) to prevent frame drop on lossless forwarding classes.

NOTE: If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the same queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless). OCX Series switches do not support the no-loss packet drop attribute and do not support lossless forwarding classes. On OCX Series switches, do not configure the no-loss packet drop attribute on forwarding classes, and do not map traffic to the default `fcoe` and `no-loss` forwarding classes (both of these default forwarding classes carry the no-loss packet drop attribute).

NOTE: On systems that do not use the ELS CLI, if you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless `fcoe` and `no-loss` forwarding classes. If you explicitly configure lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

NOTE: On systems that do not use the ELS CLI, if you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the `fcoe` and `no-loss` forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the `no-loss` option. If you do not specify the `no-loss` option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if you explicitly configure the `fcoe` forwarding class and you do not include the `no-loss` option, the `fcoe` forwarding class is lossy, not lossless.

Options

queue-number—(Switches that use different output queues for unicast and multidestination traffic) Number of the CoS unicast queue (0 through 7) or the CoS multidestination queue (8 through 11).

queue-number—(Switches that use the same output queues for unicast and multidestination traffic) Number of the CoS queue (0 through 7).

`no-loss`—Optional packet drop attribute keyword to configure the forwarding class as lossless.

Required Privilege Level

`interfaces`—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

No-loss option introduced in Junos OS Release 12.3 for the QFX Series.

rewrite-rules

IN THIS SECTION

- [Syntax \(Rewrite Rule Configuration\) | 231](#)
- [Syntax \(Rewrite Rule Association with Interface\) | 232](#)
- [Hierarchy Level \(Rewrite Rule Configuration\) | 232](#)
- [Hierarchy Level \(Rewrite Rule Association with Interface\) | 232](#)
- [Description | 232](#)
- [Options | 232](#)
- [Required Privilege Level | 232](#)
- [Release Information | 233](#)

Syntax (Rewrite Rule Configuration)

```
rewrite-rules {
  (dscp | dscp-ipv6 | ieee-802.1 | exp) rewrite-name {
    import (rewrite-name | default);
    forwarding-class class-name {
      loss-priority priority code-point (alias | bits);
    }
  }
}
```


Syntax (Rewrite Rule Association with Interface)

```
rewrite-rules {
    (dscp | dscp-ipv6 | ieee-802.1 | exp) rewrite-name;
}
```

Hierarchy Level (Rewrite Rule Configuration)

```
[edit class-of-service],
```

Hierarchy Level (Rewrite Rule Association with Interface)

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Configure rewrite rules that map traffic to code points when traffic exits the system, and apply the rewrite rules to a specific interface.

MPLS EXP rewrite rules can only be bound to logical interfaces, not to physical interfaces. You can configure up to 64 EXP rewrite rules, but you can use only 16 EXP rewrite rules on switch interfaces at any given time.

NOTE: OCX Series switches do not support MPLS, so they do not support EXP rewrite rules.

Options

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.

Release Information

Statement introduced in Junos OS Release 11.1.

EXP statement introduced in Junos OS Release 12.3 for the QFX Series.

RELATED DOCUMENTATION

Defining CoS Rewrite Rules

Configuring Rewrite Rules for MPLS EXP Classifiers

Understanding CoS Rewrite Rules

Understanding CoS MPLS EXP Classifiers and Rewrite Rules

unit

IN THIS SECTION

- [Syntax | 233](#)
- [Hierarchy Level | 234](#)
- [Description | 234](#)
- [Options | 234](#)
- [Required Privilege Level | 234](#)
- [Release Information | 234](#)

Syntax

```
unit logical-unit-number {
  classifiers {
    (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
  }
  forwarding-class class-name;
  rewrite-rules {
    (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
  }
}
```

```
}
}
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name]
```

Description

Configure a logical interface on the physical device. You must configure a logical interface to use the physical device.

NOTE: OCX Series switches do not support MPLS, so they do not support EXP classifiers and rewrite rules.

Options

logical-unit-number—Number of the logical unit.

- **Range:** 0 through 16,385

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.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

| *Assigning CoS Components to Interfaces*

Configuration Statements (Scheduling)

IN THIS CHAPTER

- [buffer-size | 235](#)
- [excess-rate | 241](#)
- [fill-level | 243](#)
- [forwarding-class | 245](#)
- [guaranteed-rate | 248](#)
- [interpolate | 251](#)
- [output-traffic-control-profile | 252](#)
- [priority \(Schedulers\) | 253](#)
- [scheduler | 256](#)
- [scheduler-map | 257](#)
- [scheduler-maps | 258](#)
- [schedulers | 260](#)
- [shaping-rate | 261](#)
- [traffic-control-profiles | 264](#)
- [transmit-rate | 268](#)

buffer-size

IN THIS SECTION

- [Syntax | 236](#)
- [Hierarchy Level | 236](#)
- [Description | 236](#)
- [Default | 238](#)

- Options | 240
- Required Privilege Level | 241
- Release Information | 241

Syntax

```
buffer-size (percent percent | remainder);
```

```
buffer-size (exact | percent percentage | remainder | temporal);
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

```
[edit class-of-service schedulers scheduler-name]
```

Description

Specify buffer size in a scheduler configuration.

On all switches, you configure the proportion of port buffers allocated to a particular output queue using the following process:

1. Configure a scheduler and set the `buffer-size` option.
2. Use a scheduler map to map the scheduler to the forwarding class that is mapped to the queue to which you want to apply the buffer size.

For example, suppose that you want to change the dedicated buffer allocation for FCoE traffic. FCoE traffic is mapped to the `fcoe` forwarding class, and the `fcoe` forwarding class is mapped to queue 3 (this is the default configuration). To use default FCoE traffic mapping, in the scheduler map configuration, map the scheduler to the `fcoe` forwarding class.

3. If you are using enhanced transmission selection (ETS) hierarchical scheduling, associate the scheduler map with the traffic control profile you want to use on the egress ports that carry FCoE traffic. If you are using direct port scheduling, skip this step.
4. If you are using ETS, associate the traffic control profile that includes the scheduler map with the desired egress ports. For this example, you associate the traffic control profile with the ports that carry FCoE traffic. If you are using port scheduling, associate the scheduler map with the desired egress ports.

Queue 3, which is mapped to the fcoe forwarding class and therefore to the FCoE traffic, receives the dedicated buffer allocation specified in the `buffer-size` statement.

NOTE: The total of all of the explicitly configured buffer size percentages for all of the queues on a port cannot exceed 100 percent.

QFX10000 Switches

On QFX10000 switches, the buffer size is the amount of time in milliseconds of port bandwidth that a queue can use to continue to transmit packets during periods of congestion, before the buffer runs out and packets begin to drop.

The switch can use up to 100 ms total (combined) buffer space for all queues on a port. A buffer-size configured as one percent is equal to 1 ms of buffer usage. A buffer-size of 15 percent (the default value for the best effort and network control queues) is equal to 15 ms of buffer usage.

The total buffer size of the switch is 4 GB. A 40-Gigabit port can use up to 500 MB of buffer space, which is equivalent to 100 ms of port bandwidth on a 40-Gigabit port. A 10-Gigabit port can use up to 125 MB of buffer space, which is equivalent to 100 ms of port bandwidth on a 10-Gigabit port. The total buffer sizes of the eight output queues on a port cannot exceed 100 percent, which is equal to the full 100 ms total buffer available to a port. The maximum amount of buffer space any queue can use is also 100 ms (which equates to a 100 percent buffer-size configuration), but if one queue uses all of the buffer, then no other queue receives buffer space.

There is no minimum buffer allocation, so you can set the buffer-size to zero (0) for a queue. However, we recommend that on queues on which you enable PFC to support lossless transport, you allocate a minimum of 5 ms (a minimum buffer-size of 5 percent). The two default lossless queues, fcoe and no-loss, have buffer-size default values of 35 ms (35 percent).

Queue buffer allocation is dynamic, shared among ports as needed. However, a queue cannot use more than its configured amount of buffer space. For example, if you are using the default CoS configuration, the best-effort queue receives a maximum of 15 ms of buffer space because the default transmit rate for the best-effort queue is 15 percent.

If a switch experiences congestion, queues continue to receive their full buffer allocation until 90 percent of the 4 GB buffer space is consumed. When 90 percent of the buffer space is in use, the amount of buffer space per port, per queue, is reduced in proportion to the configured buffer size for each queue. As the percentage of consumed buffer space rises above 90 percent, the amount of buffer space per port, per queue, continues to be reduced.

On 40-Gigabit ports, because the total buffer is 4 GB and the maximum buffer a port can use is 500 MB, up to seven 40-Gigabit ports can consume their full 100 ms allocation of buffer space. However, if an eighth 40-Gigabit port requires the full 500 MB of buffer space, then the buffer allocations are proportionally reduced because the buffer consumption is above 90 percent.

On 10-Gigabit ports, because the total buffer is 4 GB and the maximum buffer a port can use is 125 MB, up to 28 10-Gigabit ports can consume their full 100 ms allocation of buffer space. However, if a 29th 10-Gigabit port requires the full 125 MB of buffer space, then the buffer allocations are proportionally reduced because the buffer consumption is above 90 percent.

QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems

Set the dedicated buffer size of the egress queue that you bind the scheduler to in the scheduler map configuration. The switch allocates space from the global dedicated buffer pool to ports and queues in a hierarchical manner. The switch allocates an equal number of dedicated buffers to each egress port, so each egress port receives the same amount of dedicated buffer space. The amount of dedicated buffer space per port is not configurable.

However, the `buffer-size` statement allows you to control the way each port allocates its share of dedicated buffers to its queues. For example, if a port only uses two queues to forward traffic, you can configure the port to allocate all of its dedicated buffer space to those two queues and avoid wasting buffer space on queues that are not in use. We recommend that the buffer size should be the same size as the minimum guaranteed transmission rate (the `transmit-rate`).

Default

The default behavior differs on different switches.

QFX10000 Switches

If you do not configure `buffer-size` and you do not explicitly configure a queue scheduler, the default `buffer-size` is the default transmit rate of the queue. If you explicitly configure a queue scheduler, the default buffer allocations are not used. If you explicitly configure a queue scheduler, configure the `buffer-size` for each queue in the scheduler, keeping in mind that the total `buffer-size` of the queues cannot exceed 100 percent (100 ms).

[Table 44 on page 239](#) shows the default queue buffer sizes on QFX10000 switches. The default buffer size is the same as the default transmit rate for each default queue:

Table 44: Default Output Queue Buffer Sizes (QFX10000 Switches)

Queue Number	Forwarding Class	Transmit Rate	Buffer Size
0	best-effort	15%	15%
3	fcoe	35%	35%
4	no-loss	35%	35%
7	network-control	15%	15%

By default, only the queues mapped to the default forwarding classes receive buffer space from the port buffer pool. (Buffers are not wasted on queues that do not carry traffic.)

QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems

The port allocates dedicated buffers to queues that have an explicitly configured scheduler buffer size. If you do not explicitly configure a scheduler buffer size for a queue, the port serves the explicitly configured queues first. Then the port divides the remaining dedicated buffers equally among the queues that have an explicitly attached scheduler *without* an explicitly configured buffer size configuration. (If you configure a scheduler, but you do not configure the buffer size parameter, the default is equivalent to configuring the buffer size with the `remainder` option.)

If you use the default scheduler and scheduler map on a port (no explicit scheduler configuration), then the port allocates its dedicated buffer pool to queues based on the default scheduling. [Table 45 on page 239](#) shows the default queue buffer sizes. The default buffer size is the same as the default transmit rate for each default queue:

Table 45: Default Output Queue Buffer Sizes (QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems)

Queue Number	Forwarding Class	Transmit Rate	Buffer Size
0	best-effort	5%	5%
3	fcoe	35%	35%
4	no-loss	35%	35%

Table 45: Default Output Queue Buffer Sizes (QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems) (Continued)

Queue Number	Forwarding Class	Transmit Rate	Buffer Size
7	network-control	5%	5%
8	mcast	20%	20%

By default, only the queues mapped to the default forwarding classes receive buffer space from the port buffer pool. (Buffers are not wasted on queues that do not carry traffic.)

NOTE: OCX Series switches do not support lossless transport. On OCX Series switches, do not map traffic to the lossless default fcoe and no-loss forwarding classes. OCX Series default DSCP classification does not map traffic to the fcoe and no-loss forwarding classes, so by default, the OCX system does not classify traffic into those forwarding classes. (On other switches, the fcoe and no-loss forwarding classes provide lossless transport for Layer 2 traffic. OCX Series switches do not support lossless Layer 2 transport.) The active forwarding classes (best-effort, network-control, and mcast) share the unused bandwidth assigned to the fcoe and no-loss forwarding classes.

On EX Series switches except EX4300 switches, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 0, 0, 0, 0, and 5 percent, respectively. On EX4300 switches, the default scheduler transmission rate and buffer size for queues 0 through 11 are 75, 0, 0, 5, 0, 0, 0, 0, 15, 0, 0 and 5 percent, respectively, of the total available buffer.

Options

percent <i>percent</i>	Percentage of the port dedicated buffer pool allocated to the queue (or queues) mapped to the scheduler.
remainder	Remaining dedicated buffer pool after the port satisfies the needs of the explicitly configured buffers. The port divides the remaining buffers equally among the queues that are explicitly attached to a scheduler but that do not have an explicit buffer size configuration (or are configured with <code>remainder</code> as the buffer size).
exact	(Except on EX8200 standalone switches and EX8200 Virtual Chassis) Enforce the exact buffer size. When this option is configured, sharing is disabled on the queue, restricting the usage to guaranteed buffers only.

<i>percentage</i>	Buffer size as a percentage of the total buffer.
<i>remainder</i>	Remaining buffer available.
<i>temporal</i>	(EX4200 standalone switches, EX4200 Virtual Chassis, EX4300 standalone switches, EX4300 Virtual Chassis, EX8200 standalone switches, and EX8200 Virtual Chassis only) Buffer size as a temporal value.

Required Privilege Level

interface—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches

Defining CoS Schedulers and Scheduler Maps (CLI Procedure)

Defining CoS Schedulers (J-Web Procedure)

Understanding CoS Schedulers

excess-rate

IN THIS SECTION

- [Syntax | 242](#)
- [Hierarchy Level | 242](#)
- [Description | 242](#)
- [Options | 242](#)
- [Required Privilege Level | 242](#)
- [Release Information | 242](#)

Syntax

```
excess-rate percent percentage;
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles profile-name],  
[edit class-of-service schedulers scheduler-name]
```

Description

Determine the percentage of excess port bandwidth for which a queue (forwarding class) that is not a strict-high priority queue or forwarding class set (priority group) can contend. Excess bandwidth is the extra port bandwidth left after strict-high priority queues and the guaranteed minimum bandwidth requirements of other queues (as determined by each queue's transmit rate) are satisfied. With the exception of strict-high priority queues, the switch allocates extra port bandwidth to queues or to priority groups based on the configured excess rate. If you do not configure an excess rate for a queue, the default excess rate is the same as the transmit rate.

You cannot configure an excess rate on strict-high priority queues. Strict-high priority queues receive extra bandwidth based on an extra bandwidth sharing weight of "1", which is not configurable. However, the switch serves traffic on strict-high priority queues up to the configured transmit rate before it serves any other queues, so by configuring an appropriate transmit rate on a strict-high priority queue, you can guarantee strict-high priority traffic on that queue is treated in the manner you want.

Options

percent *percentage*—Percentage of the excess bandwidth to share.

- **Range:** 0 through 100 percent

Required Privilege Level

interface—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 15.1X53-D10.

RELATED DOCUMENTATION

Defining CoS Queue Schedulers for Port Scheduling

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

Understanding CoS Port Schedulers on QFX Switches

Understanding CoS Traffic Control Profiles

fill-level

IN THIS SECTION

- [QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems | 243](#)
- [QFX10000 Switches | 243](#)
- [Hierarchy Level | 243](#)
- [Description | 244](#)
- [Options | 244](#)
- [Required Privilege Level | 244](#)
- [Release Information | 245](#)

QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems

```
fill-level low-value fill-level high-value;
```

QFX10000 Switches

```
fill-level level1 level2 ... level32;
```

Hierarchy Level

```
[edit class-of-service drop-profiles profile-name interpolate]
```

Description

When configuring weighted random early detection (WRED), map the fullness of a queue to a packet *drop-probability* value. You configure the fill-level and drop-probability statements in related pairs. The pairs of fill level and drop probability values set a probability of dropping packets at a specified queue fullness value.

The first fill level is the packet drop start point. Packets do not drop until the queue fullness reaches the first fill level. The last fill level is the packet drop end point. After the queue exceeds the fullness set by the drop end point, all non-ECN packets are dropped. As the queue fills from the drop start point to the drop end point, the rate of packet drop increases in a curve pattern. The higher the queue fill level, the higher the probability of dropping packets.

On switches that support only two fill level/drop probability pairs, the two pairs are the drop start point and the drop end point. On switches that support up to 32 fill level/drop probability pairs, you can configure intermediate interpolations between the drop start point and the drop end point, which provides greater flexibility in controlling the packet drop curve.

NOTE: Do not configure the last fill level as 100 percent.

Options

low-value (switches that support only two fill level/drop probability pairs)—Fullness of the queue before packets begin to drop, expressed as a percentage. The low value must be less than the high value.

- **Range:** 0 through 100

high-value (switches that support only two fill level/drop probability pairs)—Fullness of the queue before it reaches the maximum drop probability. If the queue fills beyond the fill level high value, all packets drop. The high value must be greater than the low value.

- **Range:** 0 through 100

level1 level2 ... level32 (switches that support 32 fill level/drop probability pairs)—The queue fullness level, expressed as a percentage. Each fill level pairs with a drop probability to define the probability of a packet dropping at a specified queue fullness.

- **Range:** 0 through 100

Required Privilege Level

interface—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

forwarding-class

IN THIS SECTION

- [Classifier | 245](#)
- [Classifier \(EX Series switches\) | 246](#)
- [Rewrite Rule | 246](#)
- [Scheduler Map | 246](#)
- [Interface | 246](#)
- [Classifier Hierarchy Level | 246](#)
- [Classifier Hierarchy Level \(EX Series Switches\) | 246](#)
- [Rewrite Rule Hierarchy Level | 247](#)
- [Rewrite Rule Hierarchy Level \(EX Series Switches\) | 247](#)
- [Scheduler Map Hierarchy Level | 247](#)
- [Interface Hierarchy Level | 247](#)
- [Host-Outbound Traffic Hierarchy Level \(EX Series Switches\) | 247](#)
- [Description | 247](#)
- [Options | 248](#)
- [Required Privilege Level | 248](#)
- [Release Information | 248](#)

Classifier

```
forwarding-class class-name {
  loss-priority level {
    code-points [ aliases ] [ bit-patterns ];
  }
}
```

Classifier (EX Series switches)

```
forwarding-class class-name {
    loss-priority level {
        code-points [aliases] [6-bit-patterns];
    }
}
```

Rewrite Rule

```
forwarding-class class-name {
    loss-priority level {
        code-point [ aliases ] [ bit-patterns ];
    }
}
```

Scheduler Map

```
forwarding-class class-name {
    scheduler scheduler-name;
}
```

Interface

```
forwarding-class class-name;
```

Classifier Hierarchy Level

```
[edit class-of-service classifiers (dscp | dscp-ipv6 | ieee-802.1 | exp) classifier-name]
```

Classifier Hierarchy Level (EX Series Switches)

```
[edit class-of-service classifiers (dscp | ieee-802.1 | inet-precedence) classifier-name]
```

Rewrite Rule Hierarchy Level

```
[edit class-of-service rewrite-rules] (dscp | dscp-ipv6 | ieee-802.1) rewrite-name | exp]
```

Rewrite Rule Hierarchy Level (EX Series Switches)

```
[edit class-of-service rewrite-rules] (dscp | ieee-802.1 | inet-precedence) rewrite-rule-name]
```

Scheduler Map Hierarchy Level

```
[edit class-of-service scheduler-maps map-name]
```

Interface Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Host-Outbound Traffic Hierarchy Level (EX Series Switches)

```
[edit class-of-service host-outbound-traffic]
```

Description

Define forwarding class name and option values.

- **Classifiers**—Assign incoming traffic to the specified forwarding class based on the specified code point values and assign that traffic the specified loss priority.
- **Rewrite rules**—At the egress interface, change (rewrite) the value of the code point bits and the loss priority to specified new values for traffic assigned to the specified forwarding class, before forwarding the traffic to the next hop.
- **Scheduler maps**—Apply the specified scheduler to the specified forwarding class.
- **Interfaces**—Assign the specified forwarding class to the interface to use as a fixed classifier (all incoming traffic on the interface is classified into that forwarding class).

NOTE: OCX Series switches do not support MPLS, so they do not support EXP classifiers or rewrite rules.

Options

class-name—Name of the forwarding class.

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.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches

Defining CoS Forwarding Classes (CLI Procedure)

Defining CoS Forwarding Classes (J-Web Procedure)

Understanding CoS Forwarding Classes

guaranteed-rate

IN THIS SECTION

- [Syntax | 249](#)
- [Hierarchy Level | 249](#)
- [Description | 249](#)
- [Default | 249](#)

- Options | 250
- Required Privilege Level | 250
- Release Information | 250

Syntax

```
guaranteed-rate (rate| percent percentage);
```

Hierarchy Level

```
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
```

Description

Configure a guaranteed minimum rate of transmission for a traffic control profile. The sum of the guaranteed rates of all of the forwarding class sets (priority groups) on a port should not exceed the total port bandwidth. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the priority group (forwarding class set) can share. Extra port bandwidth is allocated among the priority groups on a port in proportion to the guaranteed rate of each priority group.

NOTE: You cannot configure a guaranteed rate for a forwarding class set (priority group) that includes strict-high priority queues. If the traffic control profile is for a forwarding class set that contains strict-high priority queues, do not configure a guaranteed rate.

Default

If you do not specify a guaranteed rate, the guaranteed rate is zero (0) and there is no minimum guaranteed bandwidth.

NOTE: If you do not configure a guaranteed rate for a traffic control profile, the queues that belong to any forwarding class set (priority group) that uses that traffic control profile cannot

have a configured transmit rate. The result is that there is no minimum guaranteed bandwidth for those queues and that those queues can be starved during periods of congestion.

Options

percent percentage—Minimum percentage of transmission capacity allocated to the forwarding class set or logical interface.

- **Range:** 1 through 100 percent

rate—Minimum transmission rate allocated to the forwarding class set or logical interface, 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 10,000,000,000 bps

Required Privilege Level

interfaces—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

Example: Configuring Minimum Guaranteed Output Bandwidth

Understanding CoS Traffic Control Profiles

output-traffic-control-profile

interpolate

IN THIS SECTION

- [QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems | 251](#)
- [QFX10000 Switches | 251](#)
- [Hierarchy Level | 251](#)
- [Description | 251](#)
- [Required Privilege Level | 252](#)
- [Release Information | 252](#)

QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems

```
interpolate {
    fill-level low-value fill-level high-value;
    drop-probability 0 drop-probability high-value;
}
```

QFX10000 Switches

```
interpolate {
    fill-level level1 level2 ... level32 drop-probability percent1 percent2 ... percent32;
}
```

Hierarchy Level

```
[edit class-of-service drop-profiles profile-name]
```

Description

Specify values for interpolating the relationship between queue fill level and drop probability for weighted random early detection (WRED) drop profiles.

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.

Release Information

Statement introduced in Junos OS Release 11.1.

output-traffic-control-profile

IN THIS SECTION

- [Syntax | 252](#)
- [Hierarchy Level | 252](#)
- [Description | 253](#)
- [Options | 253](#)
- [Required Privilege Level | 253](#)
- [Release Information | 253](#)

Syntax

```
output-traffic-control-profile profile-name;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name forwarding-class-set forwarding-class-set-name]
```

Description

Apply an output traffic scheduling and shaping profile to a forwarding class set (priority group).

Options

profile-name—Name of the traffic-control profile to apply to the specified forwarding class set.

Required Privilege Level

interfaces—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

Assigning CoS Components to Interfaces

Understanding CoS Traffic Control Profiles

priority (Schedulers)

IN THIS SECTION

- [Syntax | 254](#)
- [Hierarchy Level | 254](#)
- [Description | 254](#)
- [Options | 254](#)
- [Required Privilege Level | 255](#)
- [Release Information | 255](#)

Syntax

```
priority priority;
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

Description

Specify the packet bandwidth-scheduling priority value.

NOTE: On QFabric systems, the priority statement is valid only for Node device queue scheduling. The priority statement is not allowed for Interconnect device queue scheduling. If you map a scheduler that includes a priority configuration to a fabric forwarding class at the [edit class-of-service scheduler-map-fcset] hierarchy level, the system generates a commit error. (On the Interconnect device, fabric fc-sets are not user-definable. Only the fabric_fcset_strict_high fabric fc-set is configured with high priority, and this configuration cannot be changed.)

Options

priority—It can be one of the following:

- **low**—Scheduler has low priority.
- **medium-low**—Scheduler has medium-low priority. (QFX10000 Series switches only)
- **medium-high**—Scheduler has medium-high priority. (QFX10000 Series switches only)
- **high**—Scheduler has high priority. Assigning high priority to a queue prevents the queue from being underserved. (QFX10000 Series switches only)
- **strict-high**—Scheduler has strict high priority. On QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, you can configure only one queue as a strict-high priority queue. On QFX10000 switches, you can configure as many strict-high priority queues as you want. However, because strict-high priority traffic takes precedence over all other traffic, too much strict-high priority traffic can starve the other output queues.

Strict-high priority allocates the scheduled bandwidth to the packets on the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced.

NOTE: On QFX10000 switches, we strongly recommend that you apply a transmit rate to strict-high priority queues to prevent them from starving other queues. A transmit rate configured on a strict-high priority queue limits the amount of traffic that receives strict-high priority treatment to the amount or percentage set by the transmit rate. The switch treats traffic in excess of the transmit rate as best-effort traffic that receives bandwidth from the leftover (excess) port bandwidth pool. On strict-high priority queues, all traffic that exceeds the transmit rate shares in the port excess bandwidth pool based on the strict-high priority excess bandwidth sharing weight of “1”, which is not configurable. The actual amount of extra bandwidth that traffic exceeding the transmit rate receives depends on how many other queues consume excess bandwidth and the excess rates of those queues.

On QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, we recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. A shaping rate (shaper) sets the maximum amount of bandwidth a queue can consume. (Unlike using the transmit rate on a QFX10000 switch to limit traffic that receives strict-high priority treatment, traffic that exceeds the shaping rate is dropped, and is not treated as best-effort traffic that shares in excess bandwidth.) If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

Required Privilege Level

interface—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

medium-low and medium-low options introduced for QFX10000 Series switches in Junos OS 19.2R3.

scheduler

IN THIS SECTION

- [Syntax | 256](#)
- [Hierarchy Level | 256](#)
- [Description | 256](#)
- [Options | 256](#)
- [Required Privilege Level | 257](#)
- [Release Information | 257](#)

Syntax

```
scheduler scheduler-name;
```

Hierarchy Level

```
[edit class-of-service scheduler-maps map-name forwarding-class class-name]
```

Description

Map a scheduler to a forwarding class using a scheduler map.

NOTE: On QFX5200 only, absolute CoS rate limits for transmit rate and shaping rate do not reflect 50g and 100g interfaces. Therefore this statement does not affect those interfaces for QFX5200 in release 15.1X53-D30.

Options

scheduler-name—Name of the scheduler to map to the forwarding class.

Required Privilege Level

interfaces—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

scheduler-map

IN THIS SECTION

- [Syntax | 257](#)
- [Enhanced Transmission Selection \(ETS\) Hierarchical Scheduling | 257](#)
- [Port Scheduling | 258](#)
- [Description | 258](#)
- [Options | 258](#)
- [Required Privilege Level | 258](#)
- [Release Information | 258](#)

Syntax

```
scheduler-map map-name;
```

Enhanced Transmission Selection (ETS) Hierarchical Scheduling

```
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
```

Port Scheduling

```
[edit class-of-service interfaces interface-name]
```

Description

Associate a scheduler map with a traffic control profile.

Options

map-name—Name of the scheduler map.

Required Privilege Level

interface—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

scheduler-maps

IN THIS SECTION

- [Syntax | 259](#)
- [Hierarchy Level | 259](#)
- [Description | 259](#)
- [Options | 259](#)
- [Required Privilege Level | 259](#)
- [Release Information | 259](#)

Syntax

```
scheduler-maps {  
  map-name {  
    forwarding-class class-name scheduler scheduler-name;  
  }  
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Specify a scheduler map name to map a scheduler configuration to a forwarding class.

Options

map-name—Name of the scheduler map.

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.

Release Information

Statement introduced in Junos OS Release 11.1.

schedulers

IN THIS SECTION

- [Syntax | 260](#)
- [Hierarchy Level | 260](#)
- [Description | 261](#)
- [Options | 261](#)
- [Required Privilege Level | 261](#)
- [Release Information | 261](#)

Syntax

```
schedulers {
  scheduler-name {
    buffer-dynamic-threshold value;
    buffer-size (percent percentage | remainder);
    drop-profile-map loss-priority (low | medium-high | high) protocol protocol drop-profile drop-profile-name;
    excess-rate;
    explicit-congestion-notification;
    priority priority;
    shaping-rate (rate | percent percentage);
    transmit-rate (percent percentage) <exact>;
  }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Specify scheduler name and parameter values such as minimum bandwidth (transmit-rate), maximum bandwidth (shaping-rate), and priority (priority).

Options

NOTE: The syntax above shows all possible options. The options that appear for you can vary depending on the platform, software release, and operating system (Junos OS or Junos OS Evolved).

scheduler-name —Name of the scheduler.

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.

Release Information

Statement introduced in Junos OS Release 11.1.

buffer-dynamic-threshold option introduced in Junos OS Evolved Release 22.3R1.

shaping-rate

IN THIS SECTION

- [Syntax | 262](#)
- [Hierarchy Level | 262](#)
- [Description | 262](#)
- [Default | 263](#)
- [Options | 263](#)

- Required Privilege Level | 264
- Release Information | 264

Syntax

```
shaping-rate (rate | percent percentage);
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name],  
[edit class-of-service traffic-control-profiles profile-name]
```

NOTE: Only switches that support enhanced transmission selection (ETS) hierarchical scheduling support the traffic-control-profiles hierarchy.

Description

Configure the shaping rate. The shaping rate throttles the rate of packet transmission by setting a maximum bandwidth (rate in bits per second) or a maximum percentage of bandwidth for a queue or a forwarding class set. You specify the maximum bandwidth for a queue by using a scheduler map to associate a forwarding class (queue) with a scheduler that has a configured shaping rate.

For ETS configuration, you specify the maximum bandwidth for a forwarding class set by setting the shaping rate for a traffic control profile, then you associate the scheduler map with the traffic control profile, and then you apply the traffic control profile and a forwarding class set to an interface.

For simple port scheduling configuration, you apply the scheduler map directly to an interface (instead of indirectly through the traffic control profile as in ETS).

We recommend that you configure the shaping rate as an absolute maximum usage and not as additional usage beyond the configured transmit rate (the minimum guaranteed bandwidth for a queue) or the configured guaranteed rate (the minimum guaranteed bandwidth for a forwarding class set).

NOTE: When you set the maximum bandwidth (shaping-rate value) for a queue or for a priority group at 100 Kbps or less, the traffic shaping behavior is accurate only within +/- 20 percent of the configured shaping-rate value.

NOTE: On QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, we recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

NOTE: On QFX5200 Series switches, a granularity of 64kbps is supported for the shaping rate. Therefore, the shaping rate on queues for 100g interfaces might not be applied correctly.

NOTE: QFX10000 Series switches do not support the shaping-rate statement. However, you can configure the transmit-rate exact option to prevent a queue from consuming more bandwidth than you want the queue to consume.

On QFX10000 Series switches, we recommend that you use the transmit rate to set a limit on the amount of bandwidth that receives strict-high priority treatment on a strict-high priority queue. Traffic up to the transmit rate receives strict-high priority treatment. Traffic in excess of the transmit rate is treated as best-effort traffic that receives the strict-high priority queue excess rate weight of "1". Do not use a shaping rate to set a maximum bandwidth limit on strict-high priority queues on QFX10000 Series switches.

Default

If you do not configure a shaping rate, the default shaping rate is 100 percent (all of the available bandwidth), which is the equivalent of no rate shaping.

Options

percent *percentage*—Shaping rate as a percentage of the available interface bandwidth.

- **Range:** 1 through 100 percent

rate—Peak (maximum) 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 10,000,000,000 bps

Required Privilege Level

interface—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Queue Schedulers

Example: Configuring Queue Schedulers for Port Scheduling

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

Understanding CoS Output Queue Schedulers

Understanding CoS Port Schedulers on QFX Switches

Understanding CoS Traffic Control Profiles

traffic-control-profiles

IN THIS SECTION

- [Syntax | 265](#)
- [Hierarchy Level | 266](#)
- [Description | 266](#)
- [Options | 267](#)
- [Required Privilege Level | 267](#)

Syntax

EX Series (Except EX4600), M Series, MX Series, PTX Series, T Series

```
traffic-control-profiles profile-name {
    adjust-minimum rate;
    atm-service (cbr | rtvbr | nrtvbr);
    delay-buffer-rate (percent percentage | rate);
    excess-rate (percent percentage | proportion value);
    excess-rate-high (percent percentage | proportion value);
    excess-rate-low (percent percentage | proportion value);
    guaranteed-rate (percent percentage | rate) <burst-size bytes>;
    max-burst-size cells;
    overhead-accounting (frame-mode | cell-mode | frame-mode-bytes | cell-mode-bytes) <bytes (byte-value)>;
    peak-rate rate;
    scheduler-map map-name;
    shaping-rate (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-medium-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-medium-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-medium (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-medium-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-strict-high (percent percentage | rate) <burst-size bytes>;
    strict-priority-scheduler;
    sustained-rate rate;
}
```

QFX Series including QFabric, OCX OCX1100, EX4600, NFX Series

```
traffic-control-profiles profile-name {
    guaranteed-rate (rate | percent percentage);
    scheduler-map map-name;
```

```

    shaping-rate (rate| percent percentage);
}

```

ACX Series

```

traffic-control-profiles profile-name {
    atm-service (cbr | nrtvbr | rtvbr);
    delay-buffer-rate cps;
    max-burst-size max-burst-size;
    peak-rate peak-rate;
    sustained-rate sustained-rate;
}

```

Hierarchy Level

```

[edit class-of-service]

```

Description

ACX Series Routers

Configure traffic-shaping profiles.

NOTE: For CoS on ACX6360-OR, see the documentation for the PTX1000.

EX Series (Except EX4600), M Series, MX Series, T Series, and PTX Series Routers

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 `excess-rate` statement.

QFX Series QFabric, OCX1100, EX4600, NFX Series

Configure traffic shaping and scheduling profiles for forwarding class sets (priority groups) to implement enhanced transmission selection (ETS) or for logical interfaces.

Options

profile-name—Name of the traffic-control profile. This name is also used to specify an output traffic control profile.

The remaining statements are explained separately. See [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.

Release Information

Statement was introduced in Junos OS Release 7.6 (EX series, M series, MX series, T series, and PTX series devices).

Statement was introduced in Junos OS Release 11.1 for the QFX Series.

Statement was introduced in Junos OS Release 12.3 for ACX series routers.

Statement was introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

RELATED DOCUMENTATION

[Oversubscribing Interface Bandwidth](#)

[Understanding Scheduling on PTX Series Routers](#)

[Example: Configuring CoS Hierarchical Port Scheduling \(ETS\)](#)

[Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\)](#)

[Example: Configuring Forwarding Class Sets](#)

[Assigning CoS Components to Interfaces](#)

[output-traffic-control-profile](#)

[Understanding CoS Traffic Control Profiles](#)

transmit-rate

IN THIS SECTION

- [QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems | 268](#)
- [QFX10000 Switches | 268](#)
- [Hierarchy Level | 268](#)
- [Description | 268](#)
- [Default | 270](#)
- [Options | 272](#)
- [Required Privilege Level | 272](#)
- [Release Information | 272](#)

QFX5100, EX4600, QFX3500, and QFX3600, Switches, QFabric Systems

```
transmit-rate (rate | percent percentage);
```

QFX10000 Switches

```
transmit-rate (rate | percent percentage) <exact>;
```

Hierarchy Level

```
[edit class-of-service schedulers scheduler-name]
```

Description

On QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, the transmit rate specifies the minimum guaranteed transmission rate or percentage for a queue (forwarding class) scheduler. The queue transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share on switches that support enhanced transmission selection (ETS) hierarchical scheduling.

On QFX10000 switches, the transmit rate specifies the minimum guaranteed transmission rate or percentage for a queue (forwarding class) scheduler. The queue transmit rate also determines the amount of excess (extra) port bandwidth the queue can share if you do not explicitly configure an excess rate in the scheduler. The transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share on switches that support enhanced transmission selection (ETS) hierarchical scheduling.

On QFX10000 switch strict-high priority queues, the transmit rate limits the amount of traffic the switch treats as strict-high priority traffic. Traffic up to the transmit rate receives strict-high priority treatment. The switch treats traffic that exceeds the transmit rate as best-effort traffic that receives an excess bandwidth sharing weight of “1”; you cannot configure an excess rate on a strict-high priority queue, and unlike queues with other scheduling priorities, the switch does not use the transmit rate to determine extra bandwidth sharing for strict-high priority queues.



CAUTION: We strongly recommend that you configure a transmit rate on strict-high priority queues to limit the amount of traffic the switch treats as strict-high priority traffic on those queues. This is especially important if you configure more than one strict-high priority queue on a port. To prevent a strict-high priority queue from starving the other queues on a port, we recommend that you always configure a transmit rate, even if you only configure one strict-high priority queue.

NOTE: For ETS, the transmit-rate setting works only if you also configure the *guaranteed-rate* in the traffic control profile that is attached to the forwarding class set to which the queue belongs. If you do not configure the guaranteed rate, the minimum guaranteed rate for individual queues that you set using the transmit-rate statement does not work. The sum of all queue transmit rates in a forwarding class set should not exceed the traffic control profile guaranteed rate.

NOTE: On QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, you cannot configure a transmit rate for a strict-high priority queue. Queues (forwarding classes) with a configured transmit rate cannot be included in a forwarding class set that has a strict-high priority queue. To prevent strict-high priority queues from consuming all of the available bandwidth on these switches, we recommend that you configure a shaping rate to set a maximum amount of bandwidth for strict-high priority queues.

NOTE: For transmit rates below 1 Gbps, we recommend that you configure the transmit rate as a percentage instead of as a fixed rate. This is because the system converts fixed rates into

percentages and may round small fixed rates to a lower percentage. For example, a fixed rate of 350 Mbps is rounded down to 3 percent instead of 3.5 percent.

Default

On QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, if you do not configure the transmit rate, the default scheduler transmission rate and buffer size percentages for queues 0 through 11 are:

Table 46: Default Transmit Rates for QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems

Queue Number	Default Minimum Guaranteed Bandwidth (Transmit Rate)
0 (best-effort)	5 %
1	0
2	0
3 (fcoe)	35 %
4 (no-loss)	35 %
5	0
6	0
7 (network control)	5 %
8 (mcast)	20 %
9	0
10	0

Table 46: Default Transmit Rates for QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems (Continued)

Queue Number	Default Minimum Guaranteed Bandwidth (Transmit Rate)
11	0

NOTE: OCX Series switches do not support lossless transport. The OCX Series default DSCP classifier does not classify traffic into the default lossless fcoe and no-loss forwarding classes. The bandwidth that the default scheduler allocates to the default fcoe and no-loss forwarding classes on other switches is allocated to the default best-effort, network-control, and mcast forwarding classes on OCX Series switches.

On QFX10000 switches, if you do not configure the transmit rate, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are:

Table 47: Default Transmit Rates for QFX10000 Switches

Queue Number	Default Minimum Guaranteed Bandwidth (Transmit Rate)
0 (best-effort)	15 %
1	0
2	0
3 (fcoe)	35 %
4 (no-loss)	35 %
5	0
6	0
7 (network control)	15 %

Configure schedulers if you want to change the minimum guaranteed bandwidth and other queue characteristics.

Options

rate—Minimum transmission rate for the queue, 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:** 1000 through 10,000,000,000 bps on 10-Gigabit interfaces, 1000 through 40,000,000,000 bps on 40-Gigabit interfaces.

percent *percentage*—Minimum percentage of transmission capacity allocated to the queue.

- **Range:** 1 through 100 percent

exact—(QFX10000 switches only) Shape queues that are not strict-high priority queues to the transmit rate so that the transmit rate is the maximum bandwidth limit. Traffic that exceeds the exact transmit rate is dropped. You cannot set an excess rate on queues configured as *transmit-rate (rate / percentage)* *exact* because the purpose of setting an exact transmit rate is to set a maximum bandwidth (shaping rate) on the traffic.

NOTE: On QFX10000 switches, oversubscribing all 8 queues configured with the *transmit rate exact (shaping)* statement at the `[edit class-of-service schedulers scheduler-name]` hierarchy level might result in less than 100 percent utilization of port bandwidth.

Required Privilege Level

interfaces—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Release 11.1.

Exact option introduced in Junos OS Release 15.1X53-D10 for the QFX Series.

RELATED DOCUMENTATION

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Example: Configuring Queue Schedulers

Example: Configuring Queue Schedulers for Port Scheduling

Understanding CoS Output Queue Schedulers

Understanding CoS Port Schedulers on QFX Switches

CHAPTER 7

Operational Commands (Basic Concepts)

IN THIS CHAPTER

- [Monitoring Interfaces That Have CoS Components | 274](#)
- [show class-of-service | 276](#)
- [show class-of-service interface | 282](#)
- [show class-of-service shared-buffer | 326](#)
- [show pfe filter hw summary | 330](#)
- [show pfe next-hop | 332](#)
- [show pfe route | 339](#)
- [show pfe terse | 353](#)
- [show pfe version | 356](#)
- [show interfaces voq | 357](#)

Monitoring Interfaces That Have CoS Components

IN THIS SECTION

- [Purpose | 274](#)
- [Action | 275](#)
- [Meaning | 275](#)

Purpose

Use the monitoring functionality to display details about the physical and logical interfaces and the CoS components assigned to them.

Action

To monitor interfaces that have CoS components in the CLI, enter the command:

```
user@switch> show class-of-service interface
```

To monitor a specific interface in the CLI, enter the command:

```
user@switch> show class-of-service interface interface-name
```

Meaning

[Table 48 on page 275](#) summarizes key output fields for CoS interfaces.

Table 48: Summary of Key CoS Interfaces Output Fields

Field	Values
Physical interface	Name of a physical interface to which CoS components are assigned.
Index	Index of this interface or the internal index of a specific object.
Queues supported	Number of queues you can configure on the interface.
Queues in use	Number of queues currently configured.
Scheduler map	Name of the scheduler map associated with this interface.
Congestion-notification	Status of congestion notification (enabled or disabled). NOTE: OCX Series switches do not support congestion notification profiles.
Rewrite Input IEEE Code-point	(Fibre Channel NP_Port interfaces only) IEEE 802.1p code point (priority) the interface assigns to incoming Fibre Channel (FC) traffic when the interface encapsulates the FC traffic in Ethernet before forwarding it onto the FCoE network.
Logical Interface	Name of a logical interface on the physical interface to which CoS components are assigned.

Table 48: Summary of Key CoS Interfaces Output Fields *(Continued)*

Field	Values
Object	Category of an object—for example, classifier, scheduler-map, or rewrite.
Name	Name of the object—for example, ba-classifier.
Type	Type of the object—for example, ieee8021p for a classifier.

RELATED DOCUMENTATION

| *Assigning CoS Components to Interfaces*

show class-of-service

IN THIS SECTION

- [Syntax | 276](#)
- [Description | 277](#)
- [Required Privilege Level | 277](#)
- [Output Fields | 277](#)
- [Sample Output | 279](#)
- [Release Information | 282](#)

Syntax

```
show class-of-service
```

Description

Display the class-of-service (CoS) information.

Required Privilege Level

view

Output Fields

[Table 49 on page 277](#) lists the output fields for the `show class-of-service` command. Output fields are listed in the approximate order in which they appear.

Table 49: show class-of-service Output Fields

Field Name	Field Description	Level of Output
Forwarding class	<p>The forwarding class configuration:</p> <ul style="list-style-type: none"> Forwarding class—Name of the forwarding class. ID—Forwarding class ID. Queue—Queue number. 	All levels
Code point type	<p>The type of code-point alias:</p> <ul style="list-style-type: none"> dscp—Aliases for DiffServ code point (DSCP) values. ieee-802.1—Aliases for IEEE 802.1p values. exp—Aliases for MPLS EXP values. 	All levels
Alias	Names given to CoS values.	All levels
Bit pattern	Set of bits associated with an alias.	All levels
Classifier	Name of the classifier.	All levels
Code point	Code-point values.	All levels

Table 49: show class-of-service Output Fields (Continued)

Field Name	Field Description	Level of Output
Loss priority	Loss priority assigned to specific CoS values and aliases of the classifier.	All levels
Rewrite rule	Name of the rewrite rule if one has been configured.	All levels
Drop profile	Name of the drop profile.	All levels
Type	Type of drop profile. QFX Series supports only the discrete type of drop-profile.	All levels
Fill level	Percentage of queue buffer fullness in a drop profile at which packets begin to drop during periods of congestion.	All levels
Scheduler map	Name of the scheduler map.	All levels
Scheduler	Name of the scheduler.	All levels
Transmit rate	Transmission rate of the scheduler.	All levels
Buffer size	Delay buffer size in the queue.	All levels
Drop profiles	Drop profiles configured for the specified scheduler.	All levels
Protocol	Transport protocol corresponding to the drop profile.	All levels
Name	Name of the drop profile.	All levels
Queues supported	Number of queues that can be configured on the interface.	All levels
Queues in use	Number of queues currently configured.	All levels

Table 49: show class-of-service Output Fields (Continued)

Field Name	Field Description	Level of Output
Physical interface	Name of the physical interface.	All levels
Scheduler map	Name of the scheduler map.	All levels
Congestion-notification	Enabled if a congestion notification profile is applied to the interface; disabled if no congestion notification profile is applied to the interface. NOTE: OCX Series switches do not support congestion notification profiles.	All levels
Forwarding class set	Name of the forwarding class set (priority group). NOTE: Only on systems that support enhanced transmission selection (ETS) hierarchical port scheduling.	
Index	Internal index of an object.	All levels

Sample Output

show class-of- service

```

user@switch> show class-of-service
Forwarding class          ID      Queue
  best-effort             0        0
  fcoe                    1        3
  no-loss                 2        4
  network-control         3        7
  mcast                   8        8

Code point type: dscp
  Alias      Bit pattern
  af11      001010
  af12      001100
  ...      ...

```


Code point type: ieee-802.1

Alias	Bit pattern
af11	100
...	...

Classifier: dscp-default, Code point type: dscp, Index: 7

Code point	Forwarding class	Loss priority
000000	best-effort	low
000001	best-effort	low
...

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

Drop profile:<default-drop-profile>, Type: discrete, Index: 1

Fill level
100

Scheduler map: <default>, Index: 2

Scheduler: <default-be>, Forwarding class: best-effort, Index: 21

Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Buffer Limit: none,
Priority: low

Excess Priority: low

drop-profile-map-set-type: mark

Drop profiles:

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

Scheduler: <default-fcoe>, Forwarding class: fcoe, Index: 50

Transmit rate: 35 percent, Rate Limit: none, Buffer size: 35 percent, Buffer Limit: none,
Priority: low

Excess Priority: low

drop-profile-map-set-type: mark

Drop profiles:

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

Scheduler: <default-noloss>, Forwarding class: no-loss, Index: 51

Transmit rate: 35 percent, Rate Limit: none, Buffer size: 35 percent, Buffer Limit: none,
Priority: low

Excess Priority: low

drop-profile-map-set-type: mark

Drop profiles:

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

Scheduler: <default-nc>, Forwarding class: network-control, Index: 23

Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Buffer Limit: none,
Priority: low

Excess Priority: low

drop-profile-map-set-type: mark

Drop profiles:

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

Scheduler: <default-mcast>, Forwarding class: mcast, Index: 49

Transmit rate: 20 percent, Rate Limit: none, Buffer size: 20 percent, Buffer Limit: none,
Priority: low

Excess Priority: low

drop-profile-map-set-type: mark

Drop profiles:

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

Physical interface: xe-0/0/0, Index: 129

Queues supported: 12, Queues in use: 12

Scheduler map: <default>, Index: 2

```
Congestion-notification: Disabled

Physical interface: xe-0/0/1, Index: 130
Queues supported: 12, Queues in use: 12
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

...           ...           ...

Forwarding class set: lan-fcset, Type: normal-type, Forwarding class set index: 7
  Forwarding class                                Index
  best-effort                                     0
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

<i>Monitoring CoS Code-Point Value Aliases</i>
<i>Monitoring CoS Classifiers</i>
<i>Monitoring CoS Forwarding Classes</i>
Monitoring Interfaces That Have CoS Components 274
<i>Monitoring CoS Scheduler Maps</i>
<i>Monitoring CoS Rewrite Rules</i>

show class-of-service interface

IN THIS SECTION

- [Syntax | 283](#)
- [Description | 283](#)
- [Options | 283](#)
- [Required Privilege Level | 284](#)

- [Output Fields | 284](#)
- [Sample Output | 304](#)

Syntax

```
show class-of-service interface <comprehensive | detail> <interface-name>
```

Description

Display the logical and physical interface associations for the classifier, rewrite rules, and scheduler map objects.

NOTE: This topic lists all possible options and output fields for the `show class-of-service interface` command. Options and output fields can vary depending on the platform, software release, and operating system (Junos OS or Junos OS Evolved).

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 | (Optional) Display comprehensive quality-of-service (QoS) information about all physical and logical interfaces. |
| detail | (Optional) Display QoS and CoS information based on the interface.

If the interface <i>interface-name</i> is a physical interface, the output includes: <ul style="list-style-type: none">• 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

Output Fields

[Table 50 on page 284](#) 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 50: show class-of-service interface Output Fields

Field Name	Field Description
Physical interface	Name of a physical interface.
Index	<p>Index of this interface or the internal index of this object.</p> <p>(Enhanced subscriber management) 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.</p>

Table 50: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Dedicated Queues	Status of dedicated queues configured on an interface. (Enhanced subscriber management) This field is not displayed for enhanced subscriber management.
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. (Enhanced subscriber management) This field is not displayed for enhanced subscriber management.
Rewrite Input IEEE Code-point	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 Shaping rate 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) The name of the dynamic scheduler map object is associated with a generated UID (for example, SMAP-1_UID1002) instead of with a subscriber interface.
Scheduler map forwarding class sets	Name of the output fabric scheduler map associated with a QFabric system Interconnect device interface.
Input shaping rate	Maximum transmission rate on the input interface.

Table 50: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Input scheduler map	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) The name of the dynamic traffic control profile object is associated with a generated UID (for example, TC_PROF_100_199_SERIES_UID1006) 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	Congestion notification state, enabled or disabled.
Dedicated Buffer Profile	Name of the dedicated buffer profile associated with the interface.
Monitoring Profile Name	Name of the monitoring profile defined to monitor the peak queue length for virtual output queues (VOQs) for the interface.
Logical interface	Name of a logical interface.
Object	Category of an object: Classifier, Fragmentation-map (for LSQ interfaces only), Scheduler-map, Rewrite, Translation Table, or traffic-class-map.
Name	Name of an object.

Table 50: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Type	Type of an object: dscp, dscp-ipv6, exp, ieee-802.1, ip, inet-precedence, or ieee-802.1ad.
Link-level type	Encapsulation on the physical interface.
MTU	MTU size on the physical interface.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback.
Source filtering	Whether source filtering is enabled or disabled.
Flow control	Whether flow control is enabled or disabled.
Auto-negotiation	Whether autonegotiation is enabled or disabled.
Remote-fault	Remote fault status. <ul style="list-style-type: none"> Online—Autonegotiation is manually configured as online. Offline—Autonegotiation is manually configured as offline.

Table 50: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
Device flags	<p>The Device flags field provides information about the physical device and displays one or more of the following values:</p> <ul style="list-style-type: none"> • Down—Device has been administratively disabled. • Hear-Own-Xmit—Device receives its own transmissions. • Link-Layer-Down—The link-layer protocol has failed to connect with the remote endpoint. • Loopback—Device is in physical loopback. • Loop-Detected—The link layer has received frames that it sent, thereby detecting a physical loopback. • No-Carrier—On media that support carrier recognition, no carrier is currently detected. • No-Multicast—Device does not support multicast traffic. • Present—Device is physically present and recognized. • Promiscuous—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media. • Quench—Transmission on the device is quenched because the output buffer is overflowing. • Recv-All-Multicasts—Device is in multicast promiscuous mode and therefore provides no multicast filtering. • Running—Device is active and enabled.

Table 50: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Interface flags	<p>The Interface flags field provides information about the physical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> • Admin-Test—Interface is in test mode and some sanity checking, such as loop detection, is disabled. • Disabled—Interface is administratively disabled. • Down—A hardware failure has occurred. • Hardware-Down—Interface is nonfunctional or incorrectly connected. • Link-Layer-Down—Interface keepalives have indicated that the link is incomplete. • No-Multicast—Interface does not support multicast traffic. • No-receive No-transmit—Passive monitor mode is configured on the interface. • Point-To-Point—Interface is point-to-point. • Pop all MPLS labels from packets of depth—MPLS labels are removed as packets arrive on an interface that has the pop-all-labels statement configured. The depth value can be one of the following: <ul style="list-style-type: none"> • 1—Takes effect for incoming packets with one label only. • 2—Takes effect for incoming packets with two labels only. • [1 2]—Takes effect for incoming packets with either one or two labels. • Promiscuous—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses. • Recv-All-Multicasts—Interface is in multicast promiscuous mode and provides no multicast filtering. • SNMP-Traps—SNMP trap notifications are enabled. • Up—Interface is enabled and operational.

Table 50: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Flags	<p>The Logical interface flags field provides information about the logical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> • ACFC Encapsulation—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer). • Device-down—Device has been administratively disabled. • Disabled—Interface is administratively disabled. • Down—A hardware failure has occurred. • Clear-DF-Bit—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit. • Hardware-Down—Interface protocol initialization failed to complete successfully. • PFC—Protocol field compression is enabled for the PPP session. • Point-To-Point—Interface is point-to-point. • SNMP-Traps—SNMP trap notifications are enabled. • Up—Interface is enabled and operational.
Encapsulation	Encapsulation on the logical interface.
Admin	Administrative state of the interface (Up or Down)
Link	Status of physical link (Up or Down).
Proto	Protocol configured on the interface.
Input Filter	Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.
Output Filter	Names of any firewall filters to be evaluated when packets are transmitted on the interface, including any filters attached through activation of dynamic service.

Table 50: show class-of-service interface Output Fields (*Continued*)

Field Name	Field Description
Link flags	<p>Provides information about the physical link and displays one or more of the following values:</p> <ul style="list-style-type: none"> • ACFC—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option. • Give-Up—Link protocol does not continue connection attempts after repeated failures. • Loose-LCP—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational. • Loose-LMI—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational. • Loose-NCP—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational. • Keepalives—Link protocol keepalives are enabled. • No-Keepalives—Link protocol keepalives are disabled. • PFC—Protocol field compression is configured. The PPP session negotiates the PFC option.
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
CoS queues	Number of CoS queues configured.
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: <i>year-month-day hour:minute:second:timezone (hour:minute:second ago)</i> . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).

Table 50: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Statistics last cleared	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface.
Exclude Overhead Bytes	<p>Exclude the counting of overhead bytes from aggregate queue statistics.</p> <ul style="list-style-type: none"> • Disabled—Default configuration. Includes the counting of overhead bytes in aggregate queue statistics. • Enabled—Excludes the counting of overhead bytes from aggregate queue statistics for just the physical interface. • Enabled for hierarchy—Excludes the counting of overhead bytes from aggregate queue statistics for the physical interface as well as all child interfaces, including logical interfaces and interface sets.
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>

Table 50: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
Input errors	<p>Input errors on the interface. The labels are explained in the following list:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame terminations and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Giants—Number of frames received that are larger than the giant threshold. • Bucket Drops—Drops resulting from the traffic load exceeding the interface transmit or receive leaky bucket configuration. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. Layer 3 incomplete errors can be ignored by configuring the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • HS link FIFO overflows—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces.

Table 50: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Output errors	<p>Output errors on the interface. The labels are explained in the following list:</p> <ul style="list-style-type: none"> Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. Errors—Sum of the outgoing frame terminations and FCS errors. Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. HS link FIFO underflows—Number of FIFO underflows on the high-speed links between the ASICs responsible for handling the router interfaces. MTU errors—Number of packets whose size exceeds the MTU of the interface.
Egress queues	Total number of egress Maximum usable queues on the specified interface.
Queue counters	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> Queued packets—Number of queued packets. Transmitted packets—Number of transmitted packets. Dropped packets—Number of packets dropped by the ASIC's RED mechanism.
SONET alarms SONET defects	<p>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: SONET PHY, SONET section, SONET line, and SONET path.</p>

Table 50: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
SONET PHY	<p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET PHY field has the following subfields:</p> <ul style="list-style-type: none"> • PLL Lock—Phase-locked loop • PHY Light—Loss of optical signal
SONET section	<p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET section field has the following subfields:</p> <ul style="list-style-type: none"> • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOS—Loss of signal • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section)

Table 50: show class-of-service interface Output Fields (*Continued*)

Field Name	Field Description
SONET line	<p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET line field has the following subfields:</p> <ul style="list-style-type: none"> • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line)

Table 50: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
SONET path	<p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET path field has the following subfields:</p> <ul style="list-style-type: none"> • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload (signal) label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • ES-PFE—Errored seconds (far-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path)

Table 50: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Received SONET overhead Transmitted SONET overhead	<p>Values of the received and transmitted SONET overhead:</p> <ul style="list-style-type: none"> • C2—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P. • F1—Section user channel byte. This byte is set aside for the purposes of users. • K1 and K2—These bytes are allocated for APS signaling for the protection of the multiplex section. • J0—Section trace. This byte is defined for STS-1 number 1 of an STS-<i>N</i> 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. • S1—Synchronization status. The S1 byte is located in the first STS-1 number of an STS-<i>N</i> signal. • Z3 and Z4—Allocated for future use.
Received path trace Transmitted 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>
HDLC configuration	<p>Information about the HDLC configuration.</p> <ul style="list-style-type: none"> • Policing bucket—Configured state of the receiving policer. • Shaping bucket—Configured state of the transmitting shaper. • Giant threshold—Giant threshold programmed into the hardware. • Runt threshold—Runt threshold programmed into the hardware.
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. • PLP byte—Packet Level Protocol byte.

Table 50: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
CoS information	<p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.
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.
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.

Table 50: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
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 Packet Forwarding Engine Chassis Queues 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"> • (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"> • Low, non-TCP—Number of low-loss priority non-TCP packets dropped because of RED. • Low, TCP—Number of low-loss priority TCP packets dropped because of RED. • High, non-TCP—Number of high-loss priority non-TCP packets dropped because of RED. • High, TCP—Number of high-loss priority TCP packets dropped because of RED. • (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"> • Low—Number of low-loss priority packets dropped because of RED. • Medium-low—Number of medium-low loss priority packets dropped because of RED. • Medium-high—Number of medium-high loss priority packets dropped because of RED. • High—Number of high-loss priority packets dropped because of RED.

Table 50: 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"> • (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"> • Low, non-TCP—Number of low-loss priority non-TCP bytes dropped because of RED. • Low, TCP—Number of low-loss priority TCP bytes dropped because of RED. • High, non-TCP—Number of high-loss priority non-TCP bytes dropped because of RED. • High, TCP—Number of high-loss priority TCP bytes dropped because of RED.
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"> • None—No rate limit. • exact—Queue transmits at the configured rate.
Buffer size	Delay buffer size in the queue.
Priority	Scheduling priority configured as low or high.
Excess Priority	Priority of the excess bandwidth traffic on a scheduler: low, medium-low, medium-high, high, or none.

Table 50: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Drop profiles	<p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> • Loss priority—Packet loss priority for drop profile assignment. • Protocol—Transport protocol for drop profile assignment. • Index—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • Name—Name of the drop profile. • Type—Type of the drop profile: discrete or interpolated. • Fill Level—Percentage fullness of a queue. • Drop probability—Drop probability at this fill level.
Excess Priority	Priority of the excess bandwidth traffic on a scheduler.
Drop profiles	<p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> • Loss priority—Packet loss priority for drop profile assignment. • Protocol—Transport protocol for drop profile assignment. • Index—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • Name—Name of the drop profile. • Type—Type of the drop profile: discrete or interpolated. • Fill Level—Percentage fullness of a queue. • Drop probability—Drop probability at this fill level.

Table 50: show class-of-service interface Output Fields (*Continued*)

Field Name	Field Description
Adjustment information	<p>Display the assignment of shaping-rate adjustments on a scheduler node or queue.</p> <ul style="list-style-type: none"> Adjusting application—Application that is performing the shaping-rate adjustment. <ul style="list-style-type: none"> The adjusting application can appear as ancp LS-0, which is the Junos OS Access Node Control Profile process (ancpd) that performs shaping-rate adjustments on schedule nodes. The adjusting application can appear as DHCP, which adjusts the shaping-rate and overhead-accounting class-of-service attributes based on DSL Forum VSA conveyed in 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). The adjusting application can also appear as pppoe, 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]. 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). Adjustment type—Type of adjustment: absolute or delta. Configured shaping rate—Shaping rate configured for the scheduler node or queue. Adjustment value—Value of adjusted shaping rate. Adjustment target—Level of shaping-rate adjustment performed: node or queue. Adjustment overhead-accounting mode—Configured shaping mode: frame or cell. Adjustment overhead bytes—Number of bytes that the ANCP agent adds to or subtracts from the actual downstream frame overhead before reporting the adjusted values to CoS. Adjustment target—Level of shaping-rate adjustment performed: node or queue. Adjustment multicast index—

Sample Output

show class-of-service interface (Physical)

```

user@host> show class-of-service interface et-1/0/4
Physical interface: et-1/0/4, Index: 1098
Maximum usable queues: 8, Queues in use: 4
Exclude aggregate overhead bytes: disabled
Logical interface aggregate statistics: disabled
  Scheduler map: default, Index: 0
  Congestion-notification: Disabled
  Dedicated Buffer Profile: dbp1
  Monitoring Profile Name: XYZ

Logical interface: et-1/0/4.16386, Index: 1057

```

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>          27
  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        2726446535
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 Exclude Overhead Bytes: Disabled

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                                ID      Queue  Restricted queue  Fabric priority
Policing 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

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: Sendbroadcast-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	normal	low

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		exp-filter
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	ID	Queue	Restricted queue	Fabric priority
Policing 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				

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: Sendbroadcast-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
Policing 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	Name	Type	Index
Classifier	ipprec-compatibility	ip	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	Name	Type	Index
Rewrite	ri	ieee8021p (outer)	35392

Classifier	ci	ieee8021p	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	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Logical interface: ge-0/1/3.0, Index: 77

Object	Name	Type	Index
Rewrite	custom-exp2	exp (mpls-any)	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	Name	Type	Index
Classifier	ipprec-compatibility	ip	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	Name	Type	Index
Classifier	ipprec-compatibility	ip	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	Name	Type	Index
Classifier	ipprec-compatibility	ip	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	Name	Type	Index
Classifier	d1	dscp	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      Name      Type      Index
Classifier   ipprec-compatibility ip      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         ieee8021p 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

```

RELATED DOCUMENTATION

| *Verifying and Managing Junos OS Enhanced Subscriber Management*

show class-of-service shared-buffer

IN THIS SECTION

- [Syntax | 327](#)
- [Description | 327](#)
- [Options | 327](#)
- [Required Privilege Level | 327](#)

- [Output Fields | 327](#)
- [Sample Output | 329](#)
- [Release Information | 329](#)

Syntax

```
show class-of-service shared-buffer
<egress | ingress>
```

Description

Display the shared buffer allocation and partitioning configuration.

NOTE: Due to QFX5200 cross-point architecture, all buffer usage counters are maintained separately. When usage counters are displayed with the command `show class-of-service shared-buffer` on QFX5200, various pipe counters are displayed separately.

Options

- none** Display ingress and egress shared buffer settings.
- egress** (Optional) Display the egress shared buffer settings.
- ingress** (Optional) Display the ingress shared buffer settings.

Required Privilege Level

view

Output Fields

[Table 51 on page 328](#) describes the output fields for the `show class-of-service shared-buffer` command. Output fields are listed in the approximate order in which they appear.

Table 51: show class-of-service shared-buffer Output Fields

Field Name	Field Description
Ingress	Ingress shared buffer configuration.
Total Buffer	Total buffer space available to the ports in KB. This is the combined dedicated buffer pool and shared buffer pool.
Dedicated Buffer	Buffer space allocated to the dedicated buffer pool in KB.
Shared Buffer	Buffer space allocated to the shared buffer pool in KB.
Lossless	Buffer space allocated to the lossless traffic buffer pool in KB.
Lossless Headroom	<p>Buffer space allocated to the lossless headroom traffic buffer pool to support priority-based flow control (PFC) and Ethernet PAUSE in KB. (Ingress ports only.)</p> <p>NOTE: OCX Series switches do not support PFC.</p>
Lossy	Buffer space allocated to the lossy (best-effort) traffic buffer pool in KB.
Lossless Headroom Utilization	<p>Utilization of the ingress lossless headroom buffer pool. (These fields can help you to determine how much headroom buffer space you need to reserve to support PFC and Ethernet PAUSE for lossless flows.)</p> <p>NOTE: OCX Series switches do not support PFC.</p>
Node Device	Index number that identifies the switch. On a QFX3500 switch, this field always has a value of zero (0).
Total	Size of the lossless headroom ingress buffer pool in KB.
Used	Amount in KB of lossless headroom ingress buffer used.
Free	Amount in KB of lossless headroom ingress buffer free (unused).

Table 51: show class-of-service shared-buffer Output Fields (Continued)

Field Name	Field Description
Egress	Egress shared buffer configuration.
Multicast	Buffer space allocated to the multicast traffic buffer pool in KB. (Egress ports only.)

Sample Output

show class-of-service shared-buffer

```

user@switch> show class-of-service shared-buffer
Ingress:
  Total Buffer      : 9360.00 KB
  Dedicated Buffer  : 2158.00 KB
  Shared Buffer     : 7202.00 KB
    Lossless       : 648.18 KB
    Lossless Headroom : 3240.90 KB
    Lossy          : 3312.92 KB

  Lossless Headroom Utilization:
  Node Device      Total      Used      Free
  0                3240.90 KB  0.00 KB  3240.90 KB

Egress:
  Total Buffer      : 9360.00 KB
  Dedicated Buffer  : 2704.00 KB
  Shared Buffer     : 6656.00 KB
    Lossless       : 3328.00 KB
    Multicast      : 1264.64 KB
    Lossy          : 2063.36 KB

```

Release Information

Command introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic

Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic

Configuring Global Ingress and Egress Shared Buffers

Understanding CoS Buffer Configuration

show pfe filter hw summary

IN THIS SECTION

- [Syntax | 330](#)
- [Description | 330](#)
- [Required Privilege Level | 331](#)
- [Output Fields | 331](#)
- [Sample Output | 332](#)
- [Release Information | 332](#)

Syntax

```
show pfe filter hw summary
```

Description

Display a summary of the access control list (ACL; also known as firewall filter) ternary content-addressable memory (TCAM) hardware utilization to show the allocated, used, and free TCAM entry space.

Command supported on standalone QFX Series switches, QFX5100-only (pure QFX5100) Virtual Chassis Fabric (VCF), QFX5100-only (pure QFX5100) Virtual Chassis (VC), and QFX3500-only (pure QFX3500) VC.

Required Privilege Level

view

Output Fields

Table 52 on page 331 lists the output fields for the `show pfe filter hw summary` command. Output fields are listed in the approximate order in which they appear.

Table 52: show pfe filter hw summary Output Fields

Field Name	Field Description
Group	<p>ACL ingress and egress filter groups:</p> <ul style="list-style-type: none"> • iRACL group—ingress routing ACL filter group • iVACL group—ingress VLAN ACL filter group • iPACL group—ingress port ACL filter group • ePACL group—egress port ACL filter group • eVACL group—egress VLAN ACL filter group • eRACL group—egress routing ACL filter group • eRACL IPv6 group—egress IPv6 routing ACL filter group
Group-ID	Internal identification number of the filter group.
Allocated	Number of TCAM filter entries allocated to the filter group.
Used	Number of TCAM filter entries used by the filter group.
Free	Number of TCAM filter entries available for use by the filter group.

Sample Output

show pfe filter hw summary

user@switch> show pfe filter hw summary

Group	Group-ID	Allocated	Used	Free

> Ingress filter groups:				
iRACL group	14	512	4	508
iVACL group	13	512	2	510
iPACL group	12	256	2	254
> Egress filter groups:				
ePACL group	20	256	3	253
eVACL group	21	256	4	252
eRACL group	22	256	245	11
eRACL IPV6 group	24	256	3	253

Release Information

Command introduced in Junos OS Release 14.1X53-D10.

RELATED DOCUMENTATION

| *Planning the Number of Firewall Filters to Create*

show pfe next-hop

IN THIS SECTION

- [Syntax | 333](#)
- [Syntax \(TX Matrix and TX Matrix Plus Routers\) | 333](#)
- [Description | 333](#)

- Options | 333
- Required Privilege Level | 334
- Output Fields | 334
- Sample Output | 335
- Release Information | 339

Syntax

```
show pfe next-hop <interface interface-name>
```

Syntax (TX Matrix and TX Matrix Plus Routers)

```
show pfe next-hop <fpc slot> <interface interface-name> <lcc number>
```

Description

Display Packet Forwarding Engine next-hop information.

Options

none	Display all Packet Forwarding Engine next-hop information.
fpc <i>slot</i>	<p>(TX Matrix and TX Matrix Plus routers only) (Optional) Show the next hops for a Flexible PIC Concentrator (FPC) slot.</p> <ul style="list-style-type: none"> On a TX Matrix router, if you specify the number of a T640 router by using the <i>lccnumber</i> option (the recommended method), replace <i>slot</i> with a value from 0 through 7. Otherwise, replace <i>slot</i> with a value from 0 through 31. On a TX Matrix Plus router, if you specify the number of a T1600 router by using the <i>lccnumber</i> option (the recommended method), replace <i>slot</i> with a value from 0 through 7. Otherwise, replace <i>slot</i> with a value from 0 through 31.

- On a TX Matrix Plus router in the TXP-T1600-3D, TXP-T4000-3D, or TXP-Mixed-LCC-3D configuration, if you specify the number of a T1600 or T4000 router by using the *lccnumber* option (the recommended method), replace *slot* with a value from 0 through 7. Otherwise, replace *slot* with a value from 0 through 63.

For example, the following commands have the same result:

```
user@host> show pfe next-hop fpc 1 lcc 1
user@host> show pfe next-hop fpc 9
```

interface <i>interface-name</i>	(Optional) Display the Packet Forwarding Engine next-hop interface.
lcc <i>number</i>	(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Packet Forwarding Engine next-hop interface for a specific T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display Packet Forwarding Engine next-hop interface for the router (or line-card chassis) that is connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

Required Privilege Level

admin

Output Fields

[Table 53 on page 335](#) lists the output fields for the `show pfe next hop` command. Output fields are listed in the approximate order in which they appear.

Table 53: show pfe next-hop Output Fields

Field Name	Field Description
ID	The next-hop ID for the entry.
Type	The next-hop type for the entry.
Interface	The interface to which the next-hop entry is assigned.
Protocol	The protocol type for the next-hop entry.
Encap	Encapsulation type for the next-hop entry.
Next Hop Addr	Next-hop address for the next-hop entry.
MTU	MTU value for the nexthop entry.

Sample Output

show pfe next-hop

```
user@host> show pfe next-hop
```

```
Nexthop Info:
```

ID	Type	Interface	Protocol	Encap	Next Hop Addr	MTU
----	-----	-----	-----	-----	-----	-----
4	Mcast	-	IPv4	-	0.0.0.0	0
5	Bcast	-	IPv4	-	-	0
7	Discard	-	IPv4	-	-	0
8	MDiscard	-	IPv4	-	-	0
9	Reject	-	IPv4	-	-	0
13	Local	-	IPv4	-	192.168.4.60	0
14	Resolve	fxp0.0	IPv4	Unspecified	-	0
17	Local	-	IPv4	-	127.0.0.1	0
18	Unicast	fxp0.0	IPv4	Unspecified	192.168.4.254	0
21	Local	-	IPv4	-	11.1.0.1	0


```

22  Unicast  at-0/1/0.0      IPv4      ATM SNAP  11.1.0.2      4482
...

```

show pfe next-hop fpc (TX Matrix Router)

```

user@host> show pfe next-hop fpc 1
Slot 1
Nexthop Info:

```

ID	Type	Interface	Next Hop Addr	Protocol	Encap	MTU
5	Mcast	-	default	IPv4	-	0
6	Bcast	-	-	IPv4	-	0
8	Discard	-	-	IPv4	-	0
9	MDiscard	-	-	IPv4	-	0
13	Mcast	-	default	IPV6	-	0
17	MDiscard	-	-	IPV6	-	0
18	Reject	-	-	IPV6	-	0
24	Discard	-	-	None	-	0
68	Local	-	192.168.66.113	IPv4	-	0
69	Resolve	fxp0.0	-	IPv4	Unspecified	0
70	Unicast	fxp0.0	192.168.71.254	IPv4	Unspecified	0
256	Local	-	10.71.71.1	IPv4	-	0
257	Local	-	127.0.0.1	IPv4	-	0
258	Mcast.local..1	default		IPv4	Unspecified	0
259	Bcast.local..1	-		IPv4	Unspecified	0
261	Discard.local..1	-		IPv4	Unspecified	0
262	MDiscard.local..1	-		IPv4	Unspecified	0
269	Mcast.local..1	default		IPV6	Unspecified	0
271	Discard.local..1	-		IPV6	Unspecified	0

```

...

```

show pfe next-hop fpc (TX Matrix Plus Router)

```

user@host> show pfe next-hop fpc 0

Slot 0

```

ID	Type	Interface	Next Hop Addr	Protocol	Encap	MTU
31	Mcast	-	default	IPv4	-	0

32	Bcast	-	-	IPv4	-	0
34	Discard	-	-	IPv4	-	0
35	MDiscard	-	-	IPv4	-	0
36	Reject	-	-	IPv4	-	0
39	Mcast	-	default	IPv6	-	0
42	Discard	-	-	IPv6	-	0
43	MDiscard	-	-	IPv6	-	0
44	Reject	-	-	IPv6	-	0
49	Receive	-	-	MPLS	-	0
50	Discard	-	-	MPLS	-	0
111	Mcast	.local..1	default	IPv4	Unspecified	0
112	Bcast	.local..1	-	IPv4	Unspecified	0
114	Discard	.local..1	-	IPv4	Unspecified	0
115	MDiscard	.local..1	-	IPv4	Unspecified	0
116	Reject	.local..1	-	IPv4	Unspecified	0
119	Mcast	.local..1	default	IPv6	Unspecified	0
122	Discard	.local..1	-	IPv6	Unspecified	0
123	MDiscard	.local..1	-	IPv6	Unspecified	0
124	Reject	.local..1	-	IPv6	Unspecified	0
191	Mcast	.local..2	default	IPv4	Unspecified	0
192	Bcast	.local..2	-	IPv4	Unspecified	0
194	Discard	.local..2	-	IPv4	Unspecified	0
195	MDiscard	.local..2	-	IPv4	Unspecified	0
196	Reject	.local..2	-	IPv4	Unspecified	0
322	Local	-	10.1.0.5	IPv4	-	0
323	Resolve	bcm0.0	-	IPv4	Unspecified	0
326	Local	-	129.0.0.5	IPv4	-	0
327	Resolve	bcm0.0	-	IPv4	Unspecified	0
328	Local	-	fe80::201:ff:fe01:5	IPv6	-	0
329	Receive	bcm0.0	ff02::1:ff01:5	IPv6	Unspecified	0
330	Receive	bcm0.0	fe80::	IPv6	Unspecified	0
331	Resolve	bcm0.0	-	IPv6	Unspecified	0
332	Local	-	fec0::a:1:0:5	IPv6	-	0
333	Receive	bcm0.0	ff02::1:ff00:5	IPv6	Unspecified	0
334	Receive	bcm0.0	fec0::	IPv6	Unspecified	0
335	Resolve	bcm0.0	-	IPv6	Unspecified	0
348	Local	-	192.168.178.4	IPv4	-	0
349	Resolve	em0.0	-	IPv4	Unspecified	0
350	Unicast	em0.0	192.168.178.126	IPv4	Unspecified	0
357	Local	-	fe80::201:1ff:fe01:5	IPv6	-	0
512	Local	-	10.255.178.11	IPv4	-	0

513	Local	-	127.0.0.1	IPv4	-	0
515	Local	-	abcd::10:255:178:11	IPv6	-	0
516	Local	-	fe80::200:ff:fe00:0	IPv6	-	0
517	Local	-	127.0.0.1	IPv4	-	0
518	Mcast	.local..3	default	IPv4	Unspecified	0
519	Bcast	.local..3	-	IPv4	Unspecified	0
521	Discard	.local..3	-	IPv4	Unspecified	0
522	MDiscard	.local..3	-	IPv4	Unspecified	0
523	Reject	.local..3	-	IPv4	Unspecified	0
531	Mcast	.local..3	default	IPv6	Unspecified	0
533	Discard	.local..3	-	IPv6	Unspecified	0
534	MDiscard	.local..3	-	IPv6	Unspecified	0
535	Reject	.local..3	-	IPv6	Unspecified	0
539	Mgroup	-	-	IPv4	-	0
540	Bcast	ge-15/0/3.0	-	IPv4	Ethernet	0
541	Receive	ge-15/0/3.0	14.2.1.0	IPv4	Ethernet	0
542	Local	-	14.2.1.1	IPv4	-	0
543	Resolve	ge-15/0/3.0	-	IPv4	Ethernet	0
544	Bcast	ge-31/0/4.0	-	IPv4	Ethernet	0
545	Receive	ge-31/0/4.0	14.1.1.0	IPv4	Ethernet	0
546	Local	-	14.1.1.1	IPv4	-	0
547	Resolve	ge-31/0/4.0	-	IPv4	Ethernet	0
548	Unicast	ge-31/0/4.0	14.1.1.2	IPv4	Ethernet	0
549	Unicast	ge-15/0/3.0	14.2.1.2	IPv4	Ethernet	0
550	Bcast	ae1.0	-	IPv4	Ethernet	0
551	Receive	ae1.0	11.1.1.0	IPv4	Ethernet	0
552	Local	-	11.1.1.1	IPv4	-	0
553	Resolve	ae1.0	-	IPv4	Ethernet	0
554	Aggreg.	ae1.0	-	IPv4	Ethernet	0
555	Unicast	ge-23/0/8.0	11.1.1.2	IPv4	Ethernet	0
556	Unicast	ge-7/0/9.0	11.1.1.2	IPv4	Ethernet	0
557	Aggreg.	ae1.0	-	MPLS	Ethernet	0
558	Unicast	ge-23/0/8.0	-	MPLS	Ethernet	0
559	Unicast	ge-7/0/9.0	-	MPLS	Ethernet	0
560	Aggreg.	ae1.0	-	MPLS	Ethernet	0
561	Unicast	ge-23/0/8.0	-	MPLS	Ethernet	0
562	Unicast	ge-7/0/9.0	-	MPLS	Ethernet	0

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Routing Matrix with TXP-T1600 Configuration
Routing Matrix with TXP-T1600-3D Configuration
Routing Matrix with TXP-T4000-3D Configuration
Routing Matrix with a TXP-Mixed-LCC-3D Configuration

show pfe route

IN THIS SECTION

- [Syntax | 339](#)
- [Syntax \(EX Series Switches\) | 340](#)
- [Syntax \(QFX Series\) | 340](#)
- [Syntax \(MX Series, PTX10001-36MR, PTX10004, PTX10008, PTX10016 Devices\) | 340](#)
- [Syntax \(TX Matrix and TX Matrix Plus Routers\) | 340](#)
- [Description | 341](#)
- [Options | 341](#)
- [Required Privilege Level | 343](#)
- [Output Fields | 343](#)
- [Sample Output | 345](#)
- [Release Information | 352](#)

Syntax

```
show pfe route
<<inet6 | ip | iso> <prefix prefix> | <table table-name> <index index> <prefix prefix>>>
```

```
<mpls>
<summary>
```

Syntax (EX Series Switches)

```
show pfe route
<<inet6 | ip> <prefix prefix> | <table <table-name>  <index index> <prefix prefix>>>
<mpls>
<summary>
```

Syntax (QFX Series)

```
show pfe route
<<inet6 | ip> <prefix prefix> | <table <table-name>  <index index> <prefix prefix>>> <hw (host |
lpm | multicast)>>
<<clnp> <prefix prefix> | <table <table-name>  <index index> <prefix prefix>>>
<mpls>
<summary> <hw>
```

Syntax (MX Series, PTX10001-36MR, PTX10004, PTX10008, PTX10016 Devices)

```
show pfe route
<<inet6 | ip> <prefix prefix> | <table <table-name>  <index index> <prefix prefix>>>
<dhcp>
<mpls>
<summary>
```

Syntax (TX Matrix and TX Matrix Plus Routers)

```
show pfe route
<fpc slot>
<<inet6 | ip | iso> <prefix prefix> | <table <table-name>  <index index> <prefix prefix>>>
<lcc number>
<mpls>
<summary>
```

Description

Display the routes in the Packet Forwarding Engine forwarding table. The Packet Forwarding Engine forwards packets between input and output interfaces.

NOTE: The Routing Engine maintains a master copy of the forwarding table. It copies the forwarding table to the Packet Forwarding Engine, which is the part of the router or switch responsible for forwarding packets. To display the routes in the Routing Engine forwarding table, use the `show route forwarding table` command. For more information, see the [CLI Explorer](#).

Options

none	Display all Packet Forwarding Engine forwarding table information.
clnp	(Optional) Show International Standards Organization (ISO) connectionless-mode network protocol (CLNP) route table information.
dhcp	(Optional) Display Packet Forwarding Engine DHCP-Snooping route table information.
fpc slot	<p>(TX Matrix and TX Matrix Plus routers only) (Optional) Show the next hops for a Flexible PIC Concentrator (FPC) slot.</p> <ul style="list-style-type: none"> On a TX Matrix router, if you specify the number of a T640 router by using the <code>lcc number</code> option (the recommended method), replace <code>slot</code> with a value from 0 through 7. Otherwise, replace <code>slot</code> with a value from 0 through 31. On a TX Matrix Plus router, if you specify the number of a T1600 router by using the <code>lcc number</code> option (the recommended method), replace <code>slot</code> with a value from 0 through 7. Otherwise, replace <code>slot</code> with a value from 0 through 31. On a TX Matrix Plus router in the TXP-T1600-3D, TXP-T4000-3D, or TXP-Mixed-LCC-3D configuration, if you specify the number of a T1600 or T4000 router by using the <code>lcc number</code> option (the recommended method), replace <code>slot</code> with a value from 0 through 7. Otherwise, replace <code>slot</code> with a value from 0 through 63.

For example, the following commands have the same result:

```
user@host> show pfe route fpc 1 lcc 1
user@host> show pfe route fpc 9
```

host	(QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display host routes installed in the on-chip hardware table.
hw	(QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display routes installed in the on-chip hardware table (as opposed to displaying routes from the routing table and the PFE forwarding table before they are installed in the hardware).
index <i>index</i>	(Optional) Display table index.
inet6	(Optional) Display Packet Forwarding Engine IPv6 routes.
ip	(Optional) Display Packet Forwarding Engine IPv4 routes.
iso	(Optional) Display ISO version routing tables.
lcc <i>number</i>	<p>(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, the slot number of the T640 router (or line-card chassis) that houses the FPC. On a TX Matrix Plus router, the slot number of the router (line-card chassis) that houses the FPC.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> • 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix. • 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix. • 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix. • 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
lpm	(QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display longest prefix match (LPM) routes installed in the on-chip hardware table.
mpls	(Optional) Display Packet Forwarding Engine MPLS information.
multicast	(QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display multicast routes installed in the on-chip hardware table.
prefix <i>prefix</i>	(Optional) IPv4 or IPv6 prefix for which to show table entries.

summary (Optional) Display summary of Packet Forwarding Engine information. Additionally, displays total no of errors per route table per FPC slot on PTX series devices (PTX10001-36MR, PTX10004, PTX10008, PTX10016).

table <table-name> (Optional) Display table information.

Required Privilege Level

admin

Output Fields

[Table 54 on page 343](#) lists the output fields for the `show pfe route` command. Output fields are listed in the approximate order in which they appear.

Table 54: show pfe route Output Fields

Field Name	Field Description
Destination	Destination address for the entry.
NH IP Addr	Next-hop IP address for the entry.
Type	Next-hop type for the entry
NH ID	Next-hop ID for the entry
Encap	Encapsulation type for the next-hop entry.
Interface	Interface to which the next-hop entry is assigned.

[Table 55 on page 344](#) lists the output fields for the QFX Series `show pfe route hardware table (hw)` commands. Output fields are listed in the approximate order in which they appear.

Table 55: QFX Series, EX4600 switches, and OCX Series show pfe route Hardware Table Output Fields

Field Name	Field Description
Max	Maximum routing entries per route type.
Used	Number of routing entries consumed per route type.
Free	Number of unused routing entries per route type.
% Free	Percentage of unused routing entries per route type.
Rtt	Internal routing engine index number of the route table.
VRF	Internal hardware index number for the corresponding route table.
Destination	Destination address for the entry.
Type	<p>(show pfe route summary hw)—Route type for the entry: IPv4 or IPv6 route, and host, LPM, or multicast route.</p> <p>(show pfe route (ip inet6) hw)—Next-hop type for the entry.</p>
NH ID	Next-hop ID for the entry
Interface	Interface to which the next-hop entry is assigned.
HW NH-ID	Internal hardware index number of the next-hop.
Src-MAC-Address	Source MAC address.
Port	Port number.
Dst-MAC-Address	Destination MAC address.

Table 55: QFX Series, EX4600 switches, and OCX Series show pfe route Hardware Table Output Fields (Continued)

Field Name	Field Description
VLAN	ID of the multicast group VLAN.
GROUP	Internal hardware index number of the multicast group next-hop.
CLASS	Internal class number of the multicast group.

Sample Output

show pfe route ip

```
user@host> show pfe route ip
```

```
IPv4 Route Table 0, default.0, 0x0:
```

Destination	NH IP Addr	Type	NH ID Interface
default		Discard	8
127.0.0.1	127.0.0.1	Local	256
172.16/12	192.168.71.254	Unicast	68 fxp0.0
192.168.0/18	192.168.71.254	Unicast	68 fxp0.0
192.168.40/22	192.168.71.254	Unicast	68 fxp0.0
192.168.64/18	192.168.71.254	Unicast	68 fxp0.0
192.168.64/21		Resolve	67 fxp0.0
192.168.71.249	192.168.71.249	Local	66
192.168.220.0/30		Resolve	303 fe-0/0/0.0
192.168.220.0	192.168.220.0	Receive	301 fe-0/0/0.0
224.0.0.1		Mcast	5
255.255.255.255		Bcast	6

```
...
```

show pfe route iso

```
user@host# show pfe route iso

CLNS Route Table 0, CLNP.0, 0x0:
Destination                                Type      NH ID Interface
-----
default                                    Reject    60
47.0005.80ff.f800.0000.0108.0001.0102.5508.2159/152    Local    514
49.0001.00a0.c96b.c491/72                    Local    536
```

show pfe route lcc summary (TX Matrix Router)

```
user@host> show pfe route lcc 2 summary

Slot 0

IPv4 Route Tables:
Index      Routes      Size(b)
-----
Default    43          3081
1          4           281

MPLS Route Tables:
Index      Routes      Size(b)
-----
Default    1           68

IPv6 Route Tables:
Index      Routes      Size(b)
-----
Default    9           717
1          5           389

Slot 1

IPv4 Route Tables:
Index      Routes      Size(b)
-----
```

Default	43	3081
1	4	281

MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	68

IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	9	717
1	5	389

Slot 16

IPv4 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	41	2938
1	4	281

MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	68

IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	9	717
1	5	389

Slot 17

IPv4 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	41	2938
1	4	281

MPLS Route Tables:		
Index	Routes	Size(b)
-----	-----	-----
Default	1	68
IPv6 Route Tables:		
Index	Routes	Size(b)
-----	-----	-----
Default	9	717
1	5	389

show pfe route lcc summary (TX Matrix Plus Router)

user@host> show pfe route lcc 2 summary		
Slot 0		
IPv4 Route Tables:		
Index	Routes	Size(b)
-----	-----	-----
Default	25	2266
1	9	815
2	6	545
3	5	453
4	15	1371
5	5	453
6	13	1187
MPLS Route Tables:		
Index	Routes	Size(b)
-----	-----	-----
Default	1	88
4	5	452
IPv6 Route Tables:		
Index	Routes	Size(b)
-----	-----	-----
Default	7	697
1	13	1305
3	4	385

4	4	385
5	4	385
6	18	1833

Slot 6

IPv4 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	25	2266
1	9	815
2	6	545
3	5	453
4	15	1371
5	5	453
6	13	1187

MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	88
4	5	452

IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	7	697
1	13	1305
3	4	385
4	4	385
5	4	385
6	18	1833
...		

show pfe route summary (MX Series Router)

```
user@host> show pfe route summary
```

Slot 0

DHCP-Snooping Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	144

IPv4 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	25	2266
1	9	815
2	6	545
3	5	453
4	15	1371
5	5	453
6	13	1187

MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	88
4	5	452

IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	7	697
1	13	1305
3	4	385
4	4	385
5	4	385
6	18	1833

...

show pfe route summary hw (QFX Series, EX4600 Switches, OCX Series)

```

user@switch> show pfe route summary hw
Slot 0
Unit: 0
Profile active: l2-profile-three

```

Type	Max	Used	Free	% free

IPv4 Host	8192	103	8073	98.55
IPv4 LPM	16384	9	16369	99.91
IPv4 Mcast	4096	2	4037	98.56
IPv6 Host	4096	6	4037	98.56
IPv6 LPM(< 64)	8192	3	8185	99.91
IPv6 LPM(> 64)	256	1	255	99.61
IPv6 Mcast	2048	0	2019	98.58

show pfe route ip hw host (QFX Series)

```
user@switch> show pfe route ip hw host
```

```
Slot 0
```

```
Unit: 0
```

```
IPv4 Host entries present: 103
```

Rtt	VRF	Destination	Type	NH-ID	Interface	HW	NH-
ID	Src-MAC-Address	Port	Dst-MAC-Address				

4	3	255.255.255.255	Bcast	1695	.local.	.4	ifl 550
100003		00:00:00:01:02:03					
0	1	200.1.1.42	Unicast	1743	et-0/1/1	.0	ifl 559
100268		84:18:88:de:96:fd 53					
0	1	200.1.1.56	Unicast	1743	et-0/1/1	.0	ifl 559
100268		84:18:88:de:96:fd 53					
0	1	200.1.1.61	Unicast	1743	et-0/1/1	.0	ifl 559
100268		84:18:88:de:96:fd 53					
0	1	11.1.1.2	Unicast	1743	et-0/1/1	.0	ifl 559
100268		84:18:88:de:96:fd 53					
0	1	200.1.1.73	Unicast	1743	et-0/1/1	.0	ifl 559
100268		84:18:88:de:96:fd 53					
0	1	200.1.1.76	Unicast	1743	et-0/1/1	.0	ifl 559
100268		84:18:88:de:96:fd 53					
0	1	200.1.1.18	Unicast	1743	et-0/1/1	.0	ifl 559
100268		84:18:88:de:96:fd 53					
0	1	200.1.1.5	Unicast	1743	et-0/1/1	.0	ifl 559
100268		84:18:88:de:96:fd 53					
0	1	200.1.1.23	Unicast	1743	et-0/1/1	.0	ifl 559

100268	84:18:88:de:96:fd	53	00:00:00:21:12:23						
0	1	101.1.1.255		Bcast	1664	ae0	.0	ifl	544
100003	00:00:00:01:02:03	127	00:00:00:01:02:03						
0	1	200.1.1.40		Unicast	1743	et-0/1/1	.0	ifl	559
100268	84:18:88:de:96:fd	53	00:00:00:21:12:23						
0	1	200.1.1.58		Unicast	1743	et-0/1/1	.0	ifl	559
100268	84:18:88:de:96:fd	53	00:00:00:21:12:23.	.	.				
.	.	.							

show pfe route summary (PTX10001-36MR, PTX10004, PTX10008, PTX10016 Devices)

```
user@host> show pfe route summary
```

Slot 0

IPv4 Route Tables:

Index	Routes	Size(b)	Prefixes	Aggr	Installed	Comp(%)	Errors
-----	-----	-----	-----	-----	-----	-----	-----
Default	25	2600	22	0	22	0	2
1	0	0	0	0	0	0	0
50	5	520	5	0	5	0	1
36738	5	520	5	0	5	0	1

Release Information

Command introduced before Junos OS Release 7.4.

Command option `hw` introduced in Junos OS Release 14.1X53-D10 for the QFX Series.

Command support introduced in Junos OS Evolved Release 22.4R1 for the PTX Series.

RELATED DOCUMENTATION

Routing Matrix with TXP-T1600 Configuration
Routing Matrix with TXP-T1600-3D Configuration
Routing Matrix with TXP-T4000-3D Configuration
Routing Matrix with a TXP-Mixed-LCC-3D Configuration

show pfe terse

IN THIS SECTION

- [Syntax | 353](#)
- [Syntax \(TX Matrix and TX Matrix Plus Router\) | 353](#)
- [Syntax \(MX Series Router\) | 353](#)
- [Description | 354](#)
- [Options | 354](#)
- [Required Privilege Level | 355](#)
- [Sample Output | 355](#)
- [Release Information | 356](#)

Syntax

```
show pfe terse
```

Syntax (TX Matrix and TX Matrix Plus Router)

```
show pfe terse  
<lcc number | scc>  
<sfc number>
```

Syntax (MX Series Router)

```
show pfe terse  
<all-members>  
<local>  
<member member-id>
```

Description

Display Packet Forwarding Engine status information.

Options

none	Display brief information about the Packet Forwarding Engine.
all-members	(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for all members in the Virtual Chassis configuration.
lcc <i>number</i>	<p>(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Packet Forwarding Engine information for a T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display Packet Forwarding Engine information for the router (or line-card chassis) that is connected to a TX Matrix Plus router.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> • 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix. • 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix. • 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix. • 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
local	(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for the local Virtual Chassis member.
member <i>member-id</i>	(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for the specified member of the Virtual Chassis configuration. Replace <i>member-id</i> with a value of 0 or 1.
scc	(TX Matrix routers only) (Optional) Display Packet Forwarding Engine information for the TX Matrix router (or switch-card chassis).
sfc	(TX Matrix Plus routers only) (Optional) Display Packet Forwarding Engine information for the TX Matrix Plus router (or switch-fabric chassis).

Required Privilege Level

admin

Sample Output

show pfe terse (TX Matrix Router)

```
user@host> show pfe terse
Slot Type Slot   State  Flags Uptime
  0  SFM  Present Online 0x0bf 01:25:42
  2  SFM  Present Online 0x0bf 01:25:40
  0  FPC  Present Online 0x102 01:25:57
  1  FPC  Present Online 0x102 01:25:55
  2  FPC  Present Online 0x102 01:25:53
```

show pfe terse (TX Matrix Plus Router)

```
user@host> show pfe terse
sfc0-re0:
-----
Slot Type Slot   State  Uptime
  0  LCC  Present Online  2d 05:26

lcc0-re0:
-----
Slot Type Slot   State  Uptime
  0  GFPC Present Online  2d 05:25
  1  GFPC Present Online  2d 05:25
```

show pfe terse sfc (TX Matrix Plus Router)

```
user@host> show pfe terse sfc 0
sfc0-re0:
-----
Slot Type Slot   State  Uptime
  0  LCC  Present Online  2d 05:25
```

Release Information

Command introduced before Junos OS Release 7.4.

show pfe version

IN THIS SECTION

- [Syntax | 356](#)
- [Description | 356](#)
- [Options | 356](#)
- [Required Privilege Level | 356](#)
- [Sample Output | 357](#)
- [Release Information | 357](#)

Syntax

```
show pfe version <brief | detail>
```

Description

Display Packet Forwarding Engine version information.

Options

brief | detail Display the specified level of output.

Required Privilege Level

admin

Sample Output

show pfe version brief

```
user@host> show pfe version brief
PFED release 11.1D0 built by builder on 2010-11-11 05:16:11 UTC
```

show pfe version detail

```
user@host> show pfe version detail
PFED release 11.1D0 built by builder on 2010-11-11 05:16:11 UTC
    device01.example.com:/volume/build/junos/rpd_feb11/11.1/development/20101111.0/obj-i386/
junos/usr.sbin/pfed
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces voq

IN THIS SECTION

- [Syntax | 358](#)
- [Syntax \(Junos OS Evolved\) | 358](#)
- [Description | 358](#)
- [Options | 358](#)
- [Additional Information | 359](#)
- [Required Privilege Level | 359](#)
- [Output Fields | 359](#)
- [Sample Output | 360](#)
- [Release Information | 375](#)

Syntax

```
show interfaces voq interface-name
<forwarding-class forwarding-class-name>
<non-zero>
```

Syntax (Junos OS Evolved)

```
show interfaces voq interface-name
<forwarding-class forwarding-class-name>
<non-zero>
<source-fpc source-fpc-number>
```

Description

Display the random early detection (RED) drop statistics from all ingress Packet Forwarding Engines associated with the specified physical egress interface. In the VOQ architecture, egress output queues (shallow buffers) buffer data in virtual queues on ingress Packet Forwarding Engines. In cases of congestion, you can use this command to identify which ingress Packet Forwarding Engine is the source of RED-dropped packets contributing to congestion.

NOTE: On the PTX Series routers and QFX10000 switches, these statistics include tail-dropped packets.

Options

interface <i>interface-name</i>	Display the ingress VOQ RED drop statistics for the specified egress interface.
forwarding-class <i>forwarding-class-name</i>	Display VOQ RED drop statistics for a specified forwarding class.
non-zero	Display only non-zero VOQ RED drop statistics counters.
source-fpc <i>source-fpc-number</i>	Display VOQ RED drop statistics for the specified source FPC.

Additional Information

- On PTX Series routers, you can display VOQ statistics for only the WAN physical interface.
- VOQ statistics for aggregated physical interfaces are not supported. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can use the `show interfaces queue` command to identify the child link which is experiencing congestion and then view the VOQ statistics on the respective child link using the `show interfaces voq` command.

For information on virtual output queuing on PTX routers, see *Understanding Virtual Output Queues on PTX Series Packet Transport Routers*. For information on virtual output queueing on QFX10000 switches, see *Understanding CoS Virtual Output Queues (VOQs) on QFX10000 Switches*.

Required Privilege Level

view

Output Fields

[Table 56 on page 359](#) lists the output fields for the `show interfaces queue` command. Output fields are listed in the approximate order in which they appear.

Table 56: show interfaces voq Output Fields

Field Name	Field Description
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Queue	Egress queue number.

Table 56: show interfaces voq Output Fields (Continued)

Field Name	Field Description
Forwarding classes	Forwarding class name.
FPC number	Number of the Flexible PIC Concentrator (FPC) located on ingress.
PFE	Number of the Packet Forwarding Engine providing virtual output queues on the ingress.
RED-dropped packets	<p>Number of packets per second (pps) dropped because of random early detection (RED).</p> <p>NOTE: On the PTX Series routers, these statistics include tail-dropped packets.</p>
RED-dropped bytes	<p>Number of bytes per second dropped because of RED. The byte counts vary by interface hardware.</p> <p>NOTE: On the PTX Series routers, these statistics include tail-dropped packets.</p>

Sample Output

show interfaces voq (For a Specific Physical Interface) (PTX Series Routers)

The following example shows ingress RED-dropped statistics for the egress Ethernet interface configured on port 0 of Physical Interface Card (PIC) 0, located on the FPC in slot 7.

The sample output below shows that the cause of the congestion is ingress Packet Forwarding Engine PFE 0, which resides on FPC number 4, as denoted by the count of RED-dropped packets and RED-dropped bytes for egress queue 0, forwarding classes best-effort and egress queue 3, forwarding class network control.

```

user@host> show interfaces voq et-7/0/0
Physical interface: et-7/0/0, Enabled, Physical link is Up
  Interface index: 155, SNMP ifIndex: 699

```

Queue: 0, Forwarding classes: best-effort

FPC number: 1

PFE: 0

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 1

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 2

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 3

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

FPC number: 4

PFE: 0

RED-dropped packets : 19969426 2323178 pps

RED-dropped bytes : 2196636860 2044397464 bps

PFE: 1

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 2

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 3

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

FPC number: 6

PFE: 0

RED-dropped packets : 19969424 2321205 pps

RED-dropped bytes : 2196636640 2042660808 bps

PFE: 1

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 2

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 3

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 4

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 5

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 6

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 7

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

FPC number: 7

PFE: 0

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 1

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 2

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 3

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

Queue: 1, Forwarding classes: expedited-forwarding

FPC number: 1

PFE: 0

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 1

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 2

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 3

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

FPC number: 4

PFE: 0

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 1

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 2

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 3

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

FPC number: 6

PFE: 0

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 1

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 2

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 3

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 4

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 5

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 6

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 7

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

FPC number: 7

PFE: 0

RED-dropped packets : 0 0 pps

RED-dropped bytes	:	0	0 bps
PFE: 1			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 2			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 3			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Queue: 2, Forwarding classes: assured-forwarding

FPC number: 1

PFE: 0			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 1			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 2			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 3			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

FPC number: 4

PFE: 0			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 1			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 2			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 3			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

FPC number: 6

PFE: 0

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 1

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 2

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 3

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 4

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 5

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 6

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 7

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

FPC number: 7

PFE: 0

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 1

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 2

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 3

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

Queue: 3, Forwarding classes: network-control

FPC number: 1

PFE: 0

RED-dropped packets :	0	0 pps
-----------------------	---	-------

RED-dropped bytes	:	0	0 bps
PFE: 1			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 2			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 3			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

FPC number: 4

PFE: 0			
RED-dropped packets	:	16338670	1900314 pps
RED-dropped bytes	:	1797253700	1672276976 bps
PFE: 1			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 2			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 3			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

FPC number: 6

PFE: 0			
RED-dropped packets	:	16338698	1899163 pps
RED-dropped bytes	:	1797256780	1671263512 bps
PFE: 1			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 2			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 3			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 4			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 5			
RED-dropped packets	:	0	0 pps

```

    RED-dropped bytes      :                0                0 bps
PFE: 6
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
PFE: 7
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps

FPC number: 7
PFE: 0
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
PFE: 1
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
PFE: 2
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
PFE: 3
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps

```

show interfaces voq (For a Specific Physical Interface) (QFX10000 Switches)

The sample output below shows congestion on ingress PFE 1 on FPC number 0, and on ingress PFE 2 on FPC number 1, as denoted by the count of RED-dropped packets and RED-dropped bytes for best-effort egress queue 0.

```

user@host> show interfaces voq et-1/0/0
Physical interface: et-1/0/0, Enabled, Physical link is Up
  Interface index: 659, SNMP ifIndex: 539

Queue: 0, Forwarding classes: best-effort

FPC number: 0
PFE: 0
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
PFE: 1
    RED-dropped packets    :          411063248          16891870 pps
    RED-dropped bytes      :        52616095744        17297275600 bps

```


PFE: 2

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

FPC number: 1

PFE: 0

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 1

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 2

RED-dropped packets :	411063012	16891870 pps
RED-dropped bytes :	52616065536	17297275376 bps

Queue: 3, Forwarding classes: fcoe

FPC number: 0

PFE: 0

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 1

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 2

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

FPC number: 1

PFE: 0

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 1

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 2

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

Queue: 4, Forwarding classes: no-loss

FPC number: 0

PFE: 0

```

    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
PFE: 1
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
PFE: 2
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps

```

FPC number: 1

```

PFE: 0
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
PFE: 1
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
PFE: 2
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps

```

Queue: 7, Forwarding classes: network-control

FPC number: 0

```

PFE: 0
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
PFE: 1
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
PFE: 2
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps

```

FPC number: 1

```

PFE: 0
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
PFE: 1
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
PFE: 2

```

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

show interfaces voq et-7/0/0 (For a Specific Forwarding Class)

```
user@host> show interfaces voq et-7/0/0 forwarding-class best-effort
```

Physical interface: et-7/0/0, Enabled, Physical link is Up

Interface index: 155, SNMP ifIndex: 699

Queue: 0, Forwarding classes: best-effort

FPC number: 1

PFE: 0

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 1

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 2

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 3

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

FPC number: 4

PFE: 0

RED-dropped packets :	66604786	2321519 pps
RED-dropped bytes :	7326526460	2042936776 bps

PFE: 1

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 2

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 3

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

FPC number: 6

PFE: 0

RED-dropped packets :	66604794	371200 pps
RED-dropped bytes :	7326527340	326656000 bps

PFE: 1

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 2

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 3

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 4

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 5

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 6

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 7

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

FPC number: 7

PFE: 0

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 1

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 2

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 3

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

show interfaces voq et-5/0/12 (For a Specific Source FPC)

```
user@host> show interfaces voq et-5/0/12 source-fpc 0
Physical interface: et-5/0/12, Enabled, Physical link is Up
Interface index: 166, SNMP ifIndex: 1104
```

Queue: 0, Forwarding classes: best-effort

FPC number: 0

PFE: 0

RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

PFE: 1

RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

PFE: 2

RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

PFE: 3

RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Queue: 1, Forwarding classes: expedited-forwarding

FPC number: 0

PFE: 0

RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

PFE: 1

RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

PFE: 2

RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

PFE: 3

RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Queue: 2, Forwarding classes: assured-forwarding

FPC number: 0

PFE: 0

```

        RED-dropped packets :           0           0 pps
        RED-dropped bytes   :           0           0 bps
PFE: 1
        RED-dropped packets :           0           0 pps
        RED-dropped bytes   :           0           0 bps
PFE: 2
        RED-dropped packets :           0           0 pps
        RED-dropped bytes   :           0           0 bps
PFE: 3
        RED-dropped packets :           0           0 pps
        RED-dropped bytes   :           0           0 bps

Queue: 3, Forwarding classes: network-control

```

```

FPC number: 0
PFE: 0
        RED-dropped packets :           0           0 pps
        RED-dropped bytes   :           0           0 bps
PFE: 1
        RED-dropped packets :           0           0 pps
        RED-dropped bytes   :           0           0 bps
PFE: 2
        RED-dropped packets :           0           0 pps
        RED-dropped bytes   :           0           0 bps
PFE: 3
        RED-dropped packets :           0           0 pps
        RED-dropped bytes   :           0           0 bps

```

show interfaces voq et-5/0/12 (For a Specific Forwarding Class and Source FPC)

```

user@host> show interfaces voq et-5/0/12 forwarding-class best-effort source-fpc 5
Physical interface: et-5/0/12, Enabled, Physical link is Up
Interface index: 166, SNMP ifIndex: 1104

```

```

Queue: 0, Forwarding classes: best-effort

```

```

FPC number: 5
PFE: 0
        RED-dropped packets :           0           0 pps
        RED-dropped bytes   :           0           0 bps
PFE: 1

```

```

        RED-dropped packets :          0          0 pps
        RED-dropped bytes   :          0          0 bps
PFE: 2
        RED-dropped packets :          0          0 pps
        RED-dropped bytes   :          0          0 bps
PFE: 3
        RED-dropped packets :          0          0 pps
        RED-dropped bytes   :          0          0 bps
PFE: 4
        RED-dropped packets :          0          0 pps
        RED-dropped bytes   :          0          0 bps
PFE: 5
        RED-dropped packets :          0          0 pps
        RED-dropped bytes   :          0          0 bps
PFE: 6
        RED-dropped packets :          0          0 pps
        RED-dropped bytes   :          0          0 bps
PFE: 7
        RED-dropped packets :          0          0 pps
        RED-dropped bytes   :          0          0 bps

```

show interfaces voq et-7/0/0 (Non-Zero)

```
user@host> show interfaces voq et-7/0/0 non-zero
```

```
Physical interface: et-7/0/0, Enabled, Physical link is Up
Interface index: 155, SNMP ifIndex: 699
```

```
Queue: 0, Forwarding classes: best-effort
```

```
FPC number: 4
```

```
PFE: 0
```

```

        RED-dropped packets :          95862238          2301586 pps
        RED-dropped bytes   :        10544846180        2025396264 bps

```

```
FPC number: 6
```

```
PFE: 0
```

```

        RED-dropped packets :          95866639          2322569 pps
        RED-dropped bytes   :        10545330290        2043860728 bps

```

Queue: 3, Forwarding classes: network-control

FPC number: 4

PFE: 0

RED-dropped packets :	78433066	1899727 pps
RED-dropped bytes :	8627637260	1671760384 bps

FPC number: 6

PFE: 0

RED-dropped packets :	78436704	1900628 pps
RED-dropped bytes :	8628037440	1672553432 bps

show interfaces voq et-7/0/0 (For a Specific Forwarding Class and Non-Zero)

user@host **show interfaces voq et-7/0/0 forwarding-class best-effort non-zero**

Physical interface: et-7/0/0, Enabled, Physical link is Up

Interface index: 155, SNMP ifIndex: 699

Queue: 0, Forwarding classes: best-effort

FPC number: 4

PFE: 0

RED-dropped packets :	119540012	2322319 pps
RED-dropped bytes :	13149401320	2043640784 bps

FPC number: 6

PFE: 0

RED-dropped packets :	119540049	2322988 pps
RED-dropped bytes :	13149405390	2044229744 bps

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Understanding Virtual Output Queues on PTX Series Packet Transport Routers

Understanding CoS Virtual Output Queues (VOQs) on QFX10000 Switches

CHAPTER 8

Operational Commands (Classifiers and Rewrite Rules)

IN THIS CHAPTER

- [Monitoring CoS Classifiers | 376](#)
- [Monitoring CoS Forwarding Classes | 378](#)
- [Monitoring CoS Rewrite Rules | 382](#)
- [Monitoring CoS Code-Point Value Aliases | 384](#)
- [show class-of-service classifier | 385](#)
- [show class-of-service code-point-aliases | 389](#)
- [show class-of-service forwarding-class | 391](#)
- [show class-of-service forwarding-table | 396](#)
- [show class-of-service forwarding-table classifier | 401](#)
- [show class-of-service forwarding-table classifier mapping | 403](#)
- [show class-of-service forwarding-table rewrite-rule | 406](#)
- [show class-of-service forwarding-table rewrite-rule mapping | 408](#)
- [show class-of-service interface | 410](#)
- [show class-of-service rewrite-rule | 453](#)

Monitoring CoS Classifiers

IN THIS SECTION

- [Purpose | 377](#)
- [Action | 377](#)
- [Meaning | 377](#)

Purpose

Display the mapping of incoming CoS values to forwarding class and loss priority for each classifier.

Action

To monitor CoS classifiers in the CLI, enter the CLI command:

```
user@switch> show class-of-service classifier
```

To monitor a particular classifier in the CLI, enter the CLI command:

```
user@switch> show class-of-service classifier name classifier-name
```

To monitor a particular type of classifier in the CLI, enter the CLI command:

```
user@switch> show class-of-service classifier type classifier-type
```

Meaning

[Table 57 on page 377](#) summarizes key output fields for CoS classifiers.

Table 57: Summary of Key CoS Classifier Output Fields

Field	Values
Classifier	Name of a classifier.
Code point type	<div>Type of classifier:<ul style="list-style-type: none">dscp—All classifiers of the DSCP type.ieee-802.1—All classifiers of the IEEE 802.1 type.ieee-mcast—All classifiers of the IEEE 802.1 multicast type.<p>NOTE: QFX10000 switches do not use different classifiers for unicast and multideestination (multicast, broadcast, destination lookup fail) traffic, so multicast-specific classifiers are not supported.</p><ul style="list-style-type: none">exp—All classifiers of the MPLS exp type.<p>NOTE: OCX Series switches do not support MPLS.</p></div>

Table 57: Summary of Key CoS Classifier Output Fields (*Continued*)

Field	Values
Index	Internal index of the classifier.
Code point	DSCP or IEEE 802.1 code point value of the incoming packets, in bits. These values are used for classification.
Forwarding Class	Name of the forwarding class that the classifier assigns to an incoming packet. This class affects the forwarding and scheduling policies that are applied to the packet as it transits the switch.
Loss Priority	Loss priority value that the classifier assigns to the incoming packet based on its code point value.

Monitoring CoS Forwarding Classes

IN THIS SECTION

- [Purpose | 378](#)
- [Action | 378](#)
- [Meaning | 379](#)

Purpose

Use the monitoring functionality to view the current assignment of CoS forwarding classes to queue numbers on the system.

Action

To monitor CoS forwarding classes in the CLI, enter the following CLI command:

```
user@switch> show class-of-service forwarding-class
```

Meaning

Some switches use different forwarding classes, output queues, and classifiers for unicast and multideestination (multicast, broadcast, destination lookup fail) traffic. These switches support 12 forwarding classes and output queues, eight for unicast traffic and four for multideestination traffic.

Some switches use the same forwarding classes, output queues, and classifiers for unicast and multideestination traffic. These switches support eight forwarding classes and eight output queues.

[Table 58 on page 379](#) summarizes key output fields on switches that use different forwarding classes and output queues for unicast and multideestination traffic.

Table 58: Summary of Key CoS Forwarding Class Output Fields on Switches that Separate Unicast and Multideestination Traffic

Field	Values
Forwarding Class	<p>Names of forwarding classes assigned to queue numbers. By default, the following unicast forwarding classes are assigned to queues 0, 3, 4, and 7, respectively:</p> <ul style="list-style-type: none">• best-effort—Provides no special CoS handling of packets. Loss priority is typically not carried in a CoS value.• fcoe—Provides guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic.• no-loss—Provides guaranteed delivery for TCP lossless traffic• network-control—Packets can be delayed but not dropped. <p>By default, the following multideestination forwarding class is assigned to queue 8:</p> <ul style="list-style-type: none">• mcast—Provides no special CoS handling of packets.

Table 58: Summary of Key CoS Forwarding Class Output Fields on Switches that Separate Unicast and Multidestination Traffic *(Continued)*

Field	Values
Queue	<p>Queue number corresponding to (mapped to) the forwarding class name.</p> <p>By default, four queues (0, 3, 4, and 7) are assigned to unicast forwarding classes and one queue (8) is assigned to a multidestination forwarding class:</p> <ul style="list-style-type: none"> • Queue 0—best-effort • Queue 3—fcoe • Queue 4—no-loss • Queue 7—network-control • Queue 8—mcast
No-Loss	<p>Packet drop attribute associated with each forwarding class:</p> <ul style="list-style-type: none"> • Disabled—The forwarding class is configured for lossy transport (packets might drop during periods of congestion) • Enabled—The forwarding class is configured for lossless transport <p>NOTE: To achieve lossless transport, you must ensure that priority-based flow control (PFC) and DCBX are properly configured on the lossless priority (IEEE 802.1p code point), and that sufficient port bandwidth is reserved for the lossless traffic flows.</p> <p>OCX Series switches do not support lossless transport.</p>

NOTE: OCX Series switches do not support the default lossless forwarding classes `fcoe` and `no-loss`, and do not support the no-loss packet drop attribute used to configure lossless forwarding classes. On OCX Series switches, do not map traffic to the default `fcoe` and `no-loss` forwarding classes (both of these default forwarding classes carry the no-loss packet drop attribute), and do not configure the no-loss packet drop attribute on forwarding classes.

Table 59 on page 381 summarizes key output fields on switches that use the same forwarding classes and output queues for unicast and multidestination traffic.

Table 59: Summary of Key CoS Forwarding Class Output Fields on Switches That Do Not Separate Unicast and Multidestination Traffic

Field	Values
Forwarding Class	<p>Names of forwarding classes assigned to queue numbers. By default, the following forwarding classes are assigned to queues 0, 3, 4, and 7, respectively:</p> <ul style="list-style-type: none"> • best-effort—Provides no special CoS handling of packets. Loss priority is typically not carried in a CoS value. • fcoe—Provides guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic. • no-loss—Provides guaranteed delivery for TCP lossless traffic • network-control—Packets can be delayed but not dropped.
Queue	<p>Queue number corresponding to (mapped to) the forwarding class name.</p> <p>By default, four queues (0, 3, 4, and 7) are assigned to forwarding classes:</p> <ul style="list-style-type: none"> • Queue 0—best-effort • Queue 3—fcoe • Queue 4—no-loss • Queue 7—network-control

Table 59: Summary of Key CoS Forwarding Class Output Fields on Switches That Do Not Separate Unicast and Multidestination Traffic (*Continued*)

Field	Values
No-Loss	<p>Packet drop attribute associated with each forwarding class:</p> <ul style="list-style-type: none"> • Disabled—The forwarding class is configured for lossy transport (packets might drop during periods of congestion). • Enabled—The forwarding class is configured for lossless transport. <p>NOTE: To achieve lossless transport, you must ensure that priority-based flow control (PFC) and DCBX are properly configured on the lossless priority (IEEE 802.1p code point), and that sufficient port bandwidth is reserved for the lossless traffic flows.</p> <p>OCX Series switches do not support lossless transport.</p>

Monitoring CoS Rewrite Rules

IN THIS SECTION

- [Purpose | 382](#)
- [Action | 382](#)
- [Meaning | 383](#)

Purpose

Use the monitoring functionality to display information about CoS value rewrite rules, which are based on the forwarding class and loss priority.

Action

To monitor CoS rewrite rules in the CLI, enter the CLI command:

user@switch> **show class-of-service rewrite-rule**

To monitor a particular rewrite rule in the CLI, enter the CLI command:

user@switch> **show class-of-service rewrite-rule name *rewrite-rule-name***

To monitor a particular type of rewrite rule (for example, DSCP, DSCP IPv6, IEEE-802.1, or MPLS EXP) in the CLI, enter the CLI command:

user@switch> **show class-of-service rewrite-rule type *rewrite-rule-type***

Meaning

[Table 60 on page 383](#) summarizes key output fields for CoS rewrite rules.

Table 60: Summary of Key CoS Rewrite Rule Output Fields

Field	Values
Rewrite rule	Name of the rewrite rule.
Code point type	<p>Rewrite rule type:</p> <ul style="list-style-type: none"> dscp—For IPv4 DiffServ traffic. dscp-ipv6—For IPv6 Diffserv traffic. ieee-802.1—For Layer 2 traffic. exp—For MPLS traffic. <p>NOTE: OCX Series switches do not support MPLS.</p>
Index	Internal index for the rewrite rule.
Forwarding class	<p>Name of the forwarding class that is used to determine CoS values for rewriting in combination with loss priority.</p> <p>Rewrite rules are applied to CoS values in outgoing packets based on forwarding class and loss priority setting.</p>
Loss priority	Level of loss priority that is used to determine CoS values for rewriting in combination with forwarding class.

Table 60: Summary of Key CoS Rewrite Rule Output Fields *(Continued)*

Field	Values
Code point	Rewrite code point value.

RELATED DOCUMENTATION

| *Defining CoS Rewrite Rules*

Monitoring CoS Code-Point Value Aliases

IN THIS SECTION

- Purpose | 384
- Action | 384
- Meaning | 385

Purpose

Use the monitoring functionality to display information about the CoS code-point value aliases that the system is currently using to represent DSCP and IEEE 802.1p code point bits.

Action

To monitor CoS value aliases in the CLI, enter the CLI command:

```
user@switch> show class-of-service code-point-aliases
```

To monitor a specific type of code-point alias (DSCP, DSCP IPv6, IEEE 802.1, or MPLS EXP) in the CLI, enter the CLI command:

```
user@switch> show class-of-service code-point-aliases ieee-802.1
```

Meaning

Table 61 on page 385 summarizes key output fields for CoS value aliases.

Table 61: Summary of Key CoS Value Alias Output Fields

Field	Values
Code point type	Type of the CoS value: <ul style="list-style-type: none">dscp—Examines Layer 3 packet headers for IP packet classification.dscp-ipv6—Examines Layer 3 packet headers for IPv6 packet classification.ieee-802.1—Examines Layer 2 packet headers for packet classification.exp—Examines MPLS packet headers for packet classification. <p>NOTE: OCX Series switches do not support MPLS.</p>
Alias	Name given to a set of bits—for example, af11 is a name for bits 001010.
Bit pattern	Set of bits associated with the alias.

RELATED DOCUMENTATION

| *Defining CoS Code-Point Aliases*

show class-of-service classifier

IN THIS SECTION

- Syntax | 386
- Description | 386

- [Options | 386](#)
- [Required Privilege Level | 387](#)
- [Output Fields | 387](#)
- [Sample Output | 387](#)
- [Release Information | 389](#)

Syntax

```
show class-of-service classifier
<name name>
<type dscp | type dscp-ipv6 | type exp | type ieee-802.1 | type inet-precedence>
```

Description

For each class-of-service (CoS) classifier, display the mapping of code point value to forwarding class and loss priority.

Options

none	Display all classifiers.
name <i>name</i>	(Optional) Display named classifier.
type dscp	(Optional) Display all classifiers of the Differentiated Services code point (DSCP) type.
type dscp-ipv6	(Optional) Display all classifiers of the DSCP for IPv6 type.
type exp	(Optional) Display all classifiers of the MPLS experimental (EXP) type.
type ieee-802.1	(Optional) Display all classifiers of the ieee-802.1 type.
type inet-precedence	(Optional) Display all classifiers of the inet-precedence type.

Required Privilege Level

view

Output Fields

[Table 62 on page 387](#) describes the output fields for the `show class-of-service classifier` command. Output fields are listed in the approximate order in which they appear.

Table 62: show class-of-service classifier Output Fields

Field Name	Field Description
Classifier	Name of the classifier.
Code point type	Type of the classifier: exp (not on EX Series switch), dscp, dscp-ipv6 (not on EX Series switch), ieee-802.1, or inet-precedence.
Index	Internal index of the classifier.
Code point	Code point value used for classification
Forwarding class	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.
Loss priority	Loss priority value used for classification. For most platforms, the value is high or low. For some platforms, the value is high, medium-high, medium-low, or low.

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

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service code-point-aliases

IN THIS SECTION

- Syntax | 389
- Description | 389
- Options | 390
- Required Privilege Level | 390
- Output Fields | 390
- Sample Output | 391
- Release Information | 391

Syntax

```
show class-of-service code-point-aliases
<dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence>
```

Description

Display the mapping of class-of-service (CoS) code point aliases to corresponding bit patterns.

Options

none	Display code point aliases of all code point types.
dscp	(Optional) Display Differentiated Services code point (DSCP) aliases.
dscp-ipv6	(Optional) Display IPv6 DSCP aliases.
exp	(Optional) Display MPLS EXP code point aliases.
ieee-802.1	(Optional) Display IEEE-802.1 code point aliases.
inet-precedence	(Optional) Display IPv4 precedence code point aliases.

Required Privilege Level

view

Output Fields

[Table 63 on page 390](#) describes the output fields for the `show class-of-service code-point-aliases` command. Output fields are listed in the approximate order in which they appear.

Table 63: show class-of-service code-point-aliases Output Fields

Field Name	Field Description
Code point type	Type of the code points displayed: dscp, dscp-ipv6 (not on EX Series switch), exp (not on EX Series switch or the QFX Series), ieee-802.1, or inet-precedence (not on the QFX Series).
Alias	Alias for a bit pattern.
Bit pattern	Bit pattern for which the alias is displayed.

Sample Output

show class-of-service code-point-aliases exp

```

user@host> show class-of-service code-point-aliases exp
Code point type: exp
  Alias          Bit pattern
  af11           100
  af12           101
  be             000
  be1            001
  cs6            110
  cs7            111
  ef             010
  ef1            011
  nc1            110
  nc2            111
  
```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service forwarding-class

IN THIS SECTION

- [Syntax | 392](#)
- [Description | 392](#)
- [Required Privilege Level | 392](#)
- [Output Fields | 392](#)
- [Sample Output \(Junos OS\) | 393](#)
- [Sample Output \(Junos OS Evolved\) | 395](#)
- [Release Information | 396](#)

Syntax

```
show class-of-service forwarding-class
```

Description

Display information about forwarding classes, including the mapping of forwarding classes to queue numbers.

Required Privilege Level

view

Output Fields

[Table 64 on page 392](#) lists all possible output fields for the `show class-of-service forwarding-class` command. The output fields that appear vary depending on the platform, software release, and operating system (Junos OS or Junos OS Evolved).

Table 64: show class-of-service forwarding-class Output Fields

Field Name	Field Description
Forwarding class	Name of the forwarding class.
ID	Forwarding class identifier. For DSCP-based PFC, the forwarding class ID is assigned from (and should be the same as) the configured PFC priority for the forwarding class. See <i>Configuring DSCP-based PFC for Layer 3 Untagged Traffic</i> for details.
Queue	CoS output queue mapped to the forwarding class.
Restricted queue	Restricted queue number.
Fabric priority	Fabric priority for the forwarding class, either high or low. Determines the priority of packets entering the switch fabric.

Table 64: show class-of-service forwarding-class Output Fields (Continued)

Field Name	Field Description
Policing priority	Layer 2 policing priority, either premium or normal.
SPU priority	Services Processing Unit (SPU) priority, either high or low.
No-Loss	<p>Packet loss attribute to differentiate lossless forwarding classes from lossy forwarding classes:</p> <ul style="list-style-type: none"> • Disabled—Lossless transport is not configured on the forwarding class (packet drop attribute is drop). • Enabled—Lossless transport is configured on the forwarding class (packet drop attribute is no-loss).
PFC Priority	<p>For DSCP-based PFC, the explicitly configured PFC priority configured for the forwarding class.</p> <p>The DSCP value on which PFC is enabled maps to this priority, and this priority is used in PFC pause frames sent to the peer to request to pause traffic on the mapped DSCP value when the link becomes congested. The forwarding class ID is assigned from and should match this value in the output of this command. See <i>Configuring DSCP-based PFC for Layer 3 Untagged Traffic</i> for details.</p>

Sample Output (Junos OS)

show class-of-service forwarding-class (ACX Series, EX Series, MX Series, PTX Series)

```
user@host> show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Restricted queue	Fabric priority	Policing priority	SPU priority
best-effort	0	0	0	low	normal	low
expedited-forwarding	1	1	1	low	normal	high
assured-forwarding	2	2	2	low	normal	low
network-control	3	3	3	low	normal	low

show class-of-service forwarding-class (SRX Series)

```
user@host> show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	SPU
best-effort	0	0	normal	
low				
expedited-forwarding	1	1	normal	
low				
assured-forwarding	2	2	normal	
low				
network-control	3	3	normal	
low				

show class-of-service forwarding-class (EX8200 Switch)

```
user@switch> show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Fabric priority
best-effort	0	0	low
expedited-forwarding	1	5	low
assured-forwarding	2	1	low
network-control	3	7	low
mcast-be	4	2	low
mcast-ef	5	4	low
mcast-af	6	6	low

show class-of-service forwarding-class (QFX Series)

```
user@switch> show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	No-Loss	PFC priority
best-effort	0	0	normal	Disabled	
fcoe	1	3	normal	Enabled	
no-loss	2	4	normal	Enabled	
network-control	3	7	normal	Disabled	
mcast	8	8	normal	Disabled	

show class-of-service forwarding-class (QFX Series with DSCP-based PFC)

```
user@switch> show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	No-Loss	PFC priority
best-effort	0	0	normal	Disabled	
fcoe	1	3	normal	Enabled	
no-loss	2	4	normal	Enabled	
fc2	3	2	normal	Enabled	3
network-control	5	7	normal	Disabled	
fc1	7	1	normal	Enabled	7
mcast	8	8	normal	Disabled	

On switches that do not use different forwarding classes and output queues for unicast and multideestination (multicast, broadcast, destination lookup fail) traffic, there is no `mcast` forwarding class and there is no queue 8. (Switches that use different forwarding classes and output queues for unicast and multideestination traffic support 12 forwarding classes and output queues, of which four of each are dedicated to multideestination traffic. Switches that use the same forwarding classes and output queues for unicast and multideestination traffic support eight forwarding classes and eight output queues.)

Sample Output (Junos OS Evolved)

show class-of-service forwarding-class (ACX Series, QFX Series)

```
user@host> show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Restricted queue	Fabric priority	Policing priority	SPU priority
best-effort	0	0	0	low	normal	low
expedited-forwarding	1	1	1	low	normal	high
assured-forwarding	2	2	2	low	normal	low
network-control	3	3	3	low	normal	low

show class-of-service forwarding-class (PTX Series)

NOTE: You can see more advanced details about forwarding classes by logging into the Packet Forwarding Engine and entering the command `show class-of-service forwarding-classes`.

```
user@host> show class-of-service forwarding-class
```

Forwarding class	ID	Queue	No-Loss
assured-forwarding	2	2	disabled
best-effort	0	0	disabled
expedited-forwarding	1	1	disabled
network-control	3	3	disabled

Release Information

Command introduced in Junos OS Release 9.0.

PFC priority output field introduced for DSCP-based PFC in Junos OS Release 17.4R1 for the QFX Series.

show class-of-service forwarding-table

IN THIS SECTION

- [Syntax | 397](#)
- [Syntax \(TX Matrix and TX Matrix Plus Router\) | 397](#)
- [Description | 397](#)
- [Options | 397](#)
- [Required Privilege Level | 397](#)
- [Output Fields | 398](#)
- [Sample Output | 398](#)
- [Release Information | 401](#)

Syntax

```
show class-of-service forwarding-table
```

Syntax (TX Matrix and TX Matrix Plus Router)

```
show class-of-service forwarding-table  
<lcc number> | <sfc number>
```

Description

Display the entire class-of-service (CoS) configuration as it exists in the forwarding table. Executing this command is equivalent to executing all `show class-of-service forwarding-table` commands in succession.

Options

lcc number (TX Matrix and TX Matrix Plus router only) (Optional) On a TX Matrix router, display the forwarding table configuration for a specific T640 router (or line-card chassis) configured in a routing matrix. On a TX Matrix Plus router, display the forwarding table configuration for a specific router (or line-card chassis) configured in the routing matrix.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

sfc number (TX Matrix Plus routers only) (Optional) Display the forwarding table configuration for the TX Matrix Plus router. Replace *number* with 0.

Required Privilege Level

view

Output Fields

See the output field descriptions for show class-of-service forwarding-table commands:

- *show class-of-service forwarding-table classifier*
- *show class-of-service forwarding-table classifier mapping*
- *show class-of-service forwarding-table drop-profile*
- *show class-of-service forwarding-table fabric scheduler-map*
- *show class-of-service forwarding-table rewrite-rule*
- *show class-of-service forwarding-table rewrite-rule mapping*
- *show class-of-service forwarding-table scheduler-map*

Sample Output

show class-of-service forwarding-table

```
user@host> show class-of-service forwarding-table
Classifier table index: 9, # entries: 8, Table type: EXP
Entry #   Code point   Forwarding-class #   PLP
  0         000         0                   0
  1         001         0                   1
  2         010         1                   0
  3         011         1                   1
  4         100         2                   0
  5         101         2                   1
  6         110         3                   0
  7         111         3                   1

Table Index/
Interface  Index    Q num    Table type
sp-0/0/0.1001    66      11      IPv4 precedence
sp-0/0/0.2001    67      11      IPv4 precedence
sp-0/0/0.16383   68      11      IPv4 precedence
fe-0/0/0.0       69      11      IPv4 precedence

Interface: sp-0/0/0 (Index: 129, Map index: 2, Map type: FINAL,
Num of queues: 2):
Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
```

```

    Tx rate: 0 Kb (95%), Buffer size: 95 percent
Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
    Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
        Tx rate: 0 Kb (5%), Buffer size: 5 percent
Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

Interface: fe-0/0/0 (Index: 137, Map index: 2, Map type: FINAL,
Num of queues: 2):
    Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
        Tx rate: 0 Kb (95%), Buffer size: 95 percent
Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
    Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
        Tx rate: 0 Kb (5%), Buffer size: 5 percent
Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
Interface: fe-0/0/1 (Index: 138, Map index: 2, Map type: FINAL,
Num of queues: 2):
    Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
        Tx rate: 0 Kb (95%), Buffer size: 95 percent
Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
    Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
        Tx rate: 0 Kb (5%), Buffer size: 5 percent
Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

...

```

```
RED drop profile index: 1, # entries: 1
```

Drop		
Entry	Fullness(%)	Probability(%)
0	100	100

show class-of-service forwarding-table lcc (TX Matrix Plus Router)

```

user@host> show class-of-service forwarding-table lcc 0
lcc0-re0:
-----

```


Classifier table index: 9, # entries: 64, Table type: IPv6 DSCP

Entry #	Code point	Forwarding-class #	PLP
0	000000	0	0
1	000001	0	0
2	000010	0	0
3	000011	0	0
4	000100	0	0
5	000101	0	0
6	000110	0	0
7	000111	0	0
8	001000	0	0
9	001001	0	0
10	001010	0	0
11	001011	0	0
12	001100	0	0
13	001101	0	0
14	001110	0	0
15	001111	0	0
16	010000	0	0
17	010001	0	0
18	010010	0	0
19	010011	0	0
20	010100	0	0
21	010101	0	0
22	010110	0	0
23	010111	0	0
24	011000	0	0
25	011001	0	0
26	011010	0	0
27	011011	0	0
28	011100	0	0
29	011101	0	0
30	011110	0	0
31	011111	0	0
32	100000	0	0
33	100001	0	0
34	100010	0	0
35	100011	0	0
36	100100	0	0
37	100101	0	0
38	100110	0	0
39	100111	0	0

40	101000	0	0
41	101001	0	0
42	101010	0	0
43	101011	0	0
44	101100	0	0
45	101101	0	0
46	101110	0	0
...			

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service forwarding-table classifier

IN THIS SECTION

- Syntax | 401
- Description | 401
- Options | 402
- Required Privilege Level | 402
- Output Fields | 402
- Sample Output | 403
- Release Information | 403

Syntax

```
show class-of-service forwarding-table classifier
```

Description

Display the mapping of code point value to queue number and loss priority for each classifier as it exists in the forwarding table.

Options

This command has no options.

Required Privilege Level

view

Output Fields

[Table 65 on page 402](#) describes the output fields for the `show class-of-service forwarding-table classifier` command. Output fields are listed in the approximate order in which they appear.

Table 65: show class-of-service forwarding-table classifier Output Fields

Field Name	Field Description
Classifier table index	Index of the classifier table.
entries	Total number of entries.
Table type	Type of code points in the table: DSCP, EXP (not on the QFX Series), IEEE 802.1, IPv4 precedence (not on the QFX Series), or IPv6 DSCP.
Entry #	Entry number.
Code point	Code point value used for classification.
Forwarding-class #	Forwarding class to which the code point is assigned.
PLP	Packet loss priority value set by classification. For most platforms, the value can be 0 or 1. For some platforms, the value is 0, 1, 2, or 3. The value 0 represents low PLP. The value 1 represents high PLP. The value 2 represents medium-low PLP. The value 3 represents medium-high PLP.

Sample Output

show class-of-service forwarding-table classifier

```
user@host> show class-of-service forwarding-table classifier
Classifier table index: 62436, # entries: 64, Table type: DSCP
```

Entry #	Code point	Forwarding-class #	PLP
0	000000	0	0
1	000001	0	0
2	000010	0	0
3	000011	0	0
4	000100	0	0
5	000101	0	0
6	000110	0	0
7	000111	0	0
8	001000	0	0
9	001001	0	0
10	001010	1	1
11	001011	0	0
...			
60	111100	0	0
61	111101	0	0
62	111110	0	0
63	111111	0	0

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service forwarding-table classifier mapping

IN THIS SECTION

- [Syntax | 404](#)
- [Description | 404](#)

- Options | 404
- Required Privilege Level | 404
- Output Fields | 404
- Sample Output | 405
- Release Information | 405

Syntax

```
show class-of-service forwarding-table classifier mapping
```

Description

For each logical interface, display either the table index of the classifier for a given code point type or the queue number (if it is a fixed classification) in the forwarding table.

Options

This command has no options.

Required Privilege Level

view

Output Fields

[Table 66 on page 404](#) describes the output fields for the `show class-of-service forwarding-table classifier mapping` command. Output fields are listed in the approximate order in which they appear.

Table 66: show class-of-service forwarding-table classifier mapping Output Fields

Field Name	Field Description
Table index/ Q num	If the table type is Fixed, the number of the queue to which the interface is mapped. For all other types, this value is the classifier index number.

Table 66: show class-of-service forwarding-table classifier mapping Output Fields (Continued)

Field Name	Field Description
Interface	Name of the logical interface. This field can also show the physical interface (QFX Series).
Index	Logical interface index.
Table type	Type of code points in the table: DSCP, EXP (not on the QFX Series), Fixed, IEEE 802.1, IPv4 precedence (not on the QFX Series), or IPv6 DSCP. none if no-default option set.

Sample Output

show class-of-service forwarding-table classifier mapping

```
user@host> show class-of-service forwarding-table classifier mapping
```

Interface	Index	Table index/ Q num	Table type
so-5/0/0.0	10	62436	DSCP
so-0/1/0.0	11	62436	DSCP
so-0/2/0.0	12	1	Fixed
so-0/2/1.0	13	62436	DSCP
so-0/2/1.0	13	62437	IEEE 802.1
so-0/2/2.0	14	62436	DSCP
so-0/2/2.0	14	62438	IPv4 precedence

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service forwarding-table rewrite-rule

IN THIS SECTION

- [Syntax | 406](#)
- [Description | 406](#)
- [Options | 406](#)
- [Required Privilege Level | 406](#)
- [Output Fields | 406](#)
- [Sample Output | 407](#)
- [Release Information | 407](#)

Syntax

```
show class-of-service forwarding-table rewrite-rule
```

Description

Display mapping of queue number and loss priority to code point value for each rewrite rule as it exists in the forwarding table.

Options

This command has no options.

Required Privilege Level

view

Output Fields

[Table 67 on page 407](#) describes the output fields for the `show class-of-service forwarding-table rewrite-rule` command. Output fields are listed in the approximate order in which they appear.

Table 67: show class-of-service forwarding-table rewrite-rule Output Fields

Field Name	Field Description
Rewrite table index	Index for this rewrite rule.
# entries	Number of entries in this rewrite rule.
Table type	Type of table: DSCP , EXP (not on the QFX Series), EXP-PUSH-3 (not on the QFX Series), IEEE 802.1,IPv4 precedence (not on the QFX Series), IPv6 DSCP , or Fixed .
Q#	Queue number to which this entry is assigned.
Low bits	Code point value for low-priority loss profile.
State	State of this code point: enabled , rewritten , or disabled .
High bits	Code point value for high-priority loss profile.

Sample Output

show class-of-service forwarding-table rewrite-rule

```

user@host> show class-of-service forwarding-table rewrite-rule
Rewrite table index: 3753, # entries: 4, Table type: DSCP
Q#      Low bits  State      High bits  State
0       000111  Enabled    001010    Enabled
2       000000  Disabled   001100    Enabled
1       101110  Enabled    110111    Enabled
3       110000  Enabled    111000    Enabled

```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service forwarding-table rewrite-rule mapping

IN THIS SECTION

- [Syntax | 408](#)
- [Description | 408](#)
- [Options | 408](#)
- [Required Privilege Level | 408](#)
- [Output Fields | 408](#)
- [Sample Output | 409](#)
- [Release Information | 409](#)

Syntax

```
show class-of-service forwarding-table rewrite-rule mapping
```

Description

For each logical interface, display the table identifier of the rewrite rule map for each code point type.

Options

This command has no options.

Required Privilege Level

view

Output Fields

[Table 68 on page 409](#) describes the output fields for the `show class-of-service forwarding-table rewrite-rule mapping` command. Output fields are listed in the approximate order in which they appear.

Table 68: show class-of-service forwarding-table rewrite-rule mapping Output Fields

Field Name	Field Description
Interface	Name of the logical interface. This field can also show the physical interface (QFX Series).
Index	Logical interface index.
Table index	Rewrite table index.
Type	Type of classifier: DSCP , EXP (not on the QFX Series), EXP-PUSH-3 (not on the QFX Series), EXP-SWAP-PUSH-2 (not on the QFX Series), IEEE 802.1 , IPv4 precedence (not on the QFX Series), IPv6 DSCP , or Fixed .

Sample Output

show class-of-service forwarding-table rewrite-rule mapping

```

user@host> show class-of-service forwarding-table rewrite-rule mapping
Interface      Index  Table index  Type
so-5/0/0.0     10     3753  DSCP
so-0/1/0.0     11     3753  DSCP
so-0/2/0.0     12     3753  DSCP
so-0/2/1.0     13     3753  DSCP
so-0/2/2.0     14     3753  DSCP
so-0/2/3.0     15     3753  DSCP

```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service interface

IN THIS SECTION

- [Syntax | 410](#)
- [Description | 410](#)
- [Options | 411](#)
- [Required Privilege Level | 411](#)
- [Output Fields | 411](#)
- [Sample Output | 431](#)

Syntax

```
show class-of-service interface <comprehensive | detail> <interface-name>
```

Description

Display the logical and physical interface associations for the classifier, rewrite rules, and scheduler map objects.

NOTE: This topic lists all possible options and output fields for the `show class-of-service interface` command. Options and output fields can vary depending on the platform, software release, and operating system (Junos OS or Junos OS Evolved).

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	(Optional) Display comprehensive quality-of-service (QoS) information about all physical and logical interfaces.
detail	<p>(Optional) Display QoS and CoS information based on the interface.</p> <p>If the interface <i>interface-name</i> is a physical interface, the output includes:</p> <ul style="list-style-type: none"> • 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 <p>If the interface <i>interface-name</i> is a logical interface, the output includes:</p> <ul style="list-style-type: none"> • Brief QoS information about the logical interface • Information about filters or policers for the logical interface • CoS information about the logical interface
<i>interface-name</i>	(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

Output Fields

[Table 69 on page 412](#) 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 69: show class-of-service interface Output Fields

Field Name	Field Description
Physical interface	Name of a physical interface.
Index	<p>Index of this interface or the internal index of this object.</p> <p>(Enhanced subscriber management) 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.</p>
Dedicated Queues	<p>Status of dedicated queues configured on an interface.</p> <p>(Enhanced subscriber management) This field is not displayed for enhanced subscriber management.</p>
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	<p>Number of queues created in addition to the default queues.</p> <p>(Enhanced subscriber management) This field is not displayed for enhanced subscriber management.</p>
Rewrite Input IEEE Code-point	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 Shaping rate field is displayed for either the physical interface or the logical interface.

Table 69: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Scheduler map	<p>Name of the output scheduler map associated with this interface.</p> <p>(Enhanced subscriber management) The name of the dynamic scheduler map object is associated with a generated UID (for example, SMAP-1_UID1002) instead of with a subscriber interface.</p>
Scheduler map forwarding class sets	Name of the output fabric scheduler map associated with a QFabric system Interconnect device interface.
Input shaping rate	Maximum transmission rate on the input interface.
Input scheduler map	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	<p>Name of the associated traffic control profile.</p> <p>(Enhanced subscriber management) The name of the dynamic traffic control profile object is associated with a generated UID (for example, TC_PROF_100_199_SERIES_UID1006) instead of with a subscriber interface.</p>
Classifier	Name and type of classifiers associated with this interface.
Forwarding-class-map	Name of the forwarding map associated with this interface.
Congestion-notification	Congestion notification state, enabled or disabled.

Table 69: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Dedicated Buffer Profile	Name of the dedicated buffer profile associated with the interface.
Monitoring Profile Name	Name of the monitoring profile defined to monitor the peak queue length for virtual output queues (VOQs) for the interface.
Logical interface	Name of a logical interface.
Object	Category of an object: Classifier, Fragmentation-map (for LSQ interfaces only), Scheduler-map, Rewrite, Translation Table, or traffic-class-map.
Name	Name of an object.
Type	Type of an object: dscp, dscp-ipv6, exp, ieee-802.1, ip, inet-precedence, or ieee-802.1ad.
Link-level type	Encapsulation on the physical interface.
MTU	MTU size on the physical interface.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback.
Source filtering	Whether source filtering is enabled or disabled.
Flow control	Whether flow control is enabled or disabled.
Auto-negotiation	Whether autonegotiation is enabled or disabled.

Table 69: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Remote-fault	<p>Remote fault status.</p> <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline.
Device flags	<p>The Device flags field provides information about the physical device and displays one or more of the following values:</p> <ul style="list-style-type: none"> • Down—Device has been administratively disabled. • Hear-Own-Xmit—Device receives its own transmissions. • Link-Layer-Down—The link-layer protocol has failed to connect with the remote endpoint. • Loopback—Device is in physical loopback. • Loop-Detected—The link layer has received frames that it sent, thereby detecting a physical loopback. • No-Carrier—On media that support carrier recognition, no carrier is currently detected. • No-Multicast—Device does not support multicast traffic. • Present—Device is physically present and recognized. • Promiscuous—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media. • Quench—Transmission on the device is quenched because the output buffer is overflowing. • Recv-All-Multicasts—Device is in multicast promiscuous mode and therefore provides no multicast filtering. • Running—Device is active and enabled.

Table 69: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Interface flags	<p>The Interface flags field provides information about the physical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> • Admin-Test—Interface is in test mode and some sanity checking, such as loop detection, is disabled. • Disabled—Interface is administratively disabled. • Down—A hardware failure has occurred. • Hardware-Down—Interface is nonfunctional or incorrectly connected. • Link-Layer-Down—Interface keepalives have indicated that the link is incomplete. • No-Multicast—Interface does not support multicast traffic. • No-receive No-transmit—Passive monitor mode is configured on the interface. • Point-To-Point—Interface is point-to-point. • Pop all MPLS labels from packets of depth—MPLS labels are removed as packets arrive on an interface that has the pop-all-labels statement configured. The depth value can be one of the following: <ul style="list-style-type: none"> • 1—Takes effect for incoming packets with one label only. • 2—Takes effect for incoming packets with two labels only. • [1 2]—Takes effect for incoming packets with either one or two labels. • Promiscuous—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses. • Recv-All-Multicasts—Interface is in multicast promiscuous mode and provides no multicast filtering. • SNMP-Traps—SNMP trap notifications are enabled. • Up—Interface is enabled and operational.

Table 69: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Flags	<p>The Logical interface flags field provides information about the logical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> • ACFC Encapsulation—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer). • Device-down—Device has been administratively disabled. • Disabled—Interface is administratively disabled. • Down—A hardware failure has occurred. • Clear-DF-Bit—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit. • Hardware-Down—Interface protocol initialization failed to complete successfully. • PFC—Protocol field compression is enabled for the PPP session. • Point-To-Point—Interface is point-to-point. • SNMP-Traps—SNMP trap notifications are enabled. • Up—Interface is enabled and operational.
Encapsulation	Encapsulation on the logical interface.
Admin	Administrative state of the interface (Up or Down)
Link	Status of physical link (Up or Down).
Proto	Protocol configured on the interface.
Input Filter	Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.
Output Filter	Names of any firewall filters to be evaluated when packets are transmitted on the interface, including any filters attached through activation of dynamic service.

Table 69: show class-of-service interface Output Fields (*Continued*)

Field Name	Field Description
Link flags	<p>Provides information about the physical link and displays one or more of the following values:</p> <ul style="list-style-type: none"> • ACFC—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option. • Give-Up—Link protocol does not continue connection attempts after repeated failures. • Loose-LCP—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational. • Loose-LMI—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational. • Loose-NCP—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational. • Keepalives—Link protocol keepalives are enabled. • No-Keepalives—Link protocol keepalives are disabled. • PFC—Protocol field compression is configured. The PPP session negotiates the PFC option.
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
CoS queues	Number of CoS queues configured.
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: <i>year-month-day hour:minute:second:timezone (hour:minute:second ago)</i> . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).

Table 69: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Statistics last cleared	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface.
Exclude Overhead Bytes	<p>Exclude the counting of overhead bytes from aggregate queue statistics.</p> <ul style="list-style-type: none"> • Disabled—Default configuration. Includes the counting of overhead bytes in aggregate queue statistics. • Enabled—Excludes the counting of overhead bytes from aggregate queue statistics for just the physical interface. • Enabled for hierarchy—Excludes the counting of overhead bytes from aggregate queue statistics for the physical interface as well as all child interfaces, including logical interfaces and interface sets.
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>

Table 69: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
Input errors	<p>Input errors on the interface. The labels are explained in the following list:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame terminations and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Giants—Number of frames received that are larger than the giant threshold. • Bucket Drops—Drops resulting from the traffic load exceeding the interface transmit or receive leaky bucket configuration. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. Layer 3 incomplete errors can be ignored by configuring the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • HS link FIFO overflows—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces.

Table 69: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Output errors	<p>Output errors on the interface. The labels are explained in the following list:</p> <ul style="list-style-type: none"> Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. Errors—Sum of the outgoing frame terminations and FCS errors. Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. HS link FIFO underflows—Number of FIFO underflows on the high-speed links between the ASICs responsible for handling the router interfaces. MTU errors—Number of packets whose size exceeds the MTU of the interface.
Egress queues	Total number of egress Maximum usable queues on the specified interface.
Queue counters	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> Queued packets—Number of queued packets. Transmitted packets—Number of transmitted packets. Dropped packets—Number of packets dropped by the ASIC's RED mechanism.
SONET alarms SONET defects	<p>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: SONET PHY, SONET section, SONET line, and SONET path.</p>

Table 69: show class-of-service interface Output Fields (*Continued*)

Field Name	Field Description
SONET PHY	<p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET PHY field has the following subfields:</p> <ul style="list-style-type: none"> • PLL Lock—Phase-locked loop • PHY Light—Loss of optical signal
SONET section	<p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET section field has the following subfields:</p> <ul style="list-style-type: none"> • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOS—Loss of signal • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section)

Table 69: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
SONET line	<p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET line field has the following subfields:</p> <ul style="list-style-type: none"> • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line)

Table 69: show class-of-service interface Output Fields (*Continued*)

Field Name	Field Description
SONET path	<p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET path field has the following subfields:</p> <ul style="list-style-type: none"> • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload (signal) label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • ES-PFE—Errored seconds (far-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path)

Table 69: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Received SONET overhead Transmitted SONET overhead	<p>Values of the received and transmitted SONET overhead:</p> <ul style="list-style-type: none"> • C2—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P. • F1—Section user channel byte. This byte is set aside for the purposes of users. • K1 and K2—These bytes are allocated for APS signaling for the protection of the multiplex section. • J0—Section trace. This byte is defined for STS-1 number 1 of an STS-<i>N</i> 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. • S1—Synchronization status. The S1 byte is located in the first STS-1 number of an STS-<i>N</i> signal. • Z3 and Z4—Allocated for future use.
Received path trace Transmitted 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>
HDLC configuration	<p>Information about the HDLC configuration.</p> <ul style="list-style-type: none"> • Policing bucket—Configured state of the receiving policer. • Shaping bucket—Configured state of the transmitting shaper. • Giant threshold—Giant threshold programmed into the hardware. • Runt threshold—Runt threshold programmed into the hardware.
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. • PLP byte—Packet Level Protocol byte.

Table 69: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
CoS information	<p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.
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.
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.

Table 69: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
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 Packet Forwarding Engine Chassis Queues 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"> • (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"> • Low, non-TCP—Number of low-loss priority non-TCP packets dropped because of RED. • Low, TCP—Number of low-loss priority TCP packets dropped because of RED. • High, non-TCP—Number of high-loss priority non-TCP packets dropped because of RED. • High, TCP—Number of high-loss priority TCP packets dropped because of RED. • (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"> • Low—Number of low-loss priority packets dropped because of RED. • Medium-low—Number of medium-low loss priority packets dropped because of RED. • Medium-high—Number of medium-high loss priority packets dropped because of RED. • High—Number of high-loss priority packets dropped because of RED.

Table 69: 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"> • (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"> • Low, non-TCP—Number of low-loss priority non-TCP bytes dropped because of RED. • Low, TCP—Number of low-loss priority TCP bytes dropped because of RED. • High, non-TCP—Number of high-loss priority non-TCP bytes dropped because of RED. • High, TCP—Number of high-loss priority TCP bytes dropped because of RED.
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"> • None—No rate limit. • exact—Queue transmits at the configured rate.
Buffer size	Delay buffer size in the queue.
Priority	Scheduling priority configured as low or high.
Excess Priority	Priority of the excess bandwidth traffic on a scheduler: low, medium-low, medium-high, high, or none.

Table 69: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Drop profiles	<p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> • Loss priority—Packet loss priority for drop profile assignment. • Protocol—Transport protocol for drop profile assignment. • Index—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • Name—Name of the drop profile. • Type—Type of the drop profile: discrete or interpolated. • Fill Level—Percentage fullness of a queue. • Drop probability—Drop probability at this fill level.
Excess Priority	Priority of the excess bandwidth traffic on a scheduler.
Drop profiles	<p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> • Loss priority—Packet loss priority for drop profile assignment. • Protocol—Transport protocol for drop profile assignment. • Index—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • Name—Name of the drop profile. • Type—Type of the drop profile: discrete or interpolated. • Fill Level—Percentage fullness of a queue. • Drop probability—Drop probability at this fill level.

Table 69: show class-of-service interface Output Fields (*Continued*)

Field Name	Field Description
Adjustment information	<p>Display the assignment of shaping-rate adjustments on a scheduler node or queue.</p> <ul style="list-style-type: none"> Adjusting application—Application that is performing the shaping-rate adjustment. <ul style="list-style-type: none"> The adjusting application can appear as ancp LS-0, which is the Junos OS Access Node Control Profile process (ancpd) that performs shaping-rate adjustments on schedule nodes. The adjusting application can appear as DHCP, which adjusts the shaping-rate and overhead-accounting class-of-service attributes based on DSL Forum VSA conveyed in 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). The adjusting application can also appear as pppoe, 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]. 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). Adjustment type—Type of adjustment: absolute or delta. Configured shaping rate—Shaping rate configured for the scheduler node or queue. Adjustment value—Value of adjusted shaping rate. Adjustment target—Level of shaping-rate adjustment performed: node or queue. Adjustment overhead-accounting mode—Configured shaping mode: frame or cell. Adjustment overhead bytes—Number of bytes that the ANCP agent adds to or subtracts from the actual downstream frame overhead before reporting the adjusted values to CoS. Adjustment target—Level of shaping-rate adjustment performed: node or queue. Adjustment multicast index—

Sample Output

show class-of-service interface (Physical)

```

user@host> show class-of-service interface et-1/0/4
Physical interface: et-1/0/4, Index: 1098
Maximum usable queues: 8, Queues in use: 4
Exclude aggregate overhead bytes: disabled
Logical interface aggregate statistics: disabled
  Scheduler map: default, Index: 0
  Congestion-notification: Disabled
  Dedicated Buffer Profile: dbp1
  Monitoring Profile Name: XYZ

Logical interface: et-1/0/4.16386, Index: 1057

```

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>		27
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	2726446535
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 Exclude Overhead Bytes: Disabled

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                                ID      Queue  Restricted queue  Fabric priority
Policing 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

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: Sendbroadcast-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	normal	low

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		exp-filter
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	ID	Queue	Restricted queue	Fabric priority
Policing 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				

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: Sendbcast-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
Policing 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	Name	Type	Index
Classifier	ipprec-compatibility	ip	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	Name	Type	Index
Rewrite	ri	ieee8021p (outer)	35392

Classifier	ci	ieee8021p	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	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Logical interface: ge-0/1/3.0, Index: 77

Object	Name	Type	Index
Rewrite	custom-exp2	exp (mpls-any)	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	Name	Type	Index
Classifier	ipprec-compatibility	ip	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	Name	Type	Index
Classifier	ipprec-compatibility	ip	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	Name	Type	Index
Classifier	ipprec-compatibility	ip	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	Name	Type	Index
Classifier	d1	dscp	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      Name      Type      Index
Classifier   ipprec-compatibility ip      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         ieee8021p 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
```

RELATED DOCUMENTATION

| *Verifying and Managing Junos OS Enhanced Subscriber Management*

show class-of-service rewrite-rule

IN THIS SECTION

- [Syntax | 454](#)
- [Description | 454](#)
- [Options | 454](#)
- [Required Privilege Level | 454](#)

- [Output Fields | 455](#)
- [Sample Output | 455](#)
- [Release Information | 456](#)

Syntax

```
show class-of-service rewrite-rule
<name name>
<type type>
```

Description

Display the mapping of forwarding classes and loss priority to code point values.

Options

- | | |
|----------------------------|--|
| none | Display all rewrite rules. |
| name
<i>name</i> | (Optional) Display the specified rewrite rule. |
| type <i>type</i> | <p>(Optional) Display the rewrite rule of the specified type. The rewrite rule type can be one of the following:</p> <ul style="list-style-type: none"> • dscp—For IPv4 traffic. • dscp-ipv6—For IPv6 traffic. • exp—For MPLS traffic. • frame-relay-de—(SRX Series only) For Frame Relay traffic. • ieee-802.1—For Layer 2 traffic. • inet-precedence—For IPv4 traffic. |

Required Privilege Level

view

Output Fields

Table 70 on page 455 describes the output fields for the `show class-of-service rewrite-rule` command. Output fields are listed in the approximate order in which they appear.

Table 70: show class-of-service rewrite-rule Output Fields

Field Name	Field Description
Rewrite rule	Name of the rewrite rule.
Code point type	Type of rewrite rule: dscp , dscp-ipv6 , exp , frame-relay-de , or inet-precedence .
Forwarding class	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router or switch.
Index	Internal index for this particular rewrite rule.
Loss priority	Loss priority for rewriting.
Code point	Code point value to rewrite.

Sample Output

`show class-of-service rewrite-rule type dscp`

```

user@host> show class-of-service rewrite-rule type dscp
Rewrite rule: dscp-default, Code point type: dscp
  Forwarding class      Loss priority      Code point
  gold                  high               000000
  silver                low                110000
  silver                high               111000
  bronze                low                001010
  bronze                high               001100
  lead                  high               101110

Rewrite rule: abc-dscp-rewrite, Code point type: dscp, Index: 3245
  Forwarding class      Loss priority      Code point

```


gold	low	000111
gold	high	001010
silver	low	110000
silver	high	111000
bronze	high	001100
lead	low	101110
lead	high	110111

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| *Rewrite Rules Overview*

Operational Commands (Scheduling)

IN THIS CHAPTER

- [Monitoring CoS Scheduler Maps | 457](#)
- [show class-of-service drop-profile | 459](#)
- [show class-of-service forwarding-table | 464](#)
- [show class-of-service forwarding-table drop-profile | 469](#)
- [show class-of-service forwarding-table scheduler-map | 471](#)
- [show class-of-service interface | 475](#)
- [show class-of-service scheduler-map | 518](#)
- [show interfaces queue | 523](#)
- [show interfaces voq | 583](#)

Monitoring CoS Scheduler Maps

IN THIS SECTION

- [Purpose | 457](#)
- [Action | 458](#)
- [Meaning | 458](#)

Purpose

Use the monitoring functionality to display assignments of CoS forwarding classes to schedulers.

Action

To monitor CoS scheduler maps in the CLI, enter the CLI command:

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

To monitor a specific scheduler map in the CLI, enter the CLI command:

```
user@switch> show class-of-service scheduler-map scheduler-map-name
```

Meaning

[Table 71 on page 458](#) summarizes key output fields for CoS scheduler maps.

Table 71: Summary of Key CoS Scheduler Maps Output Fields

Field	Values
Scheduler map	Name of a scheduler map that maps forwarding classes to schedulers.
Index	Index of a specific object—scheduler maps, schedulers, or drop profiles.
Scheduler	Name of a scheduler that controls queue properties such as bandwidth and scheduling priority.
Forwarding class	Name(s) of the forwarding class(es) to which the scheduler is mapped.
Transmit rate	Guaranteed minimum bandwidth configured on the queue mapped to the scheduler. On strict-high priority queues on QFX10000 switches, defines the maximum amount of traffic on the queue that is treated as strict-high priority traffic.

Table 71: Summary of Key CoS Scheduler Maps Output Fields *(Continued)*

Field	Values
Priority	<p>Scheduling priority of traffic on a queue:</p> <ul style="list-style-type: none"> strict-high or high—Packets on a strict-high priority queue are transmitted first, before all other traffic, up to the configured maximum bandwidth (shaping rate). On QFX3500, QFX3600, EX4600, and OCX series switches, and on QFabric system, only one queue can be configured as strict-high or high priority. On QFX10000 switches, you can configure more than one strict-high priority queue. low—Packets in this queue are transmitted after packets in the strict-high queue.
Drop Profiles	Name and index of a drop profile that is mapped to a specific loss priority and protocol pair. The drop profile determines the way best effort queues drop packets during periods of congestion.
Loss Priority	Packet loss priority mapped to the drop profile. You can configure different drop profiles for low, medium-high, and high loss priority traffic.
Protocol	Transport protocol of the drop profile for the particular priority.
Name	Name of the drop profile.

show class-of-service drop-profile

IN THIS SECTION

- [Syntax | 460](#)
- [Description | 460](#)
- [Options | 460](#)

- Required Privilege Level | 460
- Output Fields | 460
- Sample Output | 461
- Release Information | 464

Syntax

```
show class-of-service drop-profile
<profile-name profile-name>
```

Description

Display data points for each class-of-service (CoS) random early detection (RED) drop profile.

Options

- none** Display all drop profiles.
- profile-name *profile-name*** (Optional) Display the specified profile only.

Required Privilege Level

view

Output Fields

[Table 72 on page 460](#) describes the output fields for the `show class-of-service drop-profile` command. Output fields are listed in the approximate order in which they appear.

Table 72: show class-of-service drop-profile Output Fields

Field Name	Field Description
Drop profile	Name of a drop profile.

Table 72: show class-of-service drop-profile Output Fields (*Continued*)

Field Name	Field Description
Type	Type of drop profile: <ul style="list-style-type: none"> • discrete (default) • interpolated (EX8200 switches, QFX Series switches, QFabric systems, EX4600 switches, OCX Series switches only)
Index	Internal index of this drop profile.
Fill Level	Percentage fullness of a queue.
Drop probability	Drop probability at this fill level.

Sample Output

show class-of-service drop-profile

```

user@host> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
      100          100
Drop profile: user-drop-profile, Type: interpolated, Index: 2989
  Fill level    Drop probability
      0           0
      1           1
      2           2
      4           4
      5           5
      6           6
      8           8
     10          10
     12          15
     14          20
     15          23
... 64 entries total

```

90	96
92	96
94	97
95	98
96	98
98	99
99	99
100	100

show class-of-service drop-profile (EX4200 Switch)

```

user@switch> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level
    100
Drop profile: dp1, Type: discrete, Index: 40496
  Fill level
    10

```

show class-of-service drop-profile (EX8200 Switch)

```

user@switch> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
    100          100
Drop profile: dp1, Type: interpolated, Index: 40496
  Fill level    Drop probability
    0            0
    1            80
    2            90
    4            90
    5            90
    6            90
    8            90
   10            90
   12            91
   14            91
   15            91
   16            91
   18            91

```

20	91
22	92
24	92
25	92
26	92
28	92
30	92
32	93
34	93
35	93
36	93
38	93
40	93
42	94
44	94
45	94
46	94
48	94
49	94
51	95
52	95
54	95
55	95
56	95
58	95
60	95
62	96
64	96
65	96
66	96
68	96
70	96
72	97
74	97
75	97
76	97
78	97
80	97
82	98
84	98
85	98
86	98
88	98

90	98
92	99
94	99
95	99
96	99
98	99
99	99
100	100
Drop profile: dp2, Type: discrete, Index: 40499	
Fill level	Drop probability
10	5
50	50

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service forwarding-table

IN THIS SECTION

- Syntax | 464
- Syntax (TX Matrix and TX Matrix Plus Router) | 465
- Description | 465
- Options | 465
- Required Privilege Level | 465
- Output Fields | 465
- Sample Output | 466
- Release Information | 469

Syntax

```
show class-of-service forwarding-table
```

Syntax (TX Matrix and TX Matrix Plus Router)

```
show class-of-service forwarding-table
<lcc number> | <sfc number>
```

Description

Display the entire class-of-service (CoS) configuration as it exists in the forwarding table. Executing this command is equivalent to executing all `show class-of-service forwarding-table` commands in succession.

Options

lcc number (TX Matrix and TX Matrix Plus router only) (Optional) On a TX Matrix router, display the forwarding table configuration for a specific T640 router (or line-card chassis) configured in a routing matrix. On a TX Matrix Plus router, display the forwarding table configuration for a specific router (or line-card chassis) configured in the routing matrix.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

sfc number (TX Matrix Plus routers only) (Optional) Display the forwarding table configuration for the TX Matrix Plus router. Replace *number* with 0.

Required Privilege Level

view

Output Fields

See the output field descriptions for `show class-of-service forwarding-table` commands:

- *show class-of-service forwarding-table classifier*

- *show class-of-service forwarding-table classifier mapping*
- *show class-of-service forwarding-table drop-profile*
- *show class-of-service forwarding-table fabric scheduler-map*
- *show class-of-service forwarding-table rewrite-rule*
- *show class-of-service forwarding-table rewrite-rule mapping*
- *show class-of-service forwarding-table scheduler-map*

Sample Output

show class-of-service forwarding-table

```
user@host> show class-of-service forwarding-table
```

```
Classifier table index: 9, # entries: 8, Table type: EXP
```

Entry #	Code point	Forwarding-class #	PLP
0	000	0	0
1	001	0	1
2	010	1	0
3	011	1	1
4	100	2	0
5	101	2	1
6	110	3	0
7	111	3	1

Table Index/			
Interface	Index	Q num	Table type
sp-0/0/0.1001	66	11	IPv4 precedence
sp-0/0/0.2001	67	11	IPv4 precedence
sp-0/0/0.16383	68	11	IPv4 precedence
fe-0/0/0.0	69	11	IPv4 precedence

```
Interface: sp-0/0/0 (Index: 129, Map index: 2, Map type: FINAL,
Num of queues: 2):
```

```
Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
```

```
Tx rate: 0 Kb (95%), Buffer size: 95 percent
```

```
Priority low
```

```
PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
```

```
Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
```

```
Tx rate: 0 Kb (5%), Buffer size: 5 percent
```

Priority low
 PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

Interface: fe-0/0/0 (Index: 137, Map index: 2, Map type: FINAL,
 Num of queues: 2):

Entry 0 (Scheduler index: 16, Forwarding-class #: 0):

Tx rate: 0 Kb (95%), Buffer size: 95 percent

Priority low

PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

Entry 1 (Scheduler index: 18, Forwarding-class #: 3):

Tx rate: 0 Kb (5%), Buffer size: 5 percent

Priority low

PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

Interface: fe-0/0/1 (Index: 138, Map index: 2, Map type: FINAL,
 Num of queues: 2):

Entry 0 (Scheduler index: 16, Forwarding-class #: 0):

Tx rate: 0 Kb (95%), Buffer size: 95 percent

Priority low

PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

Entry 1 (Scheduler index: 18, Forwarding-class #: 3):

Tx rate: 0 Kb (5%), Buffer size: 5 percent

Priority low

PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

...

RED drop profile index: 1, # entries: 1

Drop

Entry	Fullness(%)	Probability(%)
0	100	100

show class-of-service forwarding-table lcc (TX Matrix Plus Router)

user@host> show class-of-service forwarding-table lcc 0

lcc0-re0:

Classifier table index: 9, # entries: 64, Table type: IPv6 DSCP

Entry #	Code point	Forwarding-class #	PLP
0	000000	0	0
1	000001	0	0

2	000010	0	0
3	000011	0	0
4	000100	0	0
5	000101	0	0
6	000110	0	0
7	000111	0	0
8	001000	0	0
9	001001	0	0
10	001010	0	0
11	001011	0	0
12	001100	0	0
13	001101	0	0
14	001110	0	0
15	001111	0	0
16	010000	0	0
17	010001	0	0
18	010010	0	0
19	010011	0	0
20	010100	0	0
21	010101	0	0
22	010110	0	0
23	010111	0	0
24	011000	0	0
25	011001	0	0
26	011010	0	0
27	011011	0	0
28	011100	0	0
29	011101	0	0
30	011110	0	0
31	011111	0	0
32	100000	0	0
33	100001	0	0
34	100010	0	0
35	100011	0	0
36	100100	0	0
37	100101	0	0
38	100110	0	0
39	100111	0	0
40	101000	0	0
41	101001	0	0
42	101010	0	0
43	101011	0	0
44	101100	0	0

45	101101	0	0
46	101110	0	0
...			

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service forwarding-table drop-profile

IN THIS SECTION

- Syntax | 469
- Description | 469
- Options | 469
- Required Privilege Level | 470
- Output Fields | 470
- Sample Output | 470
- Release Information | 471

Syntax

```
show class-of-service forwarding-table drop-profile
```

Description

Display the data points of all random early detection (RED) drop profiles as they exist in the forwarding table.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 73 on page 470 describes the output fields for the `show class-of-service forwarding-table drop-profile` command. Output fields are listed in the approximate order in which they appear.

Table 73: show class-of-service forwarding-table drop-profile Output Fields

Field Name	Field Description
RED drop profile index	Index of this drop profile.
# entries	Number of entries in a particular RED drop profile index.
Entry	Drop profile entry number.
Fullness(%)	Percentage fullness of a queue.
Drop probability(%)	Drop probability at this fill level.

Sample Output

show class-of-service forwarding-table drop-profile

```

user@host> show class-of-service forwarding-table drop-profile
RED drop profile index: 4, # entries: 1
      Drop
Entry    Fullness(%)  Probability(%)
   0         100         100

RED drop profile index: 8742, # entries: 3
      Drop
Entry    Fullness(%)  Probability(%)
   0         10         10
   1         20         20

```

```

2          30          30

RED drop profile index: 24627, # entries: 64
      Drop
Entry    Fullness(%)  Probability(%)
  0         0           0
  1         1           1
  2         2           2
  3         4           4
...
 61        98          99
 62        99          99
 63       100         100

RED drop profile index: 25393, # entries: 64
      Drop
Entry    Fullness(%)  Probability(%)
  0         0           0
  1         1           1
  2         2           2
  3         4           4
...
 61        98          98
 62        99          99
 63       100         100
```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service forwarding-table scheduler-map

IN THIS SECTION

- Syntax | 472
- Description | 472
- Options | 472

- [Required Privilege Level | 472](#)
- [Output Fields | 472](#)
- [Sample Output | 474](#)
- [Release Information | 475](#)

Syntax

```
show class-of-service forwarding-table scheduler-map
```

Description

For each physical interface, display the scheduler map information as it exists in the forwarding table.

Options

This command has no options.

Required Privilege Level

view

Output Fields

[Table 74 on page 472](#) describes the output fields for the `show class-of-service forwarding-table scheduler-map` command. Output fields are listed in the approximate order in which they appear.

Table 74: show class-of-service forwarding-table scheduler-map Output Fields

Field Name	Field Description
Interface	Name of the physical interface.
Index	Physical interface index.

Table 74: show class-of-service forwarding-table scheduler-map Output Fields (*Continued*)

Field Name	Field Description
Map index	Scheduler map index.
Num of queues	Number of queues defined in this scheduler map.
Entry	Number of this entry in the scheduler map.
Scheduler index	Scheduler policy index.
Forwarding-class #	Forwarding class number to which this entry is applied.
Tx rate	Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword remainder , which indicates that the scheduler receives the remaining bandwidth of the interface.
Max buffer delay	Amount of transmit delay (in milliseconds) or buffer size of the queue. This amount is a percentage of the total interface buffer allocation or the keyword remainder , which indicates that the buffer is sized according to what remains after other scheduler buffer allocations.
Priority	<ul style="list-style-type: none"> • high—Queue priority is high. • low—Queue priority is low.
PLP high	Drop profile index for a high packet loss priority profile.
PLP low	Drop profile index for a low packet loss priority profile.
PLP medium-high	Drop profile index for a medium-high packet loss priority profile.
PLP medium-low	Drop profile index for a medium-low packet loss priority profile.
TCP PLP high	Drop profile index for a high TCP packet loss priority profile.

Table 74: show class-of-service forwarding-table scheduler-map Output Fields (Continued)

Field Name	Field Description
TCP PLP low	Drop profile index for a low TCP packet loss priority profile.
Policy is exact	If this line appears in the output, exact rate limiting is enabled. Otherwise, no rate limiting is enabled.

Sample Output

show class-of-service forwarding-table scheduler-map

```

user@host> show class-of-service forwarding-table scheduler-map
Interface: so-5/0/0 (Index: 9, Map index: 17638, Num of queues: 2):
  Entry 0 (Scheduler index: 6090, Forwarding-class #: 0):
    Tx rate: 0 Kb (30%), Max buffer delay: 39 bytes (0%)
    Priority low
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low:8742
    Policy is exact
  Entry 1 (Scheduler index: 38372, Forwarding-class #: 1):
    Traffic chunk: Max = 0 bytes, Min = 0 bytes
    Tx rate: 0 Kb (40%), Max buffer delay: 68 bytes (0%)
    Priority high
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742

Interface: at-6/1/0 (Index: 10, Map index: 17638, Num of queues: 2):
  Entry 0 (Scheduler index: 6090, Forwarding-class #: 0):
    Traffic chunk: Max = 0 bytes, Min = 0 bytes
    Tx rate: 0 Kb (30%), Max buffer delay: 39 bytes (0%)
    Priority high
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742
  Entry 1 (Scheduler index: 38372, Forwarding-class #: 1):
    Traffic chunk: Max = 0 bytes, Min = 0 bytes
    Tx rate: 0 Kb (40%), Max buffer delay: 68 bytes (0%)
    Priority low
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742

```

Release Information

Command introduced before Junos OS Release 7.4.

show class-of-service interface

IN THIS SECTION

- [Syntax | 475](#)
- [Description | 475](#)
- [Options | 476](#)
- [Required Privilege Level | 476](#)
- [Output Fields | 477](#)
- [Sample Output | 496](#)

Syntax

```
show class-of-service interface <comprehensive | detail> <interface-name>
```

Description

Display the logical and physical interface associations for the classifier, rewrite rules, and scheduler map objects.

NOTE: This topic lists all possible options and output fields for the `show class-of-service interface` command. Options and output fields can vary depending on the platform, software release, and operating system (Junos OS or Junos OS Evolved).

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** (Optional) Display comprehensive quality-of-service (QoS) information about all physical and logical interfaces.
- detail** (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

Output Fields

Table 75 on page 477 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 75: 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) 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. (Enhanced subscriber management) This field is not displayed for enhanced subscriber management.
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. (Enhanced subscriber management) This field is not displayed for enhanced subscriber management.
Rewrite Input IEEE Code-point	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 Shaping rate field is displayed for either the physical interface or the logical interface.

Table 75: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Scheduler map	<p>Name of the output scheduler map associated with this interface.</p> <p>(Enhanced subscriber management) The name of the dynamic scheduler map object is associated with a generated UID (for example, SMAP-1_UID1002) instead of with a subscriber interface.</p>
Scheduler map forwarding class sets	Name of the output fabric scheduler map associated with a QFabric system Interconnect device interface.
Input shaping rate	Maximum transmission rate on the input interface.
Input scheduler map	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	<p>Name of the associated traffic control profile.</p> <p>(Enhanced subscriber management) The name of the dynamic traffic control profile object is associated with a generated UID (for example, TC_PROF_100_199_SERIES_UID1006) instead of with a subscriber interface.</p>
Classifier	Name and type of classifiers associated with this interface.
Forwarding-class-map	Name of the forwarding map associated with this interface.
Congestion-notification	Congestion notification state, enabled or disabled.

Table 75: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Dedicated Buffer Profile	Name of the dedicated buffer profile associated with the interface.
Monitoring Profile Name	Name of the monitoring profile defined to monitor the peak queue length for virtual output queues (VOQs) for the interface.
Logical interface	Name of a logical interface.
Object	Category of an object: Classifier, Fragmentation-map (for LSQ interfaces only), Scheduler-map, Rewrite, Translation Table, or traffic-class-map.
Name	Name of an object.
Type	Type of an object: dscp, dscp-ipv6, exp, ieee-802.1, ip, inet-precedence, or ieee-802.1ad.
Link-level type	Encapsulation on the physical interface.
MTU	MTU size on the physical interface.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback.
Source filtering	Whether source filtering is enabled or disabled.
Flow control	Whether flow control is enabled or disabled.
Auto-negotiation	Whether autonegotiation is enabled or disabled.

Table 75: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Remote-fault	<p>Remote fault status.</p> <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline.
Device flags	<p>The Device flags field provides information about the physical device and displays one or more of the following values:</p> <ul style="list-style-type: none"> • Down—Device has been administratively disabled. • Hear-Own-Xmit—Device receives its own transmissions. • Link-Layer-Down—The link-layer protocol has failed to connect with the remote endpoint. • Loopback—Device is in physical loopback. • Loop-Detected—The link layer has received frames that it sent, thereby detecting a physical loopback. • No-Carrier—On media that support carrier recognition, no carrier is currently detected. • No-Multicast—Device does not support multicast traffic. • Present—Device is physically present and recognized. • Promiscuous—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media. • Quench—Transmission on the device is quenched because the output buffer is overflowing. • Recv-All-Multicasts—Device is in multicast promiscuous mode and therefore provides no multicast filtering. • Running—Device is active and enabled.

Table 75: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
Interface flags	<p>The Interface flags field provides information about the physical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> • Admin-Test—Interface is in test mode and some sanity checking, such as loop detection, is disabled. • Disabled—Interface is administratively disabled. • Down—A hardware failure has occurred. • Hardware-Down—Interface is nonfunctional or incorrectly connected. • Link-Layer-Down—Interface keepalives have indicated that the link is incomplete. • No-Multicast—Interface does not support multicast traffic. • No-receive No-transmit—Passive monitor mode is configured on the interface. • Point-To-Point—Interface is point-to-point. • Pop all MPLS labels from packets of depth—MPLS labels are removed as packets arrive on an interface that has the pop-all-labels statement configured. The depth value can be one of the following: <ul style="list-style-type: none"> • 1—Takes effect for incoming packets with one label only. • 2—Takes effect for incoming packets with two labels only. • [1 2]—Takes effect for incoming packets with either one or two labels. • Promiscuous—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses. • Recv-All-Multicasts—Interface is in multicast promiscuous mode and provides no multicast filtering. • SNMP-Traps—SNMP trap notifications are enabled. • Up—Interface is enabled and operational.

Table 75: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Flags	<p>The Logical interface flags field provides information about the logical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> • ACFC Encapsulation—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer). • Device-down—Device has been administratively disabled. • Disabled—Interface is administratively disabled. • Down—A hardware failure has occurred. • Clear-DF-Bit—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit. • Hardware-Down—Interface protocol initialization failed to complete successfully. • PFC—Protocol field compression is enabled for the PPP session. • Point-To-Point—Interface is point-to-point. • SNMP-Traps—SNMP trap notifications are enabled. • Up—Interface is enabled and operational.
Encapsulation	Encapsulation on the logical interface.
Admin	Administrative state of the interface (Up or Down)
Link	Status of physical link (Up or Down).
Proto	Protocol configured on the interface.
Input Filter	Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.
Output Filter	Names of any firewall filters to be evaluated when packets are transmitted on the interface, including any filters attached through activation of dynamic service.

Table 75: show class-of-service interface Output Fields (*Continued*)

Field Name	Field Description
Link flags	<p>Provides information about the physical link and displays one or more of the following values:</p> <ul style="list-style-type: none"> • ACFC—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option. • Give-Up—Link protocol does not continue connection attempts after repeated failures. • Loose-LCP—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational. • Loose-LMI—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational. • Loose-NCP—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational. • Keepalives—Link protocol keepalives are enabled. • No-Keepalives—Link protocol keepalives are disabled. • PFC—Protocol field compression is configured. The PPP session negotiates the PFC option.
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
CoS queues	Number of CoS queues configured.
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: <i>year-month-day hour:minute:second:timezone (hour:minute:second ago)</i> . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).

Table 75: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Statistics last cleared	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface.
Exclude Overhead Bytes	<p>Exclude the counting of overhead bytes from aggregate queue statistics.</p> <ul style="list-style-type: none"> • Disabled—Default configuration. Includes the counting of overhead bytes in aggregate queue statistics. • Enabled—Excludes the counting of overhead bytes from aggregate queue statistics for just the physical interface. • Enabled for hierarchy—Excludes the counting of overhead bytes from aggregate queue statistics for the physical interface as well as all child interfaces, including logical interfaces and interface sets.
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>

Table 75: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
Input errors	<p>Input errors on the interface. The labels are explained in the following list:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame terminations and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Giants—Number of frames received that are larger than the giant threshold. • Bucket Drops—Drops resulting from the traffic load exceeding the interface transmit or receive leaky bucket configuration. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. Layer 3 incomplete errors can be ignored by configuring the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • HS link FIFO overflows—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces.

Table 75: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Output errors	<p>Output errors on the interface. The labels are explained in the following list:</p> <ul style="list-style-type: none"> Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. Errors—Sum of the outgoing frame terminations and FCS errors. Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. HS link FIFO underflows—Number of FIFO underflows on the high-speed links between the ASICs responsible for handling the router interfaces. MTU errors—Number of packets whose size exceeds the MTU of the interface.
Egress queues	Total number of egress Maximum usable queues on the specified interface.
Queue counters	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> Queued packets—Number of queued packets. Transmitted packets—Number of transmitted packets. Dropped packets—Number of packets dropped by the ASIC's RED mechanism.
SONET alarms SONET defects	<p>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: SONET PHY, SONET section, SONET line, and SONET path.</p>

Table 75: show class-of-service interface Output Fields (*Continued*)

Field Name	Field Description
SONET PHY	<p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET PHY field has the following subfields:</p> <ul style="list-style-type: none"> • PLL Lock—Phase-locked loop • PHY Light—Loss of optical signal
SONET section	<p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET section field has the following subfields:</p> <ul style="list-style-type: none"> • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOS—Loss of signal • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section)

Table 75: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
SONET line	<p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET line field has the following subfields:</p> <ul style="list-style-type: none"> • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line)

Table 75: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
SONET path	<p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET path field has the following subfields:</p> <ul style="list-style-type: none"> • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload (signal) label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • ES-PFE—Errored seconds (far-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path)

Table 75: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Received SONET overhead Transmitted SONET overhead	<p>Values of the received and transmitted SONET overhead:</p> <ul style="list-style-type: none"> • C2—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P. • F1—Section user channel byte. This byte is set aside for the purposes of users. • K1 and K2—These bytes are allocated for APS signaling for the protection of the multiplex section. • J0—Section trace. This byte is defined for STS-1 number 1 of an STS-<i>N</i> 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. • S1—Synchronization status. The S1 byte is located in the first STS-1 number of an STS-<i>N</i> signal. • Z3 and Z4—Allocated for future use.
Received path trace Transmitted 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>
HDLC configuration	<p>Information about the HDLC configuration.</p> <ul style="list-style-type: none"> • Policing bucket—Configured state of the receiving policer. • Shaping bucket—Configured state of the transmitting shaper. • Giant threshold—Giant threshold programmed into the hardware. • Runt threshold—Runt threshold programmed into the hardware.
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. • PLP byte—Packet Level Protocol byte.

Table 75: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
CoS information	<p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.
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.
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.

Table 75: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
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 Packet Forwarding Engine Chassis Queues 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"> • (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"> • Low, non-TCP—Number of low-loss priority non-TCP packets dropped because of RED. • Low, TCP—Number of low-loss priority TCP packets dropped because of RED. • High, non-TCP—Number of high-loss priority non-TCP packets dropped because of RED. • High, TCP—Number of high-loss priority TCP packets dropped because of RED. • (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"> • Low—Number of low-loss priority packets dropped because of RED. • Medium-low—Number of medium-low loss priority packets dropped because of RED. • Medium-high—Number of medium-high loss priority packets dropped because of RED. • High—Number of high-loss priority packets dropped because of RED.

Table 75: 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"> • (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"> • Low, non-TCP—Number of low-loss priority non-TCP bytes dropped because of RED. • Low, TCP—Number of low-loss priority TCP bytes dropped because of RED. • High, non-TCP—Number of high-loss priority non-TCP bytes dropped because of RED. • High, TCP—Number of high-loss priority TCP bytes dropped because of RED.
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"> • None—No rate limit. • exact—Queue transmits at the configured rate.
Buffer size	Delay buffer size in the queue.
Priority	Scheduling priority configured as low or high.
Excess Priority	Priority of the excess bandwidth traffic on a scheduler: low, medium-low, medium-high, high, or none.

Table 75: show class-of-service interface Output Fields (Continued)

Field Name	Field Description
Drop profiles	<p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> • Loss priority—Packet loss priority for drop profile assignment. • Protocol—Transport protocol for drop profile assignment. • Index—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • Name—Name of the drop profile. • Type—Type of the drop profile: discrete or interpolated. • Fill Level—Percentage fullness of a queue. • Drop probability—Drop probability at this fill level.
Excess Priority	Priority of the excess bandwidth traffic on a scheduler.
Drop profiles	<p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> • Loss priority—Packet loss priority for drop profile assignment. • Protocol—Transport protocol for drop profile assignment. • Index—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • Name—Name of the drop profile. • Type—Type of the drop profile: discrete or interpolated. • Fill Level—Percentage fullness of a queue. • Drop probability—Drop probability at this fill level.

Table 75: show class-of-service interface Output Fields (*Continued*)

Field Name	Field Description
Adjustment information	<p>Display the assignment of shaping-rate adjustments on a scheduler node or queue.</p> <ul style="list-style-type: none"> Adjusting application—Application that is performing the shaping-rate adjustment. <ul style="list-style-type: none"> The adjusting application can appear as ancp LS-0, which is the Junos OS Access Node Control Profile process (ancpd) that performs shaping-rate adjustments on schedule nodes. The adjusting application can appear as DHCP, which adjusts the shaping-rate and overhead-accounting class-of-service attributes based on DSL Forum VSA conveyed in 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). The adjusting application can also appear as pppoe, 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]. 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). Adjustment type—Type of adjustment: absolute or delta. Configured shaping rate—Shaping rate configured for the scheduler node or queue. Adjustment value—Value of adjusted shaping rate. Adjustment target—Level of shaping-rate adjustment performed: node or queue. Adjustment overhead-accounting mode—Configured shaping mode: frame or cell. Adjustment overhead bytes—Number of bytes that the ANCP agent adds to or subtracts from the actual downstream frame overhead before reporting the adjusted values to CoS. Adjustment target—Level of shaping-rate adjustment performed: node or queue. Adjustment multicast index—

Sample Output

show class-of-service interface (Physical)

```

user@host> show class-of-service interface et-1/0/4
Physical interface: et-1/0/4, Index: 1098
Maximum usable queues: 8, Queues in use: 4
Exclude aggregate overhead bytes: disabled
Logical interface aggregate statistics: disabled
  Scheduler map: default, Index: 0
  Congestion-notification: Disabled
  Dedicated Buffer Profile: dbp1
  Monitoring Profile Name: XYZ

Logical interface: et-1/0/4.16386, Index: 1057

```

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>		27
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	2726446535
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 Exclude Overhead Bytes: Disabled

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                                ID      Queue  Restricted queue  Fabric priority
Policing 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

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: Sendbroadcast-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	normal	low

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		exp-filter
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	ID	Queue	Restricted queue	Fabric priority
Policing 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				

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: Sendbroadcast-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
Policing 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	Name	Type	Index
Classifier	ipprec-compatibility	ip	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	Name	Type	Index
Rewrite	ri	ieee8021p (outer)	35392

Classifier	ci	ieee8021p	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	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Logical interface: ge-0/1/3.0, Index: 77

Object	Name	Type	Index
Rewrite	custom-exp2	exp (mpls-any)	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	Name	Type	Index
Classifier	ipprec-compatibility	ip	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	Name	Type	Index
Classifier	ipprec-compatibility	ip	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	Name	Type	Index
Classifier	ipprec-compatibility	ip	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	Name	Type	Index
Classifier	d1	dscp	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      Name      Type      Index
Classifier   ipprec-compatibility ip      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          ieee8021p         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

```

RELATED DOCUMENTATION

| *Verifying and Managing Junos OS Enhanced Subscriber Management*

show class-of-service scheduler-map

IN THIS SECTION

- [Syntax | 519](#)
- [Description | 519](#)
- [Options | 519](#)
- [Required Privilege Level | 519](#)

- [Output Fields | 519](#)
- [Sample Output | 522](#)
- [Release Information | 523](#)

Syntax

```
show class-of-service scheduler-map  
<name>
```

Description

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

Options

- none** Display all scheduler maps.
- name** (Optional) Display a summary of scheduler parameters for each forwarding class to which the named scheduler is assigned.

Required Privilege Level

view

Output Fields

[Table 76 on page 520](#) describes the output fields for the `show class-of-service scheduler-map` command. Output fields are listed in the approximate order in which they appear.

Table 76: show class-of-service scheduler-map Output Fields

Field Name	Field Description
Scheduler map	<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, SMAP-1_UID1002) instead of with a subscriber interface.</p>
Index	<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>
Scheduler	Name of the scheduler.
Forwarding class	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.
Transmit rate	Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword <code>remainder</code> , which indicates that the scheduler receives the remaining bandwidth of the interface.
Rate Limit	Rate limiting configuration of the queue. Possible values are <code>none</code> , meaning no rate limiting, and <code>exact</code> , meaning the queue only transmits at the configured rate.
Maximum buffer delay	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 <code>remainder</code> to indicate that the buffer is sized according to what remains after other scheduler buffer allocations.

Table 76: show class-of-service scheduler-map Output Fields (Continued)

Field Name	Field Description
Buffer size	<p>The size of the buffer as a percent of the total buffer size for the port. The total of all of the explicitly configured buffer size percentages for all of the queues on a port cannot exceed 100 percent.</p> <p>On QFX10000 switches, the buffer size is the amount of time in milliseconds of port bandwidth that a queue can use to continue to transmit packets during periods of congestion, before the buffer runs out and packets begin to drop.</p>
Buffer dynamic threshold	The alpha value of the shared-buffer pool at the egress buffer partition for each queue.
Priority	Scheduling priority: low or high.
Excess priority	Priority of excess bandwidth: low, medium-low, medium-high, high, or none.
Explicit Congestion Notification	<p>(QFX Series, OCX Series, and EX4600 switches only) Explicit congestion notification (ECN) state:</p> <ul style="list-style-type: none"> • Disable—ECN is disabled on the specified scheduler • Enable—ECN is enabled on the specified scheduler <p>ECN is disabled by default.</p>
Adjust minimum	Minimum shaping rate for an adjusted queue, in bps.
Adjust percent	Bandwidth adjustment applied to a queue, in percent.
Drop profiles	Table displaying the assignment of drop profiles by name and index to a given loss priority and protocol pair.
Loss priority	Packet loss priority for drop profile assignment.
Protocol	Transport protocol for drop profile assignment.

Table 76: show class-of-service scheduler-map Output Fields (*Continued*)

Field Name	Field Description
Name	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

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| *Verifying and Managing Junos OS Enhanced Subscriber Management*

show interfaces queue

IN THIS SECTION

- [Syntax | 524](#)
- [Description | 524](#)
- [Options | 524](#)
- [Overhead for Layer 2 Statistics | 525](#)
- [Additional Information | 527](#)
- [Required Privilege Level | 528](#)
- [Output Fields | 528](#)
- [Sample Output | 538](#)
- [Release Information | 582](#)

Syntax

```
show interfaces queue
<aggregate>
<both-ingress-egress>
<buffer-occupancy>
<egress>
<forwarding-class forwarding-class>
<ingress>
<interface-name>
<l2-statistics>
<remaining-traffic>
<slice slice-name>
```

Description

Display class-of-service (CoS) queue information for physical interfaces.

Options

NOTE: This topic lists all possible options for the show interfaces queue command. The options that appear vary depending on the platform, software release, and operating system (Junos OS or Junos OS Evolved).

none	Show detailed CoS queue statistics for all physical interfaces.
aggregate	Display the aggregated queuing statistics of all logical interfaces that have traffic-control profiles configured.
both-ingress-egress	Display both ingress and egress queue statistics.
buffer-occupancy	Displays the peak buffer occupancy for each queue while buffer-monitor-enable is enabled at the [edit chassis fpc slot-number traffic-manager] hierarchy level.
egress	Display egress queue statistics.
forwarding-class forwarding-class	Forwarding class name for this queue. Shows detailed CoS statistics for the queue associated with the specified forwarding class.

ingress	Display ingress queue statistics.
<i>interface-name</i>	Show detailed CoS queue statistics for the specified interface. The other listed options are available for the specified interface.
l2-statistics	Display Layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles
remaining-traffic	Display the queuing statistics of all logical interfaces that do not have traffic-control profiles configured.
slice <i>slice-name</i>	Display the hierarchical CoS statistics on the interface for the specified slice.

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 77 on page 525](#). 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 77 on page 525](#).

Table 77: Layer 2 Overhead and Transmitted Packets or Byte Counts

Protocol	Fragmentation		LFI
	First fragmentation	Second to n 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:

Framerelay header	: 2 bytes
Control,NLPID	: 2 bytes
Fragmentaion header	: 2 bytes
Inner proto	: 2 bytes
HDLC flag and FCS	: 4 bytes

Fragments 2 ...n :

Framerelay 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
Framerelay header	: 2 bytes
Inner proto	: 2 bytes
HDLC flag and FCS	: 4 bytes

```

Fragments 2 ...n :
  Fragmentation 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.

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 related CoS operational mode commands, see the [CLI Explorer](#).

Required Privilege Level

view

Output Fields

[Table 78 on page 529](#) lists the output fields for the `show interfaces queue` command. Output fields are listed in the approximate order in which they appear. The output fields that appear vary depending on the platform, software release, and operating system (Junos OS or Junos OS Evolved).

Table 78: show interfaces queue Output Fields

Field Name	Field Description
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Slice	Name of the slice.
Slice index	Slice's index number, which reflects its initialization sequence.
Forwarding classes supported	Total number of forwarding classes supported on the specified interface.
Forwarding classes in use	Total number of forwarding classes in use on the specified interface.
Ingress queues supported	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues supported on the specified interface.
Ingress queues in use	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues in use on the specified interface.
Output queues supported	Total number of output queues supported on the specified interface.
Output queues in use	Total number of output queues in use on the specified interface.

Table 78: show interfaces queue Output Fields (Continued)

Field Name	Field Description
Egress queues supported	Total number of egress queues supported on the specified interface.
Egress queues in use	Total number of egress queues in use on the specified interface.
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. <ul style="list-style-type: none"> Queued packets—Number of queued packets. Transmitted packets—Number of transmitted packets. Dropped packets—Number of packets dropped by the ASIC's RED mechanism.
Burst size	(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 `show interfaces queue` command:

Queue	Queue number.
Forwarding classes	Forwarding class name.

Table 78: show interfaces queue Output Fields (*Continued*)

Field Name	Field Description
Queued Packets	<p>Number of packets queued to this queue.</p> <p>NOTE: 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 "Additional Information" on page 527.</p> <p>NOTE: This field is not supported on QFX5100, QFX5110, QFX5200, and QFX5210 switches due to hardware limitations.</p> <p>This field is not supported on EX Series switches due to hardware limitations.</p>
Queued Bytes	<p>Number of bytes queued to this queue. The byte counts vary by interface hardware. For more information, see Table 79 on page 535.</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 "Additional Information" on page 527.</p> <p>NOTE: This field is not supported on QFX5100, QFX5110, QFX5200, and QFX5210 switches due to hardware limitations.</p> <p>This field is not supported on EX Series switches due to hardware limitations.</p>
Transmitted Packets	<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 Packet Forwarding Engine Chassis Queues field) shows the prefragmentation values.</p> <p>NOTE: For Layer 2 statistics, see "Overhead for Layer 2 Statistics" on page 525</p>

Table 78: show interfaces queue Output Fields (*Continued*)

Field Name	Field Description
Transmitted Bytes	<p>Number of bytes transmitted by this queue. The byte counts vary by interface hardware. For more information, see Table 79 on page 535.</p> <p>NOTE: 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>NOTE: For Layer 2 statistics, see "Overhead for Layer 2 Statistics" on page 525</p>
Tail-dropped packets	<p>Number of packets dropped because of tail drop.</p> <p>NOTE: Starting with Junos OS 18.3R1, the Tail-dropped packets counter is supported on PTX Series Packet Transport Routers.</p>
RL-dropped packets	<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 "Additional Information" on page 527.</p> <p>NOTE: The RL-dropped packets counter is not supported on the PTX Series Packet Transport Routers, and is omitted from the output.</p>
RL-dropped bytes	<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 "Additional Information" on page 527.</p>

Table 78: show interfaces queue Output Fields (*Continued*)

Field Name	Field Description
RED-dropped packets	<p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> • (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"> • Low, non-TCP—Number of low-loss priority non-TCP packets dropped because of RED. • Low, TCP—Number of low-loss priority TCP packets dropped because of RED. • High, non-TCP—Number of high-loss priority non-TCP packets dropped because of RED. • High, TCP—Number of high-loss priority TCP packets dropped because of RED. • (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"> • Low—Number of low-loss priority packets dropped because of RED. • Medium-low—Number of medium-low loss priority packets dropped because of RED. • Medium-high—Number of medium-high loss priority packets dropped because of RED. • High—Number of high-loss priority packets dropped because of RED. <p>NOTE: 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 78: 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 Table 79 on page 535.</p> <ul style="list-style-type: none"> • (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"> • Low, non-TCP—Number of low-loss priority non-TCP bytes dropped because of RED. • Low, TCP—Number of low-loss priority TCP bytes dropped because of RED. • High, non-TCP—Number of high-loss priority non-TCP bytes dropped because of RED. • High, TCP—Number of high-loss priority TCP bytes dropped because of RED. <p>NOTE: 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 the amount of queue buffer that is in occupation at this instance. This is an indicator of the amount of data that is present in a queue at that point in time. The amount of data present is in the units of bytes.
Peak	(QFX5000 Series switches only) Displays the peak buffer occupancy for the queue while <code>buffer-monitor-enable</code> is enabled at the <code>[edit chassis fpc slot-number traffic-manager]</code> hierarchy level.
Last-packet enqueued	Starting with Junos OS Release 16.1, Last-packet enqueued output field is introduced. If <code>packet-timestamp</code> 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 79 on page 535](#) shows how the byte counts on the outbound interfaces vary depending on the interface hardware. [Table 79 on page 535](#) is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.

Table 79: 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>	-

Table 79: Byte Count by Interface Hardware *(Continued)*

Interface Hardware	Output Level	Byte Count Includes	Comments
Non-IQ PIC	Interface	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet. Transmitted: 478 bytes of Layer 3 packet. <p>T4000 routers with Type 5 FPCs :</p> <ul style="list-style-type: none"> 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. 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. <p>M Series routers:</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet. Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead. <p>PTX Series Packet Transport Routers:</p> <ul style="list-style-type: none"> Queued: The sum of the transmitted bytes and the RED dropped bytes. Transmitted: Full Layer 2 overhead (including all L2 encapsulation and CRC) + 12 inter-packet gap + 8 for the preamble. 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). 	The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.

Table 79: 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"> • Queued: 478 bytes of Layer 3 packet. • Transmitted: 478 bytes of Layer 3 packet. <p>M Series routers:</p> <ul style="list-style-type: none"> • Queued: 478 bytes of Layer 3 packet. • Transmitted: 483 bytes per packet, representing 478 bytes of Layer 3 packet + 5 bytes • RED dropped: 478 bytes per packet, representing 478 bytes of Layer 3 packet 	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.	

Table 79: Byte Count by Interface Hardware (Continued)

Interface Hardware	Output Level	Byte Count Includes	Comments
1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs	Interface	Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.	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		Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.	
8-port 1G IQ2 and IQ2-E PICs	Packet forwarding component	Queued: 478 bytes of Layer 3 packet. 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  :                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	:	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 : 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 :          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 : 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

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
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

```

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 interfaces queue ge-0/0/0 (EX2200 Switch)

```

user@switch> show interfaces queue ge-0/0/0
Physical interface: ge-0/0/0, Enabled, Physical link is Down
  Interface index: 130, SNMP ifIndex: 501
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
  Transmitted:
    Packets      :      0
    Bytes        :      0
    Tail-dropped packets :      0
Queue: 1, Forwarding classes: assured-forwarding
  Queued:
  Transmitted:
    Packets      :      0
    Bytes        :      0
    Tail-dropped packets :      0
Queue: 5, Forwarding classes: expedited-forwarding
  Queued:
  Transmitted:
    Packets      :      0
    Bytes        :      0

```

```

    Tail-dropped packets :          0
Queue: 7, Forwarding classes: network-control
  Queued:
  Transmitted:
    Packets      :          0
    Bytes        :          0
    Tail-dropped packets :          0

```

show interfaces queue xe-6/0/39 (Line Card with Oversubscribed Ports in an EX8200 Switch)

```

user@switch> show interfaces queue xe-6/0/39
Physical interface: xe-6/0/39, Enabled, Physical link is Up
  Interface index: 291, SNMP ifIndex: 1641
Forwarding classes: 16 supported, 7 in use
Ingress queues: 1 supported, 1 in use
  Transmitted:
    Packets      :      337069086018
    Bytes        :      43144843010304
    Tail-dropped packets :      8003867575
PFE chassis queues: 1 supported, 1 in use
  Transmitted:
    Packets      :          0
    Bytes        :          0
    Tail-dropped packets :          0
Forwarding classes: 16 supported, 7 in use
Egress queues: 8 supported, 7 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
  Transmitted:
    Packets      :      334481399932
    Bytes        :      44151544791024
    Tail-dropped packets :          0
Queue: 1, Forwarding classes: assured-forwarding
  Queued:
  Transmitted:
    Packets      :          0
    Bytes        :          0
    Tail-dropped packets :          0
Queue: 2, Forwarding classes: mcast-be
  Queued:
  Transmitted:

```

```

Packets          :          274948977
Bytes            :          36293264964
Tail-dropped packets :          0
Queue: 4, Forwarding classes: mcast-ef
Queued:
Transmitted:
Packets          :          0
Bytes            :          0
Tail-dropped packets :          0
Queue: 5, Forwarding classes: expedited-forwarding
Queued:
Transmitted:
Packets          :          0
Bytes            :          0
Tail-dropped packets :          0
Queue: 6, Forwarding classes: mcast-af
Queued:
Transmitted:
Packets          :          0
Bytes            :          0
Tail-dropped packets :          0
Queue: 7, Forwarding classes: network-control
Queued:
Transmitted:
Packets          :          46714
Bytes            :          6901326
Tail-dropped packets :          0

Packet Forwarding Engine Chassis Queues:
Queues: 8 supported, 7 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Transmitted:
Packets          :          739338141426
Bytes            :          94635282101928
Tail-dropped packets :          0
RED-dropped packets :          5606426444
Low              :          5606426444
High             :          0
RED-dropped bytes  :          683262846464
Low              :          683262846464
High             :          0
Queue: 1, Forwarding classes: assured-forwarding

```

Queued:

Transmitted:

Packets	:	0
Bytes	:	0
Tail-dropped packets	:	0
RED-dropped packets	:	0
Low	:	0
High	:	0
RED-dropped bytes	:	0
Low	:	0
High	:	0

Queue: 2, Forwarding classes: mcast-be

Queued:

Transmitted:

Packets	:	0
Bytes	:	0
Tail-dropped packets	:	0
RED-dropped packets	:	0
Low	:	0
High	:	0
RED-dropped bytes	:	0
Low	:	0
High	:	0

Queue: 4, Forwarding classes: mcast-ef

Queued:

Transmitted:

Packets	:	0
Bytes	:	0
Tail-dropped packets	:	0
RED-dropped packets	:	0
Low	:	0
High	:	0
RED-dropped bytes	:	0
Low	:	0
High	:	0

Queue: 5, Forwarding classes: expedited-forwarding

Queued:

Transmitted:

Packets	:	0
Bytes	:	0
Tail-dropped packets	:	0
RED-dropped packets	:	0
Low	:	0

```

    High           :           0
    RED-dropped bytes :           0
    Low            :           0
    High           :           0

```

Queue: 6, Forwarding classes: mcast-af

Queued:

Transmitted:

```

    Packets           :           0
    Bytes             :           0
    Tail-dropped packets :           0
    RED-dropped packets :           0
    Low               :           0
    High              :           0
    RED-dropped bytes  :           0
    Low               :           0
    High              :           0

```

Queue: 7, Forwarding classes: network-control

Queued:

Transmitted:

```

    Packets           :          97990
    Bytes             :         14987506
    Tail-dropped packets :           0
    RED-dropped packets :           0
    Low               :           0
    High              :           0
    RED-dropped bytes  :           0
    Low               :           0
    High              :           0

```

show interfaces queue xe-0/0/2 buffer-occupancy (QFX5000 Series switch)

```
user@switch> show interfaces queue xe-0/0/2 buffer-occupancy
```

Physical interface: xe-0/0/2, Enabled, Physical link is Up

Interface index: 689, SNMP ifIndex: 514

Forwarding classes: 12 supported, 5 in use

Egress queues: 12 supported, 5 in use

Queue: 0, Forwarding classes: fc0

Queue-depth bytes :

Peak : 1786720

Queue: 3, Forwarding classes: fcoe

Queue-depth bytes :

```

Peak          : 0
Queue: 4, Forwarding classes: no-loss
Queue-depth bytes :
Peak          : 0
Queue: 7, Forwarding classes: network-control
Queue-depth bytes :
Peak          : 416
Queue: 8, Forwarding classes: mcast
Queue-depth bytes :
Peak          : 0

```

show interfaces queue xe-3/0/3 slice slice1

```

user@router show interfaces queue xe-3/0/3 slice slice1
Physical interface: xe-3/0/3, Enabled, Physical link is Up
  Interface index: 275, SNMP ifIndex: 776
  Slice name: slice1, Slice index: 5
Forwarding classes: 16 supported, 7 in use
Egress queues: 8 supported, 7 in use
CoS scheduler resource information:
  Maximum units supported per MIC/PIC: 84
  Configured units per MIC/PIC: 1
  Maximum units allowed per port: 21
  Configured units on this port: 1
Queue: 0, Forwarding classes: BE
Queued:
  Packets          :          4          0 pps
  Bytes            :         200          0 bps
Transmitted:
  Packets          :          4          0 pps
  Bytes            :         200          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	:	120061952	
. . .			

Release Information

Command introduced before Junos OS Release 7.4.

both-ingress-egress, egress, and ingress options introduced in Junos OS Release 7.6.

l2-statistics option introduced in Junos OS Release 12.1.

buffer-occupancy statement introduced in Junos OS Release 19.1R1 for QFX5000 Series switches.

Release History Table

Release	Description
22.3R1	Starting with Junos OS 22.3R1, the slice <i>slice-name</i> option is supported on MX Series Routers.
18.3R1	Starting with Junos OS 18.3R1, the Tail-dropped packets counter is supported on PTX Series Packet Transport Routers.
16.1	Starting with Junos OS Release 16.1, Last-packet enqueued output field is introduced.

RELATED DOCUMENTATION

[Monitoring Interface Status and Traffic](#)

Monitoring Interfaces That Have CoS Components

Defining CoS Schedulers and Scheduler Maps (CLI Procedure)

Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure)

[ATM Interfaces User Guide for Routing Devices](#)

[Layer 2 Bridging, Address Learning, and Forwarding User Guide](#)

show interfaces voq

IN THIS SECTION

- [Syntax | 583](#)
- [Syntax \(Junos OS Evolved\) | 583](#)
- [Description | 583](#)
- [Options | 584](#)
- [Additional Information | 584](#)
- [Required Privilege Level | 584](#)
- [Output Fields | 584](#)
- [Sample Output | 586](#)
- [Release Information | 601](#)

Syntax

```
show interfaces voq interface-name
<forwarding-class forwarding-class-name>
<non-zero>
```

Syntax (Junos OS Evolved)

```
show interfaces voq interface-name
<forwarding-class forwarding-class-name>
<non-zero>
<source-fpc source-fpc-number>
```

Description

Display the random early detection (RED) drop statistics from all ingress Packet Forwarding Engines associated with the specified physical egress interface. In the VOQ architecture, egress output queues (shallow buffers) buffer data in virtual queues on ingress Packet Forwarding Engines. In cases of

congestion, you can use this command to identify which ingress Packet Forwarding Engine is the source of RED-dropped packets contributing to congestion.

NOTE: On the PTX Series routers and QFX10000 switches, these statistics include tail-dropped packets.

Options

interface <i>interface-name</i>	Display the ingress VOQ RED drop statistics for the specified egress interface.
forwarding-class <i>forwarding-class-name</i>	Display VOQ RED drop statistics for a specified forwarding class.
non-zero	Display only non-zero VOQ RED drop statistics counters.
source-fpc <i>source-fpc-number</i>	Display VOQ RED drop statistics for the specified source FPC.

Additional Information

- On PTX Series routers, you can display VOQ statistics for only the WAN physical interface.
- VOQ statistics for aggregated physical interfaces are not supported. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can use the `show interfaces queue` command to identify the child link which is experiencing congestion and then view the VOQ statistics on the respective child link using the `show interfaces voq` command.

For information on virtual output queuing on PTX routers, see *Understanding Virtual Output Queues on PTX Series Packet Transport Routers*. For information on virtual output queueing on QFX10000 switches, see *Understanding CoS Virtual Output Queues (VOQs) on QFX10000 Switches*.

Required Privilege Level

view

Output Fields

Table 80 on page 585 lists the output fields for the `show interfaces queue` command. Output fields are listed in the approximate order in which they appear.

Table 80: show interfaces voq Output Fields

Field Name	Field Description
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Queue	Egress queue number.
Forwarding classes	Forwarding class name.
FPC number	Number of the Flexible PIC Concentrator (FPC) located on ingress.
PFE	Number of the Packet Forwarding Engine providing virtual output queues on the ingress.
RED-dropped packets	<p>Number of packets per second (pps) dropped because of random early detection (RED).</p> <p>NOTE: On the PTX Series routers, these statistics include tail-dropped packets.</p>
RED-dropped bytes	<p>Number of bytes per second dropped because of RED. The byte counts vary by interface hardware.</p> <p>NOTE: On the PTX Series routers, these statistics include tail-dropped packets.</p>

Sample Output

show interfaces voq (For a Specific Physical Interface) (PTX Series Routers)

The following example shows ingress RED-dropped statistics for the egress Ethernet interface configured on port 0 of Physical Interface Card (PIC) 0, located on the FPC in slot 7.

The sample output below shows that the cause of the congestion is ingress Packet Forwarding Engine PFE 0, which resides on FPC number 4, as denoted by the count of RED-dropped packets and RED-dropped bytes for egress queue 0, forwarding classes best-effort and egress queue 3, forwarding class network control.

```

user@host> show interfaces voq et-7/0/0
Physical interface: et-7/0/0, Enabled, Physical link is Up
  Interface index: 155, SNMP ifIndex: 699

Queue: 0, Forwarding classes: best-effort

FPC number: 1
  PFE: 0
    RED-dropped packets :                0                0 pps
    RED-dropped bytes   :                0                0 bps
  PFE: 1
    RED-dropped packets :                0                0 pps
    RED-dropped bytes   :                0                0 bps
  PFE: 2
    RED-dropped packets :                0                0 pps
    RED-dropped bytes   :                0                0 bps
  PFE: 3
    RED-dropped packets :                0                0 pps
    RED-dropped bytes   :                0                0 bps

FPC number: 4
  PFE: 0
    RED-dropped packets :           19969426           2323178 pps
    RED-dropped bytes   :           2196636860       2044397464 bps
  PFE: 1
    RED-dropped packets :                0                0 pps
    RED-dropped bytes   :                0                0 bps
  PFE: 2
    RED-dropped packets :                0                0 pps
    RED-dropped bytes   :                0                0 bps

```

PFE: 3

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

FPC number: 6

PFE: 0

RED-dropped packets :	19969424	2321205 pps
RED-dropped bytes :	2196636640	2042660808 bps

PFE: 1

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 2

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 3

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 4

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 5

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 6

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 7

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

FPC number: 7

PFE: 0

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 1

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 2

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 3

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

Queue: 1, Forwarding classes: expedited-forwarding

FPC number: 1

PFE: 0

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 1

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 2

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 3

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

FPC number: 4

PFE: 0

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 1

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 2

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 3

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

FPC number: 6

PFE: 0

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 1

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 2

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 3

RED-dropped packets : 0 0 pps

RED-dropped bytes	:	0	0 bps
PFE: 4			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 5			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 6			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 7			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

FPC number: 7

PFE: 0			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 1			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 2			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 3			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Queue: 2, Forwarding classes: assured-forwarding

FPC number: 1

PFE: 0			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 1			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 2			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
PFE: 3			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

FPC number: 4

PFE: 0

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 1

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 2

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 3

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

FPC number: 6

PFE: 0

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 1

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 2

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 3

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 4

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 5

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 6

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

PFE: 7

RED-dropped packets : 0 0 pps

RED-dropped bytes : 0 0 bps

FPC number: 7

PFE: 0

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps
PFE: 1		
RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps
PFE: 2		
RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps
PFE: 3		
RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

Queue: 3, Forwarding classes: network-control

FPC number: 1

PFE: 0		
RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps
PFE: 1		
RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps
PFE: 2		
RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps
PFE: 3		
RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

FPC number: 4

PFE: 0		
RED-dropped packets :	16338670	1900314 pps
RED-dropped bytes :	1797253700	1672276976 bps
PFE: 1		
RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps
PFE: 2		
RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps
PFE: 3		
RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

FPC number: 6

PFE: 0

RED-dropped packets :	16338698	1899163 pps
RED-dropped bytes :	1797256780	1671263512 bps

PFE: 1

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 2

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 3

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 4

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 5

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 6

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 7

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

FPC number: 7

PFE: 0

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 1

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 2

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

PFE: 3

RED-dropped packets :	0	0 pps
RED-dropped bytes :	0	0 bps

show interfaces voq (For a Specific Physical Interface) (QFX10000 Switches)

The sample output below shows congestion on ingress PFE 1 on FPC number 0, and on ingress PFE 2 on FPC number 1, as denoted by the count of RED-dropped packets and RED-dropped bytes for best-effort egress queue 0.

```

user@host> show interfaces voq et-1/0/0
Physical interface: et-1/0/0, Enabled, Physical link is Up
  Interface index: 659, SNMP ifIndex: 539

Queue: 0, Forwarding classes: best-effort

FPC number: 0
  PFE: 0
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
  PFE: 1
    RED-dropped packets :    411063248    16891870 pps
    RED-dropped bytes   :    52616095744    17297275600 bps
  PFE: 2
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps

FPC number: 1
  PFE: 0
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
  PFE: 1
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
  PFE: 2
    RED-dropped packets :    411063012    16891870 pps
    RED-dropped bytes   :    52616065536    17297275376 bps

Queue: 3, Forwarding classes: fcoe

FPC number: 0
  PFE: 0
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
  PFE: 1
    RED-dropped packets :          0          0 pps

```

```

    RED-dropped bytes      :                0                0 bps
PFE: 2
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps

```

FPC number: 1

```

PFE: 0
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
PFE: 1
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
PFE: 2
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps

```

Queue: 4, Forwarding classes: no-loss

FPC number: 0

```

PFE: 0
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
PFE: 1
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
PFE: 2
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps

```

FPC number: 1

```

PFE: 0
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
PFE: 1
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
PFE: 2
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps

```

Queue: 7, Forwarding classes: network-control

FPC number: 0

```

PFE: 0
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 1
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 2
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps

```

FPC number: 1

```

PFE: 0
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 1
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 2
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps

```

show interfaces voq et-7/0/0 (For a Specific Forwarding Class)

```
user@host> show interfaces voq et-7/0/0 forwarding-class best-effort
```

Physical interface: et-7/0/0, Enabled, Physical link is Up

Interface index: 155, SNMP ifIndex: 699

Queue: 0, Forwarding classes: best-effort

FPC number: 1

```

PFE: 0
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 1
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 2
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 3
  RED-dropped packets :          0          0 pps

```

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

FPC number: 4

PFE: 0

RED-dropped packets	:	66604786	2321519 pps
---------------------	---	----------	-------------

RED-dropped bytes	:	7326526460	2042936776 bps
-------------------	---	------------	----------------

PFE: 1

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

PFE: 2

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

PFE: 3

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

FPC number: 6

PFE: 0

RED-dropped packets	:	66604794	371200 pps
---------------------	---	----------	------------

RED-dropped bytes	:	7326527340	326656000 bps
-------------------	---	------------	---------------

PFE: 1

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

PFE: 2

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

PFE: 3

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

PFE: 4

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

PFE: 5

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

PFE: 6

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

PFE: 7

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

FPC number: 7

```

PFE: 0
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 1
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 2
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 3
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps

```

show interfaces voq et-5/0/12 (For a Specific Source FPC)

```

user@host> show interfaces voq et-5/0/12 source-fpc 0
Physical interface: et-5/0/12, Enabled, Physical link is Up
  Interface index: 166, SNMP ifIndex: 1104

```

Queue: 0, Forwarding classes: best-effort

FPC number: 0

```

PFE: 0
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 1
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 2
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 3
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps

```

Queue: 1, Forwarding classes: expedited-forwarding

FPC number: 0

```

PFE: 0
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps

```

```

PFE: 1
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 2
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 3
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps

```

Queue: 2, Forwarding classes: assured-forwarding

```

FPC number: 0
PFE: 0
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 1
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 2
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 3
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps

```

Queue: 3, Forwarding classes: network-control

```

FPC number: 0
PFE: 0
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 1
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 2
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps
PFE: 3
  RED-dropped packets :          0          0 pps
  RED-dropped bytes   :          0          0 bps

```

show interfaces voq et-5/0/12 (For a Specific Forwarding Class and Source FPC)

```

user@host> show interfaces voq et-5/0/12 forwarding-class best-effort source-fpc 5
Physical interface: et-5/0/12, Enabled, Physical link is Up
    Interface index: 166, SNMP ifIndex: 1104

    Queue: 0, Forwarding classes: best-effort

    FPC number: 5
    PFE: 0
        RED-dropped packets :                0                0 pps
        RED-dropped bytes   :                0                0 bps
    PFE: 1
        RED-dropped packets :                0                0 pps
        RED-dropped bytes   :                0                0 bps
    PFE: 2
        RED-dropped packets :                0                0 pps
        RED-dropped bytes   :                0                0 bps
    PFE: 3
        RED-dropped packets :                0                0 pps
        RED-dropped bytes   :                0                0 bps
    PFE: 4
        RED-dropped packets :                0                0 pps
        RED-dropped bytes   :                0                0 bps
    PFE: 5
        RED-dropped packets :                0                0 pps
        RED-dropped bytes   :                0                0 bps
    PFE: 6
        RED-dropped packets :                0                0 pps
        RED-dropped bytes   :                0                0 bps
    PFE: 7
        RED-dropped packets :                0                0 pps
        RED-dropped bytes   :                0                0 bps

```

show interfaces voq et-7/0/0 (Non-Zero)

```

user@host> show interfaces voq et-7/0/0 non-zero

Physical interface: et-7/0/0, Enabled, Physical link is Up

```


Interface index: 155, SNMP ifIndex: 699

Queue: 0, Forwarding classes: best-effort

FPC number: 4

PFE: 0

RED-dropped packets :	95862238	2301586 pps
RED-dropped bytes :	10544846180	2025396264 bps

FPC number: 6

PFE: 0

RED-dropped packets :	95866639	2322569 pps
RED-dropped bytes :	10545330290	2043860728 bps

Queue: 3, Forwarding classes: network-control

FPC number: 4

PFE: 0

RED-dropped packets :	78433066	1899727 pps
RED-dropped bytes :	8627637260	1671760384 bps

FPC number: 6

PFE: 0

RED-dropped packets :	78436704	1900628 pps
RED-dropped bytes :	8628037440	1672553432 bps

show interfaces voq et-7/0/0 (For a Specific Forwarding Class and Non-Zero)

user@host **show interfaces voq et-7/0/0 forwarding-class best-effort non-zero**

Physical interface: et-7/0/0, Enabled, Physical link is Up

Interface index: 155, SNMP ifIndex: 699

Queue: 0, Forwarding classes: best-effort

FPC number: 4

PFE: 0

RED-dropped packets :	119540012	2322319 pps
RED-dropped bytes :	13149401320	2043640784 bps

FPC number: 6

PFE: 0			
RED-dropped packets	:	119540049	2322988 pps
RED-dropped bytes	:	13149405390	2044229744 bps

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

- Understanding Virtual Output Queues on PTX Series Packet Transport Routers*
- Understanding CoS Virtual Output Queues (VOQs) on QFX10000 Switches*