

Junos[®] OS for EX Series Ethernet Switches

Class of Service User Guide (EX Series Switches Except EX4600 and EX9200 Switches)

Published 2023-06-13



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.

Junos® OS for EX Series Ethernet Switches Class of Service User Guide (EX Series Switches Except EX4600 and EX9200 Switches) 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 | x

1	
_	L

Basic CoS Configuration

CoS Overview | 2

Junos OS CoS for EX Series Switches Overview | 2

Configuring CoS on EX Series Switches | 4

Understanding Junos OS CoS Components for EX Series Switches | 8

Understanding CoS Two-Color Marking | 10

Understanding CoS Queues on EX8200 Line Cards That Include Oversubscribed Ports | 11

Example: Configuring CoS on EX Series Switches | 15

Requirements | 15 Overview and Topology | 15 Configuration | 19 Verification | 35

Junos OS EZQoS and J-Web | 51

Understanding Junos OS EZQoS for CoS Configurations on EX Series Switches | 51

Configuring Junos OS EZQoS for CoS (CLI Procedure) | 52

Configuring CoS (J-Web Procedure) | 53

CoS on Interfaces | 55

Assigning CoS Components to Interfaces (CLI Procedure) | 55

Assigning CoS Components to Interfaces (J-Web Procedure) | 56

Monitoring Interfaces That Have CoS Components | 58

CoS Code-Point Aliases | 61

Understanding CoS Code-Point Aliases | 61

Defining CoS Code-Point Aliases (CLI Procedure) | 65

Defining CoS Code-Point Aliases (J-Web Procedure) | 66

Monitoring CoS Value Aliases | 67

CoS Classifiers | 70

Understanding CoS Classifiers | 70

Defining CoS Classifiers (CLI Procedure) | 74

Defining CoS Classifiers (J-Web Procedure) | 77

Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers | 80

Requirements | 81

Overview | 81

Verification | 82

Configuring and Applying IEEE 802.1ad Classifiers | 83

Configuring the IEEE 802.1p Field for CoS Host Outbound Traffic | 85

Configuring a Global Default IEEE 802.1p Value for All Host Outbound Traffic | 86

Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure) | 87

Monitoring CoS Classifiers | 88

Troubleshooting a CoS Classifier Configuration for a TCAM Space Error | 91

CoS Rewrite | 95

Understanding CoS Rewrite Rules | 95

Defining CoS Rewrite Rules (CLI Procedure) | 97

Defining CoS Rewrite Rules (J-Web Procedure) | 100

Classifiers and Rewrite Rules at the Global, Physical, and Logical Interface Levels Overview | 102

Configuring Classifiers and Rewrite Rules at the Global and Physical Interface Levels | 103

Applying Egress Interface Rewrite Rules to the IEEE 802.1p Field for All Host Outbound Traffic on the Interface | **105**

Monitoring CoS Rewrite Rules | 107

Forwarding Classes | 110

Understanding CoS Forwarding Classes | 110

Defining CoS Forwarding Classes (CLI Procedure) | 114

Defining CoS Forwarding Classes (J-Web Procedure) | 115

Monitoring CoS Forwarding Classes | 117

Flow Control | 120

Understanding Priority-Based Flow Control | 120

Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure) | 123

CoS Queue Schedulers and Scheduler Maps | 128

Understanding CoS Schedulers | 128

Defining CoS Schedulers and Scheduler Maps (CLI Procedure) | 138

Configuring a Scheduler and a Scheduler Map | 138

Assigning a Scheduler Map to Interfaces | 139

Assigning Scheduler Maps to Interfaces on EX8200 Line Cards That Include Oversubscribed Ports | 140

Defining CoS Schedulers (J-Web Procedure) | 141

Defining CoS Scheduler Maps (J-Web Procedure) | 146

Monitoring CoS Scheduler Maps | 147

Troubleshooting CoS Schedulers on a 40-port SFP+ Line Card in an EX8200 Switch | 150

Congestion Management, Tail Drop Profiles, Queue Shaping, and Explicit Congestion Notification (ECN)

Congestion Management | 154

2

Understanding CoS Congestion Management | 154

Configuring CoS Congestion Management (CLI Procedure) | 160

Configuring a Weighted Tail Drop Profile | 161

Configuring a Weighted Random Early Detection Drop Profile | 161

Tail Drop Profiles | 164

Understanding CoS Tail Drop Profiles | 164

Configuring CoS Tail Drop Profiles (CLI Procedure) | 165

Defining CoS Drop Profiles (J-Web Procedure) | 165

Monitoring CoS Drop Profiles | 167

Queue Shaping | 169

Understanding Port Shaping and Queue Shaping for CoS | 169

Configuring Shaping for CoS (CLI Procedure) | 171

Configuring Port Shaping for CoS on an EX Series Switch | **171**

Configuring Queue Shaping for CoS on an EX Series Switch | 171

Applying a Shaping Rate to Physical Interfaces Overview | 173

Configuring the Shaping Rate for Physical Interfaces | 173

Explicit Congestion Notification (ECN) | 175

Understanding CoS Explicit Congestion Notification | 175

Example: Configuring ECN | 185

Requirements | 185 Overview | 185 Configuration | 188 Verification | 191

CoS on Overlay Networks

CoS on MPLS Networks | 194

Understanding Using CoS with MPLS Networks on EX Series Switches | 194

Example: Combining CoS with MPLS on EX Series Switches | 199

Requirements | 200 Overview and Topology | 200 Configuring the Local PE Switch | 203 Configuring the Remote PE Switch | 207 Configuring the Provider Switch | 208

Verification | 210

Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS | 217

Configuring CoS | **217** Configuring an LSP Policer | **218**

Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect | 219

Configuring CoS | 220 Configuring an LSP Policer | 220

3

Configuring MPLS on EX8200 and EX4500 Provider Switches | 223

CoS on EVPN VXLANs | 225

CoS Support on EVPN VXLANs | 225

Understanding CoS on VXLAN Interfaces | 225 Configuring CoS on VXLAN Interfaces | 227 Implementing CoS on VXLAN Interfaces (Junos OS Evolved) | 230 CoS Limitations on VXLANs | 231

Configuration Statements and Operational Commands

4

Configuration Statements | 234 broadcast | 235 buffer-size | 237 class | 242 class-of-service | 244 classifiers | 247 code-point (Congestion Notification) | 249 code-point-aliases | 251 code-points | 252 congestion-notification-profile | 253 drop-probability (Fill Level) | 257 drop-profile-map | 259 dscp | 260 dscp-ipv6 | 262 ethernet (CoS for Multidestination Traffic) | 264 excess-rate (Schedulers) | 265 exp | 267 explicit-congestion-notification | 268

family | 270

flow-control | 271

forwarding-class (Forwarding Policy) | 273

forwarding-classes | 275

ieee-802.1 | 281

ieee-802.1 (Congestion Notification) | 282

import | 284

inet (CoS) | 285

inet6 (CoS Multidestination) | 287

inet-precedence | 288

input (Congestion Notification) | 290

interfaces | 291

loss-priority (Classifiers and Rewrite Rules) | 293

multi-destination | 295

policing | 297

priority (Schedulers) | 298

protocol (Drop Profiles) | 300

rewrite-rules | 301

scheduler-map | 302

scheduler-maps | 304

schedulers (CoS) | 305

shaping-rate | 307

shaping-rate (Applying to an Interface) | 309

shared-buffer | 312

transmit-rate (EX Series Switches) | 313

tri-color (EX Series Switches) | 315

unit | 318

Operational Commands | 321 show class-of-service | 322 show class-of-service classifier | 330 show class-of-service code-point-aliases | 333 show class-of-service drop-profile | 336 show class-of-service forwarding-class | 340 show class-of-service forwarding-table | 345 show class-of-service forwarding-table classifier | 350 show class-of-service forwarding-table classifier mapping | 352 show class-of-service forwarding-table drop-profile | 355 show class-of-service forwarding-table rewrite-rule | 357 show class-of-service forwarding-table rewrite-rule mapping | 359 show class-of-service forwarding-table scheduler-map | 361 show class-of-service interface | 365 show class-of-service multi-destination | 408 show class-of-service rewrite-rule | 411 show class-of-service scheduler-map | 414 show interfaces queue | 418 show pfe statistics traffic cpu | 478 show pfe statistics traffic egress-queues | 483 show pfe statistics traffic multicast | 485

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. This guide applies to all EX Series switches except the EX4600 and the EX9200 lines of switches.



Basic CoS Configuration

CoS Overview | 2 Junos OS EZQoS and J-Web | 51

CoS on Interfaces | 55

CoS Code-Point Aliases | 61

CoS Classifiers | 70

CoS Rewrite | 95

Forwarding Classes | 110

Flow Control | 120

CoS Queue Schedulers and Scheduler Maps | 128

CoS Overview

IN THIS CHAPTER

- Junos OS CoS for EX Series Switches Overview | 2
- Configuring CoS on EX Series Switches | 4
- Understanding Junos OS CoS Components for EX Series Switches | 8
- Understanding CoS Two-Color Marking | 10
- Understanding CoS Queues on EX8200 Line Cards That Include Oversubscribed Ports | 11
- Example: Configuring CoS on EX Series Switches | 15

Junos OS CoS for EX Series Switches Overview

IN THIS SECTION

- How Junos OS CoS Works | 3
- Default CoS Behavior on EX Series Switches | 4

When a network experiences congestion and delay, some packets must be dropped. Junos operating system (Junos OS) *class of service* (CoS) divides traffic into classes to which you can apply different levels of throughput and packet loss when congestion occurs. This allows packet loss to happen according to rules that you configure.

For interfaces that carry IPv4, IPv6, and MPLS traffic, you can configure Junos OS CoS features to provide multiple classes of service for different applications. CoS also allows you to rewrite the Differentiated Services code point (DSCP), IP precedence, 802.1p, or EXP CoS bits of packets egressing out of an interface, thus allowing you to tailor packets for the remote peers' network requirements. See *Understanding Using CoS with MPLS Networks on EX Series Switches* for more information about CoS for MPLS networks.

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, and schedule the transmission service level for each queue.

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

Because Juniper Networks EX Series Ethernet Switches implement CoS 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 interface of an EX Series switch. Also, each physical and *logical interface* on the switch can have custom CoS rules associated with it. When CoS is used in an MPLS network, there are some additional restrictions. See *Understanding Using CoS with MPLS Networks on EX Series Switches*.

How Junos OS CoS Works

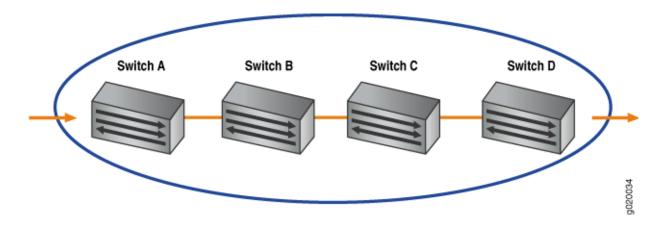
Junos OS CoS works by examining traffic entering at the edge of your network. The switches classify traffic into defined service groups to provide the special treatment of traffic across the network. For example, voice traffic can be sent across certain links, and data traffic can use other links. In addition, the data traffic streams can be serviced differently along the network path. As the traffic leaves the network at the far edge, you can rewrite the traffic to meet the policies of the targeted peer.

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 then 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.

Figure 1 on page 4 represents the network scenario of an enterprise. Switch A is receiving traffic from various network nodes such as desktop computers, servers, surveillance cameras, and VoIP telephones. As each packet enters, Switch A examines the packet's CoS settings and classifies the traffic into one of the groupings defined by the enterprise. This definition allows Switch A to prioritize resources for servicing the traffic streams it receives. Switch A might alter the CoS settings of the packets to better match the enterprise's traffic groups.

When Switch B receives the packets, it examines the CoS settings, determines the appropriate traffic groups, and processes the packets 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 rewrite the CoS bits of the packets before transmitting them.

Figure 1: Packet Flow Across the Network



Default CoS Behavior on EX Series Switches

If you do not configure any CoS settings on the switch, the software still ensures that user traffic and protocol packets are forwarded 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

Understanding Junos OS CoS Components for EX Series Switches | 8 Understanding Junos OS EZQoS for CoS Configurations on EX Series Switches | 51 Example: Configuring CoS on EX Series Switches | 15 Example: Combining CoS with MPLS on EX Series Switches

Configuring CoS on EX Series Switches

The topics in this guide 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.
- Provide congestion management with tail drop profiles, queue shaping, and congestion notification.
- Configure CoS on MPLS networks.
- 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 6 lists the primary CoS configuration tasks, includes platform limitations, and provides links to those tasks.

Table 1: CoS Configuration Tasks

CoS Configuration Task

Basic CoS Configuration:

- Configure CoS using EZQoS with templates for key traffic classes, or a browser and the J-Web interface.
- 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.
 - 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.
 - 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.
- Configure forwarding classes.
- Configure priority-based flow control to apply link-level flow control on a specific traffic class so that different types of traffic can efficiently use the same network interface card (NIC).
- Configure CoS schedulers to define the properties of output queues on EX Series switches. These properties include the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, the priority of the queue, and the drop profiles associated with the queue.
- Assign the following CoS components to physical or logical interfaces:
 - Classifiers (logical interfaces only)

Links

- "Configuring Junos OS EZQoS for CoS (CLI Procedure)" on page 52
- "Configuring CoS (J-Web Procedure)" on page 53
- "Defining CoS Code-Point Aliases (CLI Procedure)" on page 65
- "Defining CoS Classifiers (CLI Procedure)" on page 74
- (EX4300 and EX4600 only) "Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers" on page 80
- "Defining CoS Rewrite Rules (CLI Procedure)" on page 97
- "Defining CoS Forwarding Classes (CLI Procedure)" on page 114
- (EX4500 only) "Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure)" on page 123
- "Defining CoS Schedulers and Scheduler Maps (CLI Procedure)" on page 138
- "Assigning CoS Components to Interfaces (CLI Procedure)" on page 55

Table 1: CoS Configuration Tasks (Continued)

CoS Configuration Task	Links
 Forwarding classes (logical interfaces only) Scheduler maps Rewrite rules 	
 Configure congestion management mechanisms for a switch to drop arriving packets based on certain parameters when a queue is full. Based on the EX Series switch that you are using, packets are dropped depending on the priority of a packet or on both priority and drop probability of a packet. Configure a weighted tail drop profile, a simple and effective traffic congestion avoidance mechanism. When you apply this mechanism to manage congestion, packets are dropped when the output queue is full. Configure a weighted random early detection (WRED) drop profile. When the configured capacity (fill level) is reached, packets marked with a packet loss priority (PLP) of high are discarded. Configure port shaping and queue shaping to enable you to limit traffic on an interface or queue, respectively, so that you can control the amount of traffic passing through the interface or the queue. 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. 	 "Configuring CoS Congestion Management (CLI Procedure)" on page 160 "Configuring Shaping for CoS (CLI Procedure)" on page 171 (EX4300 and EX4600 only) "Example: Configuring ECN" on page 185
Configure CoS on MPLS networks to ensure better performance for low-latency applications such as VoIP and other business-critical functions.	• (EX4500 and EX8200 only) "Example: Combining CoS with MPLS on EX Series Switches" on page 199

Understanding Junos OS CoS Components for EX Series Switches

IN THIS SECTION

- Code-Point Aliases | 8
- Policers | 8
- Classifiers | 9
- Forwarding Classes | 9
- Tail Drop Profiles | 9
- Schedulers | 9
- Rewrite Rules | 10

This topic describes the Juniper Networks Junos operating system (Junos OS) class-of-service (CoS) components for Juniper Networks EX Series Ethernet Switches:

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, drop-profile maps, 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. You define policers with filters that can be associated with input interfaces.

For more information about policers, see Understanding the Use of Policers in Firewall Filters.

NOTE: You can configure policers to discard packets that exceed the rate limits. If you want to configure CoS parameters such as **loss-priority** and **forwarding-class**, you must use firewall filters.

Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. In Juniper Networks Junos operating system (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 or CoS value traffic classifiers—Examines 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, IP precedence value, and IEEE 802.1p value.
- Multifield traffic classifiers—Examines 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.

Forwarding Classes

Forwarding classes group the packets for transmission. Based on forwarding classes, you assign packets to output queues. Forwarding classes affect the forwarding, scheduling, and marking policies applied to packets as they transit a switch. By default, four categories of forwarding classes are defined: best effort, assured forwarding, expedited forwarding, and network control. For EX Series switches, 16 forwarding classes are supported, providing granular classification capability.

Tail Drop Profiles

Drop profile is a mechanism that defines parameters that allow packets to be dropped from the network. Drop profiles define the meanings of the loss priorities. When you configure drop profiles you are essentially setting the value for queue fullness. The queue fullness represents a percentage of the queue used to store packets in relation to the total amount that has been allocated for that specific queue.

Loss priorities set the priority of dropping a packet. Loss priority affects the scheduling of a packet without affecting the packet's relative ordering. You can use the loss priority setting to identify packets that have experienced congestion. Typically you mark packets exceeding some service level with a high loss priority.

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 which type of packet should be transmitted before another. You can define the priority, bandwidth, delay buffer size, and tail drop profiles to be applied to a particular queue for packet transmission.

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, thus allowing the next downstream device to classify the packet into the appropriate service group. Rewriting, or marking, outbound packets is useful when the switch is at the border of a network and must alter the CoS values to meet the policies of the targeted peer.

NOTE: Egress firewall filters can also assign forwarding class and loss priority so that the packets are rewritten based on forwarding class and loss priority.

RELATED DOCUMENTATION

Understanding CoS Code-Point Aliases 61	
Understanding CoS Classifiers 70	
Understanding CoS Forwarding Classes 110	
Understanding CoS Tail Drop Profiles 164	
Understanding CoS Schedulers 128	
Understanding CoS Two-Color Marking 10	
Understanding CoS Rewrite Rules 95	
Example: Configuring CoS on EX Series Switches 15	

Understanding CoS Two-Color Marking

Networks police traffic by limiting the input or output transmission rate of a class of traffic on the basis of user-defined criteria. Policing traffic allows you to control the maximum rate of traffic sent or received on an interface and to partition a network into multiple priority levels or classes of service.

Policers require you to apply limits to the traffic flow and set a consequence for packets that exceed these limits—usually a higher loss priority, so that packets exceeding the policer limits are discarded first.

Juniper Networks EX Series Ethernet Switches support a single-rate two-color marking type of policer, which is a simplified version of Single-Rate-Three-Color marking, defined in RFC 2697, *A Single Rate Three Color Marker*. This type of policer meters traffic based on the configured committed information rate (CIR) and committed burst size (CBS).

The single-rate two-color marker meters traffic and marks incoming packets depending on whether they are smaller than the committed burst size (CBS)—marked green—or exceed it— marked red.

The single-rate two-color marking policer operates in color-blind mode. In this mode, the policer's actions are not affected by any previous marking or metering of the examined packets. In other words, the policer is "blind? to any previous coloring a packet might have had.

RELATED DOCUMENTATION

Understanding Junos OS CoS Components for EX Series Switches | 8 Understanding the Use of Policers in Firewall Filters Configuring Policers to Control Traffic Rates (CLI Procedure)

Understanding CoS Queues on EX8200 Line Cards That Include Oversubscribed Ports

IN THIS SECTION

- Oversubscribed Ports on Line Cards | 11
- EX8200 Line Cards That Include Oversubscribed Ports | 12
- Ingress Queueing | 13
- Egress Queues | 14

Some line cards available for Juniper Networks EX8200 Ethernet Switches include oversubscribed ports that are combined in logical port groups that share bandwidth. These oversubscribed ports handle traffic differently than ports that provide continuous line-rate bandwidth. You might need to configure CoS queues differently for oversubscribed ports than for line-rate ports.

This topic describes:

Oversubscribed Ports on Line Cards

Oversubscribed ports on a line card are grouped into logical port groups. A port group collectively supports a certain bandwidth.

An EX8200 switch supports different line cards that provide line-rate and oversubscribed ports. Based on your requirement, you can choose the appropriate line card for an EX8200 switch. Line cards are field-replaceable units (FRUs) that can be installed in the line card slots in an EX8200 switch. In a line-rate EX8200 line card, each port in the line card supports the same amount of bandwidth and a single port can utilize that complete bandwidth. In an oversubscribed line card, a group of ports collectively support a certain total bandwidth and each port in that group can use either a portion or all of the available bandwidth. However, the total utilization of bandwidth by the ports in the group cannot exceed the bandwidth available for that group.

Because the port groups share bandwidth, class-of-service (CoS) ingress and egress queues are handled differently for these shared-bandwidth ports in logical port groups than they are for ports that individually support line-rate bandwidth. Some EX8200 line cards combine both port types, those that share bandwidth across port groups and those that individually support line-rate bandwidth.

EX8200 Line Cards That Include Oversubscribed Ports

Table 2 on page 12 lists EX8200 line cards that include oversubscribed ports in logical port groups.

Line Card Model	Name	Number of Oversubscribed Ports/Port Connector
EX8200-40XS	40-port SFP+	40 oversubscribed 10-gigabit SFP+ ports
EX8200-2XS-40P	40-port PoE+ with 4-port SFP and 2-port SFP+	40 oversubscribed 10/100/1000 Gigabit Ethernet ports with RJ-45 connectors, four small form-factor pluggable (SFP) ports (in which you can install 1-gigabit SFP transceivers) and two SFP+ ports
EX8200-2XS-40T	40-port RJ-45 with 4-port SFP and 2-port SFP+	40 oversubscribed 10/100/1000 Gigabit Ethernet ports with RJ-45 connectors, four SFP ports (in which you can install 1-gigabit small form- factor pluggable (SFP) transceivers) and two SFP+ ports

Table 2: EX8200 Line Cards That Include Oversubscribed Ports

Line Card Model	Name	Number of Oversubscribed Ports/Port Connector
EX8200-48PL	48-port PoE+ 20 Gbps	48 oversubscribed 10/100/1000 Gigabit Ethernet ports with RJ-45 connectors
EX8200-48TL	48-port RJ-45 20 Gbps	48 oversubscribed 10/100/1000 Gigabit Ethernet ports with RJ-45 connectors

Table 2: EX8200 Line Cards That Include Oversubscribed Ports (Continued)

Ingress Queueing

Classification of packets occurs in two phases for the oversubscribed ports in the port groups.

Preclassification of Packets and Port Ingress Queuing

Packets entering ports are forwarded to one of the ingress queues. The ingress queues schedule traffic from ports into the Packet Forwarding Engine.

The ingress queues are:

- Low-priority queue—Each interface in the line card has one low-priority queue. Traffic on these queues is scheduled using the shaped deficit weighted round-robin (SDWRR) algorithm, with each interface's queue having equal weight. On EX4300 switches, traffic is queued using the weighted deficit round-robin (WDRR) algorithm.
- High-priority queue—A set of interfaces in the line card shares a single high-priority queue. Traffic on this queue is scheduled by strict-high priority. The switch always sends critical network control packets on the high-priority queue.
- Line-rate priority queue—The packets entering line-rate ports are forwarded to this queue. Traffic on this queue is scheduled by strict priority and is always given higher priority than the traffic on the high-priority queue. This queue is used only in the following oversubscribed lines cards for an EX8200 switch:
 - EX8200-2XS-40P
 - EX8200-2XS-40T

For the purpose of port ingress queuing on oversubscribed ports, packets are classified only by behavior aggregate (BA) classification. To control the ingress queue (high priority or low priority) to which packets are sent, configure a BA classifier on the physical port and specify switch fabric priorities for the forwarding classes. On EX8200 switches, fabric priority determines the priority of packets ingressing the switch fabric. For the EX8200-40XS line card, fabric priority also determines the priority of packets ingressing the port group.

By default, the fabric priority for all forwarding classes is low. To direct packets belonging to a forwarding class to the high-priority ingress queue, set the fabric priority to high for that class.

Critical network-control packets and line-rate packets are handled differently from other packets. Instead of using the BA classifier to classify critical network-control packets, the switch always sends critical network packets to the high-priority queue. The line-rate packets are always sent to the line-rate priority queue. This difference in handling of network-control packets and line-rate packets ensures that these packets are not dropped because of congestion on the shared-bandwidth ports.

Full Classification of Packets and Fabric Ingress Queuing

When packets (apart from line-rate and critical network-control packets) from an oversubscribed port reach the Packet Forwarding Engine, it performs full packet classification, along with other actions, such as multifield (MF) classification, traffic policing, and storm control. It then schedules and queues the packets for ingressing the fabric. The fabric priority associated with the forwarding class determines whether packets are sent to the low priority or high-priority ingress queues.

Egress Queues

On EX Series switches except EX4300 switches, each interface supports eight egress CoS queues. You can map up to 16 forwarding classes to these queues. An EX4300 switch interface supports 12 egress CoS queues.

In the EX8200-40XS line card, all interfaces in a port group share a single set of eight egress queues at the Packet Forwarding Engine. Egress traffic is fanned out from the Packet Forwarding Engine queues to the corresponding queues for the individual ports. For this reason, the interfaces in a port group must share the same scheduler map configuration. If you configure different scheduler map configurations for the different interfaces in a port group, an error is logged in the system log and the default scheduler map is used for all ports in the port group.

RELATED DOCUMENTATION

Understanding Junos OS CoS Components for EX Series Switches | 8

Understanding CoS Schedulers | 128

Understanding CoS Forwarding Classes | 110

Example: Configuring CoS on EX Series Switches | 15

Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure) | 87

Example: Configuring CoS on EX Series Switches

IN THIS SECTION

- Requirements | 15
- Overview and Topology | 15
- Configuration | 19
- Verification | 35

Configure class of service (CoS) on your switch to manage traffic so that when the network experiences congestion and delay, critical applications are protected. Using CoS, you can divide traffic on your switch into classes and provide various levels of throughput and packet loss. This is especially important for traffic that is sensitive to jitter and delay, such as voice traffic.

This example shows how to configure CoS on a single EX Series switch in the network.

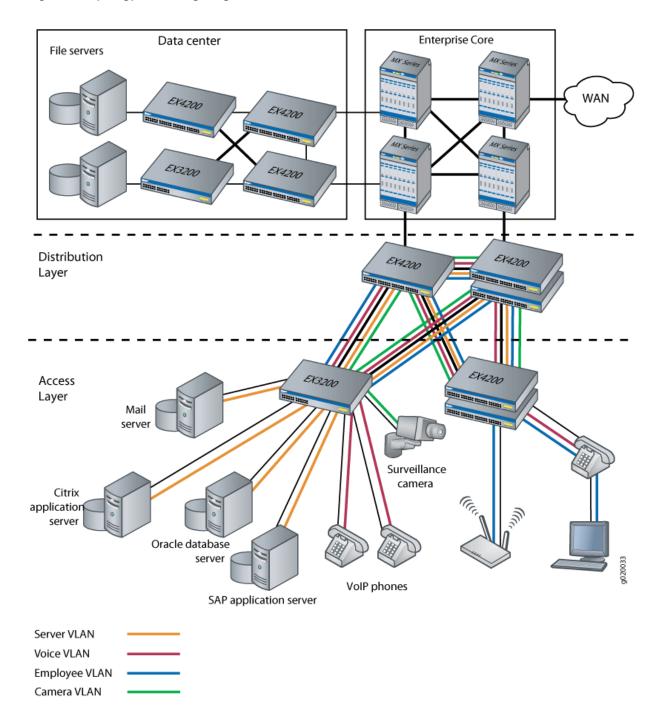
Requirements

This example uses the following hardware and software components:

- EX3200 and EX4200 switches
- Junos OS Release 9.0 or later for EX Series switches

Overview and Topology

This example uses the topology shown in Figure 2 on page 16.



The topology for this configuration example consists of EX3200 and EX4200 switches at the access layer.

The EX Series access switches are configured to support VLAN membership. On the EX3200 access layer switch, interfaces ge-0/0/0 and ge-0/0/1 are assigned to the voice VLAN (voice-vlan) for two VoIP IP phones. Switch interface ge-0/0/2 is assigned to the camera VLAN (camera-vlan) for the surveillance camera. Switch interfaces ge-0/0/3, ge-0/0/4, ge-0/0/5, and ge-0/0/6 are assigned to the server VLAN (server-vlan) for the server shosting various applications such as those provided by Citrix, Microsoft, Oracle, and SAP. The EX3200 trunk ports, ge-0/0/20 and ge-0/0/21, are assigned to the server, voice, employee, and camera VLANs and used as uplink ports to connect the distribution layer switches.

EX4200 switches are also included in the access layer to similarly connect employee and voice VLANs, although this example does not show configuration details for those switches.

Table 3 on page 17 shows the VLAN configuration components.

VLAN Name	VLAN ID	VLAN Subnet and Available IP Addresses	VLAN Description
voice-vlan	10	192.168.1.0/28 192.168.1.1 through 192.168.1.14 192.168.1.15 is the subnet's broadcast address.	Voice VLAN used for employee VoIP communication.
camera-vlan	20	192.168.1.16/28 192.168.1.17 through 192.168.1.30 192.168.1.31 is the subnet's broadcast address.	VLAN for the surveillance cameras.

Table 3: Configuration Components: VLANs

VLAN Name	VLAN ID	VLAN Subnet and Available IP Addresses	VLAN Description
server-vlan	30	192.168.1.32/28 192.168.1.33 through 192.168.1.46 192.168.1.47 is the subnet's broadcast address.	VLAN for the servers hosting enterprise applications.

Table 3: Configuration Components: VLANs (Continued)

PoE-capable ports on EX Series switches support Power over Ethernet (PoE) to provide both network connectivity and power for VoIP telephones connecting to the ports. Table 4 on page 18 shows the switch interfaces that are assigned to the VLANs and the IP addresses for devices connected to the switch ports on a 48-port switch, all ports of which are PoE-capable.

Table 4: Configuration Components: Switch Interfaces Assigned to VLANs and Devices on a 48-Port
All-PoE Switch

Interfaces	VLAN Membership	IP Addresses	Port Devices
ge-0/0/0, ge-0/0/1	voice-vlan	192.168.1.1/28 through 192.168.1.2/28	Two VoIP telephones.
ge-0/0/2	camera-vlan	192.168.1.17/28	Surveillance camera.
ge-0/0/3, ge-0/0/4, ge-0/0/5, ge-0/0/6	server-vlan	192.168.1.33/28 through 192.168.1.36/28	Four servers hosting applications such as those provided by Citrix, Microsoft, Oracle, and SAP.

NOTE: This example shows how to configure CoS on a standalone EX Series switch. This example does not consider across-the-network applications of CoS in which you might implement different configurations on ingress and egress switches to provide differentiated treatment to different classes across a set of nodes in a network.

NOTE: Although you will sometimes see schedulers configured for strict-high priority with a transmit-rate configured, that configuration is misleading because strict-high priority schedulers get unlimited bandwidth and the transmit-rate parameter has no effect on them. With this configuration, lower priority queues can suffer starvation if there is congestion. It is better that schedulers with strict-high priority have shaping-rate parameters configured, which is the correct way to limit their bandwidth.

Configuration

IN THIS SECTION

Procedure | 19

Procedure

CLI Quick Configuration

To quickly configure CoS, copy the following commands and paste them into the switch terminal window:

```
[edit]
set class-of-service forwarding-classes class app queue-num 5
set class-of-service forwarding-classes class mail queue-num 1
set class-of-service forwarding-classes class db queue-num 2
set class-of-service forwarding-classes class erp queue-num 3
set class-of-service forwarding-classes class video queue-num 4
set class-of-service forwarding-classes class best-effort queue-num 0
set class-of-service forwarding-classes class voice queue-num 6
set class-of-service forwarding-classes class network-control queue-num 7
set firewall family ethernet-switching filter voip_class term voip from source-address
192.168.1.1/28
set firewall family ethernet-switching filter voip_class term voip from source-address
192.168.1.2/28
set firewall family ethernet-switching filter voip_class term voip from protocol udp
set firewall family ethernet-switching filter voip_class term voip from source-port 2698
set firewall family ethernet-switching filter voip_class term voip then forwarding-class voice
```

loss-priority low set firewall family ethernet-switching filter voip_class term network_control from precedence [net-control internet-control] set firewall family ethernet-switching filter voip_class term network_control then forwardingclass network-control loss-priority low set firewall family ethernet-switching filter voip_class term best_effort_traffic then forwarding-class best-effort loss-priority low set interfaces ge-0/0/0 description phone1-voip-ingress-port set interfaces ge-0/0/0 unit 0 family ethernet-switching filter input voip_class set class-of-service interfaces ge-0/0/0 shaping-rate 100m set interfaces ge-0/0/1 description phone2-voip-ingress-port set interfaces ge-0/0/1 unit 0 family ethernet-switching filter input voip_class set firewall family ethernet-switching filter video_class term video from source-address 192.168.1.17/28 set firewall family ethernet-switching filter video_class term video from protocol udp set firewall family ethernet-switching filter video_class term video from source-port 2979 set firewall family ethernet-switching filter video_class term video then forwarding-class video loss-priority low set firewall family ethernet-switching filter video_class term network_control from precedence [net-control internet-control] set firewall family ethernet-switching filter video_class term network_control then forwardingclass network-control loss-priority low set firewall family ethernet-switching filter video_class term best_effort_traffic then forwarding-class best-effort loss-priority low set interfaces ge-0/0/2 description video-ingress-port set interfaces ge-0/0/2 unit 0 family ethernet-switching filter input video_class set firewall family ethernet-switching filter app_class term app from source-address 192.168.1.33/28 set firewall family ethernet-switching filter app_class term app from protocol tcp set firewall family ethernet-switching filter app_class term app from source-port [1494 2512 2513 2598 2897] set firewall family ethernet-switching filter app_class term app then forwarding-class app losspriority low set firewall family ethernet-switching filter app_class term mail from source-address 192.168.1.34/28 set firewall family ethernet-switching filter app_class term mail from protocol tcp set firewall family ethernet-switching filter app_class term mail from source-port [25 143 389 691 993 3268 3269] set firewall family ethernet-switching filter app_class term mail then forwarding-class mail loss-priority low set firewall family ethernet-switching filter app_class term db from source-address 192.168.1.35/28 set firewall family ethernet-switching filter app_class term db from protocol tcp

set firewall family ethernet-switching filter app_class term db from source-port [1521 1525 1527 1571 1810 2481] set firewall family ethernet-switching filter app_class term db then forwarding-class db losspriority low set firewall family ethernet-switching filter app_class term erp from source-address 192.168.1.36/28 set firewall family ethernet-switching filter app_class term erp from protocol tcp set firewall family ethernet-switching filter app_class term erp from source-port [3200 3300 3301 3600] set firewall family ethernet-switching filter app_class term erp then forwarding-class erp losspriority low set firewall family ethernet-switching filter app_class term network_control from precedence [net-control internet-control] set firewall family ethernet-switching filter app_class term network_control then forwardingclass network-control loss-priority low set firewall family ethernet-switching filter app_class term best_effort_traffic then forwardingclass best-effort loss-priority low set interfaces ge-0/0/3 unit 0 family ethernet-switching filter input app_class set interfaces ge-0/0/4 unit 0 family ethernet-switching filter input app_class set interfaces ge-0/0/5 unit 0 family ethernet-switching filter input app_class set interfaces ge-0/0/6 unit 0 family ethernet-switching filter input app_class set class-of-service schedulers voice-sched shaping-rate percent 10 set class-of-service schedulers voice-sched buffer-size percent 10 set class-of-service schedulers voice-sched priority strict-high set class-of-service schedulers video-sched priority low set class-of-service schedulers video-sched transmit-rate percent 15 set class-of-service schedulers app-sched buffer-size percent 10 set class-of-service schedulers app-sched priority low set class-of-service schedulers app-sched transmit-rate percent 10 set class-of-service schedulers mail-sched buffer-size percent 5 set class-of-service schedulers mail-sched priority low set class-of-service schedulers mail-sched transmit-rate percent 5 set class-of-service schedulers db-sched buffer-size percent 10 set class-of-service schedulers db-sched priority low set class-of-service schedulers db-sched transmit-rate percent 10 set class-of-service schedulers erp-sched buffer-size percent 10 set class-of-service schedulers erp-sched priority low set class-of-service schedulers erp-sched transmit-rate percent 10 set class-of-service schedulers nc-sched shaping-rate percent 5 set class-of-service schedulers nc-sched buffer-size percent 5 set class-of-service schedulers nc-sched priority strict-high set class-of-service schedulers be-sched buffer-size percent 35 set class-of-service schedulers be-sched priority low

```
set class-of-service schedulers be-sched transmit-rate percent 35
set class-of-service scheduler-maps ethernet-cos-map forwarding-class voice scheduler voice-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class video scheduler video-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class app scheduler app-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class mail scheduler mail-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class db scheduler db-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class erp scheduler erp-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class network-control scheduler
nc-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class best-effort scheduler be-
sched
set class-of-service interfaces ge-0/0/20 scheduler-map ethernet-cos-map
set class-of-service interfaces ge-0/0/21 scheduler-map ethernet-cos-map
set class-of-service schedulers voice-sched-queue-shap shaping-rate 30m
set class-of-service scheduler-maps sched-map-be forwarding-class best-effort scheduler voice-
sched-queue-shap
set class-of-service interfaces ge-0/0/2 scheduler-map sched-map-be
```

Step-by-Step Procedure

To configure and apply CoS:

1. Configure one-to-one mappings between eight forwarding classes and eight queues:

```
[edit class-of-service]
user@switch# set forwarding-classes class app queue-num 5
user@switch# set forwarding-classes class mail queue-num 1
user@switch# set forwarding-classes class db queue-num 2
user@switch# set forwarding-classes class erp queue-num 3
user@switch# set forwarding-classes class video queue-num 4
user@switch# set forwarding-classes class best-effort queue-num 0
user@switch# set forwarding-classes class voice queue-num 6
user@switch# set forwarding-classes class network-control queue-num 7
```

2. Define the firewall filter voip_class to classify the VoIP traffic:

[edit firewall]
user@switch# set family ethernet-switching filter voip_class

3. Define the term voip:

```
[edit firewall]
user@switch# set family ethernet-switching filter voip_class term voip from source-address
192.168.1.1/28
user@switch# set family ethernet-switching filter voip_class term voip from source-address
192.168.1.2/28
user@switch# set family ethernet-switching filter voip_class term voip protocol udp
user@switch# set family ethernet-switching filter voip_class term voip source-port 2698
user@switch# set family ethernet-switching filter voip_class term voip then forwarding-
class voice loss-priority low
```

4. Define the term network_control (for the voip_class filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter voip_class term network_control from
precedence [net-control internet-control]
user@switch# set family ethernet-switching filter voip_class term network_control then
forwarding-class network-control loss-priority low
```

5. Define the term best_effort_traffic with no match conditions (for the voip_class filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter voip_class term best_effort_traffic then
forwarding-class best-effort loss-priority low
```

6. Apply the firewall filter voip_class as an input filter to the interfaces for the VoIP phones:

```
[edit interfaces]
```

```
user@switch# set ge-0/0/0 description phone1-voip-ingress-port
user@switch# set ge-0/0/0 unit 0 family ethernet-switching filter input voip_class
user@switch# set ge-0/0/1 description phone2-voip-ingress-port
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input voip_class
```

7. Apply port shaping on the interface ge-0/0/0:

```
[edit]
```

user@switch# set class-of-service interfaces ge-0/0/0 shaping-rate 100m

8. Define the firewall filter video_class to classify the video traffic:

[edit firewall]
user@switch# set family ethernet-switching filter video_class

9. Define the term video:

```
[edit firewall]
user@switch# set family ethernet-switching filter video_class term video from source-
address 192.168.1.17/28
user@switch# set family ethernet-switching filter video_class term video protocol udp
user@switch# set family ethernet-switching filter video_class term video source-port 2979
user@switch# set family ethernet-switching filter video_class term video then forwarding-
class video loss-priority low
```

10. Define the term network_control (for the video_class filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter video_class term network_control from
precedence [net-control internet-control]
user@switch# set family ethernet-switching filter video_class term network_control then
forwarding-class network-control loss-priority low
```

11. Define the term best_effort_traffic with no match conditions (for the video_class filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter video_class term best_effort_traffic then
forwarding-class best-effort loss-priority low
```

12. Apply the firewall filter video_class as an input filter to the interface for the surveillance camera:

```
[edit interfaces]
user@switch# set ge-0/0/2 description video-ingress-port
user@switch# set ge-0/0/2 unit 0 family ethernet-switching filter input video_class
```

13. Define the firewall filter app_class to classify the application server traffic:

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class
```

14. Define the term app (for the app_class filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term app from source-address
192.168.1.33/28
user@switch# set family ethernet-switching filter app_class term app protocol tcp
user@switch# set family ethernet-switching filter app_class term app source-port [1494 2512
2513 2598 2897]
user@switch# set family ethernet-switching filter app_class term app then forwarding-class
app loss-priority low
```

15. Define the term mail (for the app_class filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term mail from source-address
192.168.1.34/28
user@switch# set family ethernet-switching filter app_class term mail protocol tcp
user@switch# set family ethernet-switching filter app_class term mail source-port [25 143
389 691 993 3268 3269]
user@switch# set family ethernet-switching filter app_class term mail then forwarding-class
mail loss-priority low
```

16. Define the term db (for the app_class filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term db from source-address
192.168.1.35/28
```

```
user@switch# set family ethernet-switching filter app_class term db protocol tcp
user@switch# set family ethernet-switching filter app_class term db source-port [1521 1525
1527 1571 1810 2481]
user@switch# set family ethernet-switching filter app_class term db then forwarding-class
db loss-priority low
```

17. Define the term erp (for the app_class filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term erp from source-address
192.168.1.36/28
user@switch# set family ethernet-switching filter app_class term erp protocol tcp
user@switch# set family ethernet-switching filter app_class term erp source-port [3200 3300
3301 3600]
user@switch# set family ethernet-switching filter app_class term erp then forwarding-class
erp loss-priority low
```

18. Define the term network_control (for the app_class filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term network_control from
precedence [net-control internet-control]
user@switch# set family ethernet-switching filter app_class term network_control then
forwarding-class network-control loss-priority low
```

19. Define the term best_effort_traffic (for the app_class filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term best_effort_traffic then
forwarding-class best-effort loss-priority low
```

20. Apply the firewall filter app_class as an input filter to the interfaces for the servers hosting applications:

```
[edit interfaces]
user@switch# set ge-0/0/3 unit 0 family ethernet-switching filter input app_class
user@switch# set ge-0/0/4 unit 0 family ethernet-switching filter input app_class
```

user@switch# set ge-0/0/5 unit 0 family ethernet-switching filter input app_class
user@switch# set ge-0/0/6 unit 0 family ethernet-switching filter input app_class

21. Configure schedulers:

```
[edit class-of-service]
user@switch# set schedulers voice-sched shaping-rate percent 10
```

	See Seneuarers	voree selled shaping rate percent to
user@switch#	set schedulers	voice-sched buffer-size percent 10
user@switch#	set schedulers	voice-sched priority strict-high
user@switch#	set schedulers	video-sched priority low
user@switch#	set schedulers	video-sched transmit-rate percent 15
user@switch#	set schedulers	app-sched buffer-size percent 10
user@switch#	set schedulers	app-sched priority low
user@switch#	set schedulers	app-sched transmit-rate percent 10
user@switch#	set schedulers	mail-sched buffer-size percent 5
user@switch#	set schedulers	mail-sched priority low
user@switch#	set schedulers	mail-sched transmit-rate percent 5
user@switch#	set schedulers	db-sched buffer-size percent 10
user@switch#	set schedulers	db-sched priority low
user@switch#	set schedulers	db-sched transmit-rate percent 10
user@switch#	set schedulers	erp-sched buffer-size percent 10
user@switch#	set schedulers	erp-sched priority low
user@switch#	set schedulers	erp-sched transmit-rate percent 10
user@switch#	set schedulers	nc-sched shaping-rate percent 5
user@switch#	set schedulers	nc-sched buffer-size percent 5
user@switch#	set schedulers	nc-sched priority strict-high
user@switch#	set schedulers	nc-sched transmit-rate percent 5
user@switch#	set schedulers	be-sched buffer-size percent 35
user@switch#	set schedulers	be-sched priority low
user@switch#	set schedulers	be-sched transmit-rate percent 35

22. Assign the forwarding classes to schedulers with the scheduler map ethernet-cos-map:

[edit class-of-service]

user@switch# set scheduler-maps ethernet-cos-map forwarding-class voice scheduler voicesched

user@switch# set scheduler-maps ethernet-cos-map forwarding-class video scheduler videosched

user@switch# set scheduler-maps ethernet-cos-map forwarding-class app scheduler app-sched user@switch# set scheduler-maps ethernet-cos-map forwarding-class mail scheduler mail-sched user@switch# set scheduler-maps ethernet-cos-map forwarding-class db scheduler db-sched user@switch# set scheduler-maps ethernet-cos-map forwarding-class erp scheduler erp-sched user@switch# set scheduler-maps ethernet-cos-map forwarding-class network-control scheduler nc-sched user@switch# set scheduler-maps ethernet-cos-map forwarding-class best-effort scheduler be-

23. Associate the scheduler map with the outgoing interfaces:

```
[edit class-of-service interfaces]
user@switch# set ge-0/0/20 scheduler-map ethernet-cos-map
user@switch# set ge-0/0/21 scheduler-map ethernet-cos-map
```

24. Apply queue shaping for the best-effort queue:

```
[edit]
```

sched

```
user@switch# set class-of-service schedulers voice-sched-queue-shap shaping-rate 30m
user@switch# set class-of-service scheduler-maps sched-map-be forwarding-class best-effort
scheduler voice-sched-queue-shap
user@switch# set class-of-service interfaces ge-0/0/2 scheduler-map sched-map-be
```

Results

Display the results of the configuration:

user@switch> show firewall

```
forwarding-class voice;
            loss-priority low;
       }
   }
   term network control {
       from {
            precedence [net-control internet-control];
       }
       then {
            forwarding-class network-control;
            loss-priority low;
       }
   }
    term best_effort_traffic {
        then {
            forwarding-class best-effort;
            loss-priority low;
       }
   }
}
filter video_class {
    term video {
        from {
            source-address {
                192.168.1.17/28;
            }
            protocol udp;
            source-port 2979;
       }
        then {
            forwarding-class video;
            loss-priority low;
       }
   }
   term network control {
        from {
            precedence [net-control internet-control];
       }
        then {
            forwarding-class network-control;
            loss-priority low;
       }
   }
```

```
term best_effort_traffic {
        then {
            forwarding-class best-effort;
            loss-priority low;
        }
    }
}
filter app_class {
    term app {
        from {
            source-address {
                192.168.1.33/28;
            }
            protocol tcp;
            source-port [1491 2512 2513 2598 2897];
        }
        then {
            forwarding-class app;
            loss-priority low;
        }
    }
    term mail {
        from {
            source-address {
                192.168.1.34/28;
            }
            protocol tcp;
            source-port [25 143 389 691 993 3268 3269];
        }
        then {
            forwarding-class mail;
            loss-priority low;
        }
    }
    term db {
        from {
            source-address {
                192.168.1.35/28;
            }
            protocol tcp;
            source-port [1521 1525 1527 1571 1810 2481];
        }
        then {
```

```
forwarding-class db;
            loss-priority low;
        }
    }
    term erp {
        from {
            source-address {
                192.168.1.36/28;
            }
            protocol tcp;
            source-port [3200 3300 3301 3600];
        }
        then {
            forwarding-class erp;
            loss-priority low;
        }
    }
    term network control {
        from {
            precedence [net-control internet-control];
        }
        then {
            forwarding-class network-control;
            loss-priority low;
        }
    }
    term best_effort_traffic {
        then {
            forwarding-class best-effort;
            loss-priority low;
        }
    }
}
```

user@switch# show class-of-service

forwarding-classes {
 class app queue-num 5;
 class mail queue-num 1;

}

```
class db queue-num 2;
    class erp queue-num 3;
    class video queue-num 4;
    class best-effort queue-num 0;
    class voice queue-num 6;
    class network-control queue-num 7;
}
interfaces {
    ge-0/0/0 {
        shaping-rate 100m;
    }
    ge-0/0/2 {
        scheduler-map sched-map-be;
    }
    ge-0/0/20 {
        scheduler-map ethernet-cos-map;
    }
    ge-0/0/21 {
        scheduler-map ethernet-cos-map;
    }
}
schedulers {
    voice-sched-queue-shap {
        shaping-rate 30m;
    }
    voice-sched {
        shaping-rate percent 10;
        buffer-size percent 10;
        priority strict-high;
    }
    video-sched {
        buffer-size percent 15;
        priority low;
        transmit-rate percent 15;
    }
    app-sched {
        buffer-size percent 10;
        priority low;
        transmit-rate percent 10;
    }
    mail-sched {
        buffer-size percent 5;
        priority low;
```

```
transmit-rate percent 5;
    }
    db-sched {
        buffer-size percent 10;
        priority low;
        transmit-rate percent 10;
    }
    erp-sched {
        buffer-size percent 10;
        priority low;
        transmit-rate percent 10;
   }
    nc-sched {
        shaping-rate percent 5;
        buffer-size percent 5;
        priority strict-high;
    }
    be-sched {
        buffer-size percent 35;
        priority low;
        transmit-rate percent 35;
    }
}
scheduler-maps {
    ethernet-cos-map {
        forwarding-class voice scheduler voice-sched;
        forwarding-class video scheduler video-sched;
        forwarding-class app scheduler app-sched;
        forwarding-class mail scheduler mail-sched;
        forwarding-class db scheduler db-sched;
        forwarding-class erp scheduler erp-sched;
        forwarding-class network-control scheduler nc-sched;
        forwarding-class best-effort scheduler be-sched;
    }
    sched-map-be {
        forwarding-class best-effort scheduler voice-sched-queue-shap;
```

}

}

```
user@switch# show interfaces
```

```
ge-0/0/0 {
   unit 0 {
       family ethernet {
           filter {
               input voip_class;
           }
       }
   }
}
ge-0/0/1 {
   unit 0 {
       family ethernet {
           filter {
               input voip_class;
           }
       }
   }
}
ge-0/0/2 {
    unit 0 {
       family ethernet {
           filter {
               input video_class;
           }
       }
   }
}
ge-0/0/3 {
   unit 0 {
       family ethernet {
           filter {
               input app_class;
           }
       }
    }
```

```
}
ge-0/0/4 {
    unit 0 {
        family ethernet {
            filter {
                input app_class;
            }
        }
   }
}
ge-0/0/5 {
    unit 0 {
        family ethernet {
            filter {
                input app_class;
            }
        }
    }
}
ge-0/0/6 {
    unit 0 {
        family ethernet {
            filter {
                input app_class;
            }
        }
    }
}
```

Verification

IN THIS SECTION

- Verifying That the Defined Forwarding Classes Exist and Are Mapped to Queues | 36
- Verifying That the Forwarding Classes Have Been Assigned to Schedulers | 36
- Verifying That the Scheduler Map Has Been Applied to the Interfaces | 38
- Verifying That Port Shaping Has Been Applied | 39
- Verifying That Queue Shaping Has Been Applied | 44

To confirm that the configuration is working properly, perform these tasks:

Verifying That the Defined Forwarding Classes Exist and Are Mapped to Queues

Purpose

Verify that the forwarding classes app, best-effort, db, erp, mail, network-control, video, and voice have been defined and mapped to queues.

Action

<pre>user@switch> show class-of-service forwarding-class</pre>						
Forwarding class	ID	Queue				
арр	0	5				
db	1	2				
erp	2	3				
best-effort	3	0				
mail	4	1				
voice	5	6				
video	6	4				
network-control	7	7				

Meaning

This output shows that the forwarding classes have been defined and mapped to appropriate queues.

Verifying That the Forwarding Classes Have Been Assigned to Schedulers

Purpose

Verify that the forwarding classes have been assigned to schedulers.

Action

user@switch> show class-of-service scheduler-map Scheduler map: ethernet-cos-map, Index: 2 Scheduler: voice-sched, Forwarding class: voice, Index: 22 Shaping rate: 10 percent, Rate Limit: none, Buffer size: 10 percent, Priority: Strict-high

Drop profiles: Loss priority Protocol Index Name non-TCP 1 <default-drop-profile> High TCP High 1 <default-drop-profile> Scheduler: video-sched, Forwarding class: video, Index: 22 Transmit rate: 10 percent, Rate Limit: none, Buffer size: 10 percent, Priority: low Drop profiles: Loss priority Protocol Index Name High non-TCP 1 <default-drop-profile> TCP High 1 <default-drop-profile> Scheduler: app-sched, Forwarding class: app, Index: 22 Transmit rate: 10 percent, Rate Limit: none, Buffer size: 10 percent, Priority: low Drop profiles: Loss priority Protocol Index Name High non-TCP <default-drop-profile> 1 TCP 1 <default-drop-profile> High Scheduler: mail-sched, Forwarding class: mail, Index: 22 Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Priority: low Drop profiles: Loss priority Protocol Index Name High non-TCP 1 <default-drop-profile> TCP <default-drop-profile> High 1 Scheduler: db-sched, Forwarding class: db, Index: 22 Transmit rate: 10 percent, Rate Limit: none, Buffer size: 10 percent, Priority: low Drop profiles: Loss priority Protocol Index Name <default-drop-profile> High non-TCP 1 High TCP 1 <default-drop-profile> Scheduler: erp-sched, Forwarding class: erp, Index: 22 Transmit rate: 10 percent, Rate Limit: none, Buffer size: 10 percent, Priority: low Drop profiles: Loss priority Protocol Index Name <default-drop-profile> High non-TCP 1

High TCP <default-drop-profile> 1 Scheduler: be-sched, Forwarding class: best-effort, Index: 20 Transmit rate: 35 percent, Rate Limit: none, Buffer size: 35 percent, Priority: low Drop profiles: Loss priority Protocol Index Name High non-TCP 1 <default-drop-profile> TCP <default-drop-profile> High 1 Scheduler: nc-sched, Forwarding class: network-control, Index: 22 Shaping rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Priority: Strict-high Drop profiles: Loss priority Protocol Index Name non-TCP 1 <default-drop-profile> High High TCP 1 <default-drop-profile>

Meaning

This output shows that the forwarding classes have been assigned to schedulers.

Verifying That the Scheduler Map Has Been Applied to the Interfaces

Purpose

Verify that the scheduler map has been applied to the interfaces.

Action

```
user@switch> show class-of-service interface
...
Physical interface: ge-0/0/20, Index: 149
Queues supported: 8, Queues in use: 8
Scheduler map: ethernet-cos-map, Index: 43366
Input scheduler map: <default>, Index: 3
...
Physical interface: ge-0/0/21, Index: 150
Queues supported: 8, Queues in use: 8
```

```
Scheduler map: ethernet-cos-map, Index: 15103
Input scheduler map: <default>, Index: 5
...
```

Meaning

This output includes details of the interfaces to which the scheduler map (ethernet-cos-map) has been applied (ge-0/0/20 and ge-0/0/21).

Verifying That Port Shaping Has Been Applied

Purpose

Verify that the port shaping has been applied to an interface.

Action

Following is the output before port shaping is applied to the interface ge-0/0/0, when there is egress traffic of 400 Mpbs exiting on that interface:

```
user@switch> show interfaces ge-0/0/0 extensive
Physical interface: ge-0/0/0, Enabled, Physical link is Up
 Interface index: 239, SNMP ifIndex: 548, Generation: 242
 Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto, BPDU Error: None, MAC-REWRITE
Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-
negotiation: Enabled, Remote fault: Online,
 Media type: Copper
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues
               : 8 supported, 8 maximum usable queues
                : Up 0 ms, Down 0 ms
 Hold-times
 Current address: 00:23:9c:0b:ae:8d, Hardware address: 00:23:9c:0b:ae:8d
 Last flapped : 2012-07-07 03:21:52 UTC (1d 18:02 ago)
 Statistics last cleared: 2012-07-07 23:54:34 UTC (21:29:59 ago)
 Traffic statistics:
  Input bytes :
                                     0
                                                          0 bps
  Output bytes :
                            2299853696
                                                  345934816 bps
                                     0
  Input packets:
                                                          0 pps
  Output packets:
                              17967609
                                                     337827 pps
  IPv6 transit statistics:
```

Input bytes :	6			
Output bytes :	6			
Input packets:	6			
Output packets:	6			
Input errors:				
Errors: 0, Drops:	0, Framing errors	: 0, Runts	: 0, Policed di	iscards: 0, L3 incompletes: 0, L2
channel errors: 0, L2	mismatch timeouts	: 0, FIFO e	errors: 0, Reso	purce errors: 0
Output errors:				
Carrier transition	ns: 0, Errors: 0,	Drops: 0, 0	Collisions: 0,	Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0	, MTU errors: 0, F	esource eri	rors: 0	
Egress queues: 8 su	oported, 4 in use			
Queue counters:	Queued packets	Transmitte	ed packets	Dropped packets
0 best-effort	0		18302337	0
1 assured-forw	0		0	0
5 expedited-fo	0		0	0
7 network-cont	0		0	0
Queue number:	Mapped forwardi	ng classes		
0	best-effort	-		
1	assured-forward	ling		
5	expedited-forwa	-		
7	network-control	-		
Active alarms : No	ne			
Active defects : No	ne			
MAC statistics:		Receive	Transmit	
Total octets		0	2299853696	
Total packets		0	17967609	
Unicast packets		0	17967609	
Broadcast packets		0	0	
Multicast packets		0	0	
CRC/Align errors		0	0	
FIFO errors		0	0	
MAC control frame:	S	0	0	
MAC pause frames		0	0	
Oversized frames		0		
Jabber frames		0		
Fragment frames		0		
Code violations		0		
Autonegotiation info	ormation:			
Negotiation status				
Link partner:	· · · · ·			
	ll-duplex. Flow co	ontrol: Svmr	netric, Remote	fault: OK, Link partner Speed:
1000 Mbps	, , , , , , , , , , , , , , , , , , , ,		,	
Local resolution:				

Flow control: Syn			Link Ok	K		
Packet Forwarding Engin	ne config	guration:				
Destination slot: 1						
CoS information:						
Direction : Output						
CoS transmit queue		Bandwidth		Buffer Pri	ority	Limit
	%	bps	%	usec		
0 best-effort	95	950000000	95	NA	low	none
7 network-control	5	50000000	5	NA	low	none
Interface transmit sta	tistics:	Disabled				
Logical interface ge-1,	/0/10.0 ((Index 69) (SN	MP ifInd	dex 638)(Gener	ation 1	38)
Flags: SNMP-Traps 0x0	∂ Encapsu	ulation: 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 eth-switch,	Generati	ion: 163, Rout	e table:	: 0		
Flags: Trunk-Mode						

The Traffic statistics: field in this output shows that egress traffic is ~400 Mpbs (345,934,816 bps). When a port shaping of 100 Mbps is applied to the ge-0/0/0 interface, you see the following outputs for the **show interfaces ge-0/0/0 statistics** and the **show class-of-service interface ge-0/0/0** commands:

user@switch> show interfaces ge-0/0/0 statistics
Physical interface: ge-0/0/0, Enabled, Physical link is Up
Interface index: 239, SNMP ifIndex: 548, Generation: 242
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto, BPDU Error: None, MAC-REWRITE
Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Autonegotiation: Enabled, Remote fault: Online,

Media type: Copper Device flags : Present Running Interface flags: SNMP-Traps Internal: 0x0 Link flags : None CoS queues : 8 supported, 8 maximum usable queues Hold-times : Up 0 ms, Down 0 ms Current address: 00:23:9c:0b:ae:8d, Hardware address: 00:23:9c:0b:ae:8d Last flapped : 2012-07-07 03:21:52 UTC (1d 18:10 ago) Statistics last cleared: 2012-07-07 23:54:34 UTC (21:37:58 ago) Traffic statistics: Input bytes : 0 bps 0 Output bytes : 100223104 bps 15779512832 Input packets: 0 0 pps Output packets: 123277444 97874 pps IPv6 transit statistics: Input bytes : 0 Output bytes : 0 Input packets: 0 Output packets: 0 Input errors: Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0 Output errors: Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0 Egress queues: 8 supported, 4 in use Oueue counters: Queued packets Transmitted packets Dropped packets 0 best-effort 0 57012484 123350092 1 assured-forw 0 0 0 5 expedited-fo 0 0 0 7 network-cont 0 0 0 Queue number: Mapped forwarding classes 0 best-effort 1 assured-forwarding 5 expedited-forwarding 7 network-control Active alarms : None Active defects : None MAC statistics: Transmit Receive Total octets 15779512832 0 0 Total packets 123277444 0 123277444 Unicast packets 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			
Code violations	0			
Autonegotiation informati	ion:			
Negotiation status: Com	nplete			
Link partner:				
Link mode: Full-dup	olex, Flow control:	Symmetric, Remote	fault: OK, Link partner Speed:	
1000 Mbps				
Local resolution:				
Flow control: Symme	etric, Remote fault:	Link OK		
Packet Forwarding Engine	configuration:			
Destination slot: 1				
CoS information:				
Direction : Output				
CoS transmit queue	Bandwidth	Bu	ffer Priority Limit	
	% bps	%	usec	
0 best-effort	95 9500000	95	NA low none	
7 network-control	5 500000	5	NA low none	
Interface transmit statis	stics: Disabled			
Logical interface ge-1/0/	/10.0 (Index 69) (SN	MP ifIndex 638) (Generation 138)	
Flags: SNMP-Traps 0x0 E	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		bps	
Input packets:		•		
2.1940 946166601		0		
Output packets:	0		pps pps	

```
Protocol eth-switch, Generation: 163, Route table: 0
Flags: Trunk-Mode
user@switch> show class-of-service interface ge-0/0/0
Physical interface: ge-0/0/0, Index: 165
Queues supported: 8, Queues in use: 4
Shaping rate: 10000000 bps
...
...
```

Meaning

In the output for the show interfaces ge-0/0/0 statistics command, the Traffic statistics: field shows that egress traffic is ~100 Mbps (100,223,104 bps). The output for the show class-of-service interface ge-0/0/0 command shows that the shaping rate is 100,000,000 bps, which indicates that a port shaping of 100 Mbps is applied to the ge-0/0/0 interface.

Verifying That Queue Shaping Has Been Applied

Purpose

Verify that the queue shaping has been applied to the best-effort queue.

Action

Following is the output before queue shaping is applied to the best-effort queue when there is egress traffic of 400 Mpbs exiting on that interface:

```
user@switch> show interfaces ge-0/0/2 extensive
Physical interface: ge-0/0/2, Enabled, Physical link is Up
Interface index: 239, SNMP ifIndex: 548, Generation: 242
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto, BPDU Error: None, MAC-REWRITE
Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-
negotiation: Enabled, Remote fault: Online,
Media type: Copper
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : & supported, & maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
```

Current address: 00:23:9c:0b:ae:8d, Hardware address: 00:23:9c:0b:ae:8d Last flapped : 2012-07-07 03:21:52 UTC (1d 18:02 ago) Statistics last cleared: 2012-07-07 23:54:34 UTC (21:29:59 ago) Traffic statistics: Input bytes : 0 0 bps Output bytes : 2299853696 345934816 bps Input packets: 0 0 pps Output packets: 17967609 337827 pps IPv6 transit statistics: Input bytes : 0 Output bytes : 0 Input packets: 0 Output packets: 0 Input errors: Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0 Output errors: Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0 Egress queues: 8 supported, 4 in use Queue counters: Queued packets Transmitted packets Dropped packets 0 best-effort 0 18302337 0 1 assured-forw 0 0 0 5 expedited-fo 0 0 0 7 network-cont 0 0 0 Queue number: Mapped forwarding classes 0 best-effort 1 assured-forwarding 5 expedited-forwarding 7 network-control Active alarms : None Active defects : None MAC statistics: Receive Transmit Total octets 0 2299853696 Total packets 0 17967609 Unicast packets 0 17967609 Broadcast packets 0 0 Multicast packets 0 0 0 CRC/Align errors 0 0 FIFO errors 0 MAC control frames 0 0 0 MAC pause frames 0 Oversized frames 0

Jabber frames 0 Fragment frames 0 0 Code violations Autonegotiation information: Negotiation status: Complete Link partner: Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed: 1000 Mbps Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Destination slot: 1 CoS information: Direction : Output CoS transmit queue Bandwidth Buffer Priority Limit % bps % usec 0 best-effort 95 950000000 95 NA low none 7 network-control 5 50000000 5 NA low none Interface transmit statistics: Disabled Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Input packets: 0 Output packets: 0 Local statistics: Input bytes : 0 Output bytes : 0 Input packets: 0 Output packets: 0 Transit statistics: Input bytes : 0 0 bps Output bytes : 0 0 bps Input packets: 0 0 pps Output packets: 0 0 pps Protocol eth-switch, Generation: 163, Route table: 0 Flags: Trunk-Mode

The Traffic statistics: field in this output shows that the egress traffic is ~400 Mpbs (345,934,816 bps). When a queue shaping of 30 Mbps is applied to the best-effort queue, you see the following output for the show interfaces ge-0/0/2 statistics and show class-of-service scheduler-map sched-map-be commands:

user@switch> show inte Physical interface: ge	-			
Interface index: 239	, SNMP ifIndex: 548,	Generation: 242		
Link-level type: Eth	ernet, MTU: 1514, Sp	eed: Auto, Duplex:	Auto, BPDU Error: None	e, MAC-REWRITE
Error: None, Loopback:	Disabled, Source fi	ltering: Disabled,	Flow control: Enabled,	, Auto-
negotiation: Enabled,	Remote fault: Online	,		
Media type: Copper				
Device flags : Pre	sent Running			
Interface flags: SNM	P-Traps Internal: 0x	0		
Link flags : Non	e			
CoS queues : 8 s	upported, 8 maximum	usable queues		
Hold-times : Up	0 ms, Down 0 ms			
Current address: 00:	23:9c:0b:ae:8d, Hard	ware address: 00:23	3:9c:0b:ae:8d	
Last flapped : 201	2-07-07 03:21:52 UTC	(1d 18:29 ago)		
Statistics last clea	red: 2012-07-08 21:4	6:22 UTC (00:04:56	ago)	
Traffic statistics:				
Input bytes :	0	0	bps	
Output bytes :	5376128896	30097712	bps	
Input packets:	0	0	pps	
Output packets:	42001003	29392	pps	
IPv6 transit statis	tics:			
Input bytes :	0			
Output bytes :	0			
Input packets:	0			
Output packets:	0			
Input errors:				
Errors: 0, Drops:	0, Framing errors: 0	, Runts: 0, Policed	d discards: 0, L3 incom	npletes: 0, L2
channel errors: 0, L2	mismatch timeouts: 0	, FIFO errors: 0, F	Resource errors: 0	
Output errors:				
Carrier transition	s: 0, Errors: 0, Dro	os: 0, Collisions:	0, Aged packets: 0, FI	[FO errors: 0,
HS link CRC errors: 0,	MTU errors: 0, Reso	urce errors: 0		
Egress queues: 8 sup	ported, 4 in use			
Queue counters:	Queued packets Tra	ansmitted packets	Dropped packets	
0 best-effort	0	41986978	57813642	
1 assured-forw	0	0	0	
5 expedited-fo	0	0	0	
7 network-cont	0	0	0	
Queue number:	Mapped forwarding	classes		

0 best-effort 1 assured-forwarding 5 expedited-forwarding 7 network-control Active alarms : None Active defects : None MAC statistics: Receive Transmit Total octets 0 5376128896 Total packets 0 42001003 Unicast packets 0 42001003 Unicast packets 0 42001003 Unicast packets 0 42001003 Unicast packets 0 0 Multicast packets 0 0 CRC/Align errors 0 0 MAC pause frames 0 0 CAC contol frames 0 0 Code violations 0 Autonegotiation information: Negotiation status: Complet Link partner: Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed: 1000 Mbps Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Direction : Output CoS transmit queue Bandwidth Buffer Priority Limit							
5 expedited-forwarding 7 network-control Active alarms : None Active defects : None Active defects : None MC statistics: Receive Transmit Total octets 0 5376128896 Total packets 0 42001003 Broadcast packets 0 0 ML cast packets 0 0 Multicast packets 0 0 ML cance packets 0 0 ML cance packets 0 0 ML cance packets 0 0 ML control frames 0 0 MAC pause frames 0 0 Jabber frames 0 0 Code violation 0 0 Autonegotiation information: Negotiation stutus: Complete 1 Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed: 1000 Mops Local resolution: Flags: Sinformation: Direction : Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: <td>0</td> <td>best-effort</td> <td></td> <td></td> <td></td> <td></td> <td></td>	0	best-effort					
7 network-control Active alarms: None Active defects: None MC statistics: Receive Transmit Total octets 0 5376128896 Total packets 0 42001003 Unicast packets 0 42001003 Unicast packets 0 42001003 Broadcast packets 0 0 Multicast packets 0 0 Jabber frames 0 0	1	assured-forwarding					
Active alarms : None Active defects : None MC statistics: Receive Transmit Total octets 0 5376128896 Total packets 0 42001003 Unicast packets 0 42001003 Broadcast packets 0 0 Multicast packets 0 0 Multicast packets 0 0 CRC/Align errors 0 0 MC control frames 0 0 MC control frames 0 0 MC pause frames 0 0 Oversized frames 0 Tragment frames 0 Code violations 0 Autonegotiation information: Negotiation information: Negotiation information: Negotiation information: Negotiation information: Negotiation information: Negotiation status: Complete Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed: 1000 Mpos Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Direction : Output CoS information: Direction : Output CoS informati statistics: Disabled Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Input packets: 0	5	expedited-forwarding					
Active defects : None MAC statistics: Receive Transmit Total octets 0 5376128896 Total packets 0 42001003 Unicast packets 0 0 Multicast packets 0 0 Broadcast packets 0 0 Multicast packets 0 0 Jabber frames 0 0 Code violations 0	7	network-control					
MAC statistics:ReceiveTransmitTotal octets05376128896Total packets042001003Unicast packets00Broadcast packets00Multicast packets00Multicast packets00CRC/Align errors00MAC control frames00MAC control frames00Oversized frames00Jabber frames00Code violations00Autonegotiation information:Negotiation status: CompleteLink mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed:1000 Mbps1Local resolution:Flow control: Symmetric, Remote fault: Unix OKPacket Forwarding Engine configuration:Destination slot: 1CoS transmit queueBandwidthBuffer Priority Limit%bps %usec00 best-effortrrrrNa low noneInterface transmit statistics: DisabledLogical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138)Flags: SNMP-Traps 0x0 Encapsulation: ENET2Traffic statistics:0Input bytes :0Output bytes :0Output bytes :0Input packets:0Input packets:0	Active alarms : None						
Total octets 0 5376128896 Total packets 0 42001003 Unicast packets 0 0 Broadcast packets 0 0 Multicast packets 0 0 CRCAAlign errors 0 0 MAC control frames 0 0 MAC control frames 0 0 MAC pause frames 0 0 Jabber frames 0 0 Code violations 0 0 Katonegotiation information: Negotiation status: Complete 1 Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed: 1000 Mbps Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Destination slot: 1 CoS information: Direction : Output K CoS transmit queue Bandwidth Buffer Priority Limit % <td< td=""><td>Active defects : None</td><td></td><td></td><td></td><td></td><td></td><td></td></td<>	Active defects : None						
Total packets 0 42001003 Unicast packets 0 0 Multicast packets 0 Multi	MAC statistics:	Receiv	/e	Transmit			
Unicat packets 0 42001003 Broadcast packets 0 0 Multicast packets 0 0 Multicast packets 0 0 Multicast packets 0 0 Multicast packets 0 0 FIFD errors 0 0 MAC control frames 0 0 MAC pause frames 0 0 Oversized frames 0 0 Jabber frames 0 0 Autonegotiation information: 0 0 Negotiation information: 0 0 Negotiation information: 0 0 Negotiation stuts: Complete Link partner: 1 Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed: 1000 Mbps Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: 0 Destination slot: 1 CoS transmit queue Bandwidth Buffer Priority Limit CoS transmit queue Bandwidth Buffer Priority Lim	Total octets		0 !	5376128896			
Broadcast packets 0 0 Multicast packets 0 0 CRC/Align errors 0 0 FIF0 errors 0 0 MAC control frames 0 0 MAC pause frames 0 0 Oversized frames 0 0 Jabber frames 0 0 Jabber frames 0 0 Code violations 0 0 Katonegotiation information: Negotiation status: Complete Link partner: Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed: 1000 Mbps Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Destination slot: 1 CoS transmit queue Bandwidth Buffer Priority Limit % bps % usec 0 best-effort r r r Interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Output bytes : 0	Total packets		0	42001003			
Multicast packets 0 0 CRC/Align errors 0 0 FIF0 errors 0 0 MAC control frames 0 0 MAC pause frames 0 0 Multicast frames 0 0 Oversized frames 0 0 Jabber frames 0 0 Gode violations 0 0 Fragment frames 0 0 Code violations 0 0 Autonegotiation information: Negotiation status: Complete 0 Link partner: Link partner: 1 Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed: 1000 Mbps Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Destination slot: 1 CoS information: Direction : Output User CoS transmit queue Bandwidth Buffer Priority Limit \$\$\scillatteriations : Disabled\$ User 0 Interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SMMP-Traps 0x0 Encapsulation: ENET2 Traffic statisti	Unicast packets		0	42001003			
CRC/Align errors00FIFO errors00MAC control frames00MAC pause frames00Oversized frames00Jabber frames00Code violations00Code violations00Autonegotiation information:0Negotiation status: Complete0Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed:1000 Mbps1000 mbpsLocal resolution:Flow control: Symmetric, Remote fault: Link OKPacket Forwarding Engine configuration:Destination slot: 1Cos information:Differ Priority LimitMeretion:% bps % usec00Interface transmit statistics: DisabledLogical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138)Flags: SNMP-Traps 0x0 Encapsulation: ENET2Traffic statistics:Input bytes :0Output bytes :0Input packets:0	Broadcast packets		0	0			
FIFO errors00MAC control frames00MAC pause frames00Oversized frames00Jabber frames00Jabber frames00Code violations00Autonegotiation information:0Negotiation status: Complete0Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed:1000 Mbps1000 MbpsLocal resolution:Flow control: Symmetric, Remote fault: Link OKPacket Forwarding Engine configuration:Destination slot: 1CoS information:Differ Priority LimitDirection : OutputCoS transmit queueCoS transmit queueBandwidthBuffer Priority Limit x bps%usec00best-effortrrrNAInterface transmit statistics: DisabledLogical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138)Flags: SNMP-Traps 0x0 Encapsulation: ENET2Traffic statistics:Input bytes :0Output bytes :0Input packets:0	Multicast packets		0	0			
MAC control frames 0 0 MAC pause frames 0 0 Oversized frames 0 0 Jabber frames 0 0 Jabber frames 0 0 Fragment frames 0 0 Code violations 0 0 Autonegotiation information: 0 0 Negotiation status: Complete 0 0 Link partner: Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed: 1000 Mbps Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Destination slot: 1 CoS information: Direction : Output CoS transmit queue Bandwidth Buffer Priority Limit % bps % usec 0 0 0 best-effort r r r N A low none Interface transmit statistics: Disabled Iogical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 0 0 Input packets: 0	CRC/Align errors		0	0			
MAC pause frames 0 0 Oversized frames 0 0 Jabber frames 0 0 Fragment frames 0 0 Code violations 0 0 Autonegotiation information: 0 0 Negotiation status: Complete 0 0 Link partner: 1 1 Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed: 1000 Mbps Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Destination slot: 1 CoS information: Direction : Output CoS transmit queue Bandwidth Buffer Priority Limit % bps % usec 0 best-effort r r Interface transmit statistics: Disabled 1 1 Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 0 Input packets: 0 1 1	FIFO errors		0	0			
Oversized frames0Jabber frames0Fragment frames0Code violations0Autonegotiation information:0Negotiation status: Complete1Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed:1000 MbpsLocal resolution:Flow control: Symmetric, Remote fault: Link OKPacket Forwarding Engine configuration:Destination slot: 1CoS information:Direction : OutputCoS transmit queueBandwidthBuffer Priority Limit%bps %usec0 best-effortrrrInterface transmit statistics: DisabledLogical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138)Flags: SNMP-Traps 0x0 Encapsulation: ENET2Traffic statistics:Input bytes :0Output bytes :0Input packets:0	MAC control frames		0	0			
Jabber frames 0 Fragment frames 0 Code violations 0 Autonegotiation information: Negotiation status: Complete Link partner: Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed: 1000 Mbps Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Destination slot: 1 CoS information: Direction : Output CoS transmit queue Bandwidth Buffer Priority Limit % bps % usec 0 best-effort r r r NA low none Interface transmit statistics: Disabled Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Input packets: 0	MAC pause frames		0	0			
Fragment frames 0 Code violations 0 Autonegotiation information: Negotiation status: Complete Link partner: Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed: 1000 Mbps Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Destination slot: 1 CoS information: Direction : Output CoS transmit queue Bandwidth Buffer Priority Linterface transmit statistics: Disabled Logical interface ge-1/0/10.0 (Index 69) (SNMP if Index 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Uptub tytes : 0 Input packets:	Oversized frames		0				
Code violations 0 Autonegotiation information: Negotiation status: Complete Link partner: Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed: 1000 Mbps Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Destination slot: 1 CoS information: Direction : Output CoS transmit queue Bandwidth Buffer Priority Limit % bps % 0 best-effort r Interface transmit statistics: Disabled Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0	Jabber frames		0				
Autonegotiation information: Negotiation status: Complete Link partner: Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed: 1000 Mbps Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Destination slot: 1 CoS information: Direction : Output CoS transmit queue Bandwidth Buffer Priority Limit % bps % 0 best-effort r r r Number of the statistics: Disabled Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0	Fragment frames		0				
Negotiation status: Complete Link partner: Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed: 1000 Mbps Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Destination slot: 1 CoS information: Direction : Output CoS transmit queue Bandwidth Buffer Priority Limit % bps % usec 0 best-effort r r r NA low none Interface transmit statistics: Disabled Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Input packets: 0	Code violations		0				
Link partner: Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed: 1000 Mbps Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Destination slot: 1 CoS information: Direction : Output CoS transmit queue Bandwidth Buffer Priority Limit % bps % usec 0 best-effort r r r NA low none Interface transmit statistics: Disabled Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Input packets: 0	Autonegotiation infor	mation:					
Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link partner Speed: 1000 Mbps Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Destination slot: 1 CoS information: Direction : Output CoS transmit queue Bandwidth Buffer Priority Limit % bps % usec 0 best-effort r r r NA low none Interface transmit statistics: Disabled Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Input packets: 0	Negotiation status:	Complete					
1000 Mbps Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Destination slot: 1 CoS information: Direction : Output CoS transmit queue Bandwidth Buffer Priority Limit % bps % usec 0 best-effort r r r NA low none Interface transmit statistics: Disabled Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Input packets: 0	Link partner:						
Local resolution: Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Destination slot: 1 CoS information: Direction : Output CoS transmit queue Bandwidth Buffer Priority Limit % bps % usec 0 best-effort r r r NA low none Interface transmit statistics: Disabled Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Input packets: 0	Link mode: Full	-duplex, Flow control:	Symmetr	ic, Remote faul	t: OK, Lir	nk partner	Speed:
Flow control: Symmetric, Remote fault: Link OK Packet Forwarding Engine configuration: Destination slot: 1 CoS information: Direction : Output CoS transmit queue Bandwidth Buffer Priority Limit % bps 0 best-effort r r NA low none Interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Input packets: 0	1000 Mbps						
Packet Forwarding Engine configuration: Destination slot: 1 CoS information: Direction : Output CoS transmit queue Bandwidth Buffer Priority Limit % bps % 0 best-effort r r r NA low Interface transmit statistics: Disabled Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Input packets: 0	Local resolution:						
Destination slot: 1 CoS information: Direction : Output CoS transmit queue Bandwidth Buffer Priority Limit % bps % usec 0 best-effort r r r NA low none Interface transmit statistics: Disabled Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Input packets: 0	Flow control: S	ymmetric, Remote fault	: Link Oł	<			
CoS information: Direction : Output CoS transmit queue Bandwidth Buffer Priority Limit CoS transmit queue Bandwidth Buffer Priority Limit % bps % usec Ø best-effort r r r NA low none Interface transmit statistics: Disabled Interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Output bytes : 0 Input packets: 0 Output packets: 0 Output packets: 0	Packet Forwarding Eng	ine configuration:					
Direction : Output CoS transmit queue Bandwidth Buffer Priority Limit % bps % usec 0 best-effort r r r NA low none Interface transmit statistics: Disabled Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Input packets: 0	Destination slot: 1						
CoS transmit queueBandwidthBuffer PriorityLimit%bps%usec0 best-effortrrrNAlowInterface transmit statistics:DisabledLogical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138)Flags:SNMP-Traps 0x0Encapsulation:ENET2Traffic statistics:0Output bytes :0Input packets:0	CoS information:						
% bps % usec 0 best-effort r r r NA low none Interface transmit statistics: Disabled Interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: 0 0utput bytes : 0 0 Input packets: 0 0 0 0 0	Direction : Output						
<pre>0 best-effort r r r NA low none Interface transmit statistics: Disabled Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Input packets: 0</pre>	CoS transmit queue	Bandwidt	:h	Buffer	Priority	Limit	
Interface transmit statistics: Disabled Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Input packets: 0		% br	os %	usec			
Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138) Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Input packets: 0	0 best-effort	r	r r	NA	low	none	
Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Input packets: 0	Interface transmit st	atistics: Disabled					
Flags: SNMP-Traps 0x0 Encapsulation: ENET2 Traffic statistics: Input bytes : 0 Output bytes : 0 Input packets: 0							
Traffic statistics:Input bytes :0Output bytes :0Input packets:0	Logical interface ge-	1/0/10.0 (Index 69) (S	SNMP ifInd	dex 638) (Gener	ation 138))	
Input bytes :0Output bytes :0Input packets:0	Flags: SNMP-Traps 0	x0 Encapsulation: ENET	2				
Output bytes :0Input packets:0	Traffic statistics:						
Input packets: 0	Input bytes :	0					
	Output bytes :	0					
Output packets: 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 eth-switch, Ge	eneration: 163, Route table:	0	
Flags: Trunk-Mode			

```
user@switch> show class-of-service scheduler-map sched-map-be
Scheduler map: sched-map-be, Index: 31271
 Scheduler: voice-sched-queue-shap, Forwarding class: best-effort, Index: 64106
   Transmit rate: remainder, Rate Limit: none, Buffer size: remainder,
   Buffer Limit: none, Priority: low
   Excess Priority: unspecified
   Shaping rate: 3000000 bps
   Drop profiles:
     Loss priority Protocol
                                 Index
                                          Name
     High
                     non-TCP
                                     1
                                          <default-drop-profile>
                     TCP
                                     1
                                          <default-drop-profile>
     High
```

Meaning

In the output for the show interfaces ge-0/0/2 statistics command, the Traffic statistics: field shows that the egress traffic is ~30 Mbps (30,097,712 bps). The output for the show class-of-service scheduler-map sched-map-be command, shows that a shaping rate of 30,000,000 bps (that is 30 Mbps) is applied to the best-effort queue.

RELATED DOCUMENTATION

Defining CoS Code-Point Aliases (CLI Procedure) | 65 Defining CoS Classifiers (CLI Procedure) | 74 Configuring Firewall Filters (CLI Procedure)

Junos OS EZQoS and J-Web

IN THIS CHAPTER

- Understanding Junos OS EZQoS for CoS Configurations on EX Series Switches | 51
- Configuring Junos OS EZQoS for CoS (CLI Procedure) | 52
- Configuring CoS (J-Web Procedure) | 53

Understanding Junos OS EZQoS for CoS Configurations on EX Series Switches

Junos operating system (Junos OS) EZQoS on Juniper Networks EX Series Ethernet Switches eliminates the complexities involved in configuring *class of service* (CoS) across the network. EZQoS offers templates for key traffic classes.

Junos OS CoS allows you to divide traffic into classes and offer various levels of throughput and packet loss when congestion occurs. You can use CoS to ensure that different types of traffic (voice, video, and data) get the bandwidth and consideration they need to meet user expectations and business objectives.

Configuring CoS requires careful consideration of your service needs and thorough planning and design to ensure consistency across all switches in a CoS domain. To configure CoS manually, you must define and fine-tune all CoS components such as classifiers, *rewrite rules*, forwarding classes, schedulers, and scheduler-maps and then apply these components to the interfaces. Therefore, configuring CoS can be a fairly complex and time-consuming task.

EZQoS works by automatically assigning preconfigured values to all CoS parameters based on the typical application requirements. These preconfigured values are stored in a template with a unique name. You can change the preconfigured values of these parameters to suit your particular application needs.

For using EZQoS, you must identify which switch ports are being used for a specific application (such as VoIP, video, and data) and manually apply the corresponding application-specific EZQoS template to these switch ports.

NOTE: Currently, we provide an EZQoS template for configuring CoS for VoIP.

NOTE: We recommend that you do not use the term EZQoS for defining a classifier.

RELATED DOCUMENTATION

Junos OS CoS for EX Series Switches Overview | 2

Configuring Junos OS EZQoS for CoS (CLI Procedure) | 52

Configuring Junos OS EZQoS for CoS (CLI Procedure)

You use Junos OS EZQoS on EX Series switches to eliminate the complexities involved in configuring class of service (CoS) across the network. EZQoS offers templates for key traffic classes.

When you configure EZQoS on EX Series switches, preconfigured values are assigned to all CoS parameters based on the typical application requirements. These preconfigured values are stored in a template with a unique name.

NOTE: Currently, we provide an EZQoS template for configuring CoS for VoIP applications. The EZQoS VoIP template is stored in **/etc/config/ezqos-voip.conf**.

To configure EZQoS using the CLI:

1. Load the EZQoS configuration file (/etc/config/ezqos-voip.conf):

[edit]
user@switch# load merge /etc/config/ezqos-voip.conf

2. Apply the EZQoS group (ezqos-voip):

[edit]
user@switch# set apply-groups ezqos-voip

3. Apply the DSCP classifier (ezqos-dscp-classifier) to a Gigabit Ethernet interface (ge-0/0/0):

[edit class-of-service interfaces]
user@switch# set ge-0/0/0 unit 0 classifiers dscp ezqos-dscp-classifier

4. Apply the scheduler map (ezqos-voip-sched-maps) to a Gigabit Ethernet interface (ge-0/0/1):

```
[edit class-of-service interfaces]
user@switch# set ge-0/0/1 scheduler-map ezqos-voip-sched-maps
```

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches | **15** Understanding Junos OS EZQoS for CoS Configurations on EX Series Switches | **51**

Configuring CoS (J-Web Procedure)

The Class of Service Configuration pages allow you to configure the Junos CoS components. You can configure forwarding classes for transmitting packets, define which packets are placed into each output queue, and schedule the transmission service level for each queue. After defining the CoS components you must assign classifiers to the required physical and logical interfaces.

Using the Class of Service Configuration pages, you can configure various CoS components individually or in combination to define particular CoS services.

To configure CoS components :

- 1. In the J-Web interface, select Configure>Class of Service.
- **2.** On the Class of Service Configuration page, select one of the following options depending on the CoS component that you want to define. Enter information into the pages as described in the respective table:
 - To define or edit CoS value aliases, select CoS Value Aliases .
 - To define or edit forwarding classes and assign queues, select Forwarding Classes.
 - To define or edit classifiers, select Classifiers .
 - To define or edit rewrite rules, select Rewrite Rules.

- To define or edit schedulers, select **Schedulers**.
- To define or edit virtual channel groups, select Interface Associations.
- **3.** Click **Apply** after completing configuration on any Configuration page.

RELATED DOCUMENTATION

Defining CoS Classifiers (J-Web Procedure) | 77 Defining CoS Code-Point Aliases (J-Web Procedure) | 66 Defining CoS Forwarding Classes (J-Web Procedure) | 115 Defining CoS Rewrite Rules (J-Web Procedure) | 100 Defining CoS Schedulers (J-Web Procedure) | 141 Assigning CoS Components to Interfaces (J-Web Procedure) | 56

CoS on Interfaces

IN THIS CHAPTER

- Assigning CoS Components to Interfaces (CLI Procedure) | 55
- Assigning CoS Components to Interfaces (J-Web Procedure) | 56
- Monitoring Interfaces That Have CoS Components | 58

Assigning CoS Components to Interfaces (CLI Procedure)

After you have defined the following CoS components, you must assign them to logical or physical interfaces.

- Forwarding classes—Assign only to logical interfaces.
- Classifiers-Assign only to logical interfaces.
- Scheduler maps—Assign to either physical or logical interfaces.
- Rewrite rules—Assign to either physical or logical interfaces.

You can assign a CoS component to a single interface or to multiple interfaces using wild cards.

To assign CoS components to interfaces:

• To assign CoS components to a single interface, associate a CoS component (for example a scheduler map named ethernet-cos-map) with an interface:

[edit class-of-service interfaces]
user@switch# set ge-0/0/20 scheduler-map ethernet-cos-map

 To assign a CoS component to multiple interfaces, associate a CoS component (for example, a rewrite rule named customup-rw) to all Gigabit Ethernet interfaces on the switch, use wild characters for the interface name and logical-interface (unit) number:

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

RELATED DOCUMENTATION

Assigning CoS Components to Interfaces (J-Web Procedure) | 56 Example: Configuring CoS on EX Series Switches | 15 Monitoring Interfaces That Have CoS Components | 58 Understanding Junos OS CoS Components for EX Series Switches | 8

Assigning CoS Components to Interfaces (J-Web Procedure)

NOTE: This topic applies only to the J-Web Application package.

After you have defined CoS components on an EX Series switch, you must assign them to logical or physical interfaces. You can use the J-Web interface to assign scheduler maps to physical or logical interfaces and to assign forwarding classes or classifiers to logical interfaces.

To assign CoS components to interfaces:

1. Select Configure > Class of Service > Assign to Interface.

NOTE: After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See Using the Commit Options to Commit Configuration Changes (J-Web Procedure) for details about all commit options.

- **2.** To configure an interface association, select an interface from the list and click **Edit**. For an EX8200 Virtual Chassis configuration, select the member, the FPC, and the interface from the list, and click **Edit**.
- **3.** Select one of the following:

- Associate system default scheduler map—Associates the interface with the default scheduler map.
- Select the scheduler map—Associates the interface with a configured scheduler map. Select the scheduler map from the list.

NOTE: On the 40-port SFP+ line card for EX8200 switches, you cannot commit your changes using the J-Web interface unless you assign the same scheduler map or the default scheduler map to all interfaces in a port group.

4. Click OK.

- 5. To manage a CoS assignment on a logical interface, Click one of the following options:
 - Add—Adds a CoS service to a logical interface on a specified physical interface. Enter information as described in Table 5 on page 57.
 - Edit—Modifies a CoS service assignment to a logical interface. Enter information as described in Table 5 on page 57.
 - **Delete**-Deletes the CoS service assignment to a logical interface.

Table 5: Assigning CoS Components to Logical Interfaces

Field	Function	Your Action
Unit	Specifies the name of a logical interface. Enables you to assign CoS components when you configure a logical interface on a physical interface.	Type the interface name. To assign CoS to all logical interfaces configured on this physical interface, type the wildcard character (*).
Forwarding Class	Assigns a predefined forwarding class to incoming packets on a logical interface.	To assign a forwarding class to an interface, select the forwarding class.
Classifiers	Enables you to apply classification maps to a logical interface. Classifiers assign a forwarding class and loss priority to an incoming packet based on its CoS value.	To assign a classification map to an interface, select an appropriate classifier for each CoS value type used on the interface.

Field	Function	Your Action
Rewrite Rules	Enables you to alter the CoS values in outgoing packets to meet the requirements of the targeted peer. A rewrite rule examines the forwarding class and loss priority of a packet and sets its bits to a corresponding value specified in the rule.	To assign rewrite rules to the interface, select the appropriate rewrite rule for each CoS value type used on the interface.

Table 5: Assigning CoS Components to Logical Interfaces (Continued)

RELATED DOCUMENTATION

Assigning CoS Components to Interfaces (CLI Procedure) 55			
Example: Configuring CoS on EX Series Switches 15			
Monitoring Interfaces That Have CoS Components 58			

Monitoring Interfaces That Have CoS Components

IN THIS SECTION

- Purpose | 58
- Action | **59**
- Meaning | 59

Purpose

NOTE: This topic applies only to the J-Web Application package.

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 J-Web interface, select **Monitor** > **Class of Service** > **Interface Association**.

To monitor interfaces that have CoS components in the CLI, enter the following command:

show class-of-service interface interface

Meaning

Table 6 on page 59 summarizes key output fields for CoS interfaces.

Table 6: Summary of Key CoS Interfaces Output Fields

Field	Values	Additional Information
Interface	Name of a physical interface to which CoS components are assigned.	To display names of logical interfaces configured on this physical interface, click the plus sign (+).
Scheduler Map	Name of the scheduler map associated with this interface.	
Queues Supported	Number of queues you can configure on the interface.	
Queues in Use	Number of queues currently configured.	
Logical Interface	Name of a logical interface on the physical interface to which CoS components are assigned.	
Object	Category of an object—for example, classifier, scheduler-map, or rewrite.	
Name	Name that you have given to an object —for example, ba-classifier .	

Field	Values	Additional Information
Туре	Type of an object—for example, dscp for a classifier.	
Index	Index of this interface or the internal index of a specific object.	

Table 6: Summary of Key CoS Interfaces Output Fields (Continued)

RELATED DOCUMENTATION

Assigning CoS Components to Interfaces (CLI Procedure) 55
Assigning CoS Components to Interfaces (J-Web Procedure) 56
Example: Configuring CoS on EX Series Switches 15

CoS Code-Point Aliases

IN THIS CHAPTER

- Understanding CoS Code-Point Aliases | 61
- Defining CoS Code-Point Aliases (CLI Procedure) | 65
- Defining CoS Code-Point Aliases (J-Web Procedure) | 66
- Monitoring CoS Value Aliases | 67

Understanding CoS Code-Point Aliases

IN THIS SECTION

Default Code-Point Aliases | 62

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, drop-profile maps, 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), IP precedence, and IEEE 802.1p bits to associate incoming packets with a particular CoS servicing level. On a switch, you can assign a meaningful name or alias to the CoS values and use this 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 classes and define classifiers, you can refer to the markers by alias names. You can configure user-defined classifiers in terms of alias names. If the value of an alias changes, it alters the behavior of any classifier that references it.

This topic covers:

Default Code-Point Aliases

Table 7 on page 62 shows the default mappings between the bit values and standard aliases.

CoS Value Types	Mapping

DSCP CoS Values

ef	101110
af11	001010
af12	001100
af13	001110
af21	010010
af22	010100
af23	010110
af31	011010
af32	011100

CoS Value Types	Mapping
af33	011110
af41	100010
af42	100100
af43	100110
be	000000
cs1	001000
cs2	010000
cs3	011000
cs4	100000
cs5	101000
nc1/cs6	110000
nc2/cs7	111000

Table 7: Default Code-Point Aliases (Continued)

IEEE 802.1p CoS Values

be	000
be1	001

CoS Value Types	Mapping
ef	100
ef1	101
af11	010
af12	011
nc1/cs6	110
nc2/cs7	111

Table 7: Default Code-Point Aliases (Continued)

Legacy IP Precedence CoS Values

be	000
be1	001
ef	010
ef1	011
af11	100
af12	101
nc1/cs6	110
nc2/cs7	111

RELATED DOCUMENTATION

Understanding Junos OS CoS Components for EX Series Switches | 8 Example: Configuring CoS on EX Series Switches | 15 Defining CoS Code-Point Aliases (CLI Procedure) | 65 Defining CoS Code-Point Aliases (J-Web Procedure) | 66

Defining CoS Code-Point Aliases (CLI Procedure)

You can use code-point aliases to streamline the process of configuring CoS features on your EX Series 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, drop-profile maps, and rewrite rules.

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

- dscp and dscp-ipv6-Handles incoming IPv4 and IPv6 packets, respectively.
- ieee-802.1—Handles Layer 2 CoS.
- inet-precedence—Handles incoming IPv4 packets. IP precedence mapping requires only the higher order three bits of the DSCP field.

To configure a code-point alias for a specified CoS marker type (**dscp**), assign an alias (**my1**) to the codepoint (**110001**):

[edit class-of-service code-point-aliases]
user@switch# set dscp my1 110001

The my1 alias will be applicable for incoming IPv4 packets.

RELATED DOCUMENTATION

Defining CoS Code-Point Aliases (J-Web Procedure) | 66

Example: Configuring CoS on EX Series Switches | 15

Monitoring CoS Value Aliases | 67

Understanding CoS Code-Point Aliases | 61

Defining CoS Code-Point Aliases (J-Web Procedure)

NOTE: This topic applies only to the J-Web Application package.

You can use the J-Web interface to define CoS code-point aliases on an EX Series switch. By defining aliases, you can assign meaningful names to a particular set of bit values and refer to them when configuring CoS components.

To define CoS code-point aliases:

1. Select Configure > Class of Service > CoS Value Aliases.

NOTE: After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See Using the Commit Options to Commit Configuration Changes (J-Web Procedure) for details about all commit options.

- 2. Click one of the following options:
 - Add—Adds a code-point alias. Enter information into the code point alias page as described in Table 8 on page 66.
 - Edit—Modifies an existing code-point alias. Enter information into the code point alias page as described in Table 8 on page 66.
 - Delete-Deletes an existing code-point alias.

Table 8 on page 66 describes the related fields.

Table 8: CoS Value Aliases Configuration Fields

Field	Function	Your Action
Code point name	Specifies the name for a code-point—for example, af11 or be .	Enter a name.
Code point type	Specifies a code-point type. The code-point type can be DSCP or IP precedence.	Select a value.

Field	Function	Your Action
Code point value bits	Specifies the CoS value for which an alias is defined. Changing this value alters the behavior of all classifiers that refer to this alias.	 To specify a CoS value, type it in the appropriate format: For DSCP CoS values, use the format xxxxx, where x is 1 or 0-for example, 101110. For IP precedence CoS values, use the format xxx, where x is 1 or 0-for example, 111.

Table 8: CoS Value Aliases Configuration Fields (Continued)

RELATED DOCUMENTATION

Defining CoS Code-Point Aliases (CLI Procedure) 65
Monitoring CoS Value Aliases 67
Example: Configuring CoS on EX Series Switches 15

Monitoring CoS Value Aliases

IN THIS SECTION

- Purpose | 67
- Action | **68**
- Meaning | 68

Purpose

NOTE: This topic applies only to the J-Web Application package.

Use the monitoring functionality to display information about the CoS value aliases that the system is currently using to represent DSCP, IEEE 802.1p, and IPv4 precedence bits.

Action

To monitor CoS value aliases in the J-Web interface, select **Monitor > Class of Service > CoS Value Aliases**.

To monitor CoS value aliases in the CLI, enter the following command:

show class-of-service code-point-aliases

Meaning

Table 9 on page 68 summarizes key output fields for CoS value aliases.

Table 9: Summary of Key CoS Value Alias Output Fields

Field	Values	Additional Information
CoS Value Type	 Type of the CoS value: dscp-Examines Layer 3 packet headers for IP packet classification. ieee-802.1-Examines Layer 2 packet headers for packet classification. inet-precedence-Examines Layer 3 packet headers for IP packet classification. 	To display aliases and bit patterns, click the plus sign (+).
CoS Value Alias	Name given to a set of bits—for example, af11 is a name for 001010 bits.	
CoS Value	Set of bits associated with an alias.	

RELATED DOCUMENTATION

Defining CoS Code-Point Aliases (CLI Procedure) | 65

Defining CoS Code-Point Aliases (J-Web Procedure) | 66

Example: Configuring CoS on EX Series Switches | 15

CoS Classifiers

IN THIS CHAPTER

- Understanding CoS Classifiers | 70
- Defining CoS Classifiers (CLI Procedure) | 74
- Defining CoS Classifiers (J-Web Procedure) | 77
- Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers | 80
- Configuring and Applying IEEE 802.1ad Classifiers | 83
- Configuring the IEEE 802.1p Field for CoS Host Outbound Traffic | 85
- Configuring a Global Default IEEE 802.1p Value for All Host Outbound Traffic | 86
- Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure) | 87
- Monitoring CoS Classifiers | 88
- Troubleshooting a CoS Classifier Configuration for a TCAM Space Error | 91

Understanding CoS Classifiers

IN THIS SECTION

- Behavior Aggregate Classifiers | 71
- Multifield Classifiers | 74

Packet classification associates incoming packets with a particular class-of-service (CoS) servicing level. Classifiers associate packets with a forwarding class and loss priority, and packets are associated to an output queue based on the forwarding class. You can define classifiers for the following interfaces:

 IPv4 and IPv6 traffic to network interfaces, aggregated Ethernet interfaces (also known as link aggregation groups (LAGs))

- On switches that support the ELS configuration style, inter-VLAN routing functions use an integrated routing and bridging (IRB) interface named irb
- On switches that do not support the ELS configuration style, inter-VLAN routing functions use a routed VLAN interface (RVI) named vlan

There are two general types of classifiers:

- Behavior aggregate (BA) classifiers
- Multifield (MF) classifiers

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. If the two classification results conflict, the MF classification result overrides the BA classification result.

On Juniper Networks EX8200 Ethernet Switches, you can specify BA classifiers for bridged multidestination traffic and for IP multidestination traffic. A BA classifier for multicast packets is applied to all interfaces on the EX8200 switch.

NOTE: EX8200 switches implement the on-demand allocation of memory space for ternary content addressable memory (TCAM) so that when additional TCAM space is required for CoS classifiers, it is allocated from the free TCAM space or from the unused TCAM space. An error log message is generated when you configure CoS classifiers to use memory space that exceeds the available TCAM space that includes both the free and unused space.

This topic describes:

Behavior Aggregate Classifiers

BA classifiers are based on fixed-length fields in the packet header, which makes them computationally more efficient than MF classifiers. Therefore core devices that handle high traffic volumes are normally configured to perform BA classification. The BA classifier maps packets to a forwarding class and a loss priority. The forwarding class determines the output queue for a packet. The loss priority is used by a scheduler to control packet discard during periods of congestion.

These are the following types of BA classifiers:

- dscp-Differentiated Services Code Point (DSCP) for IP DiffServ. Handles incoming IPv4 packets.
- dscp-ipv6—Handles incoming IPv6 packets.
- ieee-802.1—Handles Layer 2 CoS (IEEE 802.1p).

• inet-precedence—Handles incoming IPv4 packets. IP precedence mapping requires only the upper three bits of the DSCP field.

A BA classifier takes a specified CoS value as either the literal bit pattern or as a defined alias and attempts to match it to the type of packet arriving on the interface. If the information in the packet's header matches the specified pattern, the packet is sent to the appropriate queue, defined by the forwarding class associated with the classifier.

Default Behavior Aggregate Classification

Juniper Networks Junos operating system (Junos OS) automatically assigns implicit default BA classifiers to logical interfaces based on the type of interface. Table 10 on page 72 lists different types of interfaces and the corresponding implicit default BA classification.

Table 10: Default BA Classification

Type of Interface	Default BA Classification
Trunk and Circuit Cross-Connect (CCC) interfaces	ieee8021p-default NOTE : This BA classification for a CCC interface is applicable only for EX8200 switches.
Layer 3 interface (IPv4)	dscp-default
Layer 3 interface (IPv6)	dscp-ipv6-default
Access interface	Untrusted
Routed VLAN interface (RVI)	No default classification
MPLS	EXP NOTE : This BA classification is applicable only for EX8200 switches.

When you explicitly associate a BA classifier with a *logical interface*, you are overriding the implicit (default) BA classifier with an explicit BA classifier.

Table 11 on page 73 describes the BA classifier types you can configure on Layer 2 and Layer 3 interfaces.

Table 11: Allowed BA Classification

Type of Interface	Allowed BA Classification
Layer 2 interface	IEEE 802.1p, IP precedence, DSCP, DSCP IPv6
Layer 3 interface (IPv4)	IEEE 802.1p, IP precedence, DSCP
Layer 3 interface (IPv6)	IEEE 802.1p, IP precedence, DSCP IPv6

You cannot apply DSCP and IP precedence classifiers to the same interface. You also cannot apply IEEE 802.1p classifiers to an interface with classifiers of any other type. DSCP IPv6 classifiers can be applied to an interface with either DSCP or IP precedence classifiers, because they apply to different types of packets.

NOTE: On EX4300 switches, the three classifiers (DSCP, DSCP IPv6, and IEEE 802.1p) can coexist on an L2 interface along with a fixed classifier. BA classification takes precedence over fixed classification.

If you have not explicitly configured a classifier on a logical interface, the default classifiers are assigned and classification works as follows:

- To a logical interface configured with an IPv4 address, a DSCP classifier is assigned by default, and IPv4 and IPv6 packets are classified using the DSCP classifier.
- To logical interface configured with an IPv6 address, a DSCP IPv6 classifier is assigned by default, and IPv4 and IPv6 packets are classified using the DSCP IPv6 classifier.

NOTE: On EX8200 switches, you can configure either one classifier of type DSCP or IEEE802.1p, or you can configure one classifier each of type DSCP and IEEE802.1p.

You can configure IRB interfaces on switches that support the ELS configuration style, or routed VLAN interfaces on switches that do not support the ELS configuration style. After you do this, the User Priority (UP) bits in the incoming packets are rewritten according to the default IEEE 802.1p rewrite rule, except on EX8200 switches. On EX8200 switches, you must explicitly assign the default IEEE 802.1p rewrite rule to RVIs.

NOTE: By default, all BA classifiers classify traffic into either the best-effort forwarding class or the network-control forwarding class.

Multifield Classifiers

Multifield (MF) 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 support for DSCP or IP precedence classifiers in end-user applications. On an edge switch, an MF classifier provides the filtering functionality that scans through a variety of packet fields to determine the forwarding class for a packet. Typically, any classifier performs matching operations on the selected fields against a configured value.

RELATED DOCUMENTATION

Understanding Junos OS CoS Components for EX Series Switches | 8 Example: Configuring CoS on EX Series Switches | 15 Defining CoS Classifiers (CLI Procedure) | 74 Defining CoS Classifiers (J-Web Procedure) | 77

Defining CoS Classifiers (CLI Procedure)

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. Junos OS supports two general types of classifiers:

Behavior aggregate (BA) classifier—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, IP precedence value, or IEEE 802.1p value. EX Series switches except EX4300 switches
support two types of loss priorities: high and low. EX4300 switches support three types of loss
priorities: high, medium-high, and low.

You can configure BA classifiers for the following CoS marker types:

- dscp and dscp-ipv6–Handles incoming IPv4 and IPv6 packets, respectively.
- ieee-802.1-Handles Layer 2 CoS.
- **inet-precedence**—Handles incoming IPv4 packets. IP precedence mapping requires only the higher order three bits of the DSCP field.
- Multifield (MF) classifier—Examine multiple fields in the 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.

NOTE: Juniper Networks EX8200 Ethernet Switches implement the on-demand ternary content addressable memory (TCAM) allocation of memory so that when additional TCAM space is required for CoS, the space is allocated from the free TCAM space or from the unused TCAM space. An error log message is generated when you configure CoS classifiers beyond the available TCAM space that includes both the free and unused space.

The following example describes how to configure a BA classifier (**ba-classifier**) as the default DSCP map for handling IPv4 traffic and to apply the BA classifier to either a specific Gigabit Ethernet interface or to all the Gigabit Ethernet interfaces on the switch. The BA classifier assigns loss priorities, as shown in Table 12 on page 75, to incoming packets in the four forwarding classes.

You can use the same procedure to set MF classifiers (except that you would use firewall filter rules).

Forwarding Class	For CoS Traffic Type	ba-classifier Assignment
be	Best-effort traffic	High-priority code point: 000001
ef	Expedited-forwarding traffic	High-priority code point: 101110
af	Assured-forwarding traffic	High-priority code point: 001100
nc	Network-control traffic	High-priority code point: 110001

 Table 12: BA-classifier Loss Priority Assignments

To configure a DSCP BA classifier named **ba-classifier** as the default DSCP map:

• Associate code point 000001 with forwarding class be and loss priority high:

```
[edit class-of-service classifiers]
user@switch# set dscp ba-classifier import default forwarding-class be loss-priority high
code-points 000001
```

• Associate code point 101110 with forwarding class ef and loss priority high:

```
[edit class-of-service classifiers]
user@switch# set dscp ba-classifier forwarding-class ef loss-priority high code-points 101110
```

• Associate code point **001100** with forwarding class **af** and loss priority **high**:

[edit class-of-service classifiers]
user@switch# set dscp ba-classifier forwarding-class af loss-priority high code-points 001100

• Associate code point **110001** with forwarding class **nc** and loss priority **high**:

[edit class-of-service classifiers]
user@switch# set dscp ba-classifier forwarding-class nc loss-priority high code-points 110001

- Apply the classifier to a specific interface or to all Gigabit Ethernet interfaces on the switch.
 - To apply the classifier to a specific interface:

[edit class-of-service interfaces]
user@switch# set ge-0/0/0 unit 0 classifiers dscp ba-classifier

• To apply the classifier to all Gigabit 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 ge-* unit * classifiers dscp ba-classifier
```

NOTE: On EX8200 switches, it can take a long time to install code-point classifiers on multiple interfaces (for example, approximately 25 minutes to install 64 code-point classifiers on multiple interfaces in the order of 280 or more).

RELATED DOCUMENTATION

Defining CoS Classifiers (J-Web Procedure) | 77

Example: Configuring CoS on EX Series Switches | 15

Assigning CoS Components to Interfaces (CLI Procedure) | 55

Monitoring CoS Classifiers | 88

Understanding CoS Classifiers | 70

Troubleshooting a CoS Classifier Configuration for a TCAM Space Error | 91

Defining CoS Classifiers (J-Web Procedure)

NOTE: This topic applies only to the J-Web Application package.

You can use the J-Web interface to define CoS classifiers on an EX Series switch. Classifiers examine the CoS value or alias of an incoming packet and assign the packet a level of service by setting its forwarding class and loss priority.

To define CoS classifiers:

1. Select Configure > Class of Service > Classifiers.

NOTE: After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See Using the Commit Options to Commit Configuration Changes (J-Web Procedure) for details about all commit options.

- 2. Click one of the following options:
 - Add—Adds a classifier. Enter information into the classifier page as described in Table 13 on page 78.

- Edit—Modifies an existing classifier. Enter information into the classifier page as described in Table 13 on page 78.
- **Delete**-Deletes an existing classifier.

Table 13: Classifiers Configuration Fields

Field	Function	Your Action
Classifier Name	Specifies the name for a classifier.	To name a classifier, type the name—for example, ba-classifier .
Classifier Type	Specifies the type of classifier: dscp , ieee-802.1 , or inet-precedence .	Select a value from the list.

Field	Function	Your Action
Code Point Mapping	Sets the forwarding classes and the packet loss priorities (PLPs) for specific CoS values and aliases.	 To add a code point mapping: 1. Click Add. 2. Select the code point. 3. Select a forwarding class from the following list: expedited-forwarding—Provides low loss, low delay, low jitter, assured bandwidth, and end-to-end service. Packets can be forwarded out of sequence or dropped. best-effort—Provides no special CoS handling of packets. Typically, RED drop profile is aggressive and no loss priority is defined. assured-forwarding—Provides high assurance for packets within the specified service profile. Excess packets are dropped. network-control—Packets can be delayed but not dropped. Select the loss priority. To assign a loss priority, select one: high—Packet has a high loss priority. low—Packet has a low loss priority.

Table 13: Classifiers Configuration Fields (Continued)

RELATED DOCUMENTATION

Defining CoS Classifiers (CLI Procedure) | 74

Example: Configuring CoS on EX Series Switches | 15

Monitoring CoS Classifiers | 88

Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers

IN THIS SECTION

- Requirements | 81
- Overview | 81
- Verification | 82

Packet classification associates incoming packets with a particular CoS servicing level. Behavior aggregate (BA) classifiers examine the CoS value in the packet header to determine 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 incoming CoS value.

Beginning with Junos OS Release 17.1, EX4300 switches support multidestination classifiers. On EX4300 switches, you can apply multidestination classifiers globally or to a specific interface. If you apply multidestination classifiers both globally and to a specific interface, the classifications on the interface take precedence.

Multidestination classifiers apply to all of the switch interfaces and handle multicast, broadcast, and destination lookup fail (DLF) traffic. You cannot apply a multidestination classifier to a single interface or to a range of interfaces, except on an EX4300 switch.

Unicast and multidestination traffic must use different classifiers.

Configuring Multidestination Classifiers

Step-by-Step Procedure

To configure a multicast IEEE 802.1 BA classifier named ba-mcast-classifier:

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

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

2. Configure the classifier as a multidestination classifier:

```
[edit class-of-service]
user@switch# set multi-destination classifiers ieee-802.1 ba-mcast-classifier
```

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.

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 CoS value.
- 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 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.

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

NOTE: You must assign unicast traffic and multicast traffic to different classifiers. One classifier cannot include both unicast and multicast forwarding classes. A multidestination classifier can include only forwarding classes for multicast traffic.

The following example describes how to configure a BA classifier called ba-mcast-classifier, which is applied to all of the switch interfaces. The BA classifier assigns loss priorities, as shown in Table 14 on page 82, to incoming packets in the multidestination forwarding class.

You can also use firewall filters to set multifield classifiers.

Table 14: BA-mcast-classifier Loss Priority Assignments

Multicast Forwarding Class	Traffic Type	ba-mcast-classifier Assignment
mcast	Best-effort multicast traffic	Low loss priority code point: 000

Verification

IN THIS SECTION

- Verifying the IEEE 802.1 Multidestination Classifier | 82
- Verifying the Multidestination Classifier Configuration | 83

To verify the multidestination classifier configuration, perform these tasks:

Verifying the IEEE 802.1 Multidestination Classifier

Purpose

Verify that the classifier ba-mcast-classifier is configured as the IEEE 802.1 multidestination classifier:

Action

Verify the results of the classifier configuration using the operational mode command show configuration class-of-service multi-destination classifiers ieee-802.1:

user@switch> show configuration class-of-service multi-destination classifiers ieee-802.1
ba-mcast-classifier;

Verifying the Multidestination Classifier Configuration

Purpose

Verify that you configured the multidestination 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-mcast-classifier:

user@switch> show configuration class-of-service classifiers ieee-802.1 ba-mcast-classifier
forwarding-class mcast {
loss-priority low code-points 000;
}

Release History Table

Release	Description
17.1	Beginning with Junos OS Release 17.1, EX4300 switches support multidestination classifiers.

RELATED DOCUMENTATION

Example: Configuring Unicast Classifiers	
Defining CoS BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)	
Monitoring CoS Classifiers	
Understanding CoS Classifiers	
Understanding CoS Classifiers	

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

Configuring and Applying IEEE 802.1ad Classifiers

If you apply an IEEE 802.1 classifier to a logical interface, this classifier takes precedence and is not compatible with any other classifier type. For Juniper Networks MX Series 5G Universal Routing Platform interfaces or IQ2 PICs with IEEE 802.1ad frame formats or EX Series switches, you can set the

forwarding class and loss priority for traffic on the basis of the three IEEE 802.1p bits (three bits in either the inner virtual LAN (VLAN) tag or the outer VLAN tag) and the drop eligible indicator (DEI) bit. You can apply the default map or customize one or more of the default values.

You then apply the classifier to the interface on which you configure IEEE 802.1ad frame formats.

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the CLI User Guide.

- **1.** Define the custom IEEE 802.1ad map:
 - a. Create the classifier by specifying a name for it and defining it as an IEEE-802.1ad (DEI) classifier.

```
[edit]
user@host# edit class-of-service classifiers ieee-802.1ad dot1p_dei_class
```

b. Assign the forwarding class and loss priority to the code-point alias.

[edit class-of-service classifiers ieee-802.1ad dot1p_dei_class]
user@host# set forwarding-class best-effort loss-priority low code-points [0000 1101]

- **2.** Apply the classifier to the logical interface:
 - **a.** Specify the interface to which you want to apply the classifier.

```
[edit]
user@host# edit class-of-service interfaces ge-2/0/0 unit 0
```

b. Specify the name of the classifier you want to apply to the interface.

[edit class-of-service interfaces ge-2/0/0 unit 0] user@host# set classifiers ieee-802.1ad dot1p_dei_class

3. Verify the custom IEEE 802.1ad map configuration:

[edit]
user@host# show

class-of-service {
 classifiers {

```
ieee-802.1ad dot1p_dei_class {
    forwarding-class best-effort {
        loss-priority low code-points [ 0000 1101 ];
    }
    }
}
```

RELATED DOCUMENTATION

Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic Applying Behavior Aggregate Classifiers to Logical Interfaces

Configuring the IEEE 802.1p Field for CoS Host Outbound Traffic

This topic provides a summary of the configuration for setting the IEEE 802.1p field in the Ethernet frame header for host outbound traffic (control plane traffic). You can set a global value for the priority code point that applies to all host outbound traffic. Additionally, or alternatively, you can specify that rewrite rules are applied to all host outbound traffic on egress logical interfaces. These are rules that have been previously configured to set the IEEE 802.1p field for data traffic on those interfaces.

Configuration of 802.1p bits is supported only on the following hardware and software components:

- EX Series switches
- MX Series 5G Universal Routing Platforms

- Enhanced Queuing DPCs
- MPCs
- Junos OS Release 12.3 or later

To configure the IEEE 802.1p field settings:

- (Optional) Specify a global default value for the IEEE 802.1p field for all host outbound traffic. See *Configuring a Global Default IEEE 802.1p Value for All Host Outbound Traffic*.
- **2.** (Optional) Specify that the IEEE 802.1p rewrite rules for the egress logical interfaces are applied to all host outbound traffic on those interfaces.

See Applying Egress Interface Rewrite Rules to the IEEE 802.1p Field for All Host Outbound Traffic on the Interface.

RELATED DOCUMENTATION

Rewriting Packet Headers to Ensure Forwarding Behavior

Configuring a Global Default IEEE 802.1p Value for All Host Outbound Traffic

This topic describes how to configure a global default value for the IEEE 802.1p field for all host outbound traffic on MX Series routers and EX Series switches.

To configure a global default value for the IEEE 802.1p field:

• Specify the value.

[edit class-of-service host-outbound-traffic ieee-802.1]
user@host# set default value

For example, specify that a value of 010 is applied to all host outbound traffic:

[edit class-of-service host-outbound-traffic ieee-802.1]
user@host# set default 010

RELATED DOCUMENTATION

Configuring the IEEE 802.1p Field for CoS Host Outbound Traffic Rewriting Packet Headers to Ensure Forwarding Behavior

Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure)

EX8200 switches provide certain line cards that include oversubscribed ports. These ports are logically grouped into a port group and each port group share a certain fixed bandwidth. Because oversubscribed ports handle traffic differently than ports that provide continuous line-rate bandwidth, configuring CoS queues is different for oversubscribed ports than for line-rate ports.

Packets arriving on an oversubscribed port in a line card are directed to a high-priority, low priority, or line-rate queue. These queues are used for scheduling traffic from the port into the Packet Forwarding Engine. The fabric priority associated with the packet's forwarding class determines which queue the packet is sent to. The forwarding class of the packet in turn is determined by the behavior aggregate (BA) classifier assigned to the port. By default, the fabric priority of all forwarding classes is low. Thus all packets, with the exception of critical network packets and line-rate packets, are sent to the low-priority ingress queue by default. The critical network packets and line-rate packets do not need a BA classifier as they are always sent on the high-priority and line-rate queues, respectively.

This procedure describes how you can direct traffic into the high-priority ingress queue and thus avoid congestion at the port group.

To direct traffic to the high-priority ingress queue for a port group:

1. Create the BA classifier for the forwarding class:

[edit class-of-service]
user@switch# set classifiers classifier-type classifier-name
forwarding-class class-name loss-priority level code-points code-point

2. Assign a queue number and fabric priority to the forwarding class:

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

3. Assign the BA classifier to the physical interface:

[edit class-of-service]
user@switch# set interfaces interface-name unit 0
classifiers classifier-type classifier-name

For example, to direct voice traffic to the high-priority ingress queue for interface xe-1/0/2:

```
[edit class-of-service]
user@switch# set classifiers dscp dscp1 forwarding-class cos-voice
loss-priority low code-points ef
```

[edit class-of-service]
user@switch# set forwarding-classes class cos-voice queue-num 5 priority high

[edit class-of-service]
user@switch# set interfaces xe-1/0/2 unit 0 classifiers dscp dscp1

NOTE: You must use a BA classifier to classify traffic for ingress queuing. Multifield (MF) classification and port classification (that is, assigning a forwarding class to the interface) are not supported for classifying traffic for ingress queuing. The BA classifier must be assigned to a physical interface, not a Layer 3 tagged interface or a routed VLAN interface (RVI).

RELATED DOCUMENTATION

Understanding CoS Queues on EX8200 Line Cards That Include Oversubscribed Ports | 11

Monitoring CoS Classifiers

IN THIS SECTION

- Purpose | 89
- Action | 89

Purpose

NOTE: This topic applies only to the J-Web Application package.

Use the monitoring functionality to display the mapping of incoming CoS values to the forwarding class and loss priority for each classifier.

Action

To monitor CoS classifiers in the J-Web interface, select Monitor > Class of Service > Classifiers.

To monitor CoS classifiers in the CLI, enter the following CLI command:

show class-of-service classifier

Meaning

Table 15 on page 89 summarizes key output fields for CoS classifiers.

Table 15: Summary of Key CoS Classifier Output Fields

Field	Values	Additional Information
Classifier Name	Name of a classifier.	To display classifier assignments, click the plus sign (+).

Field	Values	Additional Information
CoS Value Type	 The classifiers are displayed by type: dscp—All classifiers of the DSCP type. ieee-802.1—All classifiers of the IEEE 802.1 type. inet-precedence—All classifiers of the IP precedence type. 	
Index	Internal index of the classifier.	
Incoming CoS Value	CoS value of the incoming packets, in bits. These values are used for classification.	
Assign to Forwarding Class	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.	
Assign to Loss Priority	Loss priority value that the classifier assigns to the incoming packet based on its CoS value.	

Table 15: Summary of Key CoS Classifier Output Fields (Continued)

RELATED DOCUMENTATION

Defining CoS Classifiers (CLI Procedure) 74	
Defining CoS Classifiers (J-Web Procedure) 77	
Example: Configuring CoS on EX Series Switches	15

Troubleshooting a CoS Classifier Configuration for a TCAM Space Error

IN THIS SECTION

- Problem | 91
- Solution | 91

Problem

Description

When a CoS classifier configuration exceeds the amount of available ternary content addressable memory (TCAM) space, the switch returns the following system log message:

<number_of_rules_being_added> rules for <filter_name> class <filter_class> will not be installed, key: <bind_point>. no space in tcam db(<shared_pool_information>)

The switch returns this message during the commit operation if the number of classifiers defined in the CoS configuration or the number of bind points (interfaces) to which classifiers are bound causes the CoS configuration to exceed the amount of available TCAM space. However, the commit operation for the CoS configuration is completed in the CLI module.

Solution

When a CoS configuration exceeds the amount of available TCAM table space, you must either define fewer classifiers or bind them to fewer interfaces, or both, so that the space requirements for the CoS configuration do not exceed the available space in TCAM.

To delete classifier definitions and bind points in a CoS configuration, and to apply a new CoS classifier definition to fewer bind points:

- **1.** Delete either the CoS classifier definition or the bind points:
 - To delete the CoS classifier definition:

• For behavioral classifiers:

[edit class-of-service]
user@switch# delete classifier dscp d1

• For multifield classifiers:

[edit]
user@switch# delete interfaces ge-3/0/2 unit 0 family ethernet-switching filter input
ipacl

This command deletes a multifield classifier defined for a port. Similarly, you can delete a multifield classifier defined for a VLAN or router.

You can also delete terms defined in a single multifield classifier:

[edit]
user@switch# delete firewall family inet filter f1 term t1

In both these examples (for behavioral and multifield classifiers), the assumption is that too many classifier definitions resulted in the error message.

• To delete the bind points:

```
[edit class-of-service]
user@switch# delete class-of-service interfaces ge-0/0/0
user@switch# delete class-of-service interfaces ge-0/0/2
user@switch# delete class-of-service interfaces ge-0/0/3
user@switch# delete class-of-service interfaces ge-0/0/4
user@switch# delete class-of-service interfaces ge-0/0/5
user@switch# delete class-of-service interfaces ge-0/0/5
user@switch# delete class-of-service interfaces ge-0/0/6
user@switch# delete class-of-service interfaces ge-0/0/7
user@switch# delete class-of-service interfaces ge-0/0/7
```

Here the assumption is that too many bind points (nine) in the configuration resulted in the error message.

2. Commit the operation:

[edit]
user@switch# commit

- **3.** Define fewer classifiers in the CoS configuration or bind classifiers to fewer interfaces, or both, so that the CoS classifier configuration does not exceed the amount of available TCAM space on the switch:
 - To define CoS classifiers:
 - For behavioral classifiers:

[edit]

user@switch# set class-of-service classifiers dscp d2 forwarding-class fc1 losspriority low code-points 000001 user@switch# set class-of-service classifiers dscp d2 forwarding-class fc2 losspriority low code-points 000010 user@switch# set class-of-service classifiers dscp d2 forwarding-class fc3 losspriority low code-points 000011 user@switch# set class-of-service classifiers dscp d2 forwarding-class fc4 losspriority low code-points 000100 user@switch# set class-of-service classifiers dscp d2 forwarding-class fc5 losspriority low code-points 000101 user@switch# set class-of-service classifiers dscp d2 forwarding-class fc6 losspriority low code-points 000110 user@switch# set class-of-service classifiers dscp d2 forwarding-class fc6 losspriority low code-points 000110 user@switch# set class-of-service classifiers dscp d2 forwarding-class fc7 losspriority low code-points 000111

• For multifield Classifiers:

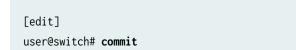
[edit] user@switch# set firewall family inet filter f1 term t1 from protocol tcp user@switch# set firewall family inet filter f1 term t1 then loss-priority high user@switch# set firewall family inet filter f1 term t1 then forwarding-class besteffort user@switch# set firewall family inet filter f1 term t2 from protocol udp user@switch# set firewall family inet filter f1 term t2 then loss-priority high user@switch# set firewall family inet filter f1 term t2 then loss-priority high user@switch# set firewall family inet filter f1 term t2 then forwarding-class assuredforwarding user@switch# set firewall family inet filter f1 term t3 from source-port ssh

user@switch# set firewall family inet filter f1 term t3 then loss-priority low user@switch# set firewall family inet filter f1 term t3 then forwarding-class fc8 user@switch#set class-of-service forwarding-classes best-effort, assured-forwarding, fc8

• To bind classifiers to fewer interfaces:

```
[edit]
user@switch# set class-of-service interfaces ge-0/0/0 unit 0 classifiers dscp d2
user@switch# set class-of-service interfaces ge-0/0/1 unit 0 classifiers dscp d2
user@switch# set class-of-service interfaces ge-0/0/2 unit 0 forwarding-class best-effort
user@switch# set class-of-service interfaces ge-0/0/3 unit 0 forwarding-class assured-
forwarding
user@switch# set class-of-service interfaces ge-0/0/4 unit 0 forwarding-class fc8
```

4. Commit the operation:



5. Check system log for an error message. If an error message is not logged, then your classifier configuration has not exceeded the TCAM space limit.

If an error message is logged, then repeat this procedure by defining fewer classifiers or binding classifiers to fewer bind points.

RELATED DOCUMENTATION

Understanding CoS Classifiers | 70

Defining CoS Classifiers (CLI Procedure) | 74

CoS Rewrite

IN THIS CHAPTER

- Understanding CoS Rewrite Rules | 95
- Defining CoS Rewrite Rules (CLI Procedure) | 97
- Defining CoS Rewrite Rules (J-Web Procedure) | 100
- Classifiers and Rewrite Rules at the Global, Physical, and Logical Interface Levels Overview | 102
- Configuring Classifiers and Rewrite Rules at the Global and Physical Interface Levels | 103
- Applying Egress Interface Rewrite Rules to the IEEE 802.1p Field for All Host Outbound Traffic on the Interface | 105
- Monitoring CoS Rewrite Rules | **107**

Understanding CoS Rewrite Rules

IN THIS SECTION

- How Rewrite Rules Work | 95
- Default Rewrite Rule | 96

As packets enter or exit a network, edge switches might be required to alter the class-of-service (CoS) settings of the packets. This topic describes how to use *rewrite rules* to alter the CoS settings. It covers:

This topic covers:

How Rewrite Rules Work

Rewrite rules set the value of the CoS bits within a packet's header. Each rewrite rule reads the current forwarding class and loss priority associated with the packet, locates the chosen CoS value from a table,

and writes this CoS value into the packet header. For rewrites to occur, rewrite rules must be explicitly assigned to an interface.

On EX Series switches, you can define rewrite rules for IPv4 and IPv6 traffic to network interfaces, aggregated Ethernet interfaces (also known as link aggregation groups (LAGs)), routed VLAN interfaces (RVIs), Layer 3 interfaces, and Layer 3 VLAN-tagged sub-interfaces. Multiple rewrite rules of different types can be assigned to a single interface.

On EX4300 switches, you cannot configure separate DSCPv4 and DSCPv6 rewrite rules on network interfaces, aggregated Ethernet interfaces, Layer 3 interfaces, and integrated routing and bridging (IRB) interfaces. If you configure a DSCPv4 rewrite rule on an interface to rewrite IPv4 traffic, then the same rewrite rule is applied to IPv6 traffic also on that interface, and vice versa. You can define only DSCPv4 rewrite rules on integrated routing and bridging (IRB) interfaces and Layer 3 VLAN-tagged logical interfaces.

In effect, the rewrite rule performs the reverse function of the behavior aggregate (BA) classifier, which is used when the packet enters the switch. As the packet leaves the switch, the final CoS action is generally the application of a rewrite rule.

You configure rewrite rules to alter CoS values in outgoing packets on the outbound interfaces of an edge switch to meet 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: When an IP precedence rewrite rule is active, bits 3, 4, and 5 of the type-of-service (ToS) byte are always reset to zero when code points are rewritten.

Default Rewrite Rule

To define a rewrite rule on an interface, you can either create your own rewrite rule and enable it on the interface or enable a default rewrite rule. See "Defining CoS Rewrite Rules (CLI Procedure)" on page 97.

Table 16 on page 97 shows the default rewrite-rule mappings. These are based on the default bit definitions of Differentiated Services code point (DSCP), IEEE 802.1p, and IP precedence values and the default forwarding classes. You can configure multiple CoS rewrite rules for DSCP, IP precedence and IEEE 802.1p.

NOTE: By default, rewrite rules are not assigned to an interface. You must explicitly assign a user-defined or system-defined rewrite rule to an interface for the rewrites to occur.

When the CoS values of a packet match the forwarding class and packet-loss-priority (PLP) values, the switch rewrites markings on the packet based on the rewrite table.

Table 16: Default Packet Header Rewrite Mappings

Map from Forwarding Class	PLP Value	Map to DSCP/IEEE 802.1p/IP Precedence Value
expedited-forwarding	low	ef
expedited-forwarding	high	ef
assured-forwarding	low	af11
assured-forwarding	high	af12 (DSCP)
best-effort	low	be
best-effort	high	be
network-control	low	nc1/cs6
network-control	high	nc2/cs7

RELATED DOCUMENTATION

Understanding Junos OS CoS Components for EX Series Switches 8
Example: Configuring CoS on EX Series Switches 15
Defining CoS Rewrite Rules (CLI Procedure) 97

Defining CoS Rewrite Rules (J-Web Procedure) | 100

Defining CoS Rewrite Rules (CLI Procedure)

You configure rewrite rules to alter CoS values in outgoing packets on the outbound interfaces of an EX Series switch to match 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 a code point, thus creating a rewrite table, and you can enable the rewrite rule on an interface. On EX Series switches except EX4300 switches, you can also enable a rewrite rule on routed VLAN interfaces (RVIs). On EX4300 switches, you can also enable rewrite rules on integrated routing and bridging (IRB) interfaces. If you need to customize a rewrite rule, you can create a customized rewrite rule using a firewall filter configuration. You can configure CoS rewrite rules for DSCP, IP precedence and IEEE 802.1p.

You can configure rewrite rules for the following CoS marker types:

- dscp and dscp-ipv6—Handles incoming IPv4 and IPv6 packets, respectively. On EX4300 switches, you cannot configure DSCP IPv4 and DSCP IPv6 rewrite rules on the same interface. If you configure a DSCP IPv4 rewrite rule on an interface to rewrite IPv4 traffic, then the same rewrite rule is applied to IPv6 traffic also on that interface, and vice versa.
- ieee-802.1—Handles Layer 2 CoS.
- **inet-precedence**—Handles incoming IPv4 packets. IP precedence mapping requires only the higher order three bits of the DSCP field.

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

To create IEEE 802.1p rewrite rules and enable them on Layer 2 interfaces:

• To create an IEEE 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 af loss-priority low code-point 010
user@switch# set ieee-802.1 customup-rw forwarding-class af loss-priority high code-point 011
user@switch# set ieee-802.1 customup-rw forwarding-class ef loss-priority low code-point 100
user@switch# set ieee-802.1 customup-rw forwarding-class ef loss-priority high code-point 100
user@switch# set ieee-802.1 customup-rw forwarding-class ef 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 110
```

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

```
[edit]
```

user@switch# set class-of-service interfaces ge-0/0/0 unit 0 rewrite-rules ieee-802.1
customup-rw

(On EX4300 switches) To enable an IEEE 802.1p rewrite rule named customup-rw on a Layer 2 interface:

[edit]

user@switch# set class-of-service interfaces ge-0/0/0 rewrite-rules ieee-802.1 customup-rw

• To enable an IEEE 802.1p rewrite rule named customup-rw on all 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 ge-* unit * rewrite-rules customup-rw
```

(On EX4300 switches) To enable an IEEE 802.1p rewrite rule named customup-rw on all Gigabit Ethernet interfaces on the switch, use wildcards for the interface name:

[edit]

user@switch# set class-of-service interfaces ge-* rewrite-rules customup-rw

RELATED DOCUMENTATION

Defining CoS Rewrite Rules (J-Web Procedure) 100	
Example: Configuring CoS on EX Series Switches 15	
Monitoring CoS Rewrite Rules 107	

Understanding CoS Rewrite Rules | 95

Defining CoS Rewrite Rules (J-Web Procedure)

NOTE: This topic applies only to the J-Web Application package.

You can use the J-Web interface to define CoS rewrite rules. Use the rewrite rules to alter the CoS values in outgoing packets to meet the requirements of the targeted peer. A rewrite rule examines the forwarding class and loss priority of a packet and sets its bits to a corresponding value specified in the rule.

To define rewrite rules:

1. Select Configure > Class of Service > Rewrite Rules.

NOTE: After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See Using the Commit Options to Commit Configuration Changes (J-Web Procedure) for details about all commit options.

- **2.** Select one of the following options:
 - Add—Adds a rewrite rule. Enter information into the rewrite rule page as described in Table 17 on page 100.
 - Edit—Modifies an existing rewrite rule. Enter information into the rewrite rule page as described in Table 17 on page 100.
 - **Delete**-Deletes an existing rewrite rule.

Table 17: Rewrite Rules Configuration Page Summary

Field	Function	Your Action
Rewrite Rule Name	Specifies the name for the rewrite rule.	To name a rule, type the name—for example, rewrite-dscps .
Rewrite rule type	Specifies the type of rewrite rule: dscp , ieee-802.1 , or inet-precedence .	Select a value from the list.

Code Point Mapping Rewrites outgoing CoS values of a packet I based on the forwarding class and loss f priority. I Allows you to remove a code point I mapping entry. I	Your Action To configure a CoS value assignment, follow these steps: To add a code point mapping: 1. Click Add. 2. Select the code point. 3. Select a forwarding class from the
based on the forwarding class and loss priority. Allows you to remove a code point mapping entry.	follow these steps: To add a code point mapping: 1. Click Add . 2. Select the code point. 3. Select a forwarding class from the
ן ג ן	 following list: expedited-forwarding—Provides low loss, low delay, low jitter, assured bandwidth, and end-to-end service. Packets can be forwarded out of sequence or dropped. best-effort—Provides no special CoS handling of packets. Typically, RED drop profile is aggressive and no loss priority is defined. assured-forwarding—Provides high assurance for packets within the specified service profile. Excess packets are dropped. network-control—Packets can be delayed but not dropped. Select the loss priority. To assign a loss priority, select one: high—Packet has a high loss priority. Iow—Packet has a low loss priority. To edit an existing code point mapping, select it and click Edit. To remove a code point mapping entry, select it and click Remove.

Table 17: Rewrite Rules Configuration Page Summary (Continued)

RELATED DOCUMENTATION

Defining CoS Rewrite Rules (CLI Procedure) | 97

Understanding CoS Rewrite Rules | 95

Monitoring CoS Rewrite Rules | 107

Example: Configuring CoS on EX Series Switches | 15

Classifiers and Rewrite Rules at the Global, Physical, and Logical Interface Levels Overview

On most ACX Series Universal Metro Routers and EX Series switches, CoS supports classification and rewrite at the global level and physical interface levels.

NOTE: The ACX6360 router does not support rewrite rules or Layer 2 (IEEE802.1p and IEEE802.1ad) classifiers.

NOTE: ACX7100 routers support classification and rewrite rules of all types (Inet-Prec/DSCP/ DSCP-v6/EXP/IEEE-802.1p/IEEE-802.1ad) at the logical interface level.

At a global level, you can define EXP classification.

At a physical interface level, you can define the following features:

- DSCP, DSCP-IPV6, and IPv4 precedence classifiers
- DSCP, DSCP-IPV6, and IPv4 precedence rewrites
- IEEE 802.1p and IEEE 802.1ad classifiers (inner and outer)
- IEEE 802.1p and IEEE 802.1ad rewrites (outer)

The IEEE 802.1ad classifier uses IEEE 802.1p and DEI bits together.

NOTE: You cannot configure both IEEE 802.1p and IEEE 802.1ad classifiers together at the physical interface level.

At a logical interface level, you can define the fixed classification and EXP rewrites.

To configure global EXP classifiers, include the **classfiers exp** *classifier-name* statement at the **[edit class-of-service system-defaults]** hierarchy level.

To configure classifiers or *rewrite rules* at the physical interface, include either the **classifiers** statement or the **rewrite-rules** statement at the **[edit class-of-service interfaces** *interface-name*] hierarchy level.

To configure fixed classifiers at the logical interface, include the **forwarding-class** *fc*] or the **rewrite-rules** statement at the **[edit class-of-service interfaces** *interface-name* **unit** *number*] hierarchy level.

To configure EXP rewrite at the logical interface, include the **[edit class-of-service interfaces** *interface-name* unit *number* rewrite-rules exp *rewrite-rule*] statement.

To display classifiers configured under **system-defaults**, enter the **show class-of-service system-defaults** command.

To display classifiers and rewrite rules bound to physical interfaces, enter the **show class-of-service interface** *interface-name* command.

RELATED DOCUMENTATION

Configuring Classifiers and Rewrite Rules at the Global and Physical Interface Levels

Configuring Classifiers and Rewrite Rules at the Global and Physical Interface Levels

On ACX Series Universal Metro Routers and EX Series switches, CoS supports classification and rewrite at the global and physical interface levels.

To configure the global EXP classifier, include the following statements at the **[edit class-of-service] system-defaults** hierarchy level.

```
[edit class-of-service]
{
    system-defaults {
        classifiers exp classifier-name
    }
}
```

CoS supports one global system default classifier of the EXP type, as shown in the following example:

```
[edit class-of-service]
{
   system-defaults {
      classifiers {
        exp exp-classf-core;
      }
   }
}
```

To configure classifiers and rewrite rules at the physical interface level, include the following statements at the **[edit class-of-service] interfaces** hierarchy level.

```
[edit class-of-service]
interfaces {
    interface-name
      classifiers dscp classifier-name
      classifiers inet-precedence classifier-name
      classifiers ieee-802.1 [vlan-tag (outer | inner)] classifier-name
      rewrite-rules dscp rewrite-name
      rewrite-rules inet-prec rewrite-name
      rewrite-rules ieee-802.1 rewrite-name
}
```

The following example shows classifiers and rewrite rules configured on physical interfaces:

```
ge-0/1/0 {
    unit 0 {
        rewrite-rules {
            exp custom-exp;
        }
    }
    classifiers {
        dscp d1;
        ieee-802.1 ci;
    }
    rewrite-rules {
        dscp default;
    }
}
```

```
ge-0/1/2 {
    classifiers {
        ieee-802.1 ci;
    }
    rewrite-rules {
        ieee-802.1 ri;
    }
}
ge-0/1/3 {
    unit 0 {
        rewrite-rules {
            exp custom-exp2;
        }
    }
}
ge-0/1/7 {
    classifiers {
        dscp d1;
    }
}
ge-0/1/8 {
    classifiers {
        dscp d1;
    }
}
```

RELATED DOCUMENTATION

Classifiers and Rewrite Rules at the Global, Physical and Logical Interface Levels Overview

Applying Egress Interface Rewrite Rules to the IEEE 802.1p Field for All Host Outbound Traffic on the Interface

This topic describes how to apply rewrite rules for egress logical interfaces to the IEEE 802.1p field for all host outbound traffic on those interfaces on MX Series routers and EX Series switches.

This task requires separately configured rewrite rules that map packet loss priority information to the code point value in the 802.1p field for data traffic on egress logical interfaces. See *Rewriting Packet Headers to Ensure Forwarding Behavior*.

To configure the rewrite rules:

- Configure the CoS rewrite rules to map the forwarding class to the desired value for the 802.1p field.
 See *Configuring Rewrite Rules*.
- 2. Associate the rewrite rules to the desired egress logical interfaces.

See Applying Rewrite Rules to Output Logical Interfaces.

3. (Optional) Configure the forwarding class for host outbound traffic. Do not configure this forwarding class if you want to use the default forwarding class assignment (input classification).

See Overriding the Input Classification.

To configure the rewrite rules to apply to the host outbound traffic IEEE 802.1p field:

• Configure the rewrite rules.

[edit class-of-service host-outbound-traffic ieee-802.1]
user@host# set rewrite-rules

NOTE: Enabling IEEE 802.1p rewrite rules for host outbound traffic on a DPC installed on an MX Series device without creating any corresponding IEEE 802.1p rewrite rules on a logical interface on the DPC causes the IEEE 802.1p code point to be automatically set to 000 for all host generated traffic that exits that logical interface.

```
[edit class-of-service]
rewrite-rules {
    ieee-802.1 rewrite_foo {
        forwarding-class network-control {
            loss-priority low code-point 101;
        }
    }
    interfaces {
        ge-1/0/0 {
            unit 100 {
               rewrite-rules {
                  ieee-802.1 rewrite_foo vlan-tag outer-and-inner;
            }
        }
    }
}
```

```
}
}
host-outbound-traffic {
   forwarding-class network-control;
}
host-outbound-traffic {
   ieee-802.1 {
      rewrite-rules;
   }
}
```

RELATED DOCUMENTATION

Configuring the IEEE 802.1p Field for CoS Host Outbound Traffic Rewriting Packet Headers to Ensure Forwarding Behavior

Monitoring CoS Rewrite Rules

IN THIS SECTION

- Purpose | 107
- Action | 108
- Meaning | 108

Purpose

NOTE: This topic applies only to the J-Web Application package.

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 J-Web interface, select Monitor > Class of Service > Rewrite Rules.

To monitor CoS rewrite rules in the CLI, enter the following command:

show class-of-service rewrite-rules

Meaning

Table 18 on page 108 summarizes key output fields for CoS rewrite rules.

Table 18: Summary of Key CoS Rewrite Rules Output Fields

Field	Values	Additional Information
Rewrite Rule Name	Names of rewrite rules.	
CoS Value Type	 Rewrite rule type: dscp–For IPv4 DiffServ traffic. exp–For MPLS traffic. ieee-802.1–For Layer 2 traffic. 	To display forwarding classes, loss priorities, and rewritten CoS values, click the plus sign (+).
Index	• inet-precedence—For IPv4 traffic. Internal index for this particular rewrite rule.	
Forwarding Class	Forwarding class that is used to determine CoS values for rewriting in combination with loss priority.	Rewrite rules are applied to CoS values in outgoing packets based on forwarding class and loss priority setting.
Loss Priority	Loss priority that is used to determine CoS values for rewriting in combination with forwarding class.	

Table 18: Summary of Key CoS Rewrite Rules Output Fields (Continued)

Field	Values	Additional Information
Rewrite CoS Value To	Value that the CoS value is rewritten to.	

RELATED DOCUMENTATION

Defining CoS Rewrite Rules (CLI Procedure) 97
Defining CoS Rewrite Rules (J-Web Procedure) 100
Example: Configuring CoS on EX Series Switches 15

Forwarding Classes

IN THIS CHAPTER

- Understanding CoS Forwarding Classes | 110
- Defining CoS Forwarding Classes (CLI Procedure) | 114
- Defining CoS Forwarding Classes (J-Web Procedure) | 115
- Monitoring CoS Forwarding Classes | 117

Understanding CoS Forwarding Classes

IN THIS SECTION

Default Forwarding Classes | 111

Class-of-Service (CoS) forwarding classes can be thought of as output queues. In effect, the result of classifying packets is the identification of an output queue for a particular packet. For a classifier to assign an output queue to a packet, it must associate the packet with one of 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.
- multicast best-effort (mcast-be)-Provides no service profile for multicast packets.

- multicast expedited forwarding (mcast-ef)-Supports high-priority multicast packets.
- multicast assured-forwarding (mcast-af)—Provides two drop profiles; high, and low, for multicast packets.
- multicast network-control (mcast-nc)—Supports high-priority multicast packets that are not delaysensitive.

NOTE: The forwarding classes multicast expedited-forwarding, multicast assured-forwarding, and multicast best-effort are applicable only to Juniper Networks EX8200 and EX4300 Ethernet Switches. The forwarding class multicast network-control is applicable only to EX4300 switches.

Juniper Networks EX Series Ethernet Switches support up to 16 forwarding classes, thus allowing granular packet classification. For example, you can configure multiple classes of expedited forwarding (EF) traffic such as EF, EF1, and EF2.

EX Series switches support up to eight output queues, except EX4300 switches that support 12 output queues. Therefore, if you configure more forwarding classes that the number of queues supported, you must map multiple forwarding classes to one or more output queues. On EX8200 *Virtual Chassis*, you can configure only eight forwarding classes and you can assign only one forwarding class to each output queue.

NOTE: On EX8200 Virtual Chassis, the queue number seven carries Virtual Chassis port (VCP) traffic and can also carry high-priority user traffic.

This topic describes:

Default Forwarding Classes

Table 19 on page 112 shows the four default forwarding classes defined for unicast traffic, and Table 20 on page 112 shows the default forwarding classes defined for multicast traffic.

NOTE: The default forwarding classes for multicast traffic are applicable only to EX8200 Virtual Chassis and EX4300 switches.

You can rename the forwarding classes associated with the queues supported on your switch. Assigning a new class name to an output queue does not alter the default classification or scheduling that is applicable to that queue. However, because CoS configurations can be quite complicated, we recommend that you avoid altering the default class names or queue number associations.

Table 19: Default Forwarding Classes for Unicast Traffic

Forwarding Class Name	Comments
best-effort (be)	The software does not apply any special CoS handling to packets with 000000 in the DiffServ field. This is a backward compatibility feature. These packets are usually dropped under congested network conditions.
expedited-forwarding (ef)	The software delivers assured bandwidth, low loss, low delay, and low delay variation (jitter) end-to-end for packets in this service class. The software accepts excess traffic in this class, but in contrast to the assured forwarding class, the out-of-profile expedited-forwarding class packets can be forwarded out of sequence or dropped.
assured-forwarding (af)	The software offers a high level of assurance that the packets are delivered as long as the packet flow from the customer stays within a certain service profile that you define. The software accepts excess traffic, but it applies a tail drop profile to determine that excess packets are dropped, and not forwarded. Two drop probabilities (low and high) are defined for this service class.
network-control (nc)	The software delivers packets in this service class with a high priority. (These packets are not delay-sensitive.) Typically, these packets represent routing protocol hello or keep alive messages. Because loss of these packets jeopardizes proper network operation, packet delay is preferable to packet discard for these packets.

Table 20: Default Forwarding Classes for Multicast Traffic on EX8200 Virtual Chassis and EX4300 Switches

Forwarding Class Name	Comments
multicast best-effort (mcast-be)	The software does not apply any special CoS handling to multicast packets. These packets are usually dropped under congested network conditions.

Table 20: Default Forwarding Classes for Multicast Traffic on EX8200 Virtual Chassis and EX4300 Switches (Continued)

Forwarding Class Name	Comments
multicast expedited-forwarding (mcast-ef)	The software delivers assured bandwidth, low loss, low delay, and low delay variation (jitter) end-to-end for multicast packets in this service class. The software accepts excess traffic in this class, but in contrast to the multicast assured forwarding class, out-of-profile multicast expedited-forwarding class packets can be forwarded out of sequence or dropped.
multicast assured-forwarding (mcast-af)	The software offers a high level of assurance that the multicast packets are delivered as long as the packet flow from the customer stays within a certain service profile that you define. The software accepts excess traffic, but it applies a tail drop profile to determine if the excess packets are dropped and not forwarded. Two drop probabilities (low and high) are defined for this service class.
multicast network-control (mcast- nc)	(EX4300 switches only) The software delivers packets in this service class with a high priority. (These packets are not delay-sensitive.) Typically, these packets represent routing protocol hello or keep alive messages. Because loss of these packets jeopardizes proper network operation, packet delay is preferable to packet discard for these packets.

The following rules govern queue assignment:

- CoS configurations that specify more queues than the switch can support are not accepted. If you commit such a configuration, the commit fails and a message displays that states the number of queues available.
- All default CoS configurations are based on queue number. The name of the forwarding class that is displayed in the default configuration for a queue number is that of the forwarding class currently associated with that queue.

RELATED DOCUMENTATION

Understanding Junos OS CoS Components for EX Series Switches | 8 Example: Configuring CoS on EX Series Switches | 15 Example: Prioritizing Snooped and Inspected Packet

Defining CoS Forwarding Classes (CLI Procedure)

Forwarding classes allow you to group packets for transmission. Based on forwarding classes, you assign packets to output queues.

By default, four categories of forwarding classes are defined: best effort, assured forwarding, expedited forwarding, and network control. EX Series switches support up to 16 forwarding classes.

You can configure forwarding classes in one of the following ways:

- Using class statement—You can configure up to 16 forwarding classes and you can map multiple forwarding classes to single queue.
- Using queue statement—You can configure up to 8 forwarding classes and you can map one forwarding class to one queue.

This example uses the class statement to configure forwarding classes.

To configure CoS forwarding classes, map the forwarding classes to queues:

```
[edit class-of-service forwarding-classes]
user@switch# set class be queue-num 0
user@switch# set class ef queue-num 1
user@switch# set class af queue-num 2
user@switch# set class nc queue-num 3
user@switch# set class ef1 queue-num 4
user@switch# set class ef2 queue-num 5
user@switch# set class af1 queue-num 6
user@switch# set class nc1 queue-num 7
```

RELATED DOCUMENTATION

Defining CoS Forwarding Classes (J-Web Procedure) 115		
Example: Configuring CoS on EX Series Switches 15		
Example: Prioritizing Snooped and Inspected Packet		
Assigning CoS Components to Interfaces (CLI Procedure) 55		

Monitoring CoS Forwarding Classes | 117

Understanding CoS Forwarding Classes | 110

Defining CoS Forwarding Classes (J-Web Procedure)

NOTE: This topic applies only to the J-Web Application package.

You can define CoS forwarding classes on an EX Series switch using the J-Web interface. Assigning a forwarding class to a queue number affects the scheduling and marking of a packet as it transits a switch.

To define forwarding classes:

1. Select Configure > Class of Service > Forwarding Classes.

NOTE: After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See Using the Commit Options to Commit Configuration Changes (J-Web Procedure) for details about all commit options.

- 2. Select one of the following options:
 - Add—Adds a forwarding class. Enter information into the forwarding class page as described in Table 21 on page 115.
 - Edit—Modifies an existing forwarding class. Enter information into the forwarding class page as described in Table 21 on page 115.
 - Delete–Deletes an existing forwarding class.

Table 21: Forwarding Classes Configuration Fields

Field	Function	Your Action
Forwarding Class Summary		

Field	Function	Your Action
Queue #	Specifies the internal queue numbers to which forwarding classes are assigned. By default, if a packet is not classified, it is assigned to the class associated with queue 0. You can have more than one forwarding class to a queue number.	To specify an internal queue number, select an integer from 0 through 11, appropriate for your platform as follows: NOTE : For EX2300 and EX2300-C switches, a maximum of eight egress queues are supported per port. To specify an internal queue number select an integer from 0 through 7.
Forwarding Class Name	Specifies the forwarding class names assigned to specific internal queue numbers. By default, four forwarding classes are assigned to queue numbers 0 (best-effort), 1 (assured-forwarding), 5 (expedited- forwarding), and 7 (network-connect).	Type the name—for example, be-class.

Table 21: Forwarding Classes Configuration Fields (Continued)

RELATED DOCUMENTATION

Defining CoS Forwarding Classes (CLI Procedure) 114	
Example: Configuring CoS on EX Series Switches 15	
Example: Prioritizing Snooped and Inspected Packet	
Monitoring CoS Forwarding Classes 117	
Assigning CoS Components to Interfaces (J-Web Procedure) 56	
Understanding CoS Forwarding Classes 110	

Monitoring CoS Forwarding Classes

IN THIS SECTION

- Purpose | **117**
- Action | **117**
- Meaning | 117

Purpose

NOTE: This topic applies only to the J-Web Application package.

View the current assignment of CoS forwarding classes to queues on the switch.

Action

To monitor CoS forwarding classes in the J-Web interface, select **Monitor > Class of Service > Forwarding Classes**.

To monitor CoS forwarding classes in the CLI, enter the following CLI command:

show class-of-service forwarding-class

Meaning

Table 22 on page 118 summarizes key output fields for CoS forwarding classes.

Table 22: Summary of Key CoS Forwarding Class Output Fields

Field	Values
Forwarding Class	Names of forwarding classes assigned to queue numbers. The following are the default forwarding classes:
	• best-effort —Provides no special CoS handling of packets. Loss priority is typically not carried in a CoS value.
	• expedited-forwarding —Provides low loss, low delay, low jitter, assured bandwidth, and end-to-end service.
	• assured-forwarding —Provides high assurance for packets within the specified service profile. Excess packets are dropped.
	• network-control —Packets can be delayed but not dropped.
	EX8200 switches have the following additional default forwarding classes:
	• mcast-be —Provides no special CoS handling of packets.
	• mcast-ef —Provides low loss, low delay, low jitter, assured bandwidth, and end-to-end service.
	• mcast-af —Provides high assurance for packets within the specified service profile. Excess packets are dropped.
	• mcast-nc —Provides multicast network-control traffic.
Queue	Queue number corresponding to the forwarding class name. The default forwarding classes are assigned as follows:
	• best-effort -0
	• expedited-forwarding-5
	• assured-forwarding-1
	• network-control-7
	• mcast-be-2
	• mcast-ef-4
	• mcast-af-6

Field	Values
Fabric Priority	(EX8200 switches only) Fabric priority for the forwarding class, either high or low . The fabric priority determines the priority of packets entering the switch fabric.

Table 22: Summary of Key CoS Forwarding Class Output Fields (Continued)

RELATED DOCUMENTATION

Defining CoS Forwarding Classes (CLI Procedure) | 114

Defining CoS Forwarding Classes (J-Web Procedure) | 115

Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure) | 87

Example: Configuring CoS on EX Series Switches | 15

Flow Control

IN THIS CHAPTER

- Understanding Priority-Based Flow Control | 120
- Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure) | 123

Understanding Priority-Based Flow Control

IN THIS SECTION

- Reliability of Packet Delivery in Standard Ethernet Networks and in Layer 2 Networks | 120
- Calculations for Buffer Requirements When Using PFC PAUSE | 121
- How PFC and Congestion Notification Profiles Work With or Without DCBX | 122

Priority-based flow control (PFC), IEEE standard 802.1Qbb, is a link-level flow control mechanism. The flow control mechanism is similar to that used by IEEE 802.3x Ethernet PAUSE, but it operates on individual priorities. Instead of pausing all traffic on a link, PFC allows you to selectively pause traffic according to its class.

This topic describes:

Reliability of Packet Delivery in Standard Ethernet Networks and in Layer 2 Networks

Standard Ethernet does not guarantee that a packet injected into the network will arrive at its intended destination. Reliability is provided by upper-layer protocols. Generally, a network path consists of multiple hops between the source and destination. A problem arises when transmitters send packets faster than receivers can accept them. When receivers run out of available buffer space to hold incoming flows, they silently drop additional incoming packets. This problem is generally resolved by upper-layer protocols that detect the drops and request retransmission.

Applications that require reliability in Layer 2 must have flow control that includes feedback from a receiver to a sender regarding buffer availability. Using IEEE 802.3x Ethernet PAUSE control frames, a receiver can generate a MAC control frame and send a PAUSE request to a sender when a specified threshold of receiver buffer has been filled to prevent buffer overflow. Upon receiving a PAUSE request, the sender stops transmission of any new packets until the receiver notifies the sender that it has sufficient buffer space to accept them again. The disadvantage of using Ethernet PAUSE is that it operates on the entire link, which might be carrying multiple traffic flows. Some traffic flows do not need flow control in Layer 2, because they are carrying applications that rely on upper-layer protocols for reliability. PFC enables you to configure Layer 2 flow control selectively for the traffic that requires it, such as Fibre Channel over Ethernet (FCoE) traffic, without impacting other traffic on the link. You can also enable PFC for other traffic types, such as iSCSI.

Calculations for Buffer Requirements When Using PFC PAUSE

The receive buffer must be large enough to accommodate all data that is received while the system is responding to a PFC PAUSE frame.

When you calculate buffer requirements, consider the following factors:

- Processing and queuing delay of the PFC PAUSE—In general, the time to detect the lack of sufficient buffer space and to transmit the PFC PAUSE is negligible. However, delays can occur if the switch detects a reduction in buffer space just as the transmitter is beginning to transmit a maximum length frame.
- Propagation delay across the media—The delay amount depends on the length and speed of the physical link.
- Response time to the PFC PAUSE frame
- Propagation delay across the media on the return path

NOTE: We recommend that you configure at least 20 percent of the buffer size for the queue that is using PFC and that you do not specify the **exact** option.

Because it is mandatory to explicitly configure a certain percentage of buffer size for PFC, you must also explicitly configure some buffer size for any other forwarding classes that you are planning to use (including the default forwarding classes and the user-defined forwarding classes). The percentage that you allocate depends on the usage of the respective classes.

How PFC and Congestion Notification Profiles Work With or Without DCBX

PFC can be applied to an interface regardless of whether the Data Center Bridging Capability Exchange protocol (DCBX) is enabled (DCBX is enabled by default for 10-Gigabit Ethernet interfaces on EX4500 CEE-enabled switches).

However, automatic control and advertisement of PFC requires DCBX:

- When DCBX is enabled—DCBX detects the data center bridging (DCB) neighbor's PFC configuration, uses autonegotiation to advertise local and peer PFC configuration, and then enables or disables PFC depending on whether the configurations are compatible or not. When PFC is enabled, it uses the congestion notification profile, which you have configured and applied to the interface.
- When DCBX is not enabled—*Class of service* (CoS) triggers PFC when the incoming frame has a User Priority (UP) field that matches the three-bit pattern specified for the congestion notification profile.

To manually control the use of PFC on the interface regardless of the configuration of the peer data center devices, you can explicitly change the configuration of DCBX on the interface to disable PFC autonegotiation. See *Disabling DCBX to Disable PFC Autonegotiation on EX Series Switches (CLI Procedure).* When PFC autonegotiation is disabled, PFC is triggered by the congestion notification profile for PFC regardless of the configuration of the DCB peer.

NOTE: PFC functions effectively only when the peer devices connected to the local interface are also using PFC and are configured compatibly with the local interface. PFC must be symmetrical —if PFC is not configured to use the same traffic class (code point) on both the local and the peer interface, it does not have any impact on the traffic.

Table 23 on page 122 shows the one-to-one mapping between the UP field of an IEEE 802.1Q tagged frame, the traffic class, and the egress queue. In addition to setting a PFC congestion notification profile on an ingress port, you must set a forwarding class to match the priority specified in the PFC congestion notification profile and to forward the frame to the appropriate queue.

Juniper Networks EX Series Ethernet Switches support up to six traffic classes and allow you to associate those classes with six different congestion notification profiles. (The switches support up to 16 forwarding classes.)

Table 23: Input for PFC Congestion Notification Profile and Mapping to Traffic Class and Egress Queue

UP Field of IEEE-802.1Q Tagged Frame	Traffic Class	Egress Queue
000	ТС 0	queue 0

UP Field of IEEE-802.1Q Tagged Frame	Traffic Class	Egress Queue
001	TC 1	queue 1
010	TC 2	queue 2
011	ТС 3	queue 3
100	TC4	queue 4
101	TC 5	queue 5

 Table 23: Input for PFC Congestion Notification Profile and Mapping to Traffic Class and Egress Queue

 (Continued)

RELATED DOCUMENTATION

Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches Example: Configuring an FCoE Transit Switch Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure) | **123**

schedulers

congestion-notification-profile

Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure)

You can configure priority-based flow control (PFC) on EX4500 switches to apply link-level flow control on a specific traffic class so that different types of traffic can efficiently use the same network interface card (NIC). You must configure PFC for all interfaces carrying Fibre Channel over Ethernet (FCoE) traffic. You can also configure PFC on interfaces carrying other traffic types, such as Internet small computer system interface (iSCSI) traffic. Using PFC is optional for traffic types other than FCoE.

NOTE:

- PFC is supported only on 10-Gigabit Ethernet interfaces.
- If you are using PFC for a non-FCoE DCBX application, use the same 802.1p code points for the PFC congestion notification profile and for the application map that is carrying that application traffic.

Data Center Bridging Capability Exchange protocol (DCBX) is enabled by default on all 10-Gigabit Ethernet interfaces on EX4500 switches. DCBX enables or disables PFC on the local interface depending on whether the PFC configuration on that interface is the same as the PFC configuration of the connected interface on the data center bridging (DCB) peer.

NOTE: When you configure PFC, we recommend that you:

- Configure at least 20 percent of the buffer for the queue that is using PFC.
- Configure an appropriate percent of the buffer for any other forwarding classes (default forwarding classes and the user-defined forwarding classes) that you are using.
- Do not specify the exact option when configuring the buffer for the queue that is using PFC.
- Configure the loss-priority statement to low for a traffic class that is using PFC.
- Verify that the PFC configurations of the local interfaces are the same as the PFC configurations of the connected interfaces on the DCB peer. See *show dcbx neighbors*.

EX Series switches support up to six congestion notification profiles for PFC.

To configure PFC:

1. Configure a congestion notification profile, specifying the name of the profile and specifying the three-bit pattern of the User Priority bits in an incoming frame that will trigger the priority-based flow control on that traffic class:

```
[edit class-of-service]
user@switch# set congestion-notification-profile profile-name input ieee-802.1 code-point
up-bits pfc
```

2. Disable standard Ethernet flow control on the interfaces that will be used for the traffic class that you have selected for PFC:

```
[edit interfaces]
user@switch# set interface-name ether-options no-flow-control
```

NOTE: You cannot apply PFC to interfaces that are using standard Ethernet flow control. You must first disable flow control on those interfaces.

3. Bind the congestion notification profile to the interfaces that will be used for the traffic class that you have selected for PFC:

```
[edit class-of-service]
user@switch# set interfaces interface-name congestion-notification-profile profile-name
```

4. Create a CoS classifier for a traffic class that will use PFC:

```
[edit class-of-service]
user@switch# set classifiers ieee-802.1 classifier-name import default
```

5. Configure this traffic class (*classifier-name*) to use a user-defined or default forwarding class with a low loss priority value and specify the 802.1p code points::

```
[edit class-of-service]
user@switch# set classifiers ieee-802.1 classifier-name forwarding-class class-name loss-
priority low code-points 3 bit-patterns
```

6. Bind the *classifier-name* classifier to all interfaces that require PFC:

```
[edit class-of-service]
user@switch# set interfaces interface-name unit logical-unit-number classifiers ieee-802.1
classifier-name
```

7. Assign the specified forwarding-class to an egress queue:

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

8. Set a scheduler for this queue, allocating at least 20 percent of the buffer to be used for FCoE traffic:

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

9. Set a scheduler to allocate buffer space for forwarding classes carrying other traffic:

NOTE: You must explicitly allocate some buffer space for the other forwarding classes. The default allocation of buffer space for forwarding classes is overridden when you manually configure the requisite amount of buffer space for the FCoE traffic.

[edit class-of-service]
user@switch# set scheduler-name buffer-size percent

10. Configure a scheduler map that associates the specified scheduler with the specified forwarding class:

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

For example:

```
[edit class-of-service]
user@switch# set scheduler-maps pfc-map forwarding-class af2 scheduler pfc-sched
user@switch# set scheduler-maps pfc-map forwarding-class best-effort scheduler default-
sched user@switch# set scheduler-maps pfc-map forwarding-class network-control scheduler
default-sched
user@switch# set scheduler-maps pfc-map forwarding-class expedited-forwarding scheduler
default-sched
```

11. Assign the scheduler map to the egress interface:

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

RELATED DOCUMENTATION

Example: Configuring an FCoE Transit Switch Understanding Priority-Based Flow Control | 120

congestion-notification-profile

CoS Queue Schedulers and Scheduler Maps

IN THIS CHAPTER

- Understanding CoS Schedulers | 128
- Defining CoS Schedulers and Scheduler Maps (CLI Procedure) | 138
- Defining CoS Schedulers (J-Web Procedure) | 141
- Defining CoS Scheduler Maps (J-Web Procedure) | 146
- Monitoring CoS Scheduler Maps | **147**
- Troubleshooting CoS Schedulers on a 40-port SFP+ Line Card in an EX8200 Switch | 150

Understanding CoS Schedulers

IN THIS SECTION

- Default Schedulers | 129
- Excess Rate | 130
- Transmission Rate | 130
- Scheduler Buffer Size | **130**
- Priority Scheduling | 131
- Scheduler Drop-Profile Maps | 132
- Scheduler Maps | 132

You use class-of-service (CoS) schedulers to define the properties of output queues on Juniper Networks EX Series Ethernet Switches. These properties include the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, the priority of the queue, and the drop profiles associated with the queue. You associate the schedulers with forwarding classes by means of scheduler maps. You can then associate each scheduler map with an interface, thereby configuring the queues, packet schedulers, and tail drop processes that operate according to this mapping.

This topic describes:

Default Schedulers

Each forwarding class has an associated scheduler priority. On EX Series switches other than Juniper Networks EX8200, EX4300, and EX3400 Ethernet Switches, only two forwarding classes—best-effort (queue 0) and network-control (queue 7)—are used in the default configuration. By default on these switches, the best-effort forwarding class (queue 0) receives 95 percent of the bandwidth and the buffer space for the output link, and the network-control forwarding class (queue 7) receives 5 percent. The default drop profile causes the buffer to fill completely and then to discard all incoming packets until it has free space.

On EX8200 switches three forwarding classes—best-effort (queue 0), multicast best-effort (queue 2), and network-control (queue 7)—are used in the default configuration. By default, the best-effort forwarding class (queue 0) receives 75 percent of the bandwidth, the multicast best-effort forwarding class (queue 2) receives 20 percent, and the network-control forwarding class (queue 7) receives 5 percent of the bandwidth and buffer space for the output link.

On EX4300 and EX 3400 switches, four forwarding classes—best-effort (queue 0), multicast best-effort (queue 8), network-control (queue 3), and multicast network-control (queue 11)—are used in the default configuration. By default, all the multicast traffic flows through the multicast best-effort queue. EX4300 and EX3400 switches support 12 queues (0–11), and the default scheduler transmission rates 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 bandwidth.

On EX Series switches other than EX4300 switches, the expedited-forwarding (queue 5) and assured-forwarding (queue 1) classes have no scheduler because no resources are assigned to queue 5 or queue 1, by default. However, you can manually configure resources to be assigned to the expedited-forwarding and assured-forwarding classes. On EX4300 switches, the expedited-forwarding (queue 1) and assured-forwarding (queue 2) classes have no scheduler because no resources are assigned to queue 1 or queue 2, by default. However, you can manually configure resources to be assigned to the expedited-forwarding and assured-forwarding classes.

Also by default, any queue can exceed the assigned bandwidth if additional bandwidth is available from other queues. When a forwarding class does not fully use the allocated transmission bandwidth, the remaining bandwidth can be used by other forwarding classes if they have a traffic load that exceeds their allocated bandwidth.

Excess Rate

Excess rate traffic determines the percentage of the excess bandwidth to share when a queue receives traffic in excess of its bandwidth allocation. By default, the excess bandwidth is shared in the ratio of the transmit rates. You can control this distribution by configuring the excess-rate statement at the [edit class-of-service schedulers *scheduler-name*] hierarchy. You can specify the excess rate sharing in percentage.

Transmission Rate

Transmission-rate control determines the actual traffic bandwidth for each forwarding class you configure. The transmission rate is specified in bits per second. Each queue is allocated some portion of the bandwidth of the interface. This bandwidth can be a fixed value, such as 1 megabit per second (Mbps), a percentage of the total available bandwidth, or the rest of the available bandwidth. In case of congestion, the configured transmission rate is guaranteed for the queue. Transmission-rate control allows you to ensure that each queue receives the bandwidth appropriate for its level of service.

Scheduler Buffer Size

To control congestion at the output stage, you can configure the delay-buffer bandwidth by using the buffer-size *configuration statement*. The delay-buffer bandwidth provides packet buffer space to absorb burst traffic up to the specified duration of delay. When the specified delay buffer becomes full, packets with 100 percent drop probability are dropped from the tail of the buffer.

On EX Series switches other than EX8200, EX4300, and EX3400 switches, the default scheduler transmission rates for queues 0 through 7 are 95, 0, 0, 0, 0, 0, 0, 0, and 5 percent, respectively, of the total available bandwidth. The default buffer-size percentages for queues 0 through 7 are 95, 0, 0, 0, 0, 0, 0, and 5 percent, respectively, of the total available buffer.

On EX8200 switches, the default scheduler transmission rates for queues 0 through 7 are 75, 0, 20, 0, 0, 0, 0, and 5 percent, respectively, of the total available bandwidth. The default buffer-size percentages for queues 0 through 7 are 75, 0, 20, 0, 0, 0, 0, and 5 percent, respectively, of the total available buffer.

On EX4300 and EX3400 switches, the default scheduler transmission rates 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. The default buffersize percentages 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.

For each scheduler on EX Series switches other than EX8200 switches, you can configure the buffer size as one of the following:

- The exact buffer size.
- A percentage of the total buffer.

• The remaining buffer available. The remainder is the buffer percentage that is not assigned to other queues. For example, if you assign 40 percent of the delay buffer to queue 0, allow queue 2 to keep the default allotment of 20 percent, allow queue 7 to keep the default allotment of 5 percent, and assign the remainder to queue 3, then queue 3 uses 35 percent of the delay buffer.

On EX8200 switches, you can configure the buffer size as a temporal value (in microseconds), percentage of the total buffer, or the remaining buffer available. You can configure the buffer size as a temporal value on Juniper Networks EX4200 and EX4300 Ethernet Switches also.

When you configure buffer size as a temporal value on EX4200 switches, if sufficient buffer size is not available in the shared pool, an error message is logged in the system log (syslog) file and the default profile is applied to the interface. After the temporal buffer space is allocated successfully, if the shared buffer size is less than the current value (which was set using the set class-of-service shared-buffer percent *value* command), the new reduced value must be greater than a sum of the existing reserved temporal buffer size and the required minimum buffer size. Otherwise, the modification to the shared-buffer configuration fails and an error message is logged in the system log.

Priority Scheduling

Priority scheduling determines the order in which an interface transmits traffic from queues, thus ensuring that queues containing important traffic are provided faster access.

Priority scheduling is accomplished through a procedure in which the scheduler examines the priority of the queue. Juniper Networks Junos operating system (Junos OS) supports two levels of transmission priority:

• Low—The scheduler determines whether the individual queue is within its defined bandwidth profile or not. This binary decision, which is re-evaluated on a regular time cycle, involves comparing the amount of data transmitted by the queue against the bandwidth allocated to it by the scheduler. If the transmitted amount is less than the allocated amount, the queue is considered to be in profile. A queue is out of profile when the amount of traffic that it transmits is larger than the queue's allocated limit. An out-of-profile queue is transmitted only if bandwidth is available. Otherwise, it is buffered.

On EX Series switches other than EX4300 switches, a queue from a set of queues is selected based on the shaped deficit weighted round robin (SDWRR) algorithm, which operates within the set. On EX4300 switches, the weighted deficit round-robin (WDRR) algorithm is used to select a queue from a set of queues.

 Strict-high—A strict-high priority queue receives preferential treatment over a low-priority queue. Unlimited bandwidth is assigned to a strict-high priority queue. On EX Series switches other than EX4300 switches, queues are scheduled according to the queue number, starting with the highest queue, 7, with decreasing priority down through queue 0. Traffic in higher-numbered queues is always scheduled prior to traffic in lower-numbered queues. In other words, if there are two highpriority queues, the queue with the higher queue number is processed first. On EX4300 switches, you can configure multiple strict-high priority queues on an interface and an EX4300 switch processes these queues in a round-robin method.

Packets in low-priority queues are transmitted only when strict-high priority queues are empty.

Scheduler Drop-Profile Maps

Drop-profile maps associate drop profiles with a scheduler. A drop-profile map sets the drop profile for a specific packet loss priority (PLP) and protocol type. The inputs for a drop-profile map are the PLP and the protocol type. The output is the drop profile.

Scheduler Maps

A scheduler map associates a specified forwarding class with a scheduler configuration. After configuring a scheduler, you must include it in a scheduler map and then associate the scheduler map with an output interface.

On EX Series switches, if you configure more than the supported number of scheduler maps on a switch or for a port group in a line card, an error is logged in the system log. On any interface in a port group on a line card or on a switch, if you configure a scheduler map that causes the number of scheduler maps for that port group to exceed the maximum number supported, the default scheduler map is bound to that interface. We recommend that you check the system log for errors after the commit operation to verify that you have not configured more than the maximum permitted number of scheduler maps.

NOTE: On EX Series switches, you cannot configure a scheduler map on an individual interface that is a member of a link aggregation group (LAG). Instead, you must configure the scheduler map on the LAG itself (that is, on the aggregated Ethernet (**ae**) interface).

Table 24 on page 132 shows the number of scheduler maps supported for each port group in a switch or line card.

Switch/Line Card	Number of Port Groups	Port Grouping Details	Number of Scheduler Maps Supported for Each Port Group
EX2200-C-12T and EX2200-C-12P switches	1	Port 0–11 and 2 uplink ports form a port group.	6

Table 24: Support for Scheduler Maps on Switches and Line Cards

Switch/Line Card	Number of Port Groups	Port Grouping Details	Number of Scheduler Maps Supported for Each Port Group
EX2200-24T and EX2200-24P switches	1	Ports 0–23 and 4 SFP uplink ports form a port group.	5
EX2200-48T and EX2200-48P switches	2	 Ports 0-23 and SFP uplink ports 0 and 1 form a port group. Ports 24-47 and SFP uplink ports 2 and 3 form a port group. 	5
EX3200-24T and EX3200-24P switches	1	 Ports 0–23 and the uplink ports form a port group. NOTE: Uplink ports include 2 SFP+ or XFP uplink ports, or 4 SFP uplink ports. 	4
EX3200-24T and EX3200-24P switches	1	 Ports 0–23 and the uplink ports form a port group. NOTE: Uplink ports include 2 SFP+ or XFP uplink ports or 4 SFP uplink ports. 	4
EX3200-48T and EX3200-48P switches	2	 Ports 0-23 and 1 SFP+ or XFP uplink port or 4 SFP uplink ports form a port group. Ports 24-47 and 1 SFP+ or XFP uplink port form a port group. 	4

Table 24: Support for Scheduler Maps on Switches and Line Cards (Continued)

Switch/Line Card	Number of Port Groups	Port Grouping Details	Number of Scheduler Maps Supported for Each Port Group
EX4200-48T and EX4200-48P switches	3	 Ports 0-23 form a port group. Ports 24-47 form a port group. 2 SFP+ or XFP uplink ports or 4 SFP uplink ports form a port group. 	4
EX4200-24T and EX4200-24P switches	2	 Ports 0-23 form a port group. 2 SFP+ or XFP uplink ports or 4 SFP uplink ports form a port group. 	4
EX4300-24Tand EX4300-24P switches	1	 Ports 0-23 ports, 4 uplink ports, and 4 ports on the real panel form a port group. NOTE: Uplink ports in the front panel contains SFP or SFP+ ports 0-3, and uplink ports in the rear panel contains QSFP+ ports 0-3. 	64

Table 24: Support for Scheduler Maps on Switches and Line Cards (Continued)

Switch/Line Card	Number of Port Groups	Port Grouping Details	Number of Scheduler Maps Supported for Each Port Group
EX4300-48T and EX4300-48P switches	1	 Ports 0-47, 4 uplink ports, and 4 ports on the real panel form a port group. NOTE: Uplink ports in the front panel contains SFP or SFP+ ports 0-3, and uplink ports in the rear panel contains QSFP+ ports 0-3. 	64
EX4500-40F switch	2	 SFP or SFP+ ports 0-19 and the first SFP or SFP+ port 0-4 form a port group. SFP or SFP+ ports 20-39 and the second SFP or SFP + uplink port 0-4 form a port group. 	4
EX4550-32F switch	1	 SFP or SFP+ ports 0-31 and the uplink ports in the front and rear panels form a port group. NOTE: Uplink ports in the front panel contains SFP, SFP+, or RJ-45 ports 0-7, and uplink ports in the rear panel contains SFP, SFP+, or RJ-45 ports 0-7. 	5
EX6200-48T (48-port RJ-45) and EX6200-48P (48-port PoE+) line cards	2	 Ports 0-23 form a port group. Ports 24-47 form a port group. 	5

Table 24: Support for Scheduler Maps on Switches and Line Cards (Continued)

Switch/Line Card	Number of Port Groups	Port Grouping Details	Number of Scheduler Maps Supported for Each Port Group
EX6200-SRE64-4XS	1	SFP+ ports 0–3 form a port group.	4
EX8200-8XS (8-port SFP+) line card	4	 SFP+ ports 0 and 1 form a port group. SFP+ ports 2 and 3 form a port group. SFP+ ports 4 and 5 form a port group. SFP+ ports 6 and 7 form a port group. 	6
EX8200-40XS (40- port SFP+) line card	8	 SFP+ ports 0-4 form a port group. SFP+ ports 5-9 form a port group. SFP+ ports 10-14 form a port group. SFP+ ports 15-19 form a port group. SFP+ ports 20-24 form a port group. SFP+ ports 25-29 form a port group. SFP+ ports 30-34 form a port group. SFP+ ports 35-39 form a port group. 	6

Table 24: Support for Scheduler Maps on Switches and Line Cards (Continued)

Switch/Line Card	Number of Port Groups	Port Grouping Details	Number of Scheduler Maps Supported for Each Port Group
EX8200-48-F (48- port SFP) and EX8200-48T (48-port RJ-45) line cards	2	 SFP or RJ-45 ports 0-23 form a port group. SFP or RJ-45 ports 24-47 form a port group. 	6
EX8200-2XS-40P 3 (40-port PoE+ with 4- port SFP and 2-port SFP+) line card	 Ports 0-19 and SFP ports 0 and 1 form a port group. Ports 20-39 and SFP ports 2 and 3 form a port group. 	5	
	• 2 SFP+ ports form a port group.	6	
EX8200-2XS-40T (40-port RJ-45 with 4-port SFP and 2-port SFP+) line card	3	 Ports 0–19, and SFP ports 0 and 1 form a port group. Ports 20–39 and SFP ports 2 and 3 form a port group. 	5
		• 2 SFP+ ports form a port group.	6
EX8200-48PL (48- port PoE+ 20 Gbps) and EX8200-48TL (48-port RJ-45 20 Gbps) line cards	2	 PoE+ or RJ-45 ports 0-23 form a port group. PoE+ or RJ-45 ports 24-47 form a port group. 	5

Table 24: Support for Scheduler Maps on Switches and Line Cards (Continued)

RELATED DOCUMENTATION

Understanding Junos OS CoS Components for EX Series Switches | 8 Example: Configuring CoS on EX Series Switches | 15 Defining CoS Schedulers and Scheduler Maps (CLI Procedure) | 138 Defining CoS Schedulers (J-Web Procedure) | 141

Defining CoS Schedulers and Scheduler Maps (CLI Procedure)

IN THIS SECTION

- Configuring a Scheduler and a Scheduler Map | 138
- Assigning a Scheduler Map to Interfaces | 139
- Assigning Scheduler Maps to Interfaces on EX8200 Line Cards That Include Oversubscribed Ports | 140

You use schedulers to define the class-of-service (CoS) properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, the priority of the queue, and the drop profiles associated with the queue.

You associate the schedulers with forwarding classes by means of scheduler maps. You can then associate each scheduler map with an interface, thereby configuring the queues and packet schedulers that operate according to this mapping.

NOTE: On EX Series switches, you cannot configure a scheduler map on an individual interface that is a member of a link aggregation group (LAG). Instead, you must configure the scheduler map on the LAG itself (that is, on the aggregated Ethernet (ae) interface).

You can associate up to four user-defined scheduler maps with an interface.

This topic describes:

Configuring a Scheduler and a Scheduler Map

You can define the properties for an output queue by configuring a scheduler. You can then define a scheduler map to associate a forwarding class with a scheduler.

To configure a scheduler and a scheduler map:

1. Create a scheduler, and assign one or more output queue properties to it:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name output-queue-properties
```

For various properties that you can define for an output queue, see the "schedulers" on page 305 hierarchy.

2. Configure a scheduler map that associates the scheduler with the forwarding class:

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

Assigning a Scheduler Map to Interfaces

After defining a scheduler map, you can assign the scheduler map to one or more interfaces. You can also assign the scheduler map to multiple interfaces by using a wildcard representation of the interface or Virtual Chassis Ports (VCPs).

Following are sample syntaxes and examples for assigning a scheduler map to a single or to multiple interfaces:

• To assign the scheduler map to one interface:

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

• To assign the scheduler map to more than one interface, you can use a wildcard representation of the interface:

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

For example, following is the configuration to assign the be-map scheduler map to all Gigabit Ethernet interfaces (ge-*):

[edit class-of-service interfaces]
user@switch# set ge-* scheduler-map be-map

• To assign the scheduler map to all VCPs:

[edit class-of-service interfaces]
user@switch# set wild-card-representation-of-vcp scheduler-map map-name

NOTE: You can assign a scheduler map to a VCP only on EX4200, EX4300 or EX4500 switches that are members of Virtual Chassis composed exclusively either of EX4200 switches, EX4300 switches or of EX4500 switches, or that are members of a mixed Virtual Chassis composed of EX4200, EX4300, and EX4500 switches.

For example, following is the configuration to assign the be-map scheduler map to all VCPs:

[edit class-of-service interfaces]
user@switch# set vcp-* scheduler-map be-map

Assigning Scheduler Maps to Interfaces on EX8200 Line Cards That Include Oversubscribed Ports

Some line cards available for Juniper Networks EX8200 Ethernet Switches include oversubscribed ports that are combined in logical port groups that share bandwidth. These oversubscribed ports handle traffic differently than ports that provide continuous line-rate bandwidth. You might need to configure CoS queues differently for oversubscribed ports than for line-rate ports. For more information about EX8200 line cards that include oversubscribed ports, see "Understanding CoS Queues on EX8200 Line Cards That Include Oversubscribed Ports" on page 11.

For interfaces on oversubscribed EX8200 line cards, you use the same procedure to configure CoS schedulers as you do for other interfaces. However, you must assign the same scheduler map to all the interfaces in a port group. When you assign a scheduler map to one interface in a port group, you do not need to assign the scheduler map to the remaining interfaces in the port group. The switch automatically uses that scheduler map for all the interfaces in the port group when you bring the interfaces up. Therefore, you do not need to assign the scheduler map to the scheduler map to the remaining interfaces in that port group.

If you assign different scheduler maps to different interfaces in a port group, you do not receive an error when you commit the configuration. Instead, an error is logged in the system log. When you bring an interface in the port group up, the default scheduler map is assigned to all interfaces in the port group. If you assign a scheduler map to an interface that is down and if that scheduler map is different from the scheduler map is used by the currently operating interfaces in the port group, then the default scheduler map is used by all interfaces in the port group, even the currently operating ones, when you bring the interface up.

To assign a scheduler map to a port group, assign a scheduler map to at least one interface in the port group:

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

Considering that the xe-0/0/2 interface is part of a port group, following is the configuration to assign a scheduler map named ef-map to a port group that contains xe-0/0/2:

[edit class-of-service interfaces]
user@switch# set xe-0/0/2 scheduler-map ef-map

RELATED DOCUMENTATION

Defining CoS Schedulers (J-Web Procedure) | 141

Example: Configuring CoS on EX Series Switches | 15

Assigning CoS Components to Interfaces (CLI Procedure) | 55

Monitoring CoS Scheduler Maps | 147

Understanding CoS Schedulers | 128

Defining CoS Schedulers (J-Web Procedure)

NOTE: This topic applies only to the J-Web Application package.

You can use the J-Web interface to define CoS schedulers on an EX Series switch. Using schedulers, you can assign attributes to queues and thereby provide congestion control for a particular class of traffic. These attributes include the amount of interface bandwidth, memory buffer size, transmit rate, and schedule priority.

To configure schedulers:

1. Select Configure > Class of Service > Schedulers.

NOTE: After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See Using the Commit Options to Commit Configuration Changes (J-Web Procedure) for details about all commit options.

- **2.** Select one of the following options:
 - Add—Adds a scheduler. Enter information into the Schedulers page as described in Table 25 on page 142.
 - Edit—Modifies an existing scheduler. Enter information into the Schedulers page as described in Table 25 on page 142.
 - Delete–Deletes an existing scheduler.

Table 25: Schedulers Configuration Page

Field	Function	Your Action
Scheduler name	Specifies the name for a scheduler.	To name a scheduler, type the name—for example, be-scheduler .
Scheduling priority	Sets the transmission priority of the scheduler, which determines the order in which an output interface transmits traffic from the queues. You can set the scheduling priority at different levels in the order of increasing priority from low to high. A high-priority queue with a high transmission rate might lock out lower- priority traffic.	 To set a priority, select one: low—Packets in this queue are transmitted last. strict-high—Packets in this queue are transmitted first. To specify no scheduling priority, select the blank check box.

Field	Function	Your Action
Buffer size	Defines the size of the delay buffer. By default, queues 0 through 11 are allotted the following percentages of the total available buffer space: • Queue 0–75 percent • Queue 1–0 percent • Queue 2–0 percent • Queue 3–5 percent • Queue 3–5 percent • Queue 5–0 percent • Queue 5–0 percent • Queue 6–0 percent • Queue 7–0 percent • Queue 8–15 percent • Queue 9–0 percent • Queue 10–0 percent • Queue 11–5 percent NOTE: A large buffer size value correlates with a greater possibility of packet delays. Such a value might not be practical for sensitive traffic such as voice or video.	<text><list-item><list-item><list-item></list-item></list-item></list-item></text>

Table 25: Schedulers Configuration Page (Continued)

Field	Function	Your Action
Shaping rate	Specifies the rate at which queues transmit packets.	• To specify shaping rate as a percentage, select Percent and type an integer from 1 through 100.
		• To specify shaping rate as a number, select Rate and enter a value.
		• To specify no shaping rate, select the blank check box.

Table 25: Schedulers Configuration Page (Continued)

Transmit rate Defines the transmission rate of a To	o define a transmit rate, select the opropriate option:
scheduler. ap The transmit rate determines the traffic bandwidth from each forwarding class you configure. By default, queues 0 through 11 are allotted the following percentages of the transmission capacity: • Queue 0–75 percent • Queue 1–0 percent • Queue 2–0 percent • Queue 3–5 percent • Queue 4–0 percent • Queue 5–0 percent • Queue 5–0 percent • Queue 6–0 percent • Queue 8–15 percent • Queue 9–0 percent • Queue 10–0 percent • Queue 10–0 percent • Queue 11–5 percent	To enforce the exact transmission rate, select Rate and enter a value. To specify the remaining transmission capacity, select Remainder Available .

Table 25: Schedulers Configuration Page (Continued)

RELATED DOCUMENTATION

Defining CoS Schedulers and Scheduler Maps (CLI Procedure) | 138

Example: Configuring CoS on EX Series Switches | 15

Monitoring CoS Scheduler Maps | 147

Defining CoS Scheduler Maps (J-Web Procedure)

NOTE: This topic applies only to the J-Web Application package.

You can use the J-Web interface to configure CoS scheduler maps on an EX Series switch.

NOTE: On EX Series switches, you cannot configure a scheduler map on an individual interface that is a member of a link aggregation group (LAG). Instead, you must configure the scheduler map on the LAG itself—that is, on the aggregated Ethernet (**ae**) interface.

To configure scheduler maps:

1. Select Configure > Class of Service > Scheduler Maps.

NOTE: After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See Using the Commit Options to Commit Configuration Changes (J-Web Procedure) for details about all commit options.

2. Select one of the following options:

- Add—Adds a scheduler map. Enter information into the scheduler map page as described in Table 26 on page 146.
- Edit—Modifies an existing scheduler map. Enter information into the scheduler map page as described in Table 26 on page 146.
- **Delete**–Deletes an existing scheduler map.

Table 26: Scheduler Maps Configuration Fields

Field	Function	Your Action
Scheduler Map Name	Specifies the name for a scheduler map.	To name a map, type the name—for example, be-scheduler-map .

Field	Function	Your Action
Scheduler Mapping	Enables you to associate a preconfigured scheduler with a forwarding class. After scheduler maps have been applied to	To associate a scheduler with a forwarding class, locate the forwarding class and select the scheduler in the box next to it.
	an interface, they affect the hardware queues and packet schedulers.	For example, for the best-effort forwarding class, select the configured scheduler from the list.

Table 26: Scheduler Maps Configuration Fields (Continued)

RELATED DOCUMENTATION

Defining CoS Schedulers (J-Web Procedure) 141	
Defining CoS Schedulers and Scheduler Maps (CLI Procedure) 138	
Example: Configuring CoS on EX Series Switches 15	
Monitoring CoS Scheduler Maps 147	

Monitoring CoS Scheduler Maps

IN THIS SECTION

- Purpose | **147**
- Action | **148**
- Meaning | 148

Purpose

NOTE: This topic applies only to the J-Web Application package.

Use the monitoring functionality to display assignments of CoS forwarding classes to schedulers.

Action

To monitor CoS scheduler maps in the J-Web interface, select **Monitor** > **Class of Service** > **Scheduler Maps**.

To monitor CoS scheduler maps in the CLI, enter the following CLI command:

show class-of-service scheduler-map

Meaning

Table 27 on page 148 summarizes key output fields for CoS scheduler maps.

Table 27: Summary of Key CoS Scheduler Maps Output Fields

Field	Values	Additional Information
Scheduler Map	Name of a scheduler map.	For details, click the plus sign (+).
Index	Index of a specific object—scheduler maps, schedulers, or drop profiles.	
Scheduler Name	Name of a scheduler.	
Forwarding Class	Forwarding classes this scheduler is assigned to.	
Transmit Rate	 Configured transmit rate of the scheduler in bits per second (bps). The rate value can be either of the following: A percentage—The scheduler receives the specified percentage of the total interface bandwidth. remainder—The scheduler receives the remaining bandwidth of the interface after bandwidth allocation to other schedulers. 	

Field	Values	Additional Information
Buffer Size	 Delay buffer size in the queue or the amount of transmit delay (in milliseconds). The buffer size can be either of the following: A percentage—The buffer is a percentage of the total buffer allocation. remainder—The buffer is sized according to what remains after other scheduler buffer allocations. 	
Priority	 Scheduling priority of a queue: strict-high—Packets in this queue are transmitted first. low—Packets in this queue are transmitted last. 	
Excess rate	The percentage of excess bandwidth traffic to share.	
Drop Profiles	Name and index of a drop profile that is assigned to a specific loss priority and protocol pair.	
Loss Priority	Packet loss priority corresponding to a drop profile.	
Protocol	Transport protocol corresponding to a drop profile.	
Drop Profile Name	Name of the drop profile.	

Table 27: Summary of Key CoS Scheduler Maps Output Fields (Continued)

Field	Values	Additional Information
Index	Index of a specific object—scheduler maps, schedulers, or drop profiles.	

Table 27: Summary of Key CoS Scheduler Maps Output Fields (Continued)

RELATED DOCUMENTATION

Defining CoS Schedulers and Scheduler Maps (CLI Procedure) | **138** Defining CoS Schedulers (J-Web Procedure) | **141** Example: Configuring CoS on EX Series Switches | **15**

Troubleshooting CoS Schedulers on a 40-port SFP+ Line Card in an EX8200 Switch

IN THIS SECTION

- Problem | 150
- Cause | **151**
- Solution | 151

Problem

Description

After you configure a scheduler map on an interface on the 40-port SFP+ line card, you notice one or both of the following:

• All packets are being dropped on a class-of-service queue configured on the interface.

• A message in the system log states that the interface is using the default scheduler map, not the scheduler map you configured. For example:

Sep 19 21:26:50 hostname cosd[907]: COSD_SCHED_MAP_GROUP_CONFLICT: Interface xe-5/0/15 cannot be bound to scheduler-map m1. It will be bound to default scheduler-map

Cause

The ports in a 40-port SFP+ line card are divided into eight groups, each group comprising five ports. The ports in a port group share 10 gigabits of bandwidth. Because the port groups share bandwidth, only one scheduler map can be active at a time in a port group. If you configure different scheduler maps for different interfaces in a port group, you do not receive an error when you commit the configuration. Instead, default scheduler map becomes the active scheduler map for all interfaces in the port group, and messages in the system log report that the default scheduler map is in use for the affected interfaces. If the default scheduler map does not define a queue, all traffic is dropped on that queue.

Solution

Check your CoS configuration for the interfaces in the port group. If you have different scheduler maps assigned to different interfaces in the port group:

- **1.** Delete the scheduler map configuration for all interfaces in the port group.
- 2. Determine the scheduler map that you want all interfaces in the port group to use.
- **3.** Assign that scheduler map to at least one interface in the port group. The remaining interfaces in the port group will adopt this scheduler map.

BEST PRACTICE: To prevent confusion and future configuration conflicts, explicitly assign the scheduler map to each interface in the port group.

4. After you commit the configuration, verify that the scheduler map is the active scheduler map for the interfaces in the port group by using the show class-of-service forwarding-table scheduler-map command.

RELATED DOCUMENTATION

40-port SFP+ Line Card in an EX8200 Switch

Defining CoS Schedulers and Scheduler Maps (CLI Procedure) | 138

Understanding CoS Queues on EX8200 Line Cards That Include Oversubscribed Ports | 11



Congestion Management, Tail Drop Profiles, Queue Shaping, and Explicit Congestion Notification (ECN)

Congestion Management | 154 Tail Drop Profiles | 164 Queue Shaping | 169 Explicit Congestion Notification (ECN) | 175

Congestion Management

IN THIS CHAPTER

- Understanding CoS Congestion Management | 154
- Configuring CoS Congestion Management (CLI Procedure) | 160

Understanding CoS Congestion Management

IN THIS SECTION

- Weighted Tail Drop Congestion Management | 154
- Weighted Random Early Detection Congestion Management | 155

A congestion in a network occurs because of various parameters and some packets must be dropped to avoid congestion and to facilitate easy flow of traffic in the network. On Juniper Networks EX Series Ethernet Switches, *class of service* (CoS) provides congestion management mechanisms for a switch to drop arriving packets based on certain parameters when a queue is full. Based on the EX Series switch that you are using, packets are dropped depending on the priority of a packet or on both priority and drop probability of a packet.

You can specify parameters at the [edit class-of-service drop-profiles] hierarchy level for dropping packets and reference the parameters in a scheduler configuration.

Weighted Tail Drop Congestion Management

A weighted tail drop (WTD) is a congestion management mechanism for packets to be dropped from the tail of the queue when the queue reaches a certain buffer capacity (that is, the fill level), and hence the name weighted tail drop. The packets that are dropped are based on priority and are those marked with a packet loss priority (PLP) of *high*. You can configure a WTD profile (a WTD mechanism) usually on edge devices in a network.

NOTE: A WTD profile is supported only on the Juniper Networks EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, and EX6200 Ethernet Switches.

When you configure a WTD profile, you are essentially setting the value for queue fullness. The queue fullness represents a percentage of the memory, known as delay-buffer bandwidth, that is used to store packets in relation to the total amount of memory that has been allocated for that specific queue. The delay-buffer bandwidth provides packet buffer space to absorb burst traffic up to the specified duration of delay. When the specified delay buffer becomes full, packets are dropped from the tail of the buffer.

By default, if you do not configure any drop profile, WTD profile is in effect and functions as the primary mechanism for managing congestion.

NOTE: The default WTD profile associated with the packets whose PLP is *low* cannot be modified. You can configure custom drop profile only for those packets whose PLP is *high*.

Weighted Random Early Detection Congestion Management

In a weighted random early detection (WRED) congestion management mechanism, random packets with a PLP of low or high are gradually dropped (based on drop probability) when the queue reaches a certain buffer capacity (that is, fill level).

NOTE: The WRED mechanism is supported only on Juniper Networks EX4300 standalone switches, EX4300 Virtual Chassis, EX4600 standalone switches, EX8200 standalone switches, EX8200 *Virtual Chassis*, and EX9200 standalone switches.

Following are the different implementations of WRED:

- Segmented Drop Profile
- Interpolated Drop Profile

From a high level, segmented drop profile is a stair-step-like drop profile, whereas interpolated drop profile is a smother (curve) drop profile. Figure 3 on page 156 and Figure 4 on page 158 show a graphical representation of segmented and interpolated drop profiles. Regardless of the implementation, a drop profile represents a graph where the x-axis represents the percentage of fill level (I) and the y-axis represents the percentage of drop probability (p). The origin (0,0) represents the drop profile in which the drop probability is 0 percent when the queue fullness is 0 percent, and the point (100,100) represents that the drop probability is 100 percent when the queue fullness is 100 percent. Although

the formation of graph lines in Figure 3 on page 156 and Figure 4 on page 158 is different, the application of the profile is the same. When a packet reaches the head of the queue, a random number between 0 and 100 is calculated. This random number is plotted against the drop profile graph using the current queue fullness of that particular queue. When the random number falls above the graph line, the packet is transmitted. When the number falls below the graph line, the packet is dropped from the network.

The following sections discuss the WRED drop profile implementations and parameters.

Segmented Drop Profile

In a segmented drop profile configuration, you can define multiple data points for fill level and drop probability. Figure 3 on page 156 shows a graphical representation of a segmented drop profile.

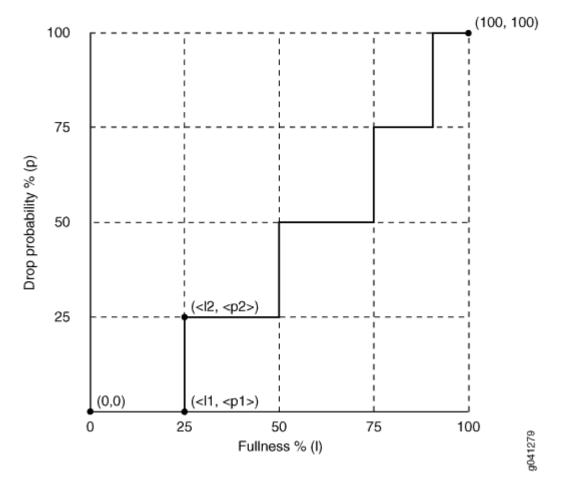


Figure 3: Graphical Representation of a Segmented Drop Profile

To create the profile's graph line, the software begins at the bottom-left corner of the graph, representing a 0 percent fill level and a 0 percent drop probability (that is the point (0,0)). The

configuration draws a line directly to the right until it reaches the first defined fill level (that is, 25 percent represented in the graph on the x-axis). The software then continues the line vertically until the first drop probability is reached (that is, 25 percent represented in the graph in the y-axis). This process is repeated for all of the defined fill levels and drop probabilities until the top-right corner of the graph is reached (that is point (100,100) in the graph).

Interpolated Drop Profile

An interpolated drop profile configuration forms a smoother graph line compared to the graph in a segmented drop profile configuration. In this method of congestion management also, a switch uses multiple drop profile values to drop incoming packets to reduce congestion in the output queue.

Following are interpolated drop profile configurations on EX Series switches:

Interpolated Drop Profile Configuration on EX Series Switches Except EX4300 Switches

An interpolated drop profile on all EX Series switches except EX4300 switches automatically generates 64 pairs of data points on the graph beginning at (0, 0) and ending at (100, 100). Along the way, the graph line intersects specific data points that you define for fullness and drop probability.

Figure 4 on page 158 shows a graphical representation of an interpolated drop profile.

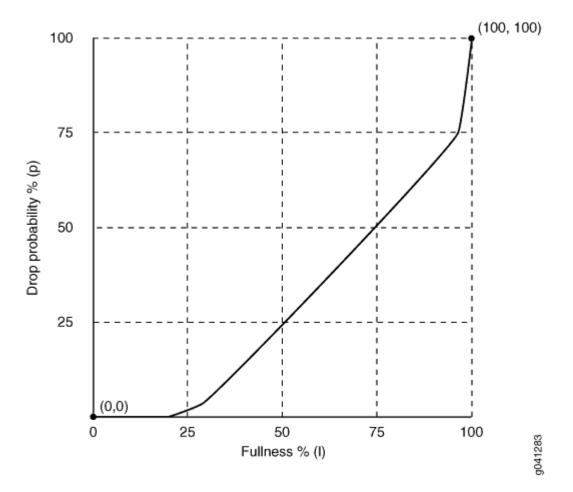


Figure 4: Graphical Representation of an Interpolated Drop Profile on EX Series Switches Except EX4300 Switches

Interpolated Drop Profile Configuration on EX4300 Switches

On EX4300 switches, you can set two queue fill levels and two drop probabilities in each drop profile. The two fill levels and the two drop probabilities create two pairs of values. The first fill level and the first drop probability create one value pair and the second fill level and the second drop probability create the second value pair.

NOTE: You can configure a maximum of 64 drop profiles on EX4300 switches.

The first fill level value specifies the percentage of queue fullness at which packets begin to drop, known as the drop start point. Until the queue reaches this level of fullness, no packets are dropped. The second fill level value specifies the percentage of queue fullness at which all packets are dropped, known as the drop end point.

The first drop probability value is always 0 (zero). This pairs with the drop start point and specifies that until the queue fullness level reaches the first fill level, no packets drop. When the queue fullness exceeds the drop start point, packets begin to drop until the queue exceeds the second fill level, when all packets drop. The second drop probability value, known as the maximum drop rate, specifies the likelihood of dropping packets when the queue fullness reaches the drop end point. As the queue fills from the drop start point to the drop end point, packets drop in a smooth, linear pattern (called an interpolated graph) as shown in Figure 5 on page 159. After the drop end point, all packets drop.

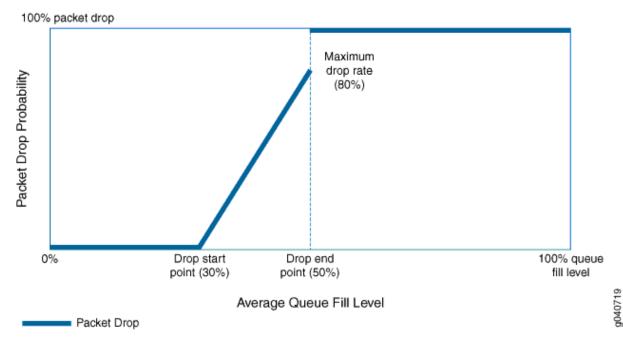


Figure 5: Tail-Drop Profile Packet Drop on EX4300 Switches

The thick line in Figure 5 on page 159 shows the packet drop characteristics for a sample tail drop profile. At the drop start point, the queue reaches a fill level of 30 percent. At the drop end point, the queue fill level reaches 50 percent, and the maximum drop rate is 80 percent.

No packets drop until the queue fill level reaches the drop start point of 30 percent. When the queue reaches the 30 percent fill level, packets begin to drop. As the queue fills, the percentage of packets dropped increases in a linear fashion. When the queue fills to the drop end point of 50 percent, the rate of packet drop has increased to the maximum drop rate of 80 percent. When the queue fill level exceeds the drop end point of 50 percent, all of the packets drop until the queue fill level drops below 50 percent.

Drop Profile Parameters

You can specify the following two values in drop profile configuration:

- Fill level—The queue fullness value, which represents a percentage of the memory used to store packets in relation to the total amount of memory allocated to the queue.
- Drop probability—The percentage value that corresponds to the likelihood that an individual packet is dropped.

RELATED DOCUMENTATION

Understanding Junos OS CoS Components for EX Series Switches | 8 Example: Configuring CoS on EX Series Switches | 15

Configuring CoS Congestion Management (CLI Procedure) | 160

Configuring CoS Congestion Management (CLI Procedure)

IN THIS SECTION

- Configuring a Weighted Tail Drop Profile | 161
- Configuring a Weighted Random Early Detection Drop Profile | 161

An effective congestion management mechanism is imperative to ensure smooth flow of traffic in a network and also to ensure minimum packet drops in the network. Class of service (CoS) provides congestion management methods that allow you to define parameters based on which packets can be dropped when the output queue is full. These parameters vary depending on the EX Series switch that you are using in a network.

You can specify parameters for dropping packets at the [edit class-of-service drop-profiles] hierarchy level and reference them in a scheduler configuration. The parameters that you can specify are fill-level and drop-probability. The first parameter defines the delay-buffer bandwidth, which provides packet buffer space to absorb burst traffic up to the specified duration of delay. When the specified delay buffer becomes full, packets with 100 percent drop probability are dropped from the head of the buffer. The second parameter represents a percentage value that correlates to the likelihood that an individual packet is dropped from the network.

Depending on the switch on which you are configuring a drop profile, you can configure either a weighted tail drop (WTD) profile or a weighted random early detection (WRED) profile.

This topic describes:

Configuring a Weighted Tail Drop Profile

A weighted tail drop (WTD) is a congestion management mechanism in which packets are dropped from the tail of the queue when the queue reaches a certain buffer capacity (that is, the fill level), and hence the name weighted tail drop. When that level is reached on EX2200, EX3200, or EX4200 Switches, packets marked with a packet loss priority (PLP) of high are prevented from entering the queue (that is, they are discarded).

To configure a WTD profile, create a drop profile name and assign a fill level:

[edit class-of-service drop-profiles]
user@switch# set profile-name fill-level percentage

Following is a sample WTD profile in which the fill level is set to 80 percent:

[edit class-of-service drop-profiles]
user@switch# set wtd-profile fill-level 80

Configuring a Weighted Random Early Detection Drop Profile

A WRED drop profile enables you to define multiple data points for fill level and drop probability so that packets are dropped at various levels of queue fullness, and for various drop probabilities. Unlike the WTD drop profile that can be defined only for packets with a PLP of high, WRED can be defined for packets with a PLP of high and also for packets with a PLP of low.

NOTE: The WRED drop profile is supported only on EX4300 standalone switches, EX4300 Virtual Chassis, EX8200 standalone switches and EX8200 Virtual Chassis.

WRED has two implementations: segmented and interpolated. From a high level, segmented is a stairstep-like drop profile, whereas interpolated is a smother (curve) drop profile. For a graphical representation of both these implementations, see "Understanding CoS Congestion Management" on page 154. Although the formation of graph lines is different for both these implementations, the application of the profile is the same. On EX Series switches except EX4300 switches, when a packet reaches the head of the queue, a random number between 0 and 100 is calculated. This random number is plotted against the drop profile using the current queue fullness of that particular queue. When the random number falls above the graph line, the packet is transmitted. When the number falls below the graph line, the packet is dropped from the network.

For information about congestion management on EX4300 switches, see "Understanding CoS Congestion Management" on page 154.

NOTE: On EX4300 switches, you cannot enable WRED on multidestination (multicast) queues. You can enable WRED only on unicast queues.

Following is the procedure to define a segmented and an interpolated drop profiles:

• To configure a segmented drop profile, specify multiple data points for fill level (1) and drop probability (p) as follows:

```
[edit class-of-service drop-profiles]
user@switch# set profile-name fill-level percentage-11 drop-probability percentage-p1
user@switch# set profile-name fill-level percentage-12 drop-probability percentage-p2
user@switch# set profile-name fill-level percentage-13 drop-probability percentage-p3
user@switch# set profile-name fill-level percentage-14 drop-probability percentage-p4
```

Following is a sample segmented drop profile:

[edit class-of-service drop-profiles] user@switch# set seg-prof fill-level 20 drop-probability 25 user@switch# set seg-prof fill-level 40 drop-probability 50 user@switch# set seg-prof fill-level 60 drop-probability 75 user@switch# set seg-prof fill-level 80 drop-probability 100

• To configure an interpolated drop profile on EX Series switches except EX4300 switches, specify multiple data points for fill level (1) and drop probability (p) using the interpolate statement as follows:

```
[edit class-of-service drop-profiles ]
user@switch# set profile-name interpolate fill-level percentage-11 drop-probability
percentage-11
user@switch# set profile-name interpolate fill-level percentage-12 drop-probability
percentage-12
user@switch# set profile-name interpolate fill-level percentage-13 drop-probability
percentage-p3
user@switch# set profile-name interpolate fill-level percentage-14 drop-probability
percentage-p4
```

Following is a sample interpolated drop profile:

[edit class-of-service drop-profiles] user@switch# set inter-prof interpolate fill-level 20 drop-probability 25 user@switch# set inter-prof interpolate fill-level 40 drop-probability 50 user@switch# set inter-prof interpolate fill-level 60 drop-probability 75 user@switch# set inter-prof interpolate fill-level 80 drop-probability 100

• To configure an interpolated drop profile EX4300 switches, specify two data points for fill level (1) and drop probability (p) by using the interpolate statement as follows:

[edit class-of-service drop-profiles]
user@switch# set profile-name interpolate fill-level percentage-11 fill-level percentage-12
drop-probability percentage-11 percentage-12

Following is a sample interpolated drop profile:

[edit class-of-service drop-profiles]
user@switch# set inter-prof interpolate fill-level 20 fill-level 80 drop-probability 25 dropprobability 100

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches | 15 Understanding CoS Congestion Management | 154

Tail Drop Profiles

IN THIS CHAPTER

- Understanding CoS Tail Drop Profiles | 164
- Configuring CoS Tail Drop Profiles (CLI Procedure) | 165
- Defining CoS Drop Profiles (J-Web Procedure) | 165
- Monitoring CoS Drop Profiles | 167

Understanding CoS Tail Drop Profiles

Tail drop profile is a congestion management mechanism that allows switch to drop arriving packets when queue buffers become full or begin to overflow.

Tail drop profiles define the meanings of the loss priorities. When you configure tail drop profiles you are essentially setting the value for queue fullness. The queue fullness represents a percentage of the memory used to store packets in relation to the total amount that has been allocated for that specific queue.

The queue fullness defines the delay-buffer bandwidth, which provides packet buffer space to absorb burst traffic up to the specified duration of delay. Once the specified delay buffer becomes full, packets with 100 percent drop probability are dropped from the tail of the buffer.

You specify drop probabilities in the drop profile section of the CoS configuration hierarchy and reference them in each scheduler configuration.

By default, if you do not configure any drop profile, tail drop profile is in effect and functions as the primary mechanism for managing congestion. In the default tail drop profile, when the fill level is 0 percent, the drop probability is 0 percent. When the fill level is 100 percent, the drop probability is 100 percent.

NOTE: The default drop profile associated with the packets whose loss priority is low cannot be modified. You can configure custom drop profile only for those packets whose loss priority is high.

RELATED DOCUMENTATION

Understanding Junos OS CoS Components for EX Series Switches | 8 Example: Configuring CoS on EX Series Switches | 15

Configuring CoS Tail Drop Profiles (CLI Procedure) | 165

Configuring CoS Tail Drop Profiles (CLI Procedure)

Tail drop is a simple and effective traffic congestion avoidance mechanism. When you apply this mechanism to manage congestion, packets are dropped when the output queue is full.

To configure CoS tail-drop profiles, create a drop profile name (be-dp) and assign a fill level (25):

[edit class-of-service drop-profiles]
user@switch# set be-dp fill-level 25

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches | **15** Understanding CoS Tail Drop Profiles | **164**

Defining CoS Drop Profiles (J-Web Procedure)

NOTE: This topic applies only to the J-Web Application package.

You can use the J-Web interface to define CoS drop profiles on EX4500 and EX8200 switches.

To configure CoS drop profiles:

1. Select Configure > Class of Service > Drop Profile.

NOTE: After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See Using the Commit Options to Commit Configuration Changes (J-Web Procedure) for details about all commit options.

- **2.** Select one of the following options:
 - Add—Adds a drop profile. Enter information into the drop profiles page as described in Table 28 on page 166.
 - Edit—Modifies an existing drop file. Enter information into the drop profiles page as described in Table 28 on page 166.
 - Delete–Deletes an existing drop profile.

Table 28: Drop Profiles Configuration parameters

Field	Function	Your Action
Drop Profile Name	Specifies the name for a drop profile.	Type the name.
Drop profile graph	Specifies the drop profile graph type	Select one: Segmented or Interpolated.
Drop profile values	Specifies values for the following two parameters of the drop profile: the queue fill level and the drop probability. The queue fill level represents a percentage of the memory used to store packets in relation to the total amount that has been allocated for that specific queue. The drop probability is a percentage value that correlates to the likelihood that an individual packet is dropped from the network.	 To add new values: 1. Click Add. 2. Enter the fill level. 3. Enter the drop probability. 4. Click OK. To edit an existing value, click Edit and modify the fill level and drop probability. To delete a value, select it and click Delete.

RELATED DOCUMENTATION

Monitoring CoS Drop Profiles | 167

Example: Configuring CoS on EX Series Switches | 15

Monitoring CoS Drop Profiles

IN THIS SECTION

- Purpose | 167
- Action | 167
- Meaning | 167

Purpose

NOTE: This topic applies only to the J-Web Application package.

Use the monitoring functionality to view data point information for each CoS random early detection (RED) drop profile on the EX8200 switch.

Action

To monitor CoS RED drop profiles in the J-Web interface, select **Monitor > Class of Service > RED Drop Profiles**.

To monitor CoS RED drop profiles in the CLI, enter the following CLI command:

show class-of-service drop-profile

Meaning

Table 29 on page 168 summarizes the key output fields for CoS RED drop profiles.

Field	Values	Additional Information
RED Drop Profile Name	Name of the RED drop profile. A drop profile consists of pairs of values between 0 and 100, one for queue buffer fill level and the other for drop probability, that determine the relationship between a buffer's fullness and the likelihood it will drop packets.	To display profile values, click the plus sign (+).
Graph RED Profile	Links to a graph of a RED curve that the system uses to determine the drop probability based on queue buffer fullness.	The x axis represents the queue buffer fill level, and the y axis represents the drop probability.
Туре	 Type of a specific drop profile: interpolated—The two coordinates (x and y) of the graph are interpolated to produce a smooth profile. segmented—The two coordinates (x and y) of the graph are represented by line fragments to produce a segmented profile. 	
Index	Internal index of this drop profile.	
Fill Level	Percentage fullness of a buffer queue. This value is the x coordinate of the RED drop profile graph.	
Drop Probability	Drop probability of a packet corresponding to a specific queue buffer fill level. This value is the y coordinate of the RED drop profile graph.	

Table 29: Summary of the Key Output Fields for CoS Red Drop Profiles

RELATED DOCUMENTATION

Defining CoS Drop Profiles (J-Web Procedure) | **165** Example: Configuring CoS on EX Series Switches | **15**

Queue Shaping

IN THIS CHAPTER

- Understanding Port Shaping and Queue Shaping for CoS | 169
- Configuring Shaping for CoS (CLI Procedure) | 171
- Applying a Shaping Rate to Physical Interfaces Overview | 173
- Configuring the Shaping Rate for Physical Interfaces | 173

Understanding Port Shaping and Queue Shaping for CoS

IN THIS SECTION

- Port Shaping | 169
- Queue Shaping | **170**

When the amount of traffic on a switch's network exceeds the maximum bandwidth, packets are lost because of congestion in the network. The excess traffic in the network must be handled carefully to ensure minimum or no data loss in the network. A class-of-service (CoS) configuration includes several parameters that classify traffic into different queues and also define packet loss priorities (PLPs) to ensure smooth transmission of data in the network. You can use these configuration parameters to control or shape traffic for a specific port on a switch or for a specific CoS queue. While port shaping defines the maximum bandwidth allocated to an interface, queue shaping defines a limit on excess-bandwidth usage for each queue.

Port Shaping

Port shaping enables you to shape the aggregate traffic through a port or channel to a rate that is less than the line rate. You can configure interfaces to shape traffic based on the rate-limited bandwidth of the total interface bandwidth. This allows you to shape the output of the interface so that the interface transmits less traffic than it is capable of transmitting. For port shaping, you specify the shaping rate as the peak rate at which traffic can pass through the interface. You specify the rate as a value in bits per second (bps) either as a decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000) and the value can range from 1000 through 160,000,000,000 bps.

By default, shaping is not configured on an interface. If you do not configure a shaping rate on an interface, the default shaping rate is 100 percent, which is the equivalent to no shaping configured for that interface.

On modern EX series switches that support Enhanced Layer 2 Software (ELS), when you configure a shaping rate on an ae interface, the traffic is equally divided among the members of the ae interface. For example, consider an interface, ae0, that consists of three interfaces: ge-0/0/0, ge-0/0/1, and ge-0/0/2. If you configure a shaping rate of X Mpbs on ae0, traffic up to the rate of X/3 Mpbs flows through each of the three interfaces. This is known as scale mode.

NOTE: On older EX swithes that don't support ELS, when you configure a shaping rate on an aggregated Ethernet (ae) interface, all members of the ae interface are shaped at the configured shaping rate. For example, consider an interface ae0 that consists of three interfaces: ge-0/0/0, ge-0/0/1, and ge-0/0/2. If you configure a shaping rate of X Mpbs on ae0, traffic up to the rate of X Mpbs flows through each of the three interfaces. Therefore, the total traffic flowing through ae0 can be at the rate of 3X Mbps. This is replicate mode.

Queue Shaping

Queue shaping throttles the rate at which queues transmit packets. For example, using queue shaping, you can rate-limit a strict-priority queue so that the strict-priority queue does not lock out (or starve) low-priority queues. Similarly, for any queue, you can configure queue shaping.

You can specify queue shaping as the maximum rate at which traffic can pass through the queue or as a percentage of the available bandwidth. On EX Series switches except EX4300 switches, you can specify the rate as a value between 3200 and 160,000,000 bps and the percentage as a value from 0 to 100 percent. On EX4300 switches, you can specify the rate as a value between 8000 and 160,000,000,000 bps and the percentage as a value from 0 to 100 percent.

RELATED DOCUMENTATION

Understanding CoS Schedulers | 128

Defining CoS Schedulers and Scheduler Maps (CLI Procedure) | 138

Configuring Shaping for CoS (CLI Procedure)

IN THIS SECTION

- Configuring Port Shaping for CoS on an EX Series Switch | 171
- Configuring Queue Shaping for CoS on an EX Series Switch | 171

Port shaping and queue shaping enable you to limit traffic on an interface or queue, respectively, so that you can control the amount of traffic passing through the interface or the queue. Port shaping enables you to shape the aggregate traffic through an interface to a rate that is less than the line rate for that interface. When you configure port shaping on an interface, you are essentially specifying a value that indicates the maximum amount of traffic that can pass through the interface. This value must be less than the maximum bandwidth for that interface. Queue shaping enables you to throttle the rate at which a queue transmits packets. When you configure queue shaping, you can specify either as the maximum rate at which traffic can pass through the queue or as a percentage of the available bandwidth.

Configuring Port Shaping for CoS on an EX Series Switch

You can configure port shaping on network interfaces, aggregated Ethernet interfaces (also known as link aggregation groups (LAGs)), and loopback interfaces.

To configure port shaping on an interface:

- 1. Ensure that the interface on which you want to configure port shaping is up and running.
- 2. Assign a shaping rate for the interface:

[edit]

user@switch# set class-of-service interfaces interface-name shaping-rate value

The value indicates the maximum amount of traffic (in bps) that can pass through the interface. This value must be less than the maximum bandwidth for that interface.

Configuring Queue Shaping for CoS on an EX Series Switch

Queue shaping enables you to restrict the rate at which queues transmit traffic. You can configure queue shaping on any queue supported by CoS on an EX Series switch that supports up to eight output queues and 16 forwarding classes. Forwarding classes can be thought of as output queues. In effect, the result of classifying packets into forwarding classes is the identification of an output queue for a particular

packet. For a classifier to assign an output queue to a packet, it must associate the packet with one of the forwarding classes discussed in "Understanding CoS Forwarding Classes" on page 110.

To configure queue shaping:

- 1. Ensure that the interface on which you want to configure queue shaping is up and running.
- **2.** Configure queue shaping:
 - a. Define a scheduler and assign a shaping rate to the scheduler:

[edit]

user@switch# set class-of-service schedulers scheduler-name shaping-rate (rate | percent percentage)

You can assign a *rate* (a value in bits per second (bps)) or a percentage value for shaping-rate.

b. Define a scheduler map and assign a forwarding class and scheduler (that you defined in the previous step) to the scheduler map:

[edit]

user@switch# setset class-of-service scheduler-maps scheduler-map-name forwarding-class
class-name scheduler scheduler-name

c. Assign the scheduler map to an interface:

[edit]

user@switch# set class-of-service interfaces interface-name scheduler-map scheduler-mapname-map-name

RELATED DOCUMENTATION

Understanding Port Shaping and Queue Shaping for CoS | 169

Understanding CoS Schedulers | 128

Example: Configuring CoS on EX Series Switches | 15

Applying a Shaping Rate to Physical Interfaces Overview

On T4000 routers with Type 5 FPCs and on EX Series switches, you can configure physical interfaces to shape traffic based on the rate-limited bandwidth of the total interface bandwidth. This allows you to shape the output of the physical interface, so that the interface transmits less traffic than it is physically capable of carrying.

If you do not configure a shaping rate on the physical interface, the default physical interface bandwidth is based on the channel bandwidth and the time slot allocation.

In general, the physical interface speed is the basis for calculating the various queue parameters for a physical interface such as delay buffer size, weighted round- robin (WRR) weight, drop profile, and so forth. However, when you apply a shaping rate by including the shaping-rate statement, the shaping rate on that physical interface becomes the basis for calculating all the queue parameters for that physical interface.

On T4000 routers with Type 5 FPCs, the shaping rate value for the physical interface must be a minimum of 292 Kbps. The maximum value of shaping rate is limited by the maximum transmission rate of the interface.

RELATED DOCUMENTATION

Configuring the Shaping Rate for Physical Interfaces

Configuring the Shaping Rate for Physical Interfaces

To configure the shaping rate on the physical interface, either include the shaping-rate statement at the [edit class-of-service interface-name] hierarchy level or include the output-traffic-control-profile statement at the [edit class-of-service interfaces *interface-name*] hierarchy level.

You can specify a peak bandwidth rate in bps, 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). For physical interfaces, the range is from 1000 through 6,400,000,000 bps.

For physical interfaces on T4000 routers with Type 5 FPCs, the shaping rate value for the physical interface must be a minimum of 292 Kbps. The maximum value of shaping-rate is limited by the maximum transmission rate of the interface.

The following are two example configurations for applying a shaping rate of 5 Gbps on a T4000 12x10 Gbps physical interface (xe-4/0/0):

Applying a shaping rate at the [edit class-of-service interfaces *interface-name*] hierarchy:

```
[edit class-of-service]
interfaces {
    xe-4/0/0 {
        shaping-rate 5g;
    }
}
```

Applying a shaping rate using traffic-control-profiles:

```
[edit class-of-service]
traffic-control-profiles {
    output {
        shaping-rate 5g;
    }
}
interfaces {
    xe-4/0/0 {
        output-traffic-control-profile output;
    }
}
```

To view the results of your configuration, issue the following show commands:

- show class-of-service interface interface-name
- show interfaces interface-name extensive

RELATED DOCUMENTATION

Applying a Shaping Rate to Physical Interfaces Overview

Explicit Congestion Notification (ECN)

IN THIS CHAPTER

- Understanding CoS Explicit Congestion Notification | 175
- Example: Configuring ECN | **185**

Understanding CoS Explicit Congestion Notification

IN THIS SECTION

- How ECN Works | 176
- WRED Drop Profile Control of ECN Thresholds | 181
- Support, Limitations, and Notes | 184

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 ECNenabled receiver. 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.

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, switches 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 switches 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 switches 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 switch might mark a packet as experiencing congestion. Whether or not a switch 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 switch, the switch 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 30 on page 177:

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

Table 30: ECN Bit Codes

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 ECNcapable by setting the DiffServ ECN field to ECT(1) (01) or ECT(0) (10). Every intermediate switch between the endpoints must have ECN enabled or it does not work.

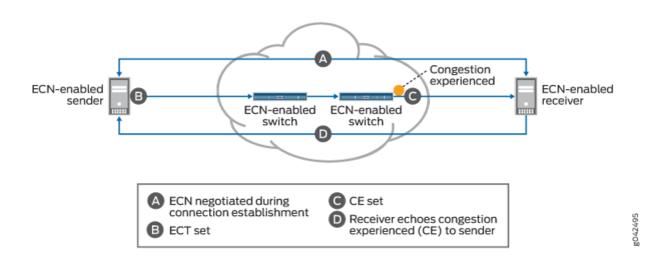
When a packet traverses a switch and experiences congestion at an output queue that uses the WRED packet drop mechanism, the switch marks the packet as experiencing congestion by setting the DiffServ ECN field to CE (11). Instead of dropping the packet (as with TCP congestion notification), the switch 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 might be dropped. See "WRED Drop Profile Control of ECN Thresholds" on page 181 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 6 on page 178 summarizes how ECN works to mitigate network congestion:

Figure 6: 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.
- **2.** After successful negotiation of ECN capability, the ECN-capable sender sends IP packets with the ECT field set to the receiver.

NOTE: All of the intermediate devices in the path between the sender and the receiver must be ECN-enabled.

3. If the WRED algorithm on a switch egress queue determines that the queue is experiencing congestion and the packet is drop eligible, the switch 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 switch), the switch forwards the packet with CE marked.

If there is no congestion at the switch egress queue, the switch 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.

On QFX5210, QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, packets that are not marked as ECN-capable (ECT, 00) are treated according to the WRED drop profile configuration and might be dropped during periods of congestion.

On QFX10000 switches, the switch uses the tail-drop algorithm to drop packets that are marked ECT (00) during periods of congestion. (When a queue fills to its maximum level of fullness, tail-drop simply drops all subsequently arriving packets until there is space in the queue to buffer more packets. All non-ECN-capable packets are treated the same.)

- **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 31 on page 179 summarizes the behavior of traffic 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
Non-ECT (00)	Does not matter	Drop (QFX5210, QFX5200, QFX5100, EX4600, QFX3500, QFX3600, QFabric systems). Tail drop occurs when queue reaches maximum fullness because no WRED drop probability is applied (QFX10000 switches).	No ECN bits marked

Table 31: 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
ECT (10 or 01)	ECN disabled	Drop	Packet dropped—no ECN bits marked
ECT (10 or 01)	ECN enabled	Do not drop. Mark packet as experiencing congestion (CE, bits 11).	Packet marked ECT (11) to indicate congestion
CE (11)	ECN disabled	Drop	Packet dropped—no ECN bits marked
CE (11)	ECN enabled	Do not drop. Packet is already marked as experiencing congestion, forward packet without changing the ECN marking.	Packet marked ECT (11) to indicate congestion

 Table 31: Traffic Behavior on ECN-Enabled Queues (Continued)

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 switches. Instead, PFC controls flows

between two PFC-enabled peer devices (for example, switches) 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.

PFC works this way: if a switch output queue fills to a certain threshold, the switch sends a PFC pause message to the connected peer device that is transmitting data. The pause message tells the transmitting switch to pause transmission of the flow. When the congestion clears, the switch sends another PFC message to tell the connected peer to resume transmission. (If the output queue of the transmitting switch also reaches a certain threshold, that switch can in turn send a PFC pause message to the connected peer that is transmitting to it. In this way, PFC can propagate a transmission pause back through the network.)

See Understanding CoS Flow Control (Ethernet PAUSE and PFC) for more information. For QFX5100 and EX4600 switches only, you can also refer to Understanding PFC Functionality Across Layer 3 Interfaces.

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 switch 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. Strict-high priority queues receive all of the port bandwidth they require up to the configured maximum bandwidth limit (scheduler transmit-rate on QFX10000 switches, and

shaping-rate on QFX5210, QFX5200, QFX5100, QFX3500, QFX3600, and EX4600 switches, and QFabric systems).

Different switches support different amounts of fill level/drop probability pairs in drop profiles. For example, QFX10000 switches support 32 fill level/drop probability pairs, so there can be as many as 30 intermediate fill level/drop probability pairs between the drop start and drop endpoints. QFX5210, QFX5200, QFX5100, QFX3500, QFX3600, and EX4600 switches, and QFabric systems support two fill level/drop probability pairs—by definition, the two pairs you configure on these switches are the drop start and drop end points.

NOTE: Do not configure the last fill level as 100 percent.

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. (At this point on QFX5210, QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, all non-ECN packets are dropped.) 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 switches 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.

Starting in Release 15.1, you can configure a WRED packet drop profile and apply it to an output queue on switches 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 switches 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

NOTE: On QFX10000 switches, when you enable a queue for ECN and apply a WRED drop profile to the queue, the WRED drop profile only sets the thresholds for marking ECN traffic as experiencing congestion (CE, 11). On ECN-enabled queues, the WRED drop profile does not set drop thresholds for non-ECT (00) traffic (traffic that is not ECN-capable). Instead, the switch uses the tail-drop algorithm on traffic is that is marked non-ECT on ECN-enabled queues during periods of congestion.

To apply a WRED drop profile to non-ECT traffic, configure a multifield (MF) classifier to assign non-ECT traffic to a different output queue that is not ECN-enabled, and then apply the WRED drop profile to that queue.

Release History Table

Release	Description
15.1	Starting in Release 15.1, you can configure a WRED packet drop profile and apply it to an output queue on switches that support port scheduling (ETS hierarchical scheduling is either not supported or not used).

RELATED DOCUMENTATION

Example: Configuring ECN

Example: Configuring ECN

IN THIS SECTION

- Requirements | 185
- Overview | 185
- Configuration | 188
- Verification | 191

This example shows how to enable explicit congestion notification (ECN) on an output queue.

Requirements

This example uses the following hardware and software components:

- One switch.
- Junos OS Release 13.2X51-D25 or later for the QFX Series or Junos OS Release 14.1X53-D20 for the OCX Series

Overview

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 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

A weighted random early detection (WRED) packet drop profile must be applied to the output queues on which ECN is enabled. ECN uses the WRED drop profile thresholds to mark packets when the output queue experiences congestion.

ECN reduces packet loss by forwarding ECN-capable packets during periods of network congestion instead of dropping those packets. (TCP notifies the network about congestion by dropping packets.) During periods of congestion, ECN marks ECN-capable packets that egress from congested queues. When the receiver receives an ECN packet that is marked as experiencing congestion, the receiver echoes the congestion state back to the sender. The sender then reduces its transmission rate to clear the congestion.

ECN is disabled by default. You can enable ECN on best-effort traffic. ECN should not be enabled on lossless traffic queues, which uses priority-based flow control (PFC) for congestion notification, and ECN should not be enabled on strict-high priority traffic queues.

To enable ECN on an output queue, you not only need to enable ECN in the queue scheduler, you also need to:

- Configure a WRED packet drop profile.
- Configure a queue scheduler that includes the WRED drop profile and enables ECN. (This example shows only ECN and drop profile configuration; you can also configure bandwidth, priority, and buffer settings in a scheduler.)
- Map the queue scheduler to a forwarding class (output queue) in a scheduler map.
- Starting in Junos OS 15.1, enhanced transmission selection (ETS) hierarchical scheduling is supported. If you are using enhanced transmission selection (ETS) hierarchical scheduling, add the forwarding class to a forwarding class set (priority group).
- If you are using ETS, associate the queue scheduler map with a traffic control profile (priority group scheduler for hierarchical scheduling).
- If you are using ETS, apply the traffic control profile and the forwarding class set to an interface. On that interface, the output queue uses the scheduler mapped to the forwarding class, as specified by the scheduler map attached to the traffic control profile. This enables ECN on the queue and applies the WRED drop profile to the queue.

If you are using port scheduling, apply the scheduler map to an interface. On that interface, the output queue uses the scheduler mapped to the forwarding class in the scheduler map, which enables ECN on the queue and applies the WRED drop profile to the queue.

Table 32 on page 186 shows the configuration components for this example.

Table 32: Components of the ECN Configuration Example

Component	Settings
Hardware	QFX Series switch

Component	Settings
Drop profile (with two fill level/ drop probability pairs)	Name: be-dp Drop start fill level: 30 percent Drop end fill level: 75 percent Drop probability at drop start (minimum drop rate): 0 percent Drop probability at drop end (maximum drop rate): 80 percent
Scheduler	Name: be-sched ECN: enabled Drop profile: be-dp Transmit rate: 25% Buffer size: 25% Priority: low
Scheduler map	Name: be-map Forwarding class: best-effort Scheduler: be-sched NOTE : By default, the best-effort forwarding class is mapped to output queue 0.
Forwarding class set (ETS only)	Name: be-pg Forwarding class: best-effort (queue 0)
Traffic control profile (ETS only)	Name: be-tcp Scheduler map: be-map
Interface (ETS only)	Name: xe-0/0/20 Forwarding class set: be-pg (Output) traffic control profile: be-tcp
Interface (port scheduling only)	Name: xe-0/0/20

Table 32: Components of the ECN Configuration Example (Continued)

NOTE: Only switches that support ETS hierarchical scheduling support forwarding class set and traffic control profile configuration. Direct port scheduling does not use the hierarchical scheduling structure.

NOTE: On QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, the WRED drop profile also controls packet drop behavior for traffic that is not ECN-capable (packets marked non-ECT, ECN bit code 00).

On QFX10000 switches, when ECN is enabled on a queue, the WRED drop profile only sets the ECN thresholds, it does not control packet drop on non-ECN packets. On ECN-enabled queues, QFX10000 switches use the tail-drop algorithm on non-ECN packets during periods of congestion. If you do not enable ECN, then the queue uses the WRED packet drop mechanism.

Configuration

IN THIS SECTION

- CLI Quick Configuration | 188
- Configuring ECN | 189

CLI Quick Configuration

To quickly configure the drop profile, scheduler with ECN enabled, and to map the scheduler to an output queue on an interface, 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.

ETS Quick Configuration

```
[edit class-of-service]
set drop-profile be-dp interpolate fill-level 30 fill-level 75 drop-probability 0 drop-
probability 80
set schedulers be-sched explicit-congestion-notification
set schedulers be-sched drop-profile-map loss-priority low protocol any drop-profile be-dp
set schedulers be-sched transmit-rate percent 25
set schedulers be-sched buffer-size percent 25
```

```
set schedulers be-sched priority low
set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
set forwarding-class-sets be-pg class best-effort
set traffic-control-profiles be-tcp scheduler-map be-map
set interfaces xe-0/0/20 forwarding-class-set be-pg output-traffic-control-profile be-tcp
```

Port Scheduling Quick Configuration (QFX10000 Switches)

[edit class-of-service]
set drop-profile be-dp interpolate fill-level 30 fill-level 75 drop-probability 0 dropprobability 80
set schedulers be-sched explicit-congestion-notification
set schedulers be-sched drop-profile-map loss-priority low protocol any drop-profile be-dp
set schedulers be-sched transmit-rate percent 25
set schedulers be-sched buffer-size percent 25
set schedulers be-sched priority low
set schedulers be-sched priority low
set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
set interfaces xe-0/0/20 scheduler-map be-map

Configuring ECN

Step-by-Step Procedure

To configure ECN:

1. Configure the WRED packet drop profile be-dp. This example uses a drop start point of 30 percent, a drop end point of 75 percent, a minimum drop rate of 0 percent, and a maximum drop rate of 80 percent:

[edit class-of-service]
user@switch# set drop-profile be-dp interpolate fill-level 30 fill-level 75 drop-probability
0 drop-probability 80

2. Create the scheduler be-sched with ECN enabled and associate the drop profile be-dp with the scheduler:

```
[edit class-of-service]
user@switch# set schedulers be-sched explicit-congestion-notification
user@switch# set schedulers be-sched drop-profile-map loss-priority low protocol any drop-
```

profile be-dp

user@switch# set be-sched transmit-rate percent 25
user be-sched transmit-rate percent 25
user@switch# set be-sched buffer-size percent 25
user@switch# set be-sched puffer-size percent 25
user@switch# set be-sched priority low

3. Map the scheduler be-sched to the best-effort forwarding class (output queue 0) using scheduler map be-map:

[edit class-of-service]
user@switch# set scheduler-maps be-map forwarding-class best-effort scheduler be-sched

4. If you are using ETS, add the forwarding class best-effort to the forwarding class set be-pg; if you are using direct port scheduling, skip this step:

[edit class-of-service]
user@switch# set forwarding-class-sets be-pg class best-effort

5. If you are using ETS, associate the scheduler map be-mapwith the traffic control profile be-tcp; if you are using direct port scheduling, skip this step:

```
[edit class-of-service]
user@switch# set traffic-control-profiles be-tcp scheduler-map be-map
```

6. If you are using ETS, associate the traffic control profile be-tcp and the forwarding class set be-pg with the interface on which you want to enable ECN on the best-effort queue:

[edit class-of-service]
user@switch# set interfaces xe-0/0/20 forwarding-class-set be-pg output-traffic-controlprofile be-tcp

If you are using direct port scheduling, associate the scheduler map be-map with the interface on which you want to enable ECN on the best-effort queue:

[edit class-of-service]
user@switch# set interfaces xe-0/0/20 scheduler-map be-map

Verification

IN THIS SECTION

Verifying That ECN Is Enabled | 191

Verifying That ECN Is Enabled

Purpose

Verify that ECN is enabled in the scheduler be-sched by showing the configuration for the scheduler map be-map.

Action

Display the scheduler map configuration using the operational mode command show class-of-service scheduler-map be-map:

```
user@switch> show class-of-service scheduler-map be-map
Scheduler map: be-map, Index: 12240
 Scheduler:be-sched, Forwarding class: best-effort, Index: 115
   Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent,
   Buffer Limit: none, Priority: low
   Excess Priority: unspecified, Explicit Congestion Notification: enable
   Drop profiles:
     Loss priority Protocol
                                 Index
                                          Name
                                  3312
     Low
                     any
                                          be-dp
     Medium-high
                                          <default-drop-profile>
                                  1
                     any
     High
                                  1
                                          <default-drop-profile>
                     any
```

Meaning

The show class-of-service scheduler-map operational command shows the configuration of the scheduler associated with the scheduler map and the forwarding class mapped to that scheduler. The output shows that:

- The scheduler associated with the scheduler map is be-sched.
- The scheduler map applies to the forwarding class best-effort (output queue 0).
- The scheduler be-sched has a transmit rate of 25 percent, a queue buffer size of 25 percent, and a drop priority of low.
- Explicit congestion notification state is enable.
- The WRED drop profile used for low drop priority traffic is be-dp.

Release History Table

Release	Description
15.1	Starting in Junos OS 15.1, enhanced transmission selection (ETS) hierarchical scheduling is supported.

RELATED DOCUMENTATION

Understanding CoS Explicit Congestion Notification



CoS on Overlay Networks

CoS on MPLS Networks | 194

CoS on EVPN VXLANs | 225

CoS on MPLS Networks

IN THIS CHAPTER

- Understanding Using CoS with MPLS Networks on EX Series Switches | 194
- Example: Combining CoS with MPLS on EX Series Switches | 199
- Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS | 217
- Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect | 219
- Configuring CoS on Provider Switches of an MPLS Network | 222
- Configuring MPLS on EX8200 and EX4500 Provider Switches | 223

Understanding Using CoS with MPLS Networks on EX Series Switches

IN THIS SECTION

- EXP Classifiers and EXP rewrite Rules | 195
- Guidelines for Using CoS Classifiers on CCCs | 195
- Using CoS Classifiers with IP over MPLS | **196**
- Setting CoS Bits in an MPLS Header | **196**
- EXP Rewrite Rules | 198
- Policer | **198**
- Schedulers | 199

You can use *class of service* (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion. See EX Series Switch Software Features Overview for a complete list of the Junos OS MPLS features that are supported on specific EX Series switches.

Juniper Networks EX Series Ethernet Switches support Differentiated Service Code Point (DSCP) or IP precedence and IEEE 802.1p CoS classifiers on the customer-edge interfaces of the ingress provider

edge (PE) switch. DSCP or IP precedence classifiers are used for Layer 3 packets. IEEE 802.1p is used for Layer 2 packets.

When a packet enters a customer-edge interface of the ingress PE switch, the switch associates the packet with a particular CoS servicing level before putting the packet onto the label-switched path (LSP). The switches within the LSP utilize the CoS value set at the ingress PE switch. The CoS value that was embedded in the classifier is translated and encoded in the MPLS header by means of the EXP or experimental bits. EX Series switches enable a default EXP classifier and a default EXP rewrite rule. For more information about EXP classifiers and EXP *rewrite rules*, see EXP Classifiers and EXP rewrite Rules.

This topic includes:

EXP Classifiers and EXP rewrite Rules

EX Series switches enable a default EXP classifier and a default EXP rewrite rule. You can configure a custom EXP classifier and a custom EXP rewrite rule if you prefer. However, the switch supports only one type of EXP classifier (default or custom) and only one EXP rewrite rule (default or custom).

You do not bind the EXP classifier or the EXP rewrite rule to individual interfaces. The switch automatically and implicitly applies the default or the custom EXP classifier and the default or the custom EXP rewrite rule to the appropriate MPLS-enabled interfaces. Because rewrite rules affect only egress interfaces, the switch applies the EXP rewrite rule only to those MPLS interfaces that are transmitting MPLS packets (not to the MPLS interfaces that are receiving the packets).

After traversing the MPLS tunnel, the traffic flows out from the egress provider edge (PE) switch. Before the traffic leaves the egress interface, the egress PE switch copies the EXP bits from the MPLS header to the most significant bits in the original IP packet--- that is, to the IP precedence bits. Note that this is the default behavior only on Juniper Networks EX8200 Ethernet Switches (standalone or *Virtual Chassis*) that are configured for MPLS.

Guidelines for Using CoS Classifiers on CCCs

When you are configuring CoS for MPLS over circuit cross-connect (CCC), there are some additional guidelines, as follows:

- You *must* explicitly bind a CoS classifier to the CCC interface on the ingress PE switch.
- You *must* use the same DSCP, IP precedence, or IEEE 802.1p classifier on CCC interfaces. However, if the CCC interfaces are on the same switch, you cannot configure both a DSCP and an IP precedence classifier on these interfaces. Thus, if you configure one CCC interface to use a DSCP classifier DSCP1, you cannot configure another CCC interface to use another DSCP classifier DSCP2. All the CCC interfaces on the switch must use the same DSCP (or IP precedence) classifier and the same IEEE 802.1p classifier.

- You *cannot* configure one CCC interface to use a DSCP classifier and another CCC interface to use an IP precedence classifier, because these classifier types overlap.
- You *can* configure one CCC interface to use a DSCP classifier and another CCC interface to use IEEE 802.1p classifier.
- You *can* configure one CCC interface to use both a DSCP and an IEEE 802.1p classifier. If you configure a CCC interface to use both these classifiers, the DSCP classifier is used for routing Layer 3 packets and the IEEE 802.1p classifier is used for routing Layer 2 packets.
- You *can* configure one CCC interface to use both an IP precedence and an IEEE 802.1p classifier. If you configure a CCC interface to use both these classifiers, the IP precedence classifier is used for routing Layer 3 packets and the IEEE 802.1p classifier is used for routing Layer 2 packets.

NOTE: These guidelines are not applicable to Juniper Networks EX8200 Ethernet Switches (standalone or Virtual Chassis).

You can define multiple DSCP, IP precedence, and IEEE 802.1p classifiers for the non-CCC interfaces on a switch.

Using CoS Classifiers with IP over MPLS

When you are configuring CoS for IP over MPLS, the customer-edge interface uses the CoS configuration for the switch as the default. You do not have to bind a classifier to the customer-edge interface in this case. There are no restrictions on using multiple DSCP, IP precedence, and IEEE 802.1p classifiers on the same switch.

- You can modify the CoS classifier for a particular interface, but it is not required.
- You can configure a DSCP classifier, DSCP1 on the first interface, another DSCP classifier, DSCP2 on the second interface, and an IP precedence classifier on a third interface, and so forth.

Setting CoS Bits in an MPLS Header

When traffic enters an LSP tunnel, the CoS bits in the MPLS header are set in one of two ways:

- The number of the output queue into which the packet was buffered and the packet loss priority (PLP) bit are written into the MPLS header and are used as the packet's CoS value. This behavior is the default, and no configuration is required. The Junos OS Class of Service User Guide for Routing Devices explains the IP CoS values, and summarizes how the CoS bits are treated.
- You set a fixed CoS value on all packets entering the LSP tunnel. A fixed CoS value means that all packets entering the LSP receive the same class of service.

The CoS value can be a decimal number from 0 through 7. This number corresponds to a 3-bit binary number. The high-order 2 bits of the CoS value select which transmit queue to use on the outbound interface card.

The low-order bit of the CoS value is treated as the PLP bit and is used to select the RED drop profile to use on the output queue. If the low-order bit is 0, the non-PLP drop profile is used, and if the low-order bit is 1, the PLP drop profile is used. It is generally expected that random early detection (RED) will more aggressively drop packets that have the PLP bit set. For more information about RED and drop profiles, see the Junos OS Class of Service User Guide for Routing Devices.

NOTE: Configuring the PLP drop profile to drop packets more aggressively (for example, setting the CoS value from 6 to 7) decreases the likelihood of traffic getting through.

Table 33 on page 197 summarizes how MPLS CoS values correspond to the transmit queue and PLP bit. Note that in MPLS, the mapping between the CoS bit value and the output queue is hard-coded. You cannot configure the mapping for MPLS; you can configure it only for IPv4 traffic flows, as described in the Junos OS Class of Service User Guide for Routing Devices.

MPLS CoS Value	Bits	Transmit Queue	PLP Bit
0	000	0	Not set
1	001	0	Set
2	010	1	Not set
3	011	1	Set
4	100	2	Not set
5	101	2	Set
6	110	3	Not set
7	111	3	Set

Table 33: MPLS CoS Values

Because the CoS value is part of the MPLS header, the value is associated with the packets only while they travel through the LSP tunnel. The value is not copied back to the IP header when the packets exit from the LSP tunnel.

NOTE: On EX8200 switches that run MPLS-based Layer 2 virtual private networks (VPNs):

- If you configure an LSP CoS, the EXP bits of the MPLS packet continue to use the same CoS values that are configured at the interface level.
- For Virtual Chassis, if the input and output interfaces are on different line cards, then the loss priority value that you configured on the first line card is not carried to the subsequent line cards. The loss priority for the outgoing traffic from the subsequent line cards is always set to low.

EXP Rewrite Rules

When traffic passes from the customer-edge interface to an MPLS interface, the DSCP, IP precedence, or IEEE 802.1p CoS classifier is translated into the EXP bits within the MPLS header. You cannot disable the default EXP rewrite rule, but you can configure your own custom EXP classifier and a custom EXP rewrite rule. You cannot bind the EXP classifier to individual MPLS interfaces; the switch applies it globally to all the MPLS-enabled interfaces on the switch.

Only one EXP rewrite rule (either default or custom) is supported on a switch. The switch applies it to all the egress interfaces on which MPLS is enabled. This is, however, not the case with EX8200 switches. With EX8200 switches, you must explicitly apply the rewrite rule on each of the egress interfaces.

Policer

Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. During periods of congestion (when the total rate of queuing packets exceeds the rate of transmission), any new packets being sent to an interface can be dropped because there is no place to store them. You can configure a policer on the ingress PE switch to prevent this:

- If you are using MPLS over CCC, you bind the policer to the LSP. You cannot bind a policer to a CCC interface.
- If you are using IP over MPLS, you bind the policer to the **inet-family** customer-edge interface. You cannot bind a policer to the LSP when you are using IP over MPLS.

NOTE: You cannot configure LSP policers on EX8200 switches.

Schedulers

The schedulers for using CoS with MPLS are the same as for the other CoS configurations on EX Series switches. Default schedulers are provided for best-effort and network-control forwarding classes. If you are using assured-forwarding, expedited-forwarding, or any custom forwarding class, we recommend that you configure a scheduler to support that forwarding class. See "Understanding CoS Schedulers" on page 128.

Example: Combining CoS with MPLS on EX Series Switches

IN THIS SECTION

- Requirements | 200
- Overview and Topology | 200
- Configuring the Local PE Switch | 203
- Configuring the Remote PE Switch | 207
- Configuring the Provider Switch | 208
- Verification | 210

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion. The CoS value is included within the MPLS label, which is passed through the network, enabling end-to-end CoS across the network.

MPLS services are often used to ensure better performance for low-latency applications such as VoIP and other business-critical functions. These applications place specific demands on a network for successful transmission. CoS gives you the ability to control the mix of bandwidth, delay, jitter, and packet loss while taking advantage of the MPLS labeling mechanism.

This example shows how to configure CoS on an MPLS network that is using a unidirectional circuit cross-connect (CCC) from the ingress provider edge (PE) switch to the egress PE switch. for the customer-edge interface of the ingress provider edge (PE) switch. It describes adding the configuration of CoS components to the ingress PE switch, the egress PE switch, and the core provider switches of the existing MPLS network. Because of the unidirectional configuration, the DSCP classifier needs to be configured only on the ingress PE switch.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 10.1 or later for EX Series switches
- Three EX Series switches

Before you configure CoS with MPLS, be sure you have:

Configured an MPLS network with two PE switches and one provider switch. See Example: Configuring MPLS on EX8200 and EX4500 Switches. This example assumes that an MPLS network has been configured using a cross circuit-connect (CCC).

Overview and Topology

IN THIS SECTION

Topology | 203

This example describes adding custom classifiers and custom rewrite rules to switches in an MPLS network that is using MPLS over CCC.

It is a unidirectional configuration. Therefore, you need to configure custom classifiers and custom rewrite rules as follows:

- On the ingress PE switch: custom DSCP classifier and custom EXP rewrite rule
- On the egress PE switch: custom EXP classifier
- On the provider switch: customer EXP classifier and custom EXP rewrite rule

NOTE: You can also configure schedulers and shapers as needed. If you are using **assuredforwarding**, **expedited-forwarding**, or other custom forwarding classes, we recommend that you configure a scheduler to support that forwarding class. See "Defining CoS Schedulers and Scheduler Maps (CLI Procedure)" on page 138.

The example creates a custom DSCP classifier (**dscp1**) on the ingress PE switch and binds this classifier to the CCC interface. It includes configuration of a policer on the ingress PE switch. The policer is applied as a filter on the label-switched path (LSP) **lsp_to_pe2_ge1**(created in Example: Configuring MPLS on EX8200 and EX4500 Switches) to ensure that the amount of traffic forwarded through the LSP never exceeds the requested bandwidth allocation.

This example creates a custom EXP rewrite rule (**exp1**) on the ingress PE switch, specifying a losspriority and code point to be used for the expedited-forwarding class as the packet travels through the LSP. The switch applies this custom rewrite rule on the core interfaces **ge-0/0/5.0** and **ge-0/0/6.0**, which are the egress interfaces for this switch.

Table 34 on page 201 shows the CoS configuration components added to the ingress PE switch.

Table 34: CoS Configuration Components on the Ingress PE Switch

Property	Settings	Description
Local PE switch hardware	EX Series switch	PE-1
Policing filter configured and applied to the LSP.	policing filter mypolicer filter myfilter	Name of the rate-limiting policer. Name of the filter, which refers to the policer
Custom DSCP classifier	dscp1	Specifies the name of the custom DSCP classifier
Custom EXP rewrite rule	e1	Name of the custom EXP rewrite rule.
Customer-edge interface	ge-0/0/1.0	Interface that receives packets from devices outside the network. The custom DSCP classifier must be specified on this CCC interface.
Core interfaces	ge-0/0/5.0 and ge-0/0/6.0	Interfaces that transmit MPLS packets to other switches within the MPLS network. The EXP rewrite rule is applied implicitly to these interfaces.

Table 35 on page 202 shows the CoS configuration components added to the egress PE switch in this example.

Table 35: CoS Configuration Components of the Egress PE Switch

Property	Settings	Description
Remote provider edge switch hardware	EX Series switch	PE-2
Custom EXP classifier	exp1	Name of custom EXP classifier
Customer-edge interface	ge-0/0/1.0	Interface that transmits packets from this network to devices outside the network. No CoS classifier is specified for this interface. A scheduler can be specified.
Core interfaces	ge-0/0/7.0 and ge-0/0/8.0	Core interfaces on PE-2 that receive MPLS packets from the provider switch. The EXP classifier is enabled by default on the switch and applied implicitly to these interfaces.

Table 36 on page 202 shows the MPLS configuration components used for the provider switch in this example.

Table 36: CoS Configuration Components of the Provider Switch

Property	Settings	Description
Provider switch hardware	EX Series switch	Transit switch within the MPLS network configuration.
Custom EXP classifier	exp1	Name of the custom EXP classifier.
Custom EXP rewrite rule	e1	Name of the custom EXP rewrite rule.

Property	Settings	Description
Core interfaces receiving packets from other MPLS switches.	ge-0/0/5.0 and ge-0/0/6.0	Interfaces that connect the provider switch to the ingress PE switch (PE-1). The EXP classifier is enabled by default on the switch and applied implicitly to these interfaces.
Core interfaces transmitting packets to other switches within the MPLS network.	ge-0/0/7.0 and ge-0/0/8.0	Interfaces that transmit packets to the egress PE (PE-2). The EXP rewrite rule is applied implicitly on these interfaces. Schedulers can also be specified and will be applied to these interfaces.

Table 36: CoS Configuration Components of the Provider Switch (Continued)

Topology

Configuring the Local PE Switch

IN THIS SECTION

Procedure | 203

Procedure

CLI Quick Configuration

To quickly configure a custom DSCP classifier, custom EXP rewrite rule, and a policer on the local PE switch, copy the following commands and paste them into the switch terminal window of PE-1:

[edit]

set class-of-service classifiers dscpset class-of-service classifiers dscp dscp1 import default set class-of-service classifiers dscp dscp1 forwarding-class expedited-forwarding loss-priority

```
low code-points 000111
set class-of-service rewrite-rules exp e1 forwarding-class expedited-forwarding loss-priority
low code-point 111
set class-of-service interfaces ge-0/0/1 unit 0 classifier dscp1
set firewall policer mypolicer if-exceeding bandwidth-limit 500m
set firewall policer mypolicer if-exceeding burst-size-limit 33553920
set firewall policer mypolicer then discard
set firewall family any filter myfilter term t1 then policer mypolicer
set protocols mpls label-switched-path lsp_to_pe2_ge1 to 127.1.1.3 policing filter myfilter
```

Step-by-Step Procedure

To configure a custom DSCP classifier, custom EXP rewrite rule, and a policer on the ingress PE switch:

1. Import the default DSCP classifier classes to the custom DSCP classifier that you are creating:

```
[edit class-of-service]
user@switch# set classifiers dscp dscp1 import default
```

2. Add the expedited-forwarding class to this custom DSCP classifier, specifying a loss priority and code point:

```
[edit class-of-service]
user@switch# set classifiers dscp dscp1 forwarding-class expedited-forwarding loss-priority
low code-points 000111
```

3. Specify the values for the custom EXP rewrite rule, e1:

```
[edit class-of-service]
user@switch# set rewrite-rules exp e1 forwarding-class expedited-forwarding loss-priority low
code-point 111
```

4. Bind the DSCP classifier to the CCC interface:

```
[edit class-of-service]
user@switch# set class-of-service interfaces ge-0/0/1 unit 0 classifier dscp1
```

5. Specify the number of bits per second permitted, on average, for the firewall policer, which will later be applied to the LSP:

```
[edit firewall]
set policer mypolicer if-exceeding bandwidth-limit 500m
```

6. Specify the maximum size permitted for bursts of data that exceed the given bandwidth limit for this policer:

```
[edit firewall policer]
set mypolicer if-exceeding burst-size-limit 33553920
```

7. Discard traffic that exceeds the rate limits for this policer:

```
[edit firewall policer]
set mypolicer then discard
```

8. To reference the policer, configure a filter term that includes the policer action:

```
[edit firewall]
user@switch# set family any filter myfilter term t1 then policer
mypolicer
```

9. Apply the filter to the LSP:

```
[edit protocols mpls]
set label-switched-path lsp_to_pe2_ge1 policing filter myfilter
```

Results

Display the results of the configuration:

[edit]
user@switch# show
class-of-service {
 classifiers {
 dscp dscp1 {

```
import default;
            forwarding-class expedited-forwarding {
                loss-priority low code-points 000111;
            }
        }
    }
    interfaces {
        ge-0/0/1 {
            unit 0 {
                classifiers {
                    dscp dscp1;
                }
            }
        }
    }
    rewrite-rules {
        exp e1 {
            forwarding-class expedited-forwarding {
                loss-priority low code-point 111;
            }
        }
    }
}
firewall {
    family any {
        filter myfilter {
            term t1 {
                then policer mypolicer;
            }
        }
    }
    policer mypolicer {
        if-exceeding {
            bandwidth-limit 500m;
            burst-size-limit 33553920;
        }
        then discard;
    }
}
```

Configuring the Remote PE Switch

IN THIS SECTION

Procedure | 207

Procedure

CLI Quick Configuration

To quickly configure a custom EXP classifier on the remote PE switch, copy the following commands and paste them into the switch terminal window of PE-2:

[edit]

```
set class-of-service classifiers exp exp1 import default
set class-of-service classifiers exp exp1 forwarding-class expedited-forwarding loss-priority
low code-points 010
```

Step-by-Step Procedure

To configure a custom EXP classifier on the egress PE switch:

1. Import the default EXP classifier classes to the custom EXP classifier that you are creating:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 import default
```

2. Add the expedited-forwarding class to this custom EXP classifier, specifying a loss priority and code point:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 forwarding-class expedited-forwarding loss-priority low
code-points 010
```

Results

Display the results of the configuration:

```
[edit]
user@switch# show
class-of-service {
    classifiers {
        exp exp1 {
            import default;
            forwarding-class expedited-forwarding {
                loss-priority low code-points 010;
            }
        }
    }
}
```

Configuring the Provider Switch

IN THIS SECTION

Procedure | 208

Procedure

CLI Quick Configuration

To quickly configure a custom EXP classifier and a custom EXP rewrite rule on the provider switch, copy the following commands and paste them into the switch terminal window of the provider switch:

```
[edit]
set class-of-service classifiers exp exp1 import default
set class-of-service classifiers exp exp1 forwarding-class expedited-forwarding loss-priority
low code-points 010
set class-of-service rewrite-rules exp e1 forwarding-class expedited-forwarding loss-priority
low code-point 111
```

Step-by-Step Procedure

To configure a custom EXP classifier and a custom EXP rewrite rule on the provider switch:

1. Import the default EXP classifier classes to the custom EXP classifier that you are creating:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 import default
```

2. Add the expedited-forwarding class to this custom EXP classifier, specifying a loss priority and code point:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 forwarding-class expedited-forwarding loss-priority low
code-points 010
```

3. Specify the values for the custom EXP rewrite rule, e1:

```
[edit class-of-service]
user@switch# set rewrite-rules exp e1 forwarding-class expedited-forwarding loss-priority low
code-point 111
```

Results

Display the results of the configuration:

```
[edit]
user@switch# show
class-of-service {
    classifiers {
        exp exp1 {
            import default;
            forwarding-class expedited-forwarding {
               loss-priority low code-points 010;
            }
        }
      }
    rewrite-rules {
```

```
exp e1 {
    forwarding-class expedited-forwarding {
        loss-priority low code-point 111;
      }
   }
}
```

Verification

IN THIS SECTION

- Verifying That the Policer Firewall Filter Is Operational | 210
- Verifying That the CoS Classifiers Are Going to the Right Queue | 211
- Verifying the CoS Forwarding Table Mapping | **215**
- Verifying the Rewrite Rules | **216**

To confirm that the configuration is working properly, perform these tasks:

Verifying That the Policer Firewall Filter Is Operational

Purpose

Verify the operational state of the policer that is configured on the ingress PE switch.

Action

user@switch> show firewall	
Filter: myfilter Policers:	
Name	Packets
mypolicer-t1	0

Meaning

This output shows that the firewall filter **mypolicer** has been created.

Verifying That the CoS Classifiers Are Going to the Right Queue

Purpose

Verify that the CoS classifiers are going to the right queue.

Action

user@swit	ch> show clas	s-of-serv	ice forwardin	g-table classifie	r
Classifie	er table index	: 7, # en ⁻	tries: 64, Ta	ble type: DSCP	
Entry #	Code point	Forwardi	ng-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	
47	101111	0	0	
48	110000	3	0	
49	110001	3	0	
50	110010	3	0	
51	110011	3	0	
52	110100	3	0	
53	110101	3	0	
54	110110	3	0	
55	110111	3	0	
56	111000	3	0	
57	111001	3	0	
58	111010	3	0	
59	111011	3	0	
60	111100	3	0	
61	111101	3	0	
62	111110	3	0	
63	111111	3	0	
Classifier	table index:	11, #	entries:	8

Classifier table index: 11, # entries: 8, Table type: IEEE 802.1 Entry # Code point Forwarding-class # PLP 0 000 0 0

1	001	0	0	
2	010	0	0	
3	011	0	0	
4	100	0	0	
5	101	0	0	
6	110	3	0	
7	111	3	0	
Classifie	er table index	: 12, # ei	entries: 8, Table type: IPv4 precedence	
Entry #	Code point	Forwardi	ing-class # PLP	
0	000	0	0	
1	001	0	0	
2	010	0	0	
3	011	0	0	
4	100	0	0	
5	101	0	0	
6	110	3	0	
7	111	3	0	
Classifie	er table index	: 16, # e	entries: 8, Table type: Untrust	
Entry #	Code point	Forwardi	ing-class # PLP	
0	000	0	0	
1	001	0	0	
2	010	0	0	
3	011	0	0	
4	100	0	0	
5	101	0	0	
6	110	0	0	
7	111	0	0	
			<pre># entries: 64, Table type: DSCP</pre>	
Entry #	Code point		ing-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	1	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
47	101111	0	0
48	110000	3	0
49	110001	3	0
50	110010	3	0
51	110011	3	0
52	110100	3	0
53	110101	3	0

54	110110	3	0
55	110111	3	0
56	111000	3	0
57	111001	3	0
58	111010	3	0
59	111011	3	0
60	111100	3	0
61	111101	3	0
62	111110	3	0
63	111111	3	0

Meaning

This output shows that a new DSCP classifier has been created, index **9346**, on the ingress PE switch (PE-1).

Verifying the CoS Forwarding Table Mapping

Purpose

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.

Action

user@switch> show class-of-service forwarding-table classifier mapping

		Table Index,	/
Interface	Index	Q num	Table type
ge-0/0/1.0	92	9346	DSCP

Meaning

The results show that the new DSCP classifier, index number 9346, is bound to interface ge-0/0/1.0.

Verifying the Rewrite Rules

Purpose

Display mapping of the queue number and loss priority to code point value for each rewrite rule as it exists in the forwarding table.

Action

user@sw:	itch>show class-of-service forwarding-table rewrite-rule
Rewrite	table index: 31, # entries: 4, Table type: DSCP
FC#	Low bits State High bits State
0	000000 Enabled 000000 Enabled
1	101110 Enabled 101110 Enabled
2	001010 Enabled 001100 Enabled
3	110000 Enabled 111000 Enabled
Rewrite	table index: 34, # entries: 4, Table type: IEEE 802.1
FC#	Low bits State High bits State
0	000 Enabled 001 Enabled
1	010 Enabled 011 Enabled
2	100 Enabled 101 Enabled
3	110 Enabled 111 Enabled
Rewrite	table index: 35, # entries: 4, Table type: IPv4 precedence
FC#	Low bits State High bits State
0	000 Enabled 000 Enabled
1	101 Enabled 101 Enabled
2	001 Enabled 001 Enabled
3	110 Enabled 111 Enabled
Rewrite	table index: 9281, # entries: 1, Table type: EXP
FC#	Low bits State High bits State
1	111 Enabled 000 Disabled

Meaning

This output shows that a new EXP classifier with the index number **9281** has been created.

Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS

IN THIS SECTION

- Configuring CoS | 217
- Configuring an LSP Policer | 218

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion. This topic describes configuring CoS components on a provider edge (PE) switch that is using IP Over MPLS.

This task describes how to create a custom DSCP classifier and a custom EXP rewrite rule on the ingress PE switch. It includes configuring a policer firewall filter and applying it to the customer-edge interface of the ingress PE switch. The policer firewall filter ensures that the amount of traffic forwarded through the MPLS tunnel never exceeds the requested bandwidth allocation.

Before you begin, configure the basic components for an MPLS network:

- Configure two PE switches. See Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect.
- Configure one or more provider switches. See Configuring MPLS on EX8200 and EX4500 Provider Switches.

Configuring CoS

To configure CoS on a provider edge switch:

1. Import the default DSCP classifier classes to the custom DSCP classifier that you are creating:

```
[edit class-of-service]
user@switch# set classifiers dscp classifier-name import default
```

2. Add a forwarding class to this custom DSCP classifier and specify a loss priority and code point:

[edit class-of-service]
user@switch# set classifiers dscp classifier-name forwarding-class forwarding-class losspriority loss-priority code-points code-point

3. Specify the values for the custom EXP rewrite rule, e1:

```
[edit class-of-service]
user@switch# set rewrite-rules exp e1 forwarding-class forwarding-class loss-priority loss-
priority code-points code-point
```

4. On EX8200 switches only, bind the custom EXP rewrite rule to the interface:

```
[edit class-of-service]
user@switch# set class-of-service interfaces interface unit unit rewrite-rules exp e1
```

Configuring an LSP Policer

To configure an LSP policer:

NOTE: You cannot configure LSP policers on EX8200 switches. EX8200 switches do not support LSP policers.

1. Specify the number of bits per second permitted, on average, for the firewall policer, which will later be applied to the customer-edge-interface:

[edit firewall]
user@switch# set policer mypolicer if-exceeding bandwidth-limit 500m

2. Specify the maximum size permitted for bursts of data that exceed the given bandwidth limit for this policer:

[edit firewall policer]
user@switch# set mypolicer if-exceeding burst-size-limit 33553920

3. Discard traffic that exceeds the rate limits for this policer:

[edit firewall policer]
user@switch# set mypolicer then discard

4. To reference the policer, configure a filter term that includes the policer action:

[edit firewall]
user@switch# set family inet filter myfilter term t1 then policer mypolicer

5. Apply the filter to the customer-edge interface:

```
[edit interfaces]
user@switch# set ge-2/0/3 unit 0 family inet address 192.168.121.1/16 policing filter myfilter
```

NOTE: You can also configure schedulers and shapers as needed. See "Defining CoS Schedulers and Scheduler Maps (CLI Procedure)" on page 138.

Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect

IN THIS SECTION

- Configuring CoS | 220
- Configuring an LSP Policer | 220

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion. This topic describes configuring CoS components on a provider edge (PE) switch that is using MPLS over circuit-cross connect (CCC).

NOTE: On EX Series switches other than EX8200 switches, if you are using MPLS over CCC, you can use only one DSCP or IP precedence classifier and only one IEEE 802.1p classifier on the CCC interfaces.

This procedure is for creating a custom DSCP classifier and a custom EXP rewrite rule on the ingress PE. It also includes enabling a policer on the label-switched path (LSP) of the ingress PE to ensure that the amount of traffic forwarded through the LSP never exceeds the requested bandwidth allocation.

This topic includes:

Configuring CoS

To configure CoS on a provider edge switch:

1. Import the default DSCP classifier classes to the custom DSCP classifier that you are creating:

```
[edit class-of-service]
user@switch# set classifiers dscp classifier-nameimport default
```

2. Add the expedited-forwarding class to this custom DSCP classifier, specifying a loss priority and code point:

```
[edit class-of-service]
user@switch# set classifiers dscp classifier-name forwarding-class forwarding-class loss-
priority loss-priority code-points code-point
```

3. Specify the values for the custom EXP rewrite rule, e1:

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

user@switch# set rewrite-rules exp e1 forwarding-class forwarding-class loss-priority losspriority code-point code-point

4. Bind the DSCP classifier to the CCC interface:

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

5. On EX8200 switches only, bind the custom EXP rewrite rule to the interface:

```
[edit class-of-service]
user@switch# set class-of-service interfaces interface unit unit rewrite-rules exp e1
```

Configuring an LSP Policer

To configure an LSP policer:

NOTE: You cannot configure LSP policers on EX8200 switches. EX8200 switches do not support LSP policers.

1. Specify the number of bits per second permitted, on average, for the policer, which will later be applied to the LSP:

```
[edit firewall]
set policer mypolicer if-exceeding bandwidth-limit 500m
```

2. Specify the maximum size permitted for bursts of data that exceed the given bandwidth limit for this policer:

[edit firewall policer]
set mypolicer if-exceeding burst-size-limit 33553920

3. Discard traffic that exceeds the rate limits for this policer:

[edit firewall policer]
set mypolicer then discard

4. To reference the policer, configure a filter term that includes the policer action:

[edit firewall]
user@switch# set family any filter myfilter term t1 then policer
mypolicer

5. Apply the filter to the LSP:

[edit protocols mpls]
set label-switched-path lsp_to_pe2_ge1 policing filter myfilter

NOTE: You can also configure schedulers and shapers as needed. See "Defining CoS Schedulers and Scheduler Maps (CLI Procedure)" on page 138.

Configuring CoS on Provider Switches of an MPLS Network

You can add class-of-service (CoS) components to your MPLS networks on EX Series switches to achieve end-to-end Differentiated Services to match your specific business requirements. The configuration of CoS components on the provider switches is the same regardless of whether the provider edge (PE) switches are using MPLS over CCC or IP over MPLS.

This task shows how to configure a custom EXP classifier and custom EXP rewrite rule on the provider switch.

1. Import the default EXP classifier classes to the custom EXP classifier that you are creating:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 import default
```

2. Add the expedited-forwarding class to this custom EXP classifier, specifying a loss priority and code point:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 forwarding-class expedited-forwarding loss-priority low
code-points 010
```

3. Specify the values for the custom EXP rewrite rule, e1:

```
[edit class-of-service]
user@switch# set rewrite-rules exp e1 forwarding-class expedited-forwarding loss-priority low
code-point 111
```

4. On EX8200 switches only, bind the custom EXP rewrite rule to the interface:

```
[edit class-of-service]
user@switch# set class-of-service interfaces ge-0/0/2 unit 0 rewrite-rules exp e1
```

NOTE: You can also configure schedulers and shapers as needed. See "Defining CoS Schedulers and Scheduler Maps (CLI Procedure)" on page 138.

Configuring MPLS on EX8200 and EX4500 Provider Switches

You can configure MPLS on EX8200 and EX4500 switches to increase transport efficiency in your network. MPLS services can be used to connect various sites to a backbone network and to ensure better performance for low-latency applications such as VoIP and other business-critical functions.

To implement MPLS on EX Series switches, you must configure at least one provider switch as a transit switch for the MPLS packets. The configuration of all the provider switches remains the same regardless of whether the provider edge (PE) switches are using circuit cross-connect (CCC) or using MPLS over IP for the customer edge interfaces. Likewise, you do not need to change the configuration of the provider switches if you implement an MPLS-based Layer 2 VPN, Layer 3 VPN, or a Layer 2 circuit configuration.

MPLS requires the configuration of a routing protocol (OSPF or IS-IS) on the core interfaces and the loopback interface of all the switches. This procedure includes the configuration of OSPF on the provider switch. For information on configuring IS-IS as the routing protocol, see *Junos OS Routing Protocols Configuration Guide*.

To configure the provider switch, complete the following tasks:

1. Enable the routing protocol (OSPF or IS-IS) on the loopback interface and on the core interfaces:

NOTE: You can use the switch address as an alternative to the loopback interface.

[edit protocols]
user@switch# set ospf area 0.0.0.0 interface lo0.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/5.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/6.0
user@switch# set ospf area 0.0.0.0 interface ae0

2. Enable traffic engineering for the routing protocol (traffic engineering must be explicitly enabled for OSPF):

[edit protocols]
user@switch# set ospf traffic-engineering

3. Enable MPLS within the protocols stanza and apply it to the core interfaces:

```
[edit protocols]
user@switch# set mpls interface ge-0/0/5.0
```

```
user@switch# set mpls interface ge-0/0/6.0
user@switch# set mpls interface ae0
```

4. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface ge-0/0/5.0
user@switch# set rsvp interface ge-0/0/6.0
user@switch# set rsvp interface ae0
```

5. Configure an IP address for the loopback interface and for the core interfaces:

[edit]

```
user@switch# set interfaces lo0 unit 0 family inet address 127.1.1.1/32
user@switch# set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
user@switch# set interfaces ae0 unit 0 family inet address 10.1.9.2/24
```

6. Configure family mpls on the logical units of the core interfaces:

[edit]

```
user@switch# set interfaces ge-0/0/5 unit 0 family mpls
user@switch# set interfaces ge-0/0/6 unit 0 family mpls
user@switch# set interfaces ae0 unit 0 family mpls
```

NOTE: You can enable **family mpls** on either individual interfaces or aggregated Ethernet interfaces. You cannot enable it on tagged VLAN interfaces.

CoS on EVPN VXLANs

IN THIS CHAPTER

CoS Support on EVPN VXLANs | 225

CoS Support on EVPN VXLANs

IN THIS SECTION

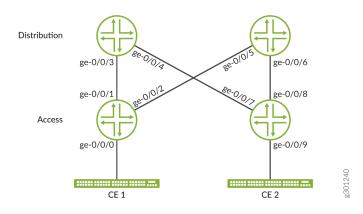
- Understanding CoS on VXLAN Interfaces | 225
- Configuring CoS on VXLAN Interfaces | 227
- Implementing CoS on VXLAN Interfaces (Junos OS Evolved) | 230
- CoS Limitations on VXLANs | 231

You can configure class of service (CoS) features on VXLAN interfaces. VXLAN traffic from different tenants traverses network boundaries over the same physical underlay network. To ensure fairness in the treatment of traffic for all tenants in the VXLAN, and to prioritize higher priority traffic, apply CoS features to the VXLAN interfaces.

Understanding CoS on VXLAN Interfaces

This section describes how classification and rewrite rules are applied to packets in a VXLAN instance. Figure 7 on page 226 shows a simple VXLAN with two leaf nodes and one spine node.

Figure 7: Classifiers and Rewrite Rules on VXLANs



Refer to Figure 7 on page 226 to understand the packet flow with DSCP/ToS fields in a VXLAN:

- **1.** CE 1 sends a packet with Layer3 DSCP/ToS bit programmed to the Leaf 1 node.
- **2.** Leaf 1 receives the original packet and appends the VXLAN header on top of the original packet. The outer VXLAN Layer3 header uses the original packet DSCP/Tos bit. You can create classifiers based on the original packet DSCP/802.1p bit. The ingress interface on the ingress leaf supports DSCP and 802.1p classifiers.
- **3.** If rewrite is configured on Leaf 1, the inner header will have the DSCP/802.1p bit set by CE 1 and the outer header will have the rewrite bit. Only DSCP rewrite rules are supported, except on QFX10000 switches where 802.1p rewrite is also supported if the underlay is tagged.
- **4.** The Spine node receives the VXLAN packet and can use ingress classification using these DSCP bits and forward the packet to the egress interface with the appropriate forwarding class.
- **5.** The Spine egress interface can rewrite these bits using rewrite rules. These Spine rewrite rules only affects the outer Layer3 DSCP field. The inner/original packet still holds the DSCP/802.1p bit that was set by CE 1.
- 6. Leaf 2 receives the packet, processes the tunnel termination, and remove the outer VXLAN header.
- 7. Leaf 2 classification and rewrite functionality works on the inner header.
- 8. The original packet arrives on CE 2.

NOTE: On the leaf nodes, if the packet is multicast, you can use multi-destination classification to create appropriate multicast classification and rewrite rules.

Configuring CoS on VXLAN Interfaces

This section shows sample configurations of classifiers and rewrite rules for the leaf and spine nodes in VXLAN using Figure 7 on page 226 as a reference. You can create schedulers as normal for the classifiers on each node.

Sample configuration of classifiers and rewrite rules on Leaf 1.

1. Create a classifier based on the *original* DSCP/ToS bits:

```
[edit class-of-service classifiers]
user@leaf1#set dscp dscp_cf forwarding-class best-effort loss-priority low code-points 100000
user@leaf1#set dscp dscp_cf forwarding-class network-control loss-priority high code-points
110000
user@leaf1#set dscp dscp_cf forwarding-class expedited-forwarding loss-priority low code-
points 011010
user@leaf1#set dscp dscp_cf forwarding-class assured-forwarding loss-priority high code-
points 001010
```

2. Apply the classier to the ingress interface:

[edit class-of-service interfaces]
user@leaf1#set ge-0/0/0 unit 0 classifiers dscp dscp_cf

3. Create a rewrite rule for the *outer* VXLAN DSCP/ToS bits:

```
[edit class-of-service rewrite-rules]
user@leaf1#set dscp dscp_rw forwarding-class best-effort loss-priority low code-points af22
user@leaf1#set dscp dscp_rw forwarding-class network-control loss-priority high code-points
af31
user@leaf1#set dscp dscp_rw forwarding-class expedited-forwarding loss-priority low code-
points af13
user@leaf1#set dscp dscp_rw forwarding-class assured-forwarding loss-priority high code-
points cs3
```

4. Apply the rewrite rule to the egress Leaf 1 interfaces:

[edit class-of-service interfaces]
user@leaf1#set ge-0/0/1 unit 0 rewrite-rules dscp dscp_rw
user@leaf1#set ge-0/0/2 unit 0 rewrite-rules dscp dscp_rw

Sample configuration of classifiers and rewrite rules on the Spine.

1. Create a classifier based on the outer VXLAN DSCP/ToS bits:

```
[edit class-of-service classifiers]
user@spine#set dscp dscp_cf forwarding-class best-effort loss-priority low code-points af22
user@spine#set dscp dscp_cf forwarding-class network-control loss-priority high code-points
af31
user@spine#set dscp dscp_cf forwarding-class expedited-forwarding loss-priority low code-
points af13
user@spine#set dscp dscp_cf forwarding-class assured-forwarding loss-priority high code-
points cs3
```

2. Apply the classier to the ingress Spine interfaces:

[edit class-of-service interfaces]
user@spine#set ge-0/0/3 unit 0 classifiers dscp dscp_cf
user@spine#set ge-0/0/5 unit 0 classifiers dscp dscp_cf

3. Create a rewrite rule for the outer VXLAN DSCP/ToS bits:

```
[edit class-of-service rewrite-rules]
user@spine#set dscp dscp_rw forwarding-class best-effort loss-priority low code-points af22
user@spine#set dscp dscp_rw forwarding-class network-control loss-priority high code-points
af31
user@spine#set dscp dscp_rw forwarding-class expedited-forwarding loss-priority low code-
points af13
user@spine#set dscp dscp_rw forwarding-class assured-forwarding loss-priority high code-
points cs3
```

4. Apply the rewrite rule to the egress Spine interfaces:

```
[edit class-of-service interfaces]
user@spine#set ge-0/0/4 unit 0 rewrite-rules dscp dscp_rw
user@spine#set ge-0/0/6 unit 0 rewrite-rules dscp dscp_rw
```

Sample configuration of classifiers and rewrite rules on Leaf 2.

1. Create a classifier based on the *original* DSCP/ToS bits, as the VXLAN header is removed at tunnel termination *before* forwarding classes are applied:

```
[edit class-of-service classifiers]
user@leaf2#set dscp dscp_cf forwarding-class best-effort loss-priority low code-points 100000
user@leaf2#set dscp dscp_cf forwarding-class network-control loss-priority high code-points
110000
user@leaf2#set dscp dscp_cf forwarding-class expedited-forwarding loss-priority low code-
points 011010
user@leaf2#set dscp dscp_cf forwarding-class assured-forwarding loss-priority high code-
points 001010
```

2. Apply the classier to the ingress Leaf 2 interfaces:

```
[edit class-of-service interfaces]
user@leaf2#set ge-0/0/7 unit 0 classifiers dscp dscp_cf
user@leaf2#set ge-0/0/8 unit 0 classifiers dscp dscp_cf
```

3. Create a rewrite rule for the *original* DSCP/ToS bits:

```
[edit class-of-service rewrite-rules]
user@leaf2#set dscp dscp_rw forwarding-class best-effort loss-priority low code-points 100000
user@leaf2#set dscp dscp_rw forwarding-class network-control loss-priority high code-points
110000
user@leaf2#set dscp dscp_rw forwarding-class expedited-forwarding loss-priority low code-
points 011010
user@leaf2#set dscp dscp_rw forwarding-class assured-forwarding loss-priority high code-
points 001010
```

4. Apply the rewrite rule to the egress Leaf 2 interface:

[edit class-of-service interfaces]
user@leaf2#set ge-0/0/9 unit 0 rewrite-rules dscp dscp_rw

To check the CoS configuration on one of the interfaces:

user@node#show class-of-service interface interface-name

To check the queue statistics on one of the interfaces:

user@node#show interfaces queue interface-name

Implementing CoS on VXLAN Interfaces (Junos OS Evolved)

CoS for EVPN VXLAN traffic is supported using a combination of classifiers, schedulers, and rewrite rules. This section describes how these components are implemented across different nodes on devices running Junos OS Evolved to apply CoS on the EVPN VXLAN traffic.

- Classification at User Network Interface (UNI)/Ingress PE Traffic classification based on IEEE 802.1p and Differentiated Services code point (DSCP) are supported on the ingress PE where the EVPN VXLAN tunnel is initiated. BA and MF classifiers can be applied to Enterprise style (EP) or Service Provider (SP) style access interfaces.
- Classification at Network Node Interface (NNI)/Egress PE Traffic classification based on IEEE 802.1p and Differentiated Services code point (DSCP) are supported on the egress PE where the EVPN VXLAN tunnel is terminated. BA classifiers can be applied to the underlying logical interface or unit. MF classifiers are not supported in tunnel terminations.
- Rewrite at NNI After the encapsulation of the VXLAN tunnel, the rewrites on the outer/tunnel header are configured using the rewrite rules on the underlying logical interface or unit. Based on the configured rewrite rules, the VXLAN traffic is classified in the Spine/Network.
 DSCP rewrites on the outer/tunnel header of VXLAN packets is supported on the NNI interface.

Rewrite rules are supported in the following EVPN VXLAN scenarios:

- Intra-VNI L2 gateway Rewrite rules are applied to both unicast and broadcast, unknown unicast, and multicast (BUM) traffic.
- Inter-VNI L3 gateway Centrally-routed bridging (CRB) and edge-routed bridging (ERB).
- EVPN Type 5 routes.
- Rewrite at UNI After the termination of the VXLAN tunnel, the rewrites on the inner headers are configured using rewrite rules on the Enterprise style (EP) or Service Provider (SP) style access interfaces. Based on the configured rewrite rules, the decapsulated packets are classified in the CE side network. The following rewrite rules are supported on the UNI interface for the decapsulated packets:
 - DSCP rewrites on the inner IPv4/IPv6 header
 - IEEE 802.1p rewrites on the inner Ethernet header (if tagged)

Rewrite rules are supported in the following EVPN VXLAN scenarios:

- Intra-VNI L2 gateway Rewrite rules are applied to both unicast and broadcast, unknown unicast, and multicast (BUM) traffic.
- Inter-VNI L3 gateway Centrally-routed bridging (CRB) and edge-routed bridging (ERB).
- EVPN Type 5 routes.
- Scheduling Traffic prioritization and bandwidth reservation are achieved by using schedulers. The schedulers are associated with a forwarding class set via classifiers.

CoS Limitations on VXLANs

The following limitations apply to PTX routers:

- DSCP rewrite rules are not supported on Integrated Routing and Bridging (IRB) (L3 gateway scenarios).
- IEEE 802.1p rewrite rules are not supported on the NNI interface.
- Explicit congestion notification (ECN) rewrites are not supported on either UNI or NNI interfaces.
- Priority-based flow control (PFC) is not supported.
- No support for CoS classification and rewrite mechanism for IPv6 or IRB underlay.

The CoS functionality on EVPN VXLAN is the same as on QFX5K platforms. All VXLAN CoS features already supported on the QFX5120 are also supported on the QFX5130 and QFX5700 platforms.

The following limitations apply to the QFX5130 and QFX5700 platforms:

- HQoS is not supported due to hardware limitations.
- Classifier, rewrite and scheduler on IRB interface is not supported.
- DOT1P rewrite and classifier on the NNI port is not supported.
- DOT1P and DSCP rewrite on the UNI port is not supported.
- DSCP rewrite on the NNI port is supported with the following limitations:
 - DSCP rewrite takes effect only after you disable TOS copy (set vxlan-disable-copy-tos-encap at [edit forwarding-options] hierarchy level) on the VXLAN encapsulation node. When TOS copy is disabled, ECN bits are not copied from the inner to the outer header, so the packet outer header will have the defined rewrite DSCP value and an ECN value of 00.
 - DSCP rewrite rewrites both the outer and the inner header. So the inner header DSCP value cannot be preserved.
- PFC configuration will cause momentary traffic drops of up to 10ms.

- DSCP IPV6 classifiers and rewrites are not supported. Use DSCP classifier and rewrite instead.
- TOS copy feature does not work for Type-5 EVPN VXLANs.

The following limitation applies to QFX10000 platforms:

• Because IRB interfaces do not support dscp rewrite rules, you can apply rewrite rules on underlying L2 interfaces. 802.1p/dscp values in a VXLAN tunneled packet are written using underlying L2 interface rules.



Configuration Statements and Operational Commands

Configuration Statements | 234

Operational Commands | 321

Configuration Statements

IN THIS CHAPTER

- broadcast | 235
- buffer-size | 237
- e class | 242
- class-of-service | 244
- classifiers | 247
- code-point (Congestion Notification) | 249
- code-point-aliases | 251
- code-points | 252
- congestion-notification-profile | 253
- drop-probability (Fill Level) | 257
- drop-profile-map | 259
- dscp | 260
- dscp-ipv6 | 262
- ethernet (CoS for Multidestination Traffic) | 264
- excess-rate (Schedulers) | 265
- exp | 267
- explicit-congestion-notification | 268
- family | **270**
- flow-control | 271
- forwarding-class (Forwarding Policy) | 273
- forwarding-classes | 275
- ieee-802.1 | 281
- ieee-802.1 (Congestion Notification) | 282
- import | 284
- inet (CoS) | 285
- inet6 (CoS Multidestination) | 287

- inet-precedence | 288
- input (Congestion Notification) | 290
- interfaces | 291
- loss-priority (Classifiers and Rewrite Rules) | 293
- multi-destination | 295
- policing | 297
- priority (Schedulers) | 298
- protocol (Drop Profiles) | 300
- rewrite-rules | 301
- scheduler-map | 302
- scheduler-maps | 304
- schedulers (CoS) | 305
- shaping-rate | 307
- shaping-rate (Applying to an Interface) | 309
- shared-buffer | 312
- transmit-rate (EX Series Switches) | 313
- tri-color (EX Series Switches) | **315**
- unit | 318

broadcast

IN THIS SECTION

- Syntax | 236
- Hierarchy Level | 236
- Description | 236
- Options | **236**
- Required Privilege Level | 236
- Release Information | 236

Syntax

broadcast forwarding-class-name;

Hierarchy Level

[edit class-of-service multi-destination family ethernet]

Description

Specify the forwarding class for the broadcast traffic belonging to the Ethernet family.

Options

forwarding-class-name – Name of the forwarding class:

- mcast-af-Default forwarding class for assured forwarding of multicast traffic.
- mcast-be-Default best-effort forwarding class for multicast traffic.
- mcast-ef-Default forwarding class for expedited forwarding of multicast traffic.

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 9.5.

RELATED DOCUMENTATION

Understanding CoS Schedulers | 128 Understanding CoS Forwarding Classes | 110

Understanding CoS Classifiers | 70

buffer-size

IN THIS SECTION

- Syntax | 237
- Hierarchy Level | 237
- Description | 237
- Default | 239
- Options | 241
- Required Privilege Level | 242
- Release Information | 242

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 noloss, 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 receives 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 37 on page 240 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:

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%

Table 37: Default Output Queue Buffer Sizes (QFX10000 Switches)

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 38 on page 241 shows the default queue buffer sizes. The default buffer size is the same as the default transmit rate for each default queue:

Queue Number	Forwarding Class	Transmit Rate	Buffer Size
0	best-effort	5%	5%
3	fcoe	35%	35%
4	no-loss	35%	35%
7	network-control	5%	5%
8	mcast	20%	20%

Table 38: Default Output Queue Buffer Sizes (QFX5100, EX4600, QFX3500, and QFX3600 Switches, and QFabric Systems)

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 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, 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, 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 remainder 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.
percentage	Buffer size as a percentage of the total buffer.
remainder	Remaining buffer available.
temporal	(EX4200 standalone switches, EX4200 Virtual Chassis, EX4300 standalong 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 | 15

Defining CoS Schedulers and Scheduler Maps (CLI Procedure) | 138

Defining CoS Schedulers (J-Web Procedure) | 141

Understanding CoS Schedulers | 128

class

IN THIS SECTION

Syntax | 243

- Hierarchy Level | 243
- Description | 243
- Options | 243
- Required Privilege Level | 244
- Release Information | 244

Syntax

class class-name queue-num queue-number priority (high | low);

Hierarchy Level

[edit class-of-service forwarding-classes]

Description

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.

On EX8200 switches, you can assign a fabric priority to a forwarding class. The fabric priority determines scheduling priority of packets ingressing the switch fabric. In addition, for interfaces on the 40-port SFP+ line card, the fabric priority determines whether packets are sent to the high or low priority queue for ingressing the port group. The primary use of this option is to prevent high priority input traffic from being dropped due to congestion on the port group of a 40-port SFP+ line card.

Options

class-name	Name of forwarding class.
priority (high low)	(Optional) (EX8200 switches only) Fabric priority.
	• Default: low

queue-num queue-number

Output queue number.

• Range: 0 through 7

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 | 15

Defining CoS Forwarding Classes (CLI Procedure) | 114

Defining CoS Forwarding Classes (J-Web Procedure) | 115

Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure) | 87

class-of-service

IN THIS SECTION

- Syntax | 245
- Hierarchy Level | 246
- Description | 246
- Default | 246
- Required Privilege Level | 247
- Release Information | 247

Syntax

```
class-of-service {
    classifiers {
        (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence) classifier-name {
            import (classifier-name | default);
            forwarding-class class-name {
                loss-priority level {
                     code-points [aliases] [6 bit-patterns];
                }
            }
        }
   }
    code-point-aliases {
        (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence) {
            alias-name bits;
        }
   }
    forwarding-classes {
        class class-name queue-num queue-number priority (high | low);
   }
    interfaces {
        interface-name {
            scheduler-map map-name;
            unit logical-unit-number {
                forwarding-class class-name;
                classifiers {
                     (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence) (classifier-name | default);
                }
            }
        }
   }
    multi-destination {
        family {
            ethernet {
                broadcast forwarding-class-name;
            }
            inet {
                classifiers {
                     (dscp | dscp-ipv6 | inet-precedence) classifier-name;
                }
            }
```

```
}
        scheduler-map map-name;
    }
    rewrite-rules {
        (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence) rewrite-name {
            import (rewrite-name | default);
            forwarding-class class-name {
                loss-priority priority code-point (alias | bits);
            }
        }
    }
    scheduler-maps {
        map-name {
            forwarding-class class-name scheduler scheduler-name;
        }
    }
    schedulers {
        scheduler-name {
            buffer-size (percent percentage | remainder);
            drop-profile-map loss-priority loss-priority protocol protocol drop-profile profile-name;
            priority priority;
            shaping-rate (rate | percent percentage);
            transmit-rate (EX Series Switches) (rate | percent percentage | remainder);
        }
    }
}
```

Hierarchy Level

[edit]

Description

Configure class-of-service (CoS) parameters on EX Series switches.

The remaining statements are explained separately. See CLI Explorer.

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 9.0.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches 15		
Defining CoS Code-Point Aliases (CLI Procedure) 65		
Defining CoS Code-Point Aliases (J-Web Procedure) 66		
Defining CoS Classifiers (CLI Procedure) 74		
Defining CoS Classifiers (J-Web Procedure) 77		
Defining CoS Forwarding Classes (CLI Procedure) 114		
Defining CoS Forwarding Classes (J-Web Procedure) 115		
Configuring CoS Tail Drop Profiles (CLI Procedure) 165		
Defining CoS Schedulers and Scheduler Maps (CLI Procedure) 138		
Defining CoS Schedulers (J-Web Procedure) 141		
Defining CoS Rewrite Rules (CLI Procedure) 97		
Defining CoS Rewrite Rules (J-Web Procedure) 100		
Assigning CoS Components to Interfaces (CLI Procedure) 55		
Assigning CoS Components to Interfaces (J-Web Procedure) 56		
Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure) 87		

classifiers

IN THIS SECTION

Syntax | 248

- Hierarchy Level | 248
- Description | 248
- Required Privilege Level | 248
- Release Information | 249

Syntax

```
classifiers {
  (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence | exp) classifier-name {
     import (classifier-name | default);
     forwarding-class class-name {
        loss-priority level {
            code-points [aliases] [6-bit-patterns];
        }
     }
   }
}
```

Hierarchy Level

```
[edit class-of-service],
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Apply a CoS aggregate behavior classifier to a logical interface. You can apply a default classifier or a custom classifier.

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.

Expanded to include EXP classifiers in Junos OS Release 10.1 for EX Series switches.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches 15	
Example: Combining CoS with MPLS on EX Series Switches	
Defining CoS Classifiers (CLI Procedure) 74	
Defining CoS Classifiers (J-Web Procedure) 77	
Assigning CoS Components to Interfaces (CLI Procedure) 55	
Assigning CoS Components to Interfaces (J-Web Procedure) 56	
Understanding CoS Classifiers 70	

code-point (Congestion Notification)

IN THIS SECTION

- Syntax | 249
- Hierarchy Level | 250
- Description | 250
- Options | **250**
- Required Privilege Level | 250
- Release Information | 250

Syntax

code-point up-bits pfc;

Hierarchy Level

[edit class-of-service congestion-notification-profile profile-name input ieee-802.1], [edit class-of-service interfaces interface-name congestion-notification-profile profile-name input ieee-802.1]

Description

Configure the IEEE 802.1p (User Priority) code point bits as input for creating the priority-based flow control (PFC) congestion notification profile, which you will associate with a particular traffic class.

Options

- pfc-PFC flow control method
- up-bits-Three-bit pattern of the User Priority field in an IEEE 802.1Q tag

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 10.4.

RELATED DOCUMENTATION

Example: Configuring an FCoE Transit Switch

Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure) | 123

code-point-aliases

IN THIS SECTION

- Syntax | 251
- Hierarchy Level | 251
- Description | 251
- Required Privilege Level | 251
- Release Information | 252

Syntax

```
code-point-aliases {
   (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence) |{
        alias-name bits;
   }
}
```

Hierarchy Level

[edit class-of-service]

Description

Define an alias for a CoS marker.

The remaining statement is 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 | 15 Defining CoS Code-Point Aliases (CLI Procedure) | 65 Defining CoS Code-Point Aliases (J-Web Procedure) | 66 Understanding CoS Code-Point Aliases | 61

code-points

IN THIS SECTION

- Syntax | 252
- Hierarchy Level | 252
- Description | 253
- Options | **253**
- Required Privilege Level | 253
- Release Information | 253

Syntax

code-points [aliases] [6 bit-patterns];

Hierarchy Level

[edit class-of-service classifiers (dscp | ieee-802.1 | inet-precedence) forwarding-class *class-name* loss-priority *level*]

Description

Specify one or more DSCP code-point aliases or bit sets for association with a forwarding class.

Options

aliases – Name of the DSCP alias.

6 bit-patterns – 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 9.0.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches | 15 Defining CoS Classifiers (CLI Procedure) | 74 Defining CoS Classifiers (J-Web Procedure) | 77

Understanding CoS Classifiers | 70

congestion-notification-profile

IN THIS SECTION

- Syntax | 254
- EX4500 and EX4550 Switches | 255
- Hierarchy Level | 255
- Description | 255
- Options | 256

- Required Privilege Level | 256
- Release Information | 257

Syntax

```
congestion-notification-profile profile-name {
    input {
        (dscp | ieee-802.1) {
            code-point [code-point-bits] {
                pfc {
                    mru mru-value;
                }
                xon value;
            }
        }
        cable-length cable-length-value;
    }
    output {
        ieee-802.1 {
            code-point [code-point-bits] {
                flow-control-queue [queue | list-of-queues];
            }
        }
    }
    pfc-watchdog {
        detection number of polling intervals;
        pfc-watchdog-action {
            drop;
            }
            poll-interval time;
        recovery time;
    }
}
```

EX4500 and EX4550 Switches

congestion-notification-profile profile-name {
 input {
 ieee-802.1 {
 code-point up-bits pfc;
 }
}

Hierarchy Level

[edit class-of-service], [edit class-of-service interfaces interface-name]

Description

Configure a congestion notification profile (CNP) to enable priority-based flow control (PFC) on traffic and apply the profile to an interface. You can apply a CNP to most interfaces, including aggregated ethernet (AE) interfaces and their individual members.

A congestion notification profile can be configured to enable PFC on incoming traffic (input stanza) that matches the following:

- A Differentiated Services code point (DSCP) value in the Layer 3 IP header (for traffic that is not VLAN-tagged).
- An IEEE 802.1 code point at Layer 2 in the VLAN header (for VLAN-tagged traffic).

A congestion notification profile can be configured to enable PFC on outgoing traffic (output stanza) specified only by an IEEE 802.1 code point at Layer 2 in the VLAN header.

NOTE: You must configure PFC for FCoE traffic. Each interface that carries FCoE traffic should be configured for PFC on the FCoE code point (usually 011).

There is no limit to the total number of congestion notification profiles you can create. However:

- You can attach a maximum of one congestion notification profile to an interface.
- DSCP-based PFC and IEEE 802.1p PFC cannot be configured under the same congestion notification profile.

NOTE: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

Options

profile-name Name of the congestion notification profile.

pfc-
watchdogEnable the Priority Flow Control (PFC) watchdog. If you do not configure any options, the
default values are used.

- pfc-watchdog-action drop—When the PFC watchdog detects that a PFC queue has stalled, it drops all queued packets and all newly arriving packets for the stalled PFC queue. This option is the default.
- poll-interval *time*—How often the PFC watchdog checks the status of PFC queues. Configure the polling interval in milliseconds.
 - Default: 100
 - Range: 100-1000
- detection *number of polling intervals*—How many polling intervals the PFC watchdog waits before it determines that a PFC queue has stalled.
 - Default: 2
 - Range: 2-10
- recovery *time*—Configure in milliseconds how long the PFC watchdog disables the affected queues before it re-enables PFC.
 - Default: 200
 - Range: 200-10,000

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control-To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Support for DSCP values introduced in Junos OS Release 17.4R1 for the QFX Series.

pfc-watchdog option introduced in Junos OS Evolved Release 20.4R1 for the PTX10008.

RELATED DOCUMENTATION

Configuring CoS PFC (Congestion Notification Profiles) Understanding CoS Flow Control (Ethernet PAUSE and PFC) Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows Understanding PFC Using DSCP at Layer 3 for Untagged Traffic Configuring DSCP-based PFC for Layer 3 Untagged Traffic PFC Watchdog

drop-probability (Fill Level)

IN THIS SECTION

- Syntax | 258
- Hierarchy Level | 258
- Description | 258
- Options | 258
- Required Privilege Level | 258
- Release Information | 258

Syntax

drop-probability percentage;

Hierarchy Level

[edit class-of-service drop-profiles profile-name fill-level percentage]

Description

(EX8200 standalone switches and EX8200 Virtual Chassis only) Drop packets at the rate of the dropprobability value when the queue fills to the percentage configured with the fill-level value. This way you can manage network congestion.

Options

percentage-The probability (expressed in percentage) for a packet to be dropped from the queue.

• 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 before Junos OS 11.4.

RELATED DOCUMENTATION

Configuring CoS Tail Drop Profiles (CLI Procedure) | **165** Understanding Junos OS CoS Components for EX Series Switches | **8**

drop-profile-map

IN THIS SECTION

- Syntax | 259
- Hierarchy Level | 259
- Description | 259
- Options | 259
- Required Privilege Level | 259
- Release Information | 260

Syntax

drop-profile-map loss-priority loss-priority protocol protocol drop-profile profile-name;

Hierarchy Level

[edit class-of-service schedulers scheduler-name]

Description

Define the loss priority value for the specified drop profile.

Options

drop-profile profile-name – Name of the drop profile.

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 | 15 Defining CoS Schedulers and Scheduler Maps (CLI Procedure) | 138 Defining CoS Schedulers (J-Web Procedure) | 141 Understanding CoS Schedulers | 128

dscp

IN THIS SECTION

- Syntax | 260
- Hierarchy Level | 261
- Description | 261
- Options | 261
- Required Privilege Level | 261
- Release Information | 261

Syntax

```
dscp classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-points [ aliases ] [ 6-bit-patterns ];
        }
    }
}
```

Hierarchy Level

[edit class-of-service classifiers], [edit class-of-service code-point-aliases], [editclass-of-service interfaces interface-name unit logical-unit-number classifiers], [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

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 15
Defining CoS Code-Point Aliases (CLI Procedure) 65
Defining CoS Code-Point Aliases (J-Web Procedure) 66
Defining CoS Classifiers (CLI Procedure) 74
Defining CoS Classifiers (J-Web Procedure) 77
Defining CoS Rewrite Rules (CLI Procedure) 97
Defining CoS Rewrite Rules (J-Web Procedure) 100
Assigning CoS Components to Interfaces (CLI Procedure) 55
Assigning CoS Components to Interfaces (J-Web Procedure) 56

dscp-ipv6

IN THIS SECTION

- Syntax | 262
- Hierarchy Level | 262
- Description | 263
- Options | 263
- Required Privilege Level | 263
- Release Information | 263

Syntax

```
dscp-ipv6 classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-points [aliases] [6-bit-patterns];
        }
    }
}
```

Hierarchy Level

[edit class-of-service classifiers], [edit class-of-service code-point-aliases], [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] [edit class-of-service rewrite-rules]

Description

Define the Differentiated Services code point (DSCP) mapping that is applied to the IPv6 packets.

Options

classifier-name—Name of the classifier.

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 10.2.

RELATED DOCUMENTATION

xample: Configuring CoS on EX Series Switches 15	
efining CoS Code-Point Aliases (CLI Procedure) 65	
efining CoS Code-Point Aliases (J-Web Procedure) 66	
efining CoS Classifiers (CLI Procedure) 74	
efining CoS Classifiers (J-Web Procedure) 77	
efining CoS Rewrite Rules (CLI Procedure) 97	
efining CoS Rewrite Rules (J-Web Procedure) 100	
ssigning CoS Components to Interfaces (CLI Procedure) 55	
ssigning CoS Components to Interfaces (J-Web Procedure) 56	
nderstanding CoS Classifiers 70	

Understanding CoS Classifiers | 70

ethernet (CoS for Multidestination Traffic)

IN THIS SECTION

- Syntax | 264
- Hierarchy Level | 264
- Description | 264
- Required Privilege Level | 265
- Release Information | 265

Syntax

```
ethernet {
    broadcast forwarding-class-name;
}
```

Hierarchy Level

[edit class-of-service multi-destination family]

Description

Specify the Ethernet broadcast traffic family.

NOTE: On EX4300 switches, including switches in a virtual chassis or virtual chassis fabric configuration, you cannot explicitly configure multi-destination (broadcast, multicast, and destination lookup fail) traffic classification. By default, all multi-destination traffic uses output queue 8. This means all multicast (multi-destination) traffic always uses Q8 and there is no control to modify the queue.

The remaining statement is 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 9.5.

RELATED DOCUMENTATION

Understanding CoS Schedulers | 128

Understanding CoS Forwarding Classes | 110

Understanding CoS Classifiers | 70

excess-rate (Schedulers)

IN THIS SECTION

- Syntax | 265
- Hierarchy Level | 266
- Description | 266
- Default | 266
- Options | 266
- Required Privilege Level | 266
- Release Information | 266

Syntax

```
excess-rate {
    percent percentage;
}
```

Hierarchy Level

[edit class-of-service schedulers scheduler-name]

Description

(EX4300 switches only) Specify the percentage of excess bandwidth traffic to share.

Default

Excess bandwidth is shared in proportion to the configured transmit rate of each queue.

Options

• percent–Percentage of the excess bandwidth to share.

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 13.2X50-D10.

RELATED DOCUMENTATION

Understanding CoS Schedulers | 128

Defining CoS Schedulers and Scheduler Maps (CLI Procedure) | 138

Defining CoS Schedulers (J-Web Procedure) | 141

Example: Configuring CoS on EX Series Switches | 15

ехр

IN THIS SECTION

- Syntax | 267
- Hierarchy Level | 267
- Description | 267
- Options | 268
- Required Privilege Level | 268
- Release Information | 268

Syntax

```
exp classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-points [aliases] [3-bit-patterns];
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service classifiers],
[edit class-of-service code-point-aliases],
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules],
[edit class-of-service rewrite-rules]
```

Description

Define the experimental bits (EXP) code point mapping that is applied to MPLS packets. You can define an exp classifier only on EX3200 switches, EX4200 and EX8200 standalone switches, and EX8200

Virtual Chassis. You can bind an exp rewrite rule on EX8200 standalone switches and EX8200 Virtual Chassis.

EX Series switches support only one EXP code mapping on the switch (either default or custom). It is applied globally and implicitly to all the MPLS-enabled interfaces on the switch. You cannot bind it or disable it on individual interfaces.

Options

classifier-name—Name of the classifier.

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 10.1.

RELATED DOCUMENTATION

Understanding Using CoS with MPLS Networks on EX Series Switches Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect Configuring MPLS on Provider Edge Switches Using IP-Over-MPLS Configuring CoS on Provider Switches of an MPLS Network

explicit-congestion-notification

IN THIS SECTION

- Syntax | **269**
- Hierarchy Level | 269

- Description | 269
- Required Privilege Level | 269
- Release Information | 270

Syntax

explicit-congestion-notification;

Hierarchy Level

[edit class-of-service schedulers scheduler-name]

Description

Enable explicit congestion notification (ECN) on the output queue (forwarding class) or output queues (forwarding classes) mapped to the scheduler. 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 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.

A weighted random early detection (WRED) packet drop profile must be applied to the output queues on which ECN is enabled. ECN uses the WRED drop profile thresholds to mark packets when the output queue experiences congestion.

ECN reduces packet loss by forwarding ECN-capable packets during periods of network congestion instead of dropping those packets. (TCP notifies the network about congestion by dropping packets.) During periods of congestion, ECN marks ECN-capable packets that egress from congested queues. When the receiver receives an ECN packet that is marked as experiencing congestion, the receiver echoes the congestion state back to the sender. The sender then reduces its transmission rate to clear the congestion.

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 13.2X51.

Statement supported on SRX380, SRX300, SRX320, SRX340, SRX345, vSRX3.0, and NFX devices in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

Example: Configuring ECN

Understanding CoS Explicit Congestion Notification

family

IN THIS SECTION

- Syntax | 270
- Hierarchy Level | 271
- Description | 271
- Required Privilege Level | 271
- Release Information | 271

Syntax

```
family {
    ethernet {
        broadcast forwarding-class-name;
    }
    inet {
        classifiers{
            (dscp | inet-precedence) classifier-name;
        }
    }
}
```

Hierarchy Level

[edit class-of-service multi-destination]

Description

Specify the multidestination traffic family.

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 9.5.

RELATED DOCUMENTATION

Understanding CoS Schedulers | 128

Understanding CoS Forwarding Classes | 110

Understanding CoS Classifiers | 70

flow-control

IN THIS SECTION

- Syntax | 272
- Hierarchy Level | 272
- Description | 272
- Default | **273**
- Required Privilege Level | 273

Release Information | 273

Syntax

(flow-control | no-flow-control);

Hierarchy Level

[edit interfaces *interface-name* ether-options]

Description

Explicitly enable or disable symmetric Ethernet PAUSE flow control, which regulates the flow of packets from the switch to the remote side of the connection by pausing all traffic flows on a link during periods of network congestion. Symmetric flow control means that Ethernet PAUSE is enabled in both directions. The interface generates and sends Ethernet PAUSE messages when the receive buffers fill to a certain threshold and the interface responds to PAUSE messages received from the connected peer.

You can configure asymmetric flow control by including the configured-flow-control statement at the [edit interfaces *interface-name* ether-options hierarchy level. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.

NOTE: Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual "lanes" of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).

Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.

- flow-control—Enable flow control; flow control is useful when the remote device is a Gigabit Ethernet switch.
- no-flow-control—Disable flow control.

Default

Whether flow control is enabled or disabled by default depends on the platform. Run show interfaces *interface-name* to see whether flow control is enabled on an interface.

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

configured-flow-control

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches

Understanding CoS Flow Control (Ethernet PAUSE and PFC)

Junos OS Network Interfaces Library for Routing Devices

forwarding-class (Forwarding Policy)

IN THIS SECTION

- Syntax | **274**
- Hierarchy Level | 274
- Description | 274
- Options | 274
- Required Privilege Level | 275
- Release Information | 275

Syntax

```
forwarding-class class-name {
    discard;
    lsp-next-hop [ lsp-regular-expression ];
    next-hop [ next-hop-name];
    non-labelled-next-hop;
    non-lsp-next-hop;
    match-next-hop-forwarding-class;
}
```

Hierarchy Level

[edit class-of-service forwarding-policy next-hop-map map-name]
[edit class-of-service forwarding-policy class class-name classification-override]

Description

Define forwarding class name and associated next hops.

Options

class-name		Name of the forwarding class.
non-labelled-next-hop		Match any non-labelled next hop.
match-next-hop- class	-forwarding-	Match any pre-defined forwarding-class or a forwarding-class that is configured using CLI.
transport-class color <i>color</i>	· · · · · · · · · · · · · · · · · · ·	
		option is available only at [edit class-of-service forwarding-policy ap <i>next-hop-map-name</i> forwarding-class <i>forwarding-class-name</i>] <i>v</i> el.

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 before Junos OS Release 7.4.

Statement introduced for QFX10000 Series switches in Junos OS Release 17.1R1.

non-labelled-next-hop option introduced in Junos OS Release 19.1R1 for all platforms.

match-next-hop-forwarding-class option introduced in Junos OS Evolved Release 21.3R1 for PTX10008 platform.

RELATED DOCUMENTATION

Overriding the Input Classification forwarding-class-default (Forwarding Policy)

forwarding-classes

IN THIS SECTION

- SRX Series | 276
- QFX Series and OCX Series | 276
- EX Series (Except EX4300) | 276
- EX4300 | 277
- M320, MX Series, T Series, and PTX Series | 277
- Hierarchy Level | 277
- Description | 277
- Options | 279
- Required Privilege Level | 280
- Release Information | 280

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 multidestination (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 multidestination forwarding classes to a multidestination queue (8 through 11). The queue to which you map a forwarding class determines if the forwarding class is a unicast or multidestination forwarding class.

Switches that use the same forwarding classes for unicast and multidestination 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 multidestination) 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 <i>class-name</i>	Define t	he forwarding class name.	
queue-num <i>queue-</i> <i>number</i>	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 | 281
- Hierarchy Level | 281
- Description | 281
- Options | **282**
- Required Privilege Level | 282
- Release Information | 282

Syntax

```
ieee-802.1 classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-points [ aliases ] [ 6 bit-patterns ];
        }
    }
}
```

Hierarchy Level

[edit class-of-service classifiers], [edit class-of-service code-point-aliases], [editclass-of-service interfaces interface-name unit logical-unit-number classifiers], [edit class-of-service rewrite-rules]

Description

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

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 15
Defining CoS Classifiers (CLI Procedure) 74
Defining CoS Classifiers (J-Web Procedure) 77
Defining CoS Code-Point Aliases (CLI Procedure) 65
Defining CoS Code-Point Aliases (J-Web Procedure) 66
Defining CoS Rewrite Rules (CLI Procedure) 97
Defining CoS Rewrite Rules (J-Web Procedure) 100
Understanding CoS Classifiers 70
Understanding CoS Rewrite Rules 95

ieee-802.1 (Congestion Notification)

IN THIS SECTION

- Syntax | 283
- Hierarchy Level | 283
- Description | 283

- Required Privilege Level | 283
- Release Information | 283

Syntax

```
ieee-802.1 {
    code-point up-bits pfc ;
}
```

Hierarchy Level

[edit class-of-service congestion-notification-profile profile-name], [edit class-of-service interfaces interface-name congestion-notification-profile profile-name]

Description

Set an association between the traffic class and the congestion notification profile.

The remaining statement is explained separately. See CLI Explorer.

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 10.4.

RELATED DOCUMENTATION

Example: Configuring an FCoE Transit Switch Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure) | 123

import

IN THIS SECTION

- Syntax | 284
- Hierarchy Level | 284
- Description | 284
- Options | 284
- Required Privilege Level | 285
- Release Information | 285

Syntax

import (classifier-name | default);

Hierarchy Level

[edit class-of-service classifiers (dscp | ieee-802.1 | inet-precedence) classifier-name],

[edit class-of-service rewrite-rules (dscp | ieee-802.1 | inet-precedence) rewrite-name]

Description

Specify a default or previously defined classifier.

Options

classifier-name – Name of the classifier mapping configured at the [edit class-of-service classifiers] hierarchy level.

default-Default classifier mapping.

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 | 15

Defining CoS Classifiers (CLI Procedure) | 74

Defining CoS Classifiers (J-Web Procedure) | 77

Defining CoS Rewrite Rules (CLI Procedure) | 97

Defining CoS Rewrite Rules (J-Web Procedure) | 100

Understanding CoS Classifiers | 70

Understanding CoS Rewrite Rules | 95

inet (CoS)

IN THIS SECTION

- Syntax | 286
- Hierarchy Level | 286
- Description | 286
- Required Privilege Level | 286
- Release Information | 286

Syntax

```
inet {
    classifiers {
        (dscp | inet-precedence) classifier-name ;
    }
}
```

Hierarchy Level

[edit class-of-service multi-destination family]

Description

Specify the IP multicast family.

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

Option inet introduced in Junos OS Release 9.5 for EX Series switches.

The remaining statements are explained separately. See CLI Explorer.

RELATED DOCUMENTATION

Understanding CoS Schedulers | 128

Understanding CoS Forwarding Classes | 110

Understanding CoS Classifiers | 70

IN THIS SECTION

- Syntax | 287
- Hierarchy Level | 287
- Description | 287
- Required Privilege Level | 287
- Release Information | 288

Syntax

```
inet6 {
    classifiers {
        dscp-ipv6 classifier-name;
    }
}
```

Hierarchy Level

[edit class-of-service multi-destination family]

Description

(EX8200 standalone switches and EX8200 Virtual Chassis only) Specify the IPv6 multicast family.

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

Option inet6 introduced in Junos OS Release before Junos OS 11.4 for EX Series switches.

RELATED DOCUMENTATION

Understanding CoS Schedulers | 128 Understanding CoS Forwarding Classes | 110

Understanding CoS Classifiers | 70

inet-precedence

IN THIS SECTION

- Syntax | 288
- Hierarchy Level | 289
- Description | 289
- Options | 289
- Required Privilege Level | 289
- Release Information | 289

Syntax

```
inet-precedence classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
        loss-priority level {
            code-points [ aliases ] [ 6-bit-patterns ];
        }
    }
}
```

Hierarchy Level

[edit class-of-service classifiers], [edit class-of-service code-point-aliases], [editclass-of-service interfaces interface-name unit logical-unit-number classifiers], [edit class-of-service rewrite-rules]

Description

Apply an IPv4 precedence rewrite rule.

Options

classifier-name—Name of the classifier.

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 15
Defining CoS Classifiers (CLI Procedure) 74
Defining CoS Classifiers (J-Web Procedure) 77
Defining CoS Code-Point Aliases (CLI Procedure) 65
Defining CoS Code-Point Aliases (J-Web Procedure) 66
Defining CoS Rewrite Rules (CLI Procedure) 97
Defining CoS Rewrite Rules (J-Web Procedure) 100
Understanding CoS Classifiers 70
Understanding CoS Rewrite Rules 95

input (Congestion Notification)

IN THIS SECTION

- Syntax | 290
- Hierarchy Level | 290
- Description | 290
- Required Privilege Level | 290
- Release Information | 291

Syntax

```
input {
    ieee-802.1 {
        code-point up-bits pfc ;
    }
}
```

Hierarchy Level

```
[edit class-of-service congestion-notification-profile profile-name],
[edit class-of-service interfaces interface-name congestion-notification-profile profile-name]
```

Description

Identify the three-bit pattern of the User Priority field that triggers the priority-based congestion notification profile for a specified traffic class.

The remaining statements are explained separately. See CLI Explorer.

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 10.4.

RELATED DOCUMENTATION

Example: Configuring an FCoE Transit Switch

Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure) | 123

interfaces

IN THIS SECTION

- Syntax | 291
- Hierarchy Level | 292
- Description | 292
- Options | **292**
- Required Privilege Level | 292
- Release Information | 292

Syntax

```
forwarding-class class-name;
classifiers {
    (dscp | ieee-802.1 | inet-precedence) (classifier-name | default);
    }
}
}
```

Hierarchy Level

[edit class-of-service]

Description

Configure interface-specific class-of-service (CoS) properties for incoming packets.

Options

interface-name Name of the interface.

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 Classifiers (CLI Procedure) Defining CoS Classifiers (J-Web Procedure) Defining CoS Forwarding Classes (CLI Procedure) Defining CoS Forwarding Classes (J-Web Procedure) Defining CoS Schedulers and Scheduler Maps (CLI Procedure)

Defining CoS Schedulers (J-Web Procedure)

Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure) | 123

loss-priority (Classifiers and Rewrite Rules)

IN THIS SECTION

- Syntax | 293
- Hierarchy Level | 293
- Description | 293
- Options | 294
- Required Privilege Level | 294
- Release Information | 294

Syntax

```
loss-priority level {
    code-points [aliases] [6-bit-patterns / 3-bit-patterns];
}
```

Hierarchy Level

```
[edit class-of-service classifiers (dscp | ieee-802.1 | inet-precedence | exp) classifier-name
forwarding-class class-name],
[edit class-of-service rewrite-rules (dscp | ieee-802.1 | inet-precedence | exp) rewrite-rule-name
forwarding-class class-name]
```

Description

Specify packet loss priority value for a specific set of code-point aliases and bit patterns.

Options

level—Can be one of the following:

- high—Packet has high loss priority.
- medium-high— (On EX3200, EX4200, EX4300, and EX4500 switches only) Code points to classify to loss priority medium-high.
- low—Packet has low loss priority.
- medium-low (On EX3200, EX4200, EX4300, and EX4500 switches only) Code points to classify to loss priority medium-low.

NOTE: The EX4300 hardware supports only three levels of loss priority — high, medium-high, and low. You can configure a loss priority of medium-low on the EX4300, but the hardware will convert it to medium-high.

The remaining statement is 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.

Statement expanded to apply to EXP classifiers in Junos OS Release 10.1 for EX Series switches.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches 15	
Defining CoS Classifiers (CLI Procedure) 74	
Defining CoS Classifiers (J-Web Procedure) 77	
Defining CoS Rewrite Rules (CLI Procedure) 97	
Defining CoS Rewrite Rules (J-Web Procedure) 100	
Understanding CoS Classifiers 70	
Understanding CoS Rewrite Rules 1 95	

multi-destination

IN THIS SECTION

- EX4300 | 295
- EX8200 | 295
- Hierarchy Level | 296
- Description | 296
- Required Privilege Level | 296
- Release Information | 296

EX4300

```
multi-desitnation
    classifiers
        (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence | exp) classifier-name {
        {
        }
    }
}
```

EX8200

```
dscp-ipv6 classifier-name;
    }
    scheduler-map map-name;
}
```

Hierarchy Level

[edit class-of-service]
[edit class-of-service interfaces interface-name] (EX4300 only)

Description

(EX8200 standalone switches, EX8200 Virtual Chassis, and EX4300 only) Define the CoS configuration for multidestination traffic.

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 9.5.

RELATED DOCUMENTATION

Understanding CoS Schedulers | 128

Understanding CoS Forwarding Classes | 110

Understanding CoS Classifiers | 70

policing

IN THIS SECTION

- Syntax | 297
- Hierarchy Level | 297
- Description | 297
- Options | 297
- Required Privilege Level | 298
- Release Information | 298

Syntax

policing (filter filter-name | no-automatic-policing);

Hierarchy Level

[edit protocols mpls label-switched-path lsp-name]
[edit interfaces interface-id unit number-of-logical-unit family inet address ip-address]

Description

Apply a rate-limiting policer as the specified policing filter:

- To the LSP for MPLS over CCC.
- To the customer-edge interface for IP over MPLS.

Options

filter filter-name-Specify the name of the policing filter.

no-automatic-policing-Disable automatic policing on this LSP.

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 10.1 .

RELATED DOCUMENTATION

Configuring Policers to Control Traffic Rates (CLI Procedure) Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS

priority (Schedulers)

IN THIS SECTION

- Syntax | 298
- Hierarchy Level | 299
- Description | 299
- Options | **299**
- Required Privilege Level | 299
- Release Information | 299

Syntax

priority priority;

Hierarchy Level

[edit class-of-service schedulers scheduler-name]

Description

Specify packet-scheduling priority value.

Options

priority—It can be one of the following:

- low–Scheduler has low priority.
- **strict-high**—Scheduler has strictly high priority.

NOTE: On EX4400 switches, applying strict-high priority schedulers to queues 0 through 3 also applies strict-high priority to queues 8 through 11. Therefore, Juniper recommends applying strict-high priority schedulers only to queues 4 through 7.

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 | 15

Defining CoS Schedulers and Scheduler Maps (CLI Procedure) | 138

Defining CoS Schedulers (J-Web Procedure) | 141

Understanding CoS Schedulers | 128

protocol (Drop Profiles)

IN THIS SECTION

- Syntax | 300
- Hierarchy Level | 300
- Description | 300
- Options | 300
- Required Privilege Level | 300
- Release Information | 301

Syntax

protocol protocol drop-profile profile-name;

Hierarchy Level

[edit class-of-service schedulers scheduler-name]

Description

Specify the protocol type for the specified drop profile.

Options

drop-profile profile-name – Name of the drop profile.

protocol –Type of protocol. It can be:

• **any**-Accept any protocol type.

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 | **15** Configuring CoS Tail Drop Profiles (CLI Procedure) | **165** Understanding CoS Tail Drop Profiles | **164**

rewrite-rules

IN THIS SECTION

- Syntax | 301
- Hierarchy Level | 302
- Description | 302
- Required Privilege Level | 302
- Release Information | 302

Syntax

```
rewrite-rules {
    (dscp | dscp-ipv6 | exp |ieee-802.1 | inet-precedence ) rewrite-name {
        import ( default | rewrite-name);
        forwarding-class class-name {
            loss-priority level code-point (alias | bits);
        }
    }
}
```

Hierarchy Level

[edit class-of-service]

Description

Specify a rewrite-rules mapping for the traffic that passes through all queues on the interface.

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.

Statement expanded for use with global EXP classifiers in Junos OS Release 10.1 for EX Series switches.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches | 15 Defining CoS Rewrite Rules (CLI Procedure) | 97 Defining CoS Rewrite Rules (J-Web Procedure) | 100 Understanding CoS Rewrite Rules | 95 Understanding Using CoS with MPLS Networks on EX Series Switches

scheduler-map

IN THIS SECTION

Syntax | 303

- Hierarchy Level | 303
- Description | 303
- Options | **303**
- Required Privilege Level | 303
- Release Information | 303

Syntax

scheduler-map map-name;

Hierarchy Level

[edit class-of-service interfaces],
[edit class-of-service multi-destination]

Description

Associate a scheduler map name with an interface or with a multidestination traffic configuration.

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 9.0.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches 15	
Assigning CoS Components to Interfaces (CLI Procedure) 55	
Assigning CoS Components to Interfaces (J-Web Procedure) 56	
Understanding CoS Schedulers 128	
Understanding CoS Classifiers 70	

scheduler-maps

IN THIS SECTION

- Syntax | 304
- Hierarchy Level | 304
- Description | 305
- Options | 305
- Required Privilege Level | 305
- Release Information | 305

Syntax

```
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
```

Hierarchy Level

[edit class-of-service]

Description

Specify a scheduler map name and associate it with the scheduler configuration and forwarding class.

Options

map-name – Name of the scheduler map.

The remaining statement is 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 15
Defining CoS Forwarding Classes (CLI Procedure) 114
Defining CoS Forwarding Classes (J-Web Procedure) 115
Understanding CoS Schedulers 128
Understanding CoS Forwarding Classes 110

schedulers (CoS)

IN THIS SECTION

- Syntax | 306
- Hierarchy Level | 306
- Description | 306
- Options | **306**

- Required Privilege Level | 306
- Release Information | 307

Syntax

```
schedulers {
    scheduler-name {
        buffer-size (percent percentage | remainder);
        drop-profile-map loss-priority loss-priority protocol protocol drop-profile profile-name;
        excess-rate (percent percentage);
        explicit-congestion-notification;
        priority priority;
        shaping-rate (rate | percent percentage);
        transmit-rate (rate | percent percentage | remainder);
    }
}
```

Hierarchy Level

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

Description

Specify scheduler name and parameter values.

Options

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 9.0.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches | 15 Defining CoS Schedulers and Scheduler Maps (CLI Procedure) | 138 Defining CoS Schedulers (J-Web Procedure) | 141 Understanding CoS Schedulers | 128

shaping-rate

IN THIS SECTION

- Syntax | 307
- Hierarchy Level | 307
- Description | 308
- Default | 308
- Options | 308
- Required Privilege Level | 308
- Release Information | 308

Syntax

shaping-rate (percent percentage | rate);

Hierarchy Level

[edit class-of-service schedulers scheduler-name]

Description

Configure shaping rate to throttle the rate at which queues transmit packets.

We recommend that you configure the shaping rate as an absolute maximum usage and not as additional usage beyond the configured transmit rate.

Default

If you do not include this statement, the default shaping rate is 100 percent, which is the same as no shaping at all.

Options

percent percentage – Shaping rate as a percentage of the available interface bandwidth.

• Range: 0 through 100 percent

rate—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

• Range: 3200 through 32,000,000,000 bps

(EX4300 switches only) 8000 through 160,000,000,000 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 9.3.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches | **15** Understanding Junos OS CoS Components for EX Series Switches | **8**

shaping-rate (Applying to an Interface)

IN THIS SECTION

- Syntax | 309
- Hierarchy Level | 309
- Description | 309
- Default | 310
- Options | 311
- Required Privilege Level | 311
- Release Information | 311

Syntax

shaping-rate rate;

Hierarchy Level

```
[edit class-of-service interfaces interface-name],
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

For logical interfaces on which you configure packet scheduling, configure traffic shaping by specifying the amount of bandwidth to be allocated to the logical interface. Applying a shaping rate can help ensure that higher-priority services do not starve lower-priority services.

For physical interfaces, by default, shaping is not configured and traffic can be up to the line rate for that interface. Port shaping enables you to control the amount of traffic passing through a physical interface. Port shaping enables you to shape the aggregate traffic through an interface to a rate that is less than the line rate for that interface. This can be useful to reduce downstream congestion.

NOTE: For platforms that support both logical and physical interface shaping, logical and physical interface traffic shaping rates are mutually exclusive. This means you can include the shaping-rate statement at the [edit class-of-service interfaces *interface-name*] hierarchy level or at the [edit class-of-service interface-name unit *logical-unit-number*] hierarchy level, but not at both.

NOTE: For MX Series routers and for EX Series switches, the shaping rate value for the physical interface at the [edit class-of-service interfaces *interface-name*] hierarchy level must be a minimum of 160 Kbps. If the value is less than the sum of the logical interface guaranteed rates, you cannot apply the shaping rate to a physical interface.

For PTX Series routers, the shaping rate value for the physical interface at the [edit class-ofservice interfaces *interface-name*] hierarchy level must be a minimum of 1 Gbps and an incremental granularity of 0.1 percent of the physical interface speed after that (for example, 10 Mbps increments on a 10 Gbps interface).

On EX4650, QFX5110, QFX5120, QFX5200, QFX5210 Series switches, when you configure a shaping rate on an aggregated Ethernet (ae) interface, all members of the ae interface are shaped at the configured shaping rate. For example, consider an interface ae0 that consists of three interfaces: xe-0/0/0, xe-0/0/1, and xe-0/0/2. If you configure a shaping rate of X Mpbs on ae0, traffic up to the rate of X Mpbs flows through each of the three interfaces. Therefore, the total traffic flowing through ae0 can be at the rate of 3X Mbps.

Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the shaping-rate statement at the [edit class-of-service traffic-control-profiles] hierarchy level. With this configuration approach, you can independently control the delay-buffer rate, as described in *Oversubscribing Interface Bandwidth*.

Default

If you do not include this statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number*] hierarchy level, the default logical interface bandwidth is the average of unused bandwidth for the number of logical interfaces that require default bandwidth treatment. If you do not include this statement at the [edit class-of-service interfaces *interface-name*] hierarchy level, the default physical interface bandwidth for the number of physical interfaces that require default bandwidth for the number of physical interfaces that require default bandwidth for the number of physical interfaces that require default bandwidth for the number of physical interfaces that require default bandwidth for the number of physical interfaces that require default bandwidth treatment.

Options

- *rate* Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).
 - Range: 1000 through 6,400,000,000,000 bps.

NOTE: For all MX Series and EX Series interfaces, the rate can be from 65,535 to 6,400,000,000,000 bps.

For all PTX Series interfaces, the rate can be from 1,000,000,000 to 160,000,000 bps in increments of 0.1 percent of the interface speed.

Required Privilege Level

interface—To view this statement in the configuration.

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

Release Information

Statement introduced before Junos OS Release 7.4.

[edit class-of-service interfaces interface-name] hierarchy level added in Junos OS Release 7.5.

RELATED DOCUMENTATION

Applying Scheduler Maps Overview Configuring Virtual LAN Queuing and Shaping on PTX Series Routers CoS Port Shaping

shared-buffer

IN THIS SECTION

- Syntax | 312
- Hierarchy Level | 312
- Description | 312
- Options | 312
- Required Privilege Level | 312
- Release Information | 313

Syntax

shared-buffer percent percentage

Hierarchy Level

[edit class-of-service],

Description

Configure the buffer allocation for the shared buffer pool.

Options

percent *percentage* Size of the shared buffer as a percentage of the buffer allocated to the shared buffer pool.

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 10.1.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches | 15

Understanding Junos OS CoS Components for EX Series Switches | 8

transmit-rate (EX Series Switches)

IN THIS SECTION

- Syntax | 313
- Hierarchy Level | 313
- Description | 314
- Default | 314
- Options | 314
- Required Privilege Level | 314
- Release Information | 314

Syntax

transmit-rate (rate | percent percentage | remainder);

Hierarchy Level

[edit class-of-service schedulers scheduler-name]

Description

Specify the transmit rate or percentage for a scheduler.

Default

If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 0, 0, 0, 0, and 5 percent.

Options

rate	Transmission rate, in bps. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).
	• Range: 3200 through 160,000,000 bps
	(EX4300 switches only) 8000 through 160,000,000,000 bps
percent <i>percentage</i>	Percentage of transmission capacity. A percentage of zero drops all packets in the queue.
	Range: 0 through 100 percent
remainder	Remaining rate available

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 15
Defining CoS Schedulers and Scheduler Maps (CLI Procedure) 138
Defining CoS Schedulers (J-Web Procedure) 141
Understanding CoS Schedulers 128

tri-color (EX Series Switches)

IN THIS SECTION

- Syntax | 315
- Hierarchy Level | 318
- Description | 318
- Options | 318
- Required Privilege Level | 318
- Release Information | 318

Syntax

```
tri-color {
    classifiers {
        (dscp | dscp-ipv6 |exp |ieee-802.1 | inet-precedence) classifier-name {
            forwarding-class (Forwarding Policy) (class-name | assured-forwarding | best-effort |
expedited-forwarding | network-control) {
                loss-priority (high | low | medium-high | medium-low) {
                    code-points [aliases] [6 bit-patterns];
                }
            }
            import (classifier-name | default);
        }
    }
    code-point-aliases {
        (dscp | dscp-ipv6 | exp |ieee-802.1 | inet-precedence) {
            alias-name bits;
        }
    }
    drop-profiles {
        profile-name {
            fill-level percentage drop-probability percentage;
            interpolate {
                drop-probability [values];
                fill-level [values]
            }
```

```
}
   }
    forwarding-classes {
        class class-name {
            priority (high | low);
            queue-num queue-number;
        }
        queue queue-number;
   }
    host-outbound-traffic {
        forwarding-class class-name;
        dscp-code-point value;
   }
    interfaces {
        interface-name {
            congestion-notification-profile profile-name {
                input {
                     ieee-802.1 {
                         code-point up-bits pfc;
                         }
                    }
                }
            }
            scheduler-map map-name;
            shaping-rate;
            unit (logical-unit-number | * ) {
                classifiers {
                     (dscp | dscp-ipv6|ieee-802.1 | inet-precedence) (classifier-name | default);
                }
                forwarding-class (Forwarding Policy) (class-name | assured-forwarding | best-effort
| expedited-forwarding | network-control) ; {
                     rewrite-rules {
                         (dscp | dscp-ipv6|ieee-802.1 | inet-precedence) (rewrite-rule-name |
default);
                    }
                     classifiers {
                         (dscp | dscp-ipv6|ieee-802.1 | inet-precedence) classifier-name;
                     }
                }
                rewrite-rules {
                     (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) rewrite-rule-name;
                }
            }
```

```
}
    }
    rewrite-rules {
        (dscp | dscp-ipv6 | exp |ieee-802.1 | inet-precedence) rewrite-rule-name {
            forwarding-class (Forwarding Policy) (class-name | assured-forwarding | best-effort |
expedited-forwarding | network-control) {
                loss-priority (high | low | medium-high | medium-low) code-point (alias | bits);
            }
            import (rewrite-rule-name | default);
        }
    }
    scheduler-maps {
        map-name {
            forwarding-class (Forwarding Policy) (class-name | assured-forwarding | best-effort |
expedited-forwarding | network-control) {
                scheduler scheduler-name;
            }
        }
    }
    schedulers {
        scheduler-name {
            buffer-size (exact | percent percentage | remainder | temporal);
            drop-profile-map {
                loss-priority (high | low);
                protocol any;
            }
            priority (low | strict-high);
            shaping-rate (rate | percent percentage);
            transmit-rate (EX Series Switches) (rate | percent percentage | remainder);
        }
    }
    shared-buffer {
        percent percentage;
    }
    traceoptions {
        file (file-name | files files | match match | no-world-readable | size size | world-
readable);
        flag ( all | asynch | chassis-scheduler | cos-adjustment | dynamic | hardware-database |
init | parse | performance-monitor | process | restart | route-socket | show | snmp | util);
        no-remote-trace;
    }
}
```

Hierarchy Level

[edit class-of-service]

Description

(EX4500 and EX4550 switches only) Enable tricolor marking.

Options

percent *percentage*—Size of the shared buffer as a percentage of the buffer allocated to the shared buffer pool.

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 11.4 .

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches | **15** Understanding Junos OS CoS Components for EX Series Switches | **8**

unit

IN THIS SECTION

- Syntax | 319
- Hierarchy Level | 319
- Description | 319



- Required Privilege Level | 319
- Release Information | 320

Syntax

```
unit logical-unit-number {
   forwarding-class class-name;
   classifiers {
      (dscp | ieee-802.1 | inet-precedence) (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 be able to use the physical device.

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 9.0.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches | 15

Assigning CoS Components to Interfaces (CLI Procedure) | 55

Assigning CoS Components to Interfaces (J-Web Procedure) | 56

Operational Commands

IN THIS CHAPTER

- show class-of-service | 322
- show class-of-service classifier | 330
- show class-of-service code-point-aliases | 333
- show class-of-service drop-profile | 336
- show class-of-service forwarding-class | 340
- show class-of-service forwarding-table | 345
- show class-of-service forwarding-table classifier | 350
- show class-of-service forwarding-table classifier mapping | 352
- show class-of-service forwarding-table drop-profile | 355
- show class-of-service forwarding-table rewrite-rule | 357
- show class-of-service forwarding-table rewrite-rule mapping | 359
- show class-of-service forwarding-table scheduler-map | 361
- show class-of-service interface | 365
- show class-of-service multi-destination | 408
- show class-of-service rewrite-rule | 411
- show class-of-service scheduler-map | 414
- show interfaces queue | 418
- show pfe statistics traffic cpu | 478
- show pfe statistics traffic egress-queues | 483
- show pfe statistics traffic multicast | 485

show class-of-service

IN THIS SECTION

- Syntax | 322
- Description | 322
- Required Privilege Level | 322
- Output Fields | 322
- Sample Output | 325
- Release Information | 329

Syntax

show class-of-service

Description

Display the class-of-service (CoS) information.

Required Privilege Level

view

Output Fields

Table 39 on page 323 lists the output fields for the show class-of-service command. Output fields are listed in the approximate order in which they appear.

Table 39: show class-of-service Output Fields

Field Name	Field Description	Level of Output
Forwarding class	 The forwarding class configuration: Forwarding class—Name of the forwarding class. ID—Forwarding class ID. Queue—Queue number. Fabric Priority—(EX8200 switches only) Fabric priority: either high or low. The fabric priority determines which CoS ingress queues packets are sent to. 	All levels
Code point type	 The type of code-point alias: dscp-Aliases for DiffServ code point (DSCP) values. ieee-802.1-Aliases for IEEE 802.1p values. inet-precedence-Aliases for IP precedence values. exp-Aliases for experimental (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
Loss priority	Loss priority assigned to specific CoS values and aliases of the classifier.	All levels
Rewrite rule	Name of the rewrite-rule.	All levels
Drop profile	Name of the drop profile.	All levels

Field Name	Field Description	Level of Output
Туре	Type of drop profile. EX Series switches support only the discrete type of drop profile.	All levels
Fill level	Percentage of queue buffer fullness of <i>high</i> packets beyond which <i>high</i> packets are dropped.	All levels
Scheduler	Name of the scheduler.	All levels
Transmit rate	Transmission rate of the scheduler.	All levels
Excess rate	Percentage of excess bandwidth traffic to share.	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
Physical interface	Name of the physical interface.	All levels
Scheduler map	Name of the scheduler map.	All levels
Index	Internal index of a specific object.	All levels

Table 39: show class-of-service Output Fields (Continued)

Sample Output

show class-of- service

user@switch> show c	lass-of-service			
Forwarding class		ID	Queue	
best-effort		0	0	
expedited-forward	ing	1	5	
assured-forwardin		2	1	
network-control	-	3	7	
Code point type: ds	ср			
Alias	Bit pattern			
af11	001010			
af12	001100			
Code point type: ie	ee-802.1			
Alias	Bit pattern			
af11	010			
Code point type: in	et-precedence			
Alias	Bit pattern			
af11	001			
Classifier: dscp-de		ype: ds	scp, Index:	7
Code point	Forwarding class			Loss priority
000000	best-effort			low
000001	best-effort			low
Classifier: ieee802		oint typ	be: ieee-80	2.1, Index: 11
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	network-control			low

	111	network-control		low
C	Classifier: ipprec-d	efault, Code poin	t type: inet-precede	ence, Index: 12
	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	network-control		low
	111	network-control		low
C	Classifier: ieee8021	p-untrust, Code p	oint 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
R	Rewrite rule: dscp-de	efault, Code poin	t type: dscp, Index:	: 27
	Forwarding class		Loss priority	Code point
	best-effort		low	000000
	best-effort		high	000000
	expedited-forwardi	ng	low	101110
	expedited-forwardi	ng	high	101110
	assured-forwarding		low	001010
	assured-forwarding		high	001100
	network-control		low	110000
	network-control		high	111000
R	Rewrite rule: ieee80	21p-default, Code	point type: ieee-80	02.1, Index: 30
	Forwarding class		Loss priority	Code point
	best-effort		low	000
	best-effort		high	001
	expedited-forwardi	ng	low	100
	expedited-forwardi	ng	high	101
	assured-forwarding		low	010
	assured-forwarding		high	011

```
network-control
                                      low
                                                           110
                                      high
  network-control
                                                           111
Rewrite rule: ipprec-default, Code point type: inet-precedence, Index: 31
 Forwarding class
                                      Loss priority
                                                           Code point
 best-effort
                                                           000
                                      low
 best-effort
                                      high
                                                           000
 expedited-forwarding
                                      low
                                                           101
 expedited-forwarding
                                      high
                                                           101
 assured-forwarding
                                      low
                                                           001
 assured-forwarding
                                      high
                                                           001
 network-control
                                      low
                                                          110
 network-control
                                      high
                                                           111
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: 20
   Transmit rate: 95 percent, Rate Limit: none, Buffer size: 95 percent,
   Priority: low
   Drop profiles:
     Loss priority Protocol
                                  Index
                                           Name
      High
                      non-TCP
                                      1
                                           <default-drop-profile>
                      TCP
      High
                                      1
                                           <default-drop-profile>
 Scheduler: <default-nc>, Forwarding class: network-control, Index: 22
   Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent,
   Priority: low
   Drop profiles:
      Loss priority
                     Protocol
                                  Index
                                           Name
      High
                      non-TCP
                                      1
                                           <default-drop-profile>
      High
                      TCP
                                      1
                                           <default-drop-profile>
Physical interface: ge-0/0/0, Index: 129
Queues supported: 8, Queues in use: 4
 Scheduler map: <default>, Index: 2
Physical interface: ge-0/0/1, Index: 130
Queues supported: 8, Queues in use: 4
 Scheduler map: <default>, Index: 2
```

Fabric priority: low Scheduler: <defaul Drop profiles:</defaul 		ndex: 23	
Loss priority	Protocol	Index	Name
High	non-TCP	1	<default-drop-profile></default-drop-profile>
High	ТСР	1	<default-drop-profile></default-drop-profile>
Fabric priority: hig Scheduler: <defaul Drop profiles:</defaul 		ndex: 23	
Loss priority	Protocol	Index	Name
High	non-TCP	1	<default-drop-profile></default-drop-profile>
High	TCP	1	<default-drop-profile></default-drop-profile>

show class-of-service rewrite-rule

<pre>user@switch> show class-of-service rewrite-rule</pre>				
Rewrite rule: dscp-default, Code po	oint type: dscp, Inde	x: 31		
Forwarding class	Loss priority	Code point		
best-effort	low	000000		
best-effort	high	000000		
expedited-forwarding	low	101110		
expedited-forwarding	high	101110		
fw-class	low	001010		
fw-class	high	001100		
network-control	low	110000		
network-control	high	111000		
Rewrite rule: exp-default, Code poi	int type: exp, Index:	33		
Forwarding class	Loss priority	Code point		
best-effort	low	000		
best-effort	high	001		
expedited-forwarding	low	010		
expedited-forwarding	high	011		
fw-class	low	100		
fw-class	high	101		
network-control	low	110		

network-control	high	111
Rewrite rule: ieee8021p-default, C	ode point type: ieee	-802.1, Index: 34
Forwarding class	Loss priority	Code point
best-effort	low	000
best-effort	high	001
expedited-forwarding	low	010
expedited-forwarding	high	011
fw-class	low	100
fw-class	high	101
network-control	low	110
network-control	high	111
Rewrite rule: ipprec-default, Code	point type: inet-pro	ecedence, Index: 35
Forwarding class	Loss priority	Code point
best-effort	low	000
best-effort	high	000
expedited-forwarding	low	101
expedited-forwarding	high	101
fw-class	low	001
fw-class	high	001
network-control	low	110
network-control	high	111

Release Information

Command introduced in Junos OS Release 9.0.

EXP classifiers added in Junos OS Release 10.1 for EX Series switches.

RELATED DOCUMENTATION

Example: Configuring CoS on EX Series Switches 15	
Monitoring CoS Value Aliases 67	
Monitoring CoS Classifiers 88	
Monitoring CoS Forwarding Classes 117	
Monitoring CoS Scheduler Maps 147	
Monitoring CoS Rewrite Rules 107	

show class-of-service classifier

IN THIS SECTION

- Syntax | 330
- Description | 330
- Options | 330
- Required Privilege Level | 331
- Output Fields | 331
- Sample Output | 332
- Release Information | 333

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 40 on page 331 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 40: 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: iee	ee802.1-default, Code point ty	pe: 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	1 network-control high				
Classifier: use	ers-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: ieee	8021p-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: ieee	8021p-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: i	eee-mcast, Code point type: i	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 | 334
- Description | 334
- Options | **334**
- Required Privilege Level | 334
- Output Fields | 334
- Sample Output | 335
- Release Information | 335

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.	
ехр	(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 41 on page 334 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 41: 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).

Field Name	Field Description
Alias	Alias for a bit pattern.
Bit pattern	Bit pattern for which the alias is displayed.

Table 41: show class-of-service code-point-aliases Output Fields (Continued)

Sample Output

show class-of-service code-point-aliases exp

us	ser@host>	show class-of-service	code-point-aliases exp
Сс	ode 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 drop-profile

IN THIS SECTION

- Syntax | 336
- Description | 336
- Options | 336
- Required Privilege Level | 336
- Output Fields | 337
- Sample Output | 337
- Release Information | 340

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 42 on page 337 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 42: show class-of-service drop-profile Output Fields

Field Name	Field Description	
Drop profile	Name of a drop profile.	
Туре	 Type of drop profile: 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 entri	es 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</default-drop-profile>			
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	
129114911591169118912091	
14 91 15 91 16 91 18 91 20 91	
15 91 16 91 18 91 20 91	
16 91 18 91 20 91	
18 91 20 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-class

IN THIS SECTION

- Syntax | 341
- Description | 341
- Required Privilege Level | 341
- Output Fields | 341
- Sample Output (Junos OS) | 342
- Sample Output (Junos OS Evolved) | 344
- Release Information | 345

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 43 on page 341 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).

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</i> <i>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.	

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	 Packet loss attribute to differentiate lossless forwarding classes from lossy forwarding classes: 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	For DSCP-based PFC, the explicitly configured PFC priority configured for the forwarding class. 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.

Table 43: show class-of-service forwarding-class Output Fields (Continued)

Sample Output (Junos OS)

show class-of-service forwarding-class (ACX Series, EX Series, MX Series, PTX Series)

user@host> show class-o	of-se	rvice f	orwarding-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

user@host> show class-of-service forwarding-class Forwarding class ID Queue Policing priority SPU priority best-effort 0 0 normal low expedited-forwarding normal 1 1 low 2 2 assured-forwarding normal low network-control 3 3 normal low

show class-of-service forwarding-class (SRX Series)

show class-of-service forwarding-class (EX8200 Switch)

user@switch> <pre>show class-of-service</pre>	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	

user@switch> show cl	ass-of-ser	vice for	warding-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	

show class-of-service forwarding-class (QFX Series with DSCP-based PFC)

On switches that do not use different forwarding classes and output queues for unicast and multidestination (multicast, broadcast, destination lookup fail) traffic, there is no meast forwarding class and there is no queue 8. (Switches that use different forwarding classes and output queues for unicast and multidestination traffic support 12 forwarding classes and output queues, of which four of each are dedicated to multidestination traffic. Switches that use the same forwarding classes and output queues for unicast for unicast and multidestination 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-o	f-se	rvice f	orwarding-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 forward	ling-class	
Forwarding class	I	D Queu	ie No-Loss
assured-forwarding	2	2 2	disabled
best-effort	e	0 0	disabled
expedited-forwarding	g 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 | 346
- Syntax (TX Matrix and TX Matrix Plus Router) | 346
- Description | 346
- Options | 346
- Required Privilege Level | 346
- Output Fields | 347
- Sample Output | 347
- Release Information | 350

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

Icc (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(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> shc	w class-	of-service	forwarding-table
Classifier tab	le index	: 9, # entr	ies: 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 Inde	ex/
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 (1	ndex: 129,	Map index: 2, Map type: FINAL,
Num of queues	: 2):		
Entry 0 (Sch	eduler i	ndex: 16, F	orwarding-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
                         Probability(%)
Entry
           Fullness(%)
```

show class-of-service forwarding-table lcc (TX Matrix Plus Router)

100

100

0

```
user@host> show class-of-service forwarding-table lcc 0
lcc0-re0:
```

Classifie	r table inde>	(: 9, # en	tries: 64	, Table	type:	IPv6 DSCP
Entry #	Code point	Forwardi	ng-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 | 350
- Description | 350
- Options | 351
- Required Privilege Level | 351
- Output Fields | 351
- Sample Output | 352
- Release Information | 352

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 44 on page 351 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.

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 \emptyset or 1. For some platforms, the value is \emptyset , 1, 2, or 3. The value \emptyset represents low PLP. The value 1 represents high PLP. The value 2 represents medium-low PLP. The value 3 represents medium-high PLP.

Table 44: show class-of-service forwarding-table classifier Output Fields

Sample Output

show class-of-service forwarding-table classifier

user@hos [.]	t> show class -	of-servic	e forwardi	ng-table	classifier
Classifi	er table inde>	<: 62436,	<pre># entries:</pre>	64, Tabl	le type: DSCP
Entry #	Code point	Forwardi	ng-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 | 353
- Description | 353



- Required Privilege Level | 353
- Output Fields | 353
- Sample Output | 354
- Release Information | 354

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 45 on page 353 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 45: 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.

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.

Table 45: show class-of-service forwarding-table classifier mapping Output Fields (Continued)

Sample Output

show class-of-service forwarding-table classifier mapping

user@host> <pre>show class-of-service forwarding-table classifier mapping</pre>							
	Table index/						
Interface	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 drop-profile

IN THIS SECTION

- Syntax | 355
- Description | 355
- Options | 355
- Required Privilege Level | 355
- Output Fields | 355
- Sample Output | 356
- Release Information | 357

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 46 on page 356 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.

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.

Table 46: show class-of-service forwarding-table drop-profile Output Fields

Sample Output

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

		f-service forwardi 4, # entries: 1	ng-table drop-profile
KED UI OP	profile index.		
_		Drop	
Entry	Fullness(%)	Probability(%)	
0	100	100	
RED drop	profile index:	8742, # entries:	3
		Drop	
Entry	Fullness(%)	<pre>Probability(%)</pre>	
0	10	10	
1	20	20	
2	30	30	
RED drop	profile index:	24627, # entries:	64
		Drop	
Entry	Fullness(%)	<pre>Probability(%)</pre>	
0	0	0	
1	1	1	
2	2	2	
3	- 4	- 4	
5	т	4	

61	98	99	
62	99	99	
63	100	100	
RED drop	<pre>profile index:</pre>	25393, # entries:	64
		Drop	
Entry	Fullness(%)	<pre>Probability(%)</pre>	
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 rewrite-rule

IN THIS SECTION

- Syntax | 358
- Description | 358
- Options | **358**
- Required Privilege Level | 358
- Output Fields | 358
- Sample Output | 359
- Release Information | 359

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 47 on page 358 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.

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.	

Field Name	Field Description
State	State of this code point: enabled , rewritten , or disabled .
High bits	Code point value for high-priority loss profile.

Table 47: show class-of-service forwarding-table rewrite-rule Output Fields (Continued)

Sample Output

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

user@ho	ost> show	class-of-se	ervice forwa	arding-table rewrite-rule
Rewrite	e table i	ndex: 3753,	<pre># entries:</pre>	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 | 360
- Description | 360
- Options | **360**
- Required Privilege Level | 360
- Output Fields | 360

- Sample Output | 361
- Release Information | 361

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 48 on page 360 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 48: 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.

Field Name	Field Description
Туре	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 .

Table 48: show class-of-service forwarding-table rewrite-rule mapping Output Fields (Continued)

Sample Output

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

user@host> sł	now class-	of-service	e forwarding-table rewrite-rule mapping
Interface	Index	Table ind	dex 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 forwarding-table scheduler-map

IN THIS SECTION

- Syntax | 362
- Description | 362
- Options | **362**
- Required Privilege Level | 362
- Output Fields | 362

- Sample Output | 364
- Release Information | 364

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 49 on page 362 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 49: show class-of-service forwarding-table scheduler-map Output Fields

Field Name	Field Description
Interface	Name of the physical interface.
Index	Physical interface index.
Map index	Scheduler map index.

Field Name	Field Description
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	 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.
TCP PLP low	Drop profile index for a low TCP packet loss priority profile.

Table 49: show class-of-service forwarding-table scheduler-map Output Fields (Continued)

Field Name	Field Description
Policy is exact	If this line appears in the output, exact rate limiting is enabled. Otherwise, no rate limiting is enabled.

Table 49: show class-of-service forwarding-table scheduler-map Output Fields (Continued)

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 | 365
- Description | 365
- Options | **366**
- Required Privilege Level | 366
- Output Fields | 366
- Sample Output | 386

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:
	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 <i>interface-name</i> 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 367 describes the output fields for the show class-of-service interface command. Output fields are listed in the approximate order in which they appear.

Field Name Field Description Name of a physical interface. Physical interface Index of this interface or the internal index of this object. Index (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. Status of dedicated queues configured on an interface. Dedicated Queues (Enhanced subscriber management) This field is not displayed for enhanced subscriber management. Maximum usable Number of queues you can configure on the interface. queues Maximum number of queues you can use. Maximum usable queues Number of queues created in addition to the default queues. Total non-default queues created (Enhanced subscriber management) This field is not displayed for enhanced subscriber management. IEEE 802.1p code point (priority) rewrite value. Incoming traffic from the Fibre Channel Rewrite Input IEEE (FC) SAN is classified into the forwarding class specified in the native FC interface Code-point (NP_Port) fixed classifier and uses the priority specified as the IEEE 802.1p rewrite value. Maximum transmission rate on the physical interface. You can configure the shaping rate 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.

Field Description
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.
Name of the output fabric scheduler map associated with a QFabric system Interconnect device interface.
Maximum transmission rate on the input interface.
Name of the input scheduler map associated with this interface.
Name of the scheduler map associated with the packet forwarding component queues.
Name and type of the rewrite rules associated with this interface.
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.
Name and type of classifiers associated with this interface.
Name of the forwarding map associated with this interface.
Congestion notification state, enabled or disabled.

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 intrface.
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.
Туре	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.

Field Name	Field Description
Remote-fault	 Remote fault status. Online—Autonegotiation is manually configured as online. Offline—Autonegotiation is manually configured as offline.
Device flags	 The Device flags field provides information about the physical device and displays one or more of the following values: 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.

Field Name	Field Description		
Interface flags	The Interface flags field provides information about the physical interface and displays one or more of the following values:		
	• 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:		
	 1—Takes effect for incoming packets with one label only. 		
	• 2-Takes effect for incoming packets with two labels only.		
	• [12]—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.		

rical interface and	

372

Table 50: show class-of-service interface Output Fields (Continued)

Field Description

Field Name

Flags	The Logical interface flags field provides information about the logical interface and displays one or more of the following values:
	• 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.

CId55-	or-service interface Output Fields (Continued)
	Field Description
	Provides information about the physical link and displays one or more of the following values:
	• 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.
	Current interface hold-time up and hold-time down, in milliseconds.

Date, time, and how long ago the interface went from down to up. The format is Last flapped: *year-month-day hour:minute: second: timezone (hour: minute: second* ago). For

example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).

Table 50: show class-of-service interface Output Fields (Continued)

Number of CoS queues configured.

Field Name

Link flags

Hold-times

CoS queues

Last flapped

Field Name	Field Description
Statistics last cleared	 Number and rate of bytes and packets received and transmitted on the physical interface. 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	 Exclude the counting of overhead bytes from aggregate queue statistics. 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	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.

Field Name	Field Description
Input errors	 Input errors on the interface. The labels are explained in the following list: 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-13-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 FIF0 overflows—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces.

-of-service interface Output Fields (Continued)	
Field Description	
Output errors on the interface. The labels are explained in the following list:	
• Carrier transitions—Number of times the interface has gone from down to up number does not normally increment quickly, increasing only when the cable unplugged, the far-end system is powered down and up, or another problem the number of carrier transitions increments quickly (perhaps once every 10 the cable, the far-end system, or the PIC is malfunctioning.	e is occurs. If
• Errors—Sum of the outgoing frame terminations and FCS errors.	
• Drops—Number of packets dropped by the output queue of the I/O Manager the interface is saturated, this number increments once for every packet that dropped by the ASIC's RED mechanism.	
• Aged packets—Number of packets that remained in shared packet SDRAM so the system automatically purged them. The value in this field should never in If it does, it is most likely a software bug or possibly malfunctioning hardware	ncrement.
• HS link FIF0 underflows—Number of FIFO underflows on the high-speed link between the ASICs responsible for handling the router interfaces.	(S

Table 50: show class-of-service interface

Field Name

Output errors

Egress queues	Total number of egress Maximum usable queues on the specified interface.
Queue counters	 CoS queue number and its associated user-configured forwarding class name. 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	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.

Field Name	Field Description
SONET PHY	 Counts of specific SONET errors with detailed information. 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 0K indicates a problem. The SONET PHY field has the following subfields: PLL Lock—Phase-locked loop PHY Light—Loss of optical signal
SONET section	 Counts of specific SONET errors with detailed information. 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 0K indicates a problem. The SONET section field has the following subfields: BIP-B1-Bit interleaved parity for SONET section overhead SEF-Severely errored framing L0S-Loss of signal L0F-Loss of frame ES-S-Errored seconds (section) SEF-S-Severely errored framing seconds (section)

Field Name	Field Description
SONET line	 Active alarms and defects, plus counts of specific SONET errors with detailed information. 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 0K indicates a problem. The SONET line field has the following subfields: 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) UAS-L—Unavailable seconds (near-end line) SES-LFE—Errored seconds (far-end line) UAS-LFE—Unavailable seconds (far-end line) UAS-LFE—Unavailable seconds (far-end line)

Field Name	Field Description
SONET path	 Active alarms and defects, plus counts of specific SONET errors with detailed information. 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 0K indicates a problem. The SONET path field has the following subfields: BIP-B3—Bit interleaved parity for SONET section overhead REI-P—Remote error indication L0P-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) SES-PFE—Errored seconds (far-end STS path) SES-PFE—Severely errored seconds (far-end STS path) UAS-PFE—Unavailable seconds (far-end STS path) UAS-PFE—Unavailable seconds (far-end STS path)

Field Name	Field Description
Received SONET overhead Transmitted SONET overhead	 Values of the received and transmitted SONET overhead: 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	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.
HDLC configuration	 Information about the HDLC configuration. 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	 Information about the configuration of the Packet Forwarding Engine: Destination slot—FPC slot number. PLP byte—Packet Level Protocol byte.

face.
ed user-configured forwarding

Field Description

Field Name

CoS information Information about the CoS queue for the physical interface.			
	• 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.		

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	 Number of packets dropped because of random early detection (RED). (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: 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: Low—Number of low-loss priority 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: 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.				

Field Name	Field Description			
RED-dropped bytes	 Number of bytes dropped because of RED. The byte counts vary by PIC type. (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: 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	 Rate limiting configuration of the queue. Possible values are : 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.			

-of-service interface Output Fields <i>(Continued)</i>
Field Description
Display the assignment of drop profiles.
• 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 inter

Field Name

Drop profiles

Excess Priority	Priority of the excess bandwidth traffic on a scheduler.
Drop profiles	Display the assignment of drop profiles.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.

Field Name	Field Description
Adjustment	Display the assignment of shaping-rate adjustments on a scheduler node or queue.
information	• Adjusting application—Application that is performing the shaping-rate adjustment.
	• 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 specifie 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 subtract 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	Туре	Index		
	Scheduler-map	<default></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)

<pre>user@host> show class-of-service interface ge-6/2/0</pre>
Physical interface: ge-6/2/0, Index: 175
Maximum usable queues: 4, Queues in use: 4
Scheduler map: <default>, Index: 2</default>
Input scheduler map: <default>, Index: 3</default>
Chassis scheduler map: <default-chassis>, Index: 4</default-chassis>

show class-of-service interface (ANCP)

	ass-of-service interface pp e: pp0.1073741842, Index: 34		
Object	Name	Туре	Index
Traffic-control-pr	ofile TCP-CVLAN	Output	12408
Classifier	dscp-ipv6-compatibilit	ty dscp-ipv6	9
Classifier	ipprec-compatibility	ip	13
Adjustment t Configured s Adjustment v Adjustment o	ication: ancp LS-0 ype: absolute haping rate: 4000000 alue: 11228000 verhead-accounting mode: Fra verhead bytes: 50 arget: node	ame Mode	

show class-of-service interface (PPPoE Interface)

us	er@host> show class-o	f-service interface pp0.1			
Lo	gical interface: pp0.	1, Index: 85			
	Object	Name	Туре	Index	
	Traffic-control-pro	file tcp-pppoe.o.pp0.1	Output	2726446535	
	Classifier	<pre>ipprec-compatibility</pre>	ip	13	
	Adjusting application Adjustment type: Adjustment value: Adjustment overhe Adjustment target	absolute 5000000 ad-accounting mode: cell			

show class-of-service interface (DHCP Interface)

user@host> show class- Logical interface: pp0	o f-service interface demux0 1, Index: 85	.1		
Object	Name	Туре	Index	
Traffic-control-pro	ofile 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-servi	ce interface xe-4/0/0		
Physical interface: xe-4/0/0,	Index: 153		
Maximum usable queues: 8,	Queues in use: 4		
Shaping rate: 500000000	bps		
Scheduler map: <default></default>	, Index: 2		
Congestion-notification:	Disabled		
Logical interface: xe-4/	0/0.0, Index: 77		
Object	Name	Туре	Index
Classifier	ipprec-compatibility	ip	13

show class-of-service interface detail

user@host> <pre>show class-of-service interface ge-0/3/0 detail</pre>
<pre>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</pre>
Physical interface: ge-0/3/0, Index: 138 Maximum usable queues: 4, Queues in use: 5 Shaping rate: 50000 bps Scheduler map: interface-schedular-map, Index: 58414 Input shaping rate: 10000 bps Input scheduler map: schedular-map, Index: 15103 Chassis scheduler map: <default-chassis>, Index: 4 Congestion-notification: Disabled</default-chassis>
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 inet up up mpls Logical interface: ge-0/3/0.0, Index: 68 Object Name Type Index 33 Rewrite exp-default exp (mpls-any) Classifier exp-default 10 exp Classifier ipprec-compatibility 13 ip 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 Index Object Name Туре 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 pps 0 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 0 Input packets: 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 Dropped packets Queue counters: Queued packets Transmitted 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: Transmit Receive Total octets 0 0

Packet Forwarding Engine Destination slot: 0 CoS information: Direction : Output CoS transmit queue 2 ef2 Direction : Input CoS transmit queue	% 39 %	Bandwidth bps 19500 Bandwidth bps	% Ø %	usec 120	Priority high Priority	
Packet Forwarding Engine Destination slot: 0 CoS information: Direction : Output CoS transmit queue 2 ef2 Direction : Input		bps 19500		usec 120	high	Limit none
Packet Forwarding Engine Destination slot: 0 CoS information: Direction : Output CoS transmit queue 2 ef2		bps		usec	-	
Packet Forwarding Engine Destination slot: 0 CoS information: Direction : Output CoS transmit queue		bps		usec	-	
Packet Forwarding Engine Destination slot: 0 CoS information: Direction : Output	C/		٥/		Priority	Limit
Packet Forwarding Engine Destination slot: 0 CoS information: Direction : Output				P 00		
Packet Forwarding Engine Destination slot: 0 CoS information:						
Packet Forwarding Engine Destination slot: 0						
Packet Forwarding Engine						
	e configur	ation:				
F1 (1 0		te fault: Lin	k OK			
Local resolution:						
Link mode: Full-de	uplex, Flo	w control: Sy	mmetric/	Asymmetric,	Remote fa	ult: OK
Link partner:						
Negotiation status: Co	omplete					
Autonegotiation informa	tion:					
CAM destination filte	rs: 0, CAM	source filte	rs: 0			
Output packet error co	ount			0		
Output packet pad cou	nt			0		
Output packet count				0		
Input SA rejects		0				
Input DA rejects		0				
Input packet rejects		0				
Input packet count		0				
Filter statistics:						
Code violations		0				
VLAN tagged frames		0				
Fragment frames		0				
Jabber frames		0				
Oversized frames		0				
MAC pause frames		0		0		
MAC control frames		0		0		
FIFO errors		0		0		
CRC/Align errors		0		0		
		0		0		
Multicast packets		0		0		
Unicast packets Broadcast packets Multicast packets				0		

Interface index: 138, SNMP ifIndex: 601

Forwarding classes: 16 supported, 5 in use

Ingress queues: 4 su			
Queue: 0, Forwarding	; classes: af3		
Queued:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
	ckets : Not Availab	le	
RED-dropped pack		0	0 pps
RED-dropped byte		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 pac	ckets : Not Availab	le	
RED-dropped pack	kets :	0	0 pps
RED-dropped byte	es :	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 pac	ckets : Not Availab	le	
RED-dropped pack	cets :	0	0 pps
RED-dropped byte	es :	0	0 bps
Queue: 3, Forwarding	g classes: ef1		
Queued:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped pac	kets : Not Availab	le	
RED-dropped pack	æts :	0	0 pps
RED-dropped byte	es :	0	0 bps
Forwarding classes:	16 supported, 5 in	use	
Egress queues: 4 sup	ported, 5 in use		

Queue: 0, Forwarding cla	sses: 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 cla	sses: 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 cla	sses: 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 cla	sses: 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 ir	n use			
Queue: 0, Forwarding clas	sses: 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 clas	ses: 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 clas	sses: 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 clas	sses: ef1			
Queued:				
Packets	:	108546		0 pps
Bytes	:	12754752	37	76 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-schedular-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 < default-drop-profile> Low any 1 Medium low < default-drop-profile> 1 any Medium high any 1 < default-drop-profile> 1 < default-drop-profile> High any 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: schedular-map Scheduler map: schedular-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
                                  40582
      Low
                      any
                                           green
      Medium low
                      any
                                      1
                                           < default-drop-profile>
      Medium high
                                      1
                                           < default-drop-profile>
                      any
                                  18928
     High
                                           yellow
                      any
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
                                      1
                                           < default-drop-profile>
                      any
      Medium low
                                      1
                                           < default-drop-profile>
                      any
      Medium high
                                      1
                                           < default-drop-profile>
                      any
      High
                                      1
                                           < default-drop-profile>
                      any
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
                                      1
                                           < default-drop-profile>
      Low
                      any
      Medium low
                                      1
                                           < default-drop-profile>
                      any
      Medium high
                                      1
                                           < default-drop-profile>
                      any
                                           < default-drop-profile>
     High
                                      1
                      any
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
                                      1
                                           < default-drop-profile>
      Low
                      any
      Medium low
                                      1
                                           < default-drop-profile>
                      any
      Medium high
                                      1
                                           < default-drop-profile>
                      any
                                           < default-drop-profile>
      High
                      any
                                      1
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
                                      1
                                           < default-drop-profile>
                      any
      Medium low
                                      1
                                            < default-drop-profile>
                      any
      Medium high
                      any
                                      1
                                            < default-drop-profile>
                                      1
                                            < default-drop-profile>
      High
                      any
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
                                                                          high
                                                 3
normal
  af1
                                         4
                                                            0
                                                                          low
                                                 4
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 0 Output packets: 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 0 Output packets: 0 pps Protocol inet, MTU: 1500, Generation: 172, Route table: 0 Flags: Sendbcast-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:
                              0
    Input unicast bytes:
    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
                          inet filter-in-ge-0/3/0.0-i
                      up
                up
                           mpls
                                                      exp-filter
                Admin Link Proto Input Policer
Interface
                                                       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:
                                                  Packets
Name
```

p1-ge-0/3/0.0-inet-i

		60		
-	ce: ge-0/3/0.0, Inde			
Object	Name	Ту	ре	Index
Rewrite	exp-defau	lt ex	p (mpls-any)	33
Rewrite rule: exp	-default, Code point	type: exp, Inde	x: 33	
Forwarding class	S	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	Ту	pe	Index
Classifier	exp-defau	-		10
010001110	onp dor da		P	
Classifier: exp-de	efault, Code point t	ype: exp, Index:	10	
Code point	Forwarding class		Loss priority	
000	af3		low	
001	af3		high	
			-	
000 001 010 011			low high low high	

0

011	af2			high	
100	ef2			low	
101	ef2			high	
110	ef1			low	
111	ef1			high	
Object		Name	Туре		Index
Classifier		ipprec-compatibility	ip		13

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13

	1 37 1 31	,
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
                                        0
  af3
                                                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
                                        0
    Input packets:
     Output packets:
                                        0
    Local statistics:
     Input bytes :
                                        0
                                        0
     Output bytes :
    Input packets:
                                        0
     Output packets:
                                        0
    Transit statistics:
     Input bytes :
                                        0
                                                             0 bps
                                        0
     Output bytes :
                                                             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
                          mpls
                up
                      up
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		Т	уре	Index
Classifier	ipprec-com	patibilit	y i	р	13
Classifier: ipprec-o	compatibility, Code	e point ty	pe: i	net-precedence, In	dex: 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	e 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- Maximum usable queues: Scheduler map: <defau Congestion-notificati</defau 	4, Queues in use: 4 lt>, Index: 2		
Logical interface: at	-0/0/1.0, Index: 71		
Logical interface: at	-0/0/1.32767, Index: 72	2	
Physical interface: ge-	0/1/0, Index: 146		
Maximum usable queues:			
Scheduler map: <defau< td=""><td>lt>, Index: 2</td><td></td><td></td></defau<>	lt>, Index: 2		
Congestion-notificati	on: Disabled		
Object	Name	Туре	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	Туре	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	5	
Physical interface: ge-	0/1/1 Index: 147		
Maximum usable queues:			
Scheduler map: <defau< td=""><td></td><td></td><td></td></defau<>			
Congestion-notificati			
Object	Name	Туре	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:			
Scheduler map: <defau< td=""><td></td><td></td><td></td></defau<>			
Congestion-notificati			
Object	Name	Туре	Index
Rewrite	ri	ieee8021p (outer)	35392

Classifier	ci	ieee8021p	583	
Physical interfac	e: ge-0/1/3, Index: 149			
-	eues: 8, Queues in use: 5			
Scheduler map:	<default>, Index: 2</default>			
Congestion-noti	fication: Disabled			
Object	Name	Туре	Index	
Classifier	ipprec-compatibility	ip	13	
Logical interfa	ce: ge-0/1/3.0, Index: 77			
Object	Name	Туре	Index	
Rewrite	custom-exp2	exp (mpls-any)	53581	
Physical interfac	e: ge-0/1/4, Index: 150			
Maximum usable qu	eues: 8, Queues in use: 5			
Scheduler map:	<default>, Index: 2</default>			
Congestion-noti	fication: Disabled			
Object	Name	Туре	Index	
Classifier	ipprec-compatibility	ip	13	
Physical interfac	e: ge-0/1/5, Index: 151			
Maximum usable qu	eues: 8, Queues in use: 5			
	<default>, Index: 2</default>			
-	fication: Disabled			
Object	Name	Туре	Index	
Classifier	ipprec-compatibility	ip	13	
Physical interfac	e: ge-0/1/6, Index: 152			
Maximum usable qu	eues: 8, Queues in use: 5			
Scheduler map:	<default>, Index: 2</default>			
Congestion-noti	fication: Disabled			
Object	Name	Туре	Index	
Classifier	ipprec-compatibility	ip	13	
Physical interfac	e: ge-0/1/7, Index: 153			
Maximum usable qu	eues: 8, Queues in use: 5			
Scheduler map:	<default>, Index: 2</default>			
-	fication: Disabled			
Object	Name	Туре	Index	
Classifier	d1	dscp	11331	
	e: ge-0/2/0, Index: 154			
Maximum usable qu	eues: 8, Queues in use: 5			

	<default>, Index: 2</default>		
-	fication: Disabled	_	
Object	Name	Туре	Index
Classifier	ipprec-compatibility	ip	13
Physical interfac	e: ge-0/2/1, Index: 155		
Maximum usable qu	eues: 8, Queues in use: 5		
Scheduler map:	<default>, Index: 2</default>		
Congestion-noti	fication: Disabled		
Object	Name	Туре	Index
Classifier	ipprec-compatibility	ip	13
Logical interfa	ce: ge-0/2/1.0, Index: 78		
Logical interfa	ce: ge-0/2/1.32767, Index: 7	9	
Physical interfac	e: xe-0/3/0, Index: 156		
Maximum usable qu	eues: 8, Queues in use: 5		
Scheduler map:	<default>, Index: 2</default>		
Congestion-noti	fication: Disabled		
Object	Name	Туре	Index
Classifier	ipprec-compatibility	ip	13
Logical interfa	ce: xe-0/3/0.0, Index: 80		
Physical interfac	e: xe-0/3/1, Index: 157		
Maximum usable qu	eues: 8, Queues in use: 5		
Scheduler map:	<default>, Index: 2</default>		
Congestion-noti	fication: Disabled		
Object	Name	Туре	Index
Classifier	<pre>ipprec-compatibility</pre>	ip	13
Logical interfa	ce: xe-0/3/1.0, Index: 81		
ç	,		
[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	Туре	Index	
Traffic-control-prof	file TC_PROF_100_199_SE	RIES_UID1006 Output	4294967312	
Scheduler-map	SMAP-1_UID1002	Output	4294967327	
Rewrite-Output	ieee-rewrite	ieee8021p	60432	
Rewrite-Output	rule1	ip	50463	
Adjusting applic	cation: PPPoE IA tags			
Adjustment type: absolute				
Configured sha	Configured shaping rate: 11000000			
Adjustment val	Adjustment value: 5000000			
Adjustment tar	get: node			
Adjusting applic	ation: ucac			
Adjustment typ	be: delta			
Configured sha	aping rate: 5000000			
Adjustment val	ue: 100000			
Adjustment tar	get: node			

RELATED DOCUMENTATION

Verifying and Managing Junos OS Enhanced Subscriber Management

show class-of-service multi-destination

IN THIS SECTION

- Syntax | 409
- Description | 409
- Options | 409
- Required Privilege Level | 409

- Output Fields | 409
- Sample Output | 410
- Release Information | 410

Syntax

show class-of-service multi-destination

Description

For each class-of-service (CoS) multidestination classifier, display the classifier type.

Options

none Display all multidestination classifiers.

Required Privilege Level

view

Output Fields

Table 51 on page 409 describes the output fields for the show class-of-service multi-destination command. Output fields are listed in the approximate order in which they appear.

Table 51: show class-of-service multi-destination Output Fields

Field Name	Field Description
Family ethernet	Family to which the classifier belongs.
Classifier Name	Name of the classifier.

Field Name	Field Description
Classifier Type	Type of the classifier: dscp or ieee-802.1.
Classifier Index	Internal index of the classifier.

Table 51: show class-of-service multi-destination Output Fields (Continued)

Sample Output

show class-of-service multi-destination

	user@switch>	show class	-of-service multi-destination	
Family ethernet:				
Classifier Name	Classif	ier Type	Classifier Index	
ba-mcast-classi	fier ieee-	802.1	62376	

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

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

Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers

Understanding CoS Classifiers

Understanding CoS Classifiers

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

show class-of-service rewrite-rule

IN THIS SECTION

- Syntax | **411**
- Description | 411
- Options | **411**
- Required Privilege Level | 412
- Output Fields | 412
- Sample Output | **413**
- Release Information | 413

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>	(Optional) Display the rewrite rule of the specified type. The rewrite rule type can be one of the following:
	• dscp—For IPv4 traffic

- **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 52 on page 412 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.

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-servic	e rewrite-rule type dsc	p
Rewrite rule: dscp-default, Cod	e 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

show class-of-service scheduler-map

IN THIS SECTION

- Syntax | 414
- Description | 414
- Options | **414**
- Required Privilege Level | 414
- Output Fields | 415
- Sample Output | 417
- Release Information | 418

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 53 on page 415 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 53: show class-of-service scheduler-map Output Fields

Field Name	Field Description
Scheduler map	Name of the scheduler map. (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.
Index	Index of the indicated object. Objects having indexes in this output include scheduler maps, schedulers, and drop profiles. (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.
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 remainder, which indicates that the scheduler receives the remaining bandwidth of the interface.
Rate Limit	Rate limiting configuration of the queue. Possible values are none, meaning no rate limiting, and exact, 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 remainder to indicate that the buffer is sized according to what remains after other scheduler buffer allocations.

Field Name	Field Description
Buffer size	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. 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.
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	 (QFX Series, OCX Series, and EX4600 switches only) Explicit congestion notification (ECN) state: Disable—ECN is disabled on the specified scheduler Enable—ECN is enabled on the specified scheduler ECN is disabled by default.
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 53: show class-of-service scheduler-map Output Fields (Continued)

Table 53: 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
                   non-TCP
    Low
                                8724
                                         aa-drop-profile
                   TCP
    Low
                                9874
                                         bb-drop-profile
    High
                   non-TCP
                                8833
                                         cc-drop-profile
                                         dd-drop-profile
    High
                   TCP
                                8484
 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
                   TCP
                                9874
                                         bb-drop-profile
    Low
    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: non	e, 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 | 419
- Description | 419
- Options | 419
- Overhead for Layer 2 Statistics | 420
- Additional Information | 422
- Required Privilege Level | 423
- Output Fields | 423
- Sample Output | 433
- Release Information | 477

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 <i>slot-number</i> 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.
interface-name	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 54 on page 420. 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 54 on page 420.

Protocol	Fragmentation		LFI
	First fragmentation	Second to <i>n</i> fragmentations	
	Bytes	Bytes	
MLPPP (Long)	13	12	8
MLPPP (short)	11	10	8
MLFR (FRF15)	12	10	8
MFR (FRF16)	10	8	-
MCMLPPP(Long)	13	12	-
MCMLPPP(Short)	11	10	-

Table 54: Layer 2 Overhead and Transmitted Packets or Byte Counts

```
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 :

Fragmentaion header : 2 bytes

Framerelay header : 2 bytes

HDLC flag and FCS : 4 bytes
```

Overhead with LFI

MLPPP(Long & short sequence):			
Outer PPP header HDLC flag and FCS		4 bytes 4 bytes	
MLFR (FRF15):			
Framerelay 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 55 on page 424 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 55: 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.

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. 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.

Field Name	Field Description
Queued Packets	 Number of packets queued to this queue. 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. 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 422. NOTE: This field is not supported on QFX5100, QFX5110, QFX5200, and QFX5210 switches due to hardware limitations. This field is not supported on EX Series switches due to hardware limitations.
Queued Bytes	 Number of bytes queued to this queue. The byte counts vary by interface hardware. For more information, see Table 56 on page 430. 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 422. NOTE: This field is not supported on QFX5100, QFX5110, QFX5200, and QFX5210 switches due to hardware limitations. This field is not supported on EX Series switches due to hardware limitations.
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. NOTE: For Layer 2 statistics, see "Overhead for Layer 2 Statistics" on page 420

Field Name	Field Description
Transmitted Bytes	Number of bytes transmitted by this queue. The byte counts vary by interface hardware. For more information, see Table 56 on page 430.
	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.
	NOTE: For Layer 2 statistics, see "Overhead for Layer 2 Statistics" on page 420
Tail-dropped packets	Number of packets dropped because of tail drop.
	NOTE : Starting with Junos OS 18.3R1, the Tail-dropped packets counter is supported on PTX Series Packet Transport Routers.
RL-dropped packets	Number of packets dropped due to rate limiting.
	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 422.
	NOTE : The RL-dropped packets counter is not supported on the PTX Series Packet Transport Routers, and is omitted from the output.
RL-dropped bytes	Number of bytes dropped due to rate limiting.
	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 422.

Field Name	Field Description
RED-dropped packets	 Number of packets dropped because of random early detection (RED). (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: 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. 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: 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-low–Number of medium-low loss priority packets dropped because of RED. Medium-low–Number of medium-low loss priority packets dropped because of RED. Medium-low–Number of medium-low loss priority packets dropped because of RED. Medium-low–Number of medium-low loss priority packets dropped because of RED. Medium-low–Number of medium-low loss priority packets dropped because of RED. MoTE: 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.

Field Name	Field Description
RED-dropped bytes	 Number of bytes dropped because of RED. The byte counts vary by interface hardware. For more information, see Table 56 on page 430. (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: 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. 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.
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) Diplays the peak buffer occupancy for the queue while buffer-monitor-enable is enabled at the [edit chassis fpc <i>slot-number</i> traffic-manager] hierarchy level.
Last-packet enqueued	Starting with Junos OS Release 16.1, Last-packet enqueued output field is introduced. If packet-timestamp 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 56 on page 430 shows how the byte counts on the outbound interfaces vary depending on the interface hardware. Table 56 on page 430 is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.

Interface Hardware	Output Level	Byte Count Includes	Comments
Gigabit Ethernet IQ and IQE PICs	Interface	Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes	The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type. For RED dropped, 6 bytes are added for the source MAC address.
	Packet forwarding component	Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet	-

Table 56: Byte Count by Interface Hardware

Interface Hardware	Output Level	Byte Count Includes	Comments
Non-IQ PIC	Interface	 T Series, TX Series, T1600, and MX Series routers: Queued: 478 bytes of Layer 3 packet. Transmitted: 478 bytes of Layer 3 packet. T4000 routers with Type 5 FPCs : 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 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. M Series routers: Queued: 478 bytes of Layer 3 packet. Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead. PTX Series Packet Transport Routers: 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. RED dropped: Full Layer 2 overhead (including all L2 encapsulation and CRC) + 12 inter-packet gap + 8 for the preamble. 	The Layer 2 overhead is 14 bytes for non- VLAN traffic and 18 bytes for VLAN traffic.

Table 56: Byte Count by Interface Hardware (Continued)

Interface Hardware	Output Level	Byte Count Includes	Comments
IQ and IQE PICs with a SONET/SDH interface	Interface	Queued: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes Transmitted: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes RED dropped: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes	The additional 4 bytes are for the Layer 2 Point-to-Point Protocol (PPP) header.
	Packet forwarding component	Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet Transmitted: 486 bytes per packet, representing 478 bytes of Layer 3 packet + 8 bytes	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	 T Series, TX Series, T1600, and MX Series routers: Queued: 478 bytes of Layer 3 packet. Transmitted: 478 bytes of Layer 3 packet. M Series routers: 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 of Layer 3 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 56: Byte Count by Interface Hardware (Continued)

Interface Hardware	Output Level	Byte Count Includes	Comments
1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs 4-port 1G IQ2 and IQ2-E PICs	Interface	Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC. Transmitted: 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.
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.	-

Table 56: Byte Count by Interface Hardware (Continued)

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.

u	ser@host> show interfac	es qu	eue ge-4/2/0			
Ρ	Physical interface: ge-4/2/0, Enabled, Physical link is Up					
	Interface index: 203, S	SNMP	ifIndex: 1054			
F	orwarding classes: 16 s	uppor	ted, 4 in use			
E	gress queues: 8 support	ed, 4	in use			
Q	ueue: 0, Forwarding cla	sses:	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	5
RED-dropped packets	:	0	0 pps	6
Low	:	0	0 pps	6
Medium-low	:	0	0 pps	3
Medium-high	:	0	0 pps	6
High	:	0	0 pps	6
RED-dropped bytes	:	0	0 bps	3
Low	:	0	0 bps	3
Medium-low	:	0	0 bps	6
Medium-high	:	0	0 bps	6
High	:	0	0 bps	6
Queue: 1, Forwarding cla	asses: ex	pedited-forwarding		
Queued:				
Packets	:	0	0 pps	3
Bytes	:	0	0 bps	3

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,	-	vsical link is Up				
Interface index: 158, SN	MP ifIndex: 3	3 Forwarding classes: 8	supported	, 8 in use		
Output queues: 8 support	ed, 8 in use					
Queue: 0, Forwarding cla	sses: 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 cla	sses: af1					
Queued:						
Packets	:	42603765	595484	pps		
Bytes	:	5453281920	609776496	bps		
Transmitted:						
Packets	:	42603765	595484	pps		

Bytes	:	5453281920	6097	76496	bps
Tail-dropped packets	:	0		0	pps
RED-dropped packets	:	0		0	pps
RED-dropped bytes	:	0		0	bps
Queue: 2, Forwarding clas	sses: 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 clas	sses: 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 clas	sses: 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 clas	sses: 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 clas	sses: af1	2		
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 clas	sses: nc2	2		
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 inte					
	o ,	ed, Physical link is Up			
Interface index:					
Forwarding classes:	8 supported, 8	in use			
Output queues: 8 su	pported, 8 in u	se			
Queue: 0, Forwardin	g classes: be				
Queued:					
Packets	:	13	0 pps		
Bytes	:	622	0 bps		
Transmitted:					
Packets	:	13	0 pps		
Bytes	:	622	0 bps		
Tail-dropped pa	ckets :	0	0 pps		
RED-dropped pac	kets :	0	0 pps		
RED-dropped bytes : 0 0 bps					
Queue: 1, Forwardin	g classes: af1				
Queued:					

Packets	:	1725947945	3721	78 pps
Bytes	:	220921336960	3811104	32 bps
Transmitted:				
Packets	:	1725947945	3721	78 pps
Bytes	:	220921336960	3811104	32 bps
Tail-dropped packets	:	0		0 pps
RED-dropped packets	:	0		0 pps
RED-dropped bytes	:	0		0 bps
Queue: 2, Forwarding clas	sses: 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 clas	sses: nc			
Queued:				
Packets	:	571		0 pps
Bytes	:	49318	3	36 bps
Transmitted:				
Packets	:	571		0 pps
Bytes	:	49318	3	36 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	Physical interface: ge-2/2/9, Enabled, Physical link is Up							
Interface index: 238,	SNMP ifIndex	: 71						
Forwarding classes: 16 s	supported, 4	in use						
Ingress queues: 4 suppor	rted, 4 in us	e						
Queue: 0, Forwarding cla	asses: best-e	ffort						
Queued:								
Packets	:	148450735	947295	pps				
Bytes	:	8016344944	409228848	bps				
Transmitted:								

Packets	:	76397439	487512	pps
Bytes	:	4125461868	210602376	bps
Tail-dropped packets	: Not Availab	le		
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 cla	sses: expedite	d-forwarding		
Queued:				
Packets	:	0	0	pps
Bytes	:	0	0	bps
Transmitted:				
Packets	:	0	0	pps
Bytes	:	0	0	bps
Tail-dropped packets	: Not Availab	le		
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 clas	ses: assured-	forwarding		
Queued:				
Packets	:	410278257	473940	pps
Bytes	: 2	2156199518	204742296	bps
Transmitted:				
Packets	:	4850003	4033	pps
Bytes	:	261900162	1742256	bps
Tail-dropped packets	: Not Availab	le		
RED-dropped packets	:	405425693	469907	pps
Low	:	405425693	469907	
Medium-low	:	0	0	pps
Medium-high	:	0		pps

	High	: 0	0 pps
	RED-dropped bytes	: 21892988124 20300004	0 bps
	Low	: 21892988124 20300004	0 bps
	Medium-low	: 0	0 bps
	Medium-high	: 0	0 bps
	High	: 0	0 bps
Q	ueue: 3, Forwarding clas	ses: 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
F	orwarding classes: 16 su	pported, 4 in use	
E	gress queues: 4 supporte	d, 4 in use	
Q	ueue: 0, Forwarding clas	ses: best-effort	
	Queued:		
	Packets	: 76605230 48537	6 pps
	Bytes	: 5209211400 26404456	0 bps
	Transmitted:		
	Packets	: 76444631 48433	6 pps
	Bytes	: 5198235612 26347880	0 bps
	Tail-dropped packets	: Not Available	
	RED-dropped packets	: 160475 104	0 pps
	Low	: 160475 104	0 pps
	Medium-low	: 0	0 pps
	Medium-high	: 0	0 pps
	High	: 0	0 pps
	RED-dropped bytes	: 10912300 56576	0 bps
	Low		0 bps
	Medium-low	: 0	0 bps
	Medium-high	: 0	0 bps
	-		

High	:	0	0 bps	
Queue: 1, Forwarding c	lasses: ex	pedited-forwarding		
Queued:				
Packets	:	0	0 pps	
Bytes	:	0	0 bps	
Transmitted:				
Packets	:	0	0 pps	
Bytes	:	0	0 bps	
Tail-dropped packe	ts : Not A	vailable		
RED-dropped packet	s :	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 c	lasses: as	sured-forwarding		
Queued:				
Packets	:	4836136	3912 pps	
Bytes	:	333402032	2139056 bps	
Transmitted:				
Packets	:	3600866	1459 pps	
Bytes	:	244858888	793696 bps	
Tail-dropped packe	ts : Not A	vailable		
RED-dropped packet	s :	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 c	lasses: ne	etwork-control		
Queued:				
Packets	:	0	0 pps	
Bytes	:	0	0 bps	
Transmitted:				
Packets	:	0	0 pps	

Bytes	:	0	0	bps
Tail-dropped packets	: Not A	vailable		
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		Queues:		
Queues: 4 supported, 4 i				
Queue: 0, Forwarding cla	sses: be	est-effort		
Queued:				
Packets	:	77059796	486384	
Bytes	:	3544750624	178989576	bps
Transmitted:				
Packets	:	77059797	486381	
Bytes	:	3544750670	178988248	bps
Tail-dropped packets	:	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 cla	sses: e>	pedited-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
	:	0		pps
Low	:	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 cla	sses: assured	l-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		bps
Low	:	0		bps
Medium-low	:	0		bps
Medium-high	:	0	0	bps
High	:	0	0	bps
Queue: 3, Forwarding cla	sses: network	-control		
Queued:				
Packets	:	0	0	pps
Bytes	:	0	0	bps
Transmitted:				
Packets	:	0	0	pps
Bytes	:	0		bps
Tail-dropped packets	:	0		pps
RED-dropped packets		0		pps
Low	:	0		pps
Medium-low	:	0		pps
Medium-high	:	0		pps
High	:	0		pps
RED-dropped bytes	:	0		bps
Low	:	0		bps
Medium-low	:	0		bps
				•

Medium-high	:	0	0 bps	
High	:	0	0 bps	

show interfaces queue (Gigabit Ethernet IQ2 PIC)

user@host> show int	terfaces queue	e 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	d, 4 in use					
Queue: 0, Forwardir	ng classes: be	est-effort				
Queued:						
Packets	:	418390039	10	pps		
Bytes	:	38910269752	7440	bps		
Transmitted:						
Packets	:	418390039	10	pps		
Bytes	:	38910269752	7440	bps		
Tail-dropped pa	ackets : Not A	Available				
RED-dropped pac	ckets :	0	0	pps		
RED-dropped byt	tes :	0	0	bps		
Queue: 1, Forwardir	ng classes: ex	xpedited-forwarding				
Queued:						
Packets	:	0	0	pps		
Bytes	:	0	0	bps		
Transmitted:						
Packets	:	0	0	pps		
Bytes	:	0	0	bps		
Tail-dropped pa	ackets : Not A	Available				
RED-dropped pac	ckets :	0	0	pps		
RED-dropped byt	tes :	0	0	bps		
Queue: 2, Forwardir	ng classes: as	ssured-forwarding				
Queued:						
Packets	:	0	0	pps		
Bytes	:	0	0	bps		
Transmitted:						
Packets	:	0	0	pps		
Bytes	:	0	0	bps		
Tail-dropped pa	ackets : Not A	Available				
RED-dropped pac	ckets :	0	0	pps		
RED-dropped byt	tes :	0	0	bps		
Queue: 3, Forwardir	ng classes: ne	etwork-control				
Queued:						

	Packets	: 70	55	1 pps
	Bytes	: 4515	52 5	12 bps
	Transmitted:			
	Packets	: 70	55	1 pps
	Bytes	: 4515	52 5	12 bps
	Tail-dropped packets	: Not Available		
	RED-dropped packets	:	0	0 pps
	RED-dropped bytes	:	0	0 bps
Fo	rwarding classes: 16 s	upported, 4 in use Egr	ess queues: 4 supporte	d, 4 in use
Qu	eue: 0, Forwarding cla	sses: best-effort		
	Queued:			
	Packets	: 10	31	0 pps
	Bytes	: 1432	92	0 bps
	Transmitted:			
	Packets	: 10	31	0 pps
	Bytes	: 1432	92	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
Qu	eue: 1, Forwarding cla	ses: expedited-forwar	ding	
	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
Qu	eue: 2, Forwarding cla	sses: assured-forwardi	ng	
	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 cla	sses:	network-control		
Queued:				
Packets	:	77009		pps
Bytes	:	6894286	7888	bps
Transmitted:				
Packets	:	77009	11	pps
Bytes	:	6894286	7888	bps
Tail-dropped packets	: No	ot 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	Chas	ssis Queues:		
Queues: 4 supported, 4 i	n use			
Queue: 0, Forwarding cla	sses:	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		bps
Low, non-TCP	:	0	0	bps
Low, TCP	:	0		bps
High, non-TCP	:	0	0	bps
High, TCP	:	0		bps
Queue: 1, Forwarding cla	sses:	expedited-forwardin		
Queued:				
Packets	:	0	0	pps
Bytes	:	0		bps
Transmitted:				
Packets	:	0	0	pps
Bytes	:	0		bps
Tail-dropped packets	:	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 clas	sses: assured-forwardir	ıg	
Queued:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:		0 bps
Tail-dropped packets	:		0 pps
RED-dropped packets	:		0 pps
Low, non-TCP	:		0 pps
Low, TCP	:		0 pps
High, non-TCP	:		0 pps
High, TCP	:		0 pps
RED-dropped bytes	:		0 bps
Low, non-TCP	:		0 bps
Low, TCP	:		0 bps
High, non-TCP	:		0 bps
High, TCP	:		0 bps
Queue: 3, Forwarding clas	sses: network-control		
Queued:			
Packets	: 9438	36 1	2 pps
Bytes	: 1375679		8 bps
Transmitted:			
Packets	: 9438	36 1	2 pps
Bytes	: 1375679		8 bps
Tail-dropped packets			0 pps
	:		0 pps
Low, non-TCP	:		0 pps
Low, TCP	:		0 pps
High, non-TCP	:		0 pps
High, TCP	:		0 pps
RED-dropped bytes	:		0 bps
Low, non-TCP	:		0 bps
	•	Ŭ	~ ~p>

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 interface	es queue ge-6/2/0 bot	h-ingress-egress	
Physical interface: ge-6/			
Interface index: 175, S	SNMP ifIndex: 121		
Forwarding classes: 8 sup	oported, 4 in use		
Ingress queues: 4 support	ted, 4 in use		
Queue: 0, Forwarding clas	sses: best-effort		
Queued:			
Packets	: Not Available		
Bytes	:	0	0 bps
Transmitted:			
Packets	:	254	0 pps
Bytes	: 16	274	0 bps
Tail-dropped packets	: Not Available		
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Queue: 1, Forwarding clas	sses: expedited-forwa	rding	
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 clas	sses: assured-forward	ing	
Queued:			
Packets	: Not Available		
Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Queue: 3, Forwarding cla	sses: network-control		
Queued:			
Packets	: Not Available		
Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	: Not Available		
····· ··········	:	0	0 pps
11 3	:	0	0 bps
Forwarding classes: 8 su			
Egress queues: 4 support			
Queue: 0, Forwarding cla	sses: best-effort		
Queued:			
Packets	: Not Available		
Bytes	:	0	0 bps
Transmitted:			
Packets	:	3	0 pps
Bytes	: 1	26	0 bps
Tail-dropped packets	: Not Available		
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Queue: 1, Forwarding cla	sses: expedited-forwar	rding	
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 cla	sses: assured-forward	ing	
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 cla	sses: network-control		

Queued:		
Packets	: Not Available	
Bytes	: 0	0 bps
Transmitted:		
Packets	: 0	0 pps
Bytes	: 0	0 bps
Tail-dropped packets	: Not Available	
	: 0	0 pps
RED-dropped bytes	: 0	0 bps
Packet Forwarding Engine	Chassis Queues:	
Queues: 4 supported, 4 ir	use	
Queue: 0, Forwarding clas	ses: 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 clas	ses: 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 clas	ses: 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 clas	ses: network-control	
Queued:		

Packets	:	9397	0 pps	;
Bytes	:	3809052	232 bps	i
Transmitted:				
Packets	:	9397	0 pps	i
Bytes	:	3809052	232 bps	i
Tail-dropped packets	:	0	0 pps	;
RED-dropped packets	:	0	0 pps	;
RED-dropped bytes	:	0	0 bps	i

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, F	Physical link is Up			
Interface index: 175, SNMP ifIndex: 121					
Forwarding classes: 8 s	supported, 4 in us	se			
Ingress queues: 4 suppo	orted, 4 in use				
Queue: 0, Forwarding cl	lasses: best-effor	^t			
Queued:					
Packets	: Not Available	9			
Bytes	:	0	0 bps		
Transmitted:					
Packets	:	288	0 pps		
Bytes	:	18450	0 bps		
Tail-dropped packet	ts : Not Available	9			
RED-dropped packets	6 :	0	0 pps		
RED-dropped bytes	:	0	0 bps		
Queue: 1, Forwarding c	lasses: expedited-	forwarding			
Queued:					
Packets	: Not Available	9			
Bytes	:	0	0 bps		
Transmitted:					
Packets	:	0	0 pps		
Bytes	:	0	0 bps		
Tail-dropped packet	ts : Not Available	2			
RED-dropped packets	5 :	0	0 pps		
RED-dropped bytes	:	0	0 bps		
Queue: 2, Forwarding cl	lasses: assured-fo	prwarding			
Queued:					
Packets	: Not Available	2			
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 clas	sses: 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, Phys:	ical link is Up		
Interface index: 175,	SNMP ifIndex: 121			
Forwarding classes: 8 s	upported, 4 in use			
Egress queues: 4 suppor	ted, 4 in use			
Queue: 0, Forwarding cl	asses: 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 cl	asses: expedited-for	warding		
Queued:				
Packets	: Not Available			
Bytes	:	0	0 bps	
Transmitted:				
Packets	:	0	0 pps	
Bytes	:	0	0 bps	
Tail-dropped packet	s : Not Available			

RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Queue: 2, Forwarding cla	sses: assured-fo	orwarding	
Queued:			
Packets	: Not Available	2	
Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	: Not Available	2	
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Queue: 3, Forwarding cla	sses: network-co	ontrol	
Queued:			
Packets	: Not Available	2	
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 i	n use		
Queue: 0, Forwarding cla	sses: best-effor	t	
Queued:			
Packets	:	80564692	0 pps
Bytes	: 33	83717100	0 bps
Transmitted:			
Packets	:	80564692	0 pps
Bytes	: 33	83717100	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Queue: 1, Forwarding cla	sses: expedited-	forwarding	
Queued:			
Packets	:	80564685	0 pps
Bytes	: 33	83716770	0 bps
Transmitted:			
Packets	:	80564685	0 pps
Bytes	: 33	83716770	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 clas	sses: network-con	trol			
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	suppor	ted, 4 in use			
Ingress queues: 4 supp	orted,	4 in use			
Queue: 0, Forwarding o	lasses:	best-effort			
Queued:					
Packets	:	110208969	472875	pps	
Bytes	:	5951284434	204282000	bps	
Transmitted:					
Packets	:	110208969	472875	pps	
Bytes	:	5951284434	204282000	bps	
Tail-dropped packe	ts : No	t Available			
RED-dropped packet	.s :	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 clas	ses: expedited-forward	ding	
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 clas	ses: assured-forwardi	ng	
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 clas	ses: network-control		

Queued:				
Packets	:	0	0	pps
Bytes	:	0	0	bps
Transmitted:				
Packets	:	0	0	pps
Bytes	:	0	0	bps
Tail-dropped packet	s : Not Avail	able		
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 suppor	ted, 4 in use	!		
Queue: 0, Forwarding cl	asses: best-e	ffort		
Queued:				
Packets	:	109355853	471736	pps
Bytes	:	7436199152	256627968	bps
Transmitted:				
Packets	:	109355852	471736	pps
Bytes	:	7436198640	256627968	bps
Tail-dropped packet	s : Not Avail	able		
RED-dropped packets	:	0	0	pps
Low	:	0	0	pps
Medium-low	:	0	0	pps
Medium-high	:	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 cl	asses: expedi	ted-forwarding		
Queued:				
Packets	:	0	0	pps
Bytes	:	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 clas	sses: assured-forwardin	ng	
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 clas	sses: 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 s				-	
Egress queues: 8 support	ed, 8 in use				
Queue: 0, Forwarding cla	sses: 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 cla	sses: 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 cla	sses: 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 cla	sses: 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 clas	sses: 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 clas	sses: 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 cla	sses: 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 clas	sses: CLASS_S_O	UTPUT, G	ETS		
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		bps
Medium-high	:		0		bps
High	:		0		bps
-					
Packet Forwarding Engine	Chassis Queues	:			
Queues: 8 supported, 8 in					
Queue: 0, Forwarding clas					
Queued:					
Packets	:	37136	5 23	8620	pps
Bytes	:	1559733			bps
Transmitted:					
Packets	:	37136	5 23	3620	pps
Bytes	:	1559733			bps
Tail-dropped packets			0		pps
RED-dropped packets			0		pps
neb al opped packets	•		v	0	PP3

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 cla	sses: 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
	:	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 cla	sses: 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
	:	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
			I

Medium-high	:	0	0 bps
High	:	0	0 bps
Queue: 3, Forwarding cla	sses: CONTROL		
Queued:			
Packets	:	32843	0 pps
Bytes	: 20	541754	56 bps
Transmitted:			
Packets	:	32843	0 pps
Bytes	: 20	541754	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 cla	sses: CLASS_B_OUTPU	JT	
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 cla	sses: CLASS_C_OUTPU	JT	
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 clas	sses: 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 clas	sses: 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)

ser@switch> show in	terfaces queue xe-0	/0/15		
Physical interface	e: xe-0/0/15, Enabl	ed, Physical link is Up		
Interface index	: 49165, SNMP ifInd	ex: 539		
Forwarding classes	s: 12 supported, 8	in use		
Egress queues: 12	supported, 8 in us	e		
Queue: 0, Forward:	ing classes: best-e	ffort		
Queued:				
Packets	:	0	0 pps	
Bytes	:	0	0 bps	
Transmitted:				
Packets	:	0	0 pps	
Bytes	:	0	0 bps	
Tail-dropped	packets : Not Avail	able		
Total-dropped	packets:	0	0 pps	
Total-dropped	bytes :	0	0 bps	
Queue: 3, Forward	ing classes: fcoe			
Queued:				
Packets	:	0	0 pps	
Bytes	:	0	0 bps	
Transmitted:				
Packets	:	0	0 pps	
Bytes	:	0	0 bps	
Tail-dropped	packets : Not Avail	able		
Total-dropped	packets:	0	0 pps	
Total-dropped	bytes :	0	0 bps	
Queue: 4, Forward	ing classes: no-los	S		
Queued:				
Packets	:	0	0 pps	
Bytes	:	0	0 bps	
Transmitted:				
Packets	:	0	0 pps	
Bytes	:	0	0 bps	

	Tail-dropped packets	s : Not Available		
	Total-dropped packet	ts:	0	0 pps
	Total-dropped bytes	:	0	0 bps
Qu	ueue: 7, Forwarding cla	asses: network-control		
	Queued:			
	Packets	:	0	0 pps
	Bytes	:	0	0 bps
	Transmitted:			
	Packets	:	0	0 pps
	Bytes	:	0	0 bps
	Tail-dropped packets	s : Not Available		
	Total-dropped packet	ts:	0	0 pps
	Total-dropped bytes	:	0	0 bps
Qu	ueue: 8, Forwarding cla	asses: mcast		
	Queued:			
	Packets	:	0	0 pps
	Bytes	:	0	0 bps
	Transmitted:			
	Packets	:	0	0 pps
	Bytes	:	0	0 bps
	Tail-dropped packets	s : Not Available		
	Total-dropped packet	ts:	0	0 pps
	Total-dropped bytes	:	0	0 bps

show interfaces queue I2-statistics (Isq 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, F	orwarding classes: be			
Queued:				
Packets	:	1	0 pps	
Bytes	:	1001	0 bps	
Transmitt	ed:			
Packets	:	5	0 pps	
Bytes	:	1062	0 bps	
Tail-dr	opped packets :	0	0 pps	
RED-dro	pped packets :	0	0 pps	
RED-dro	oped bytes :	0	0 bps	

Queue: 1, Forwarding clas: Queued: Packets	ses: et		
-	:		
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 class	ses: 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 class	ses: 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 cla	sses: ef			
Queued:				
Packets	:	0	0	pps
Bytes	:	0		bps
Transmitted:				
Packets	:	0	0	pps
Bytes	:	0		bps
Tail-dropped packets	:	0		pps
RL-dropped packets	:	0		pps
RL-dropped bytes	:	0		bps
	:	0		pps
Low	:	0		pps
Medium-low	:	0		pps
Medium-high	:	0		pps
High	:	0		pps
RED-dropped bytes	:	0		bps
Low	:	0		bps
Medium-low	:	0		bps
Medium-high	:	0		bps
High	:	0		bps
Queue: 2, Forwarding cla	sses: af			
Queued:				
Packets	:	0	0	pps
Bytes	:	ů 0		bps
Transmitted:			Ũ	

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	s :	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 cla	asses: nc		
Queued:			
Packets	:	22231	482 pps
Bytes	:	11849123	2057600 bps
Transmitted:			
Packets	:	22231	482 pps
Bytes	:	11849123	2057600 bps
Tail-dropped packets	s :	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	support	ed, 4 in use	
Egress queues: 8 suppor	ted, 4	in use	
Queue: 0, Forwarding cl	asses: l	best-effort	
Queued:			
Packets	:	16	0 pps
Bytes	:	1896	0 bps
Transmitted:			
Packets	:	16	0 pps
Bytes	:	1896	0 bps
Tail-dropped packet	s :	0	0 pps
RL-dropped packets	:	0	0 pps
-	:	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 cl	asses: (expedited-forwarding	
Queued:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packet	s :	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	: 3276	8	
Queue: 2, Forwarding clas	sses: assured-forwardin	g	
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	: 3276	8	
Queue: 3, Forwarding clas	sses: 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
		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)

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 Bytes : 0 Tail-dropped packets : 0 Queue: 5, Forwarding classes: expedited-forwarding Queued:	user@switch> show interfa	ces queue ge-0/0/0	
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 Bytes : 0 Tail-dropped packets : 0 Queue: 5, Forwarding classes: expedited-forwarding Queued:	Physical interface: ge-0/	0/0, Enabled, Physical	l link is Down
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 Bytes : 0 Tail-dropped packets : 0 Queue: 5, Forwarding classes: expedited-forwarding Queued:	Interface index: 130, S	NMP ifIndex: 501	
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 Bytes : 0 Tail-dropped packets : 0 Queue: 5, Forwarding classes: expedited-forwarding Queued:	Forwarding classes: 16 su	pported, 4 in use	
Queued: Transmitted: Packets : 0 Bytes : 0 Tail-dropped packets : 0 Queue: 1, Forwarding classes: assured-forwarding Queued: Transmitted: Packets : 0 Bytes : 0 Fackets : 0 Queue: 5, Forwarding classes: expedited-forwarding Queued:	Egress queues: 8 supporte	d, 4 in use	
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:	Queue: 0, Forwarding clas	ses: best-effort	
Packets:0Bytes:0Tail-dropped packets :0Queue: 1, Forwarding classes: assured-forwarding Queued:0Queued:Transmitted:Packets:0Bytes:0Tail-dropped packets :0Queue: 5, Forwarding classes: expedited-forwarding Queued:	Queued:		
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:		
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:	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:	Bytes	:	0
Queued: Transmitted: Packets : 0 Bytes : 0 Tail-dropped packets : 0 Queue: 5, Forwarding classes: expedited-forwarding Queued:	Tail-dropped packets	:	0
Transmitted: Packets : 0 Bytes : 0 Tail-dropped packets : 0 Queue: 5, Forwarding classes: expedited-forwarding Queued:	Queue: 1, Forwarding clas	ses: assured-forwardi	ng
Packets : 0 Bytes : 0 Tail-dropped packets : 0 Queue: 5, Forwarding classes: expedited-forwarding Queued:	Queued:		
Bytes : 0 Tail-dropped packets : 0 Queue: 5, Forwarding classes: expedited-forwarding Queued:	Transmitted:		
Tail-dropped packets : 0 Queue: 5, Forwarding classes: expedited-forwarding Queued:	Packets	:	0
Queue: 5, Forwarding classes: expedited-forwarding Queued:	Bytes	:	0
Queued:	Tail-dropped packets	:	0
	Queue: 5, Forwarding clas	ses: expedited-forward	ding
	Queued:		
Transmitted:	Transmitted:		
Packets : 0	Packets	:	0
Bytes : 0	Bytes	:	0

Tail-dropped packets :0Queue: 7, Forwarding classes: network-controlQueued:Transmitted:Packets:Bytes: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
                        :
                                 43144843010304
   Bytes
                        :
   Tail-dropped packets :
                                     8003867575
PFE chassis queues: 1 supported, 1 in use
 Transmitted:
   Packets
                                              0
                        :
   Bytes
                        :
                                              0
                                              0
   Tail-dropped packets :
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
                        :
                                              0
   Bytes
                         :
   Tail-dropped packets :
                                              0
 Queue: 2, Forwarding classes: mcast-be
 Queued:
 Transmitted:
```

Packets 274948977 : 36293264964 Bytes : Tail-dropped packets : 0 Queue: 4, Forwarding classes: mcast-ef Queued: Transmitted: Packets 0 : Bytes 0 : 0 Tail-dropped packets : 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 : 0 Tail-dropped packets : Queue: 7, Forwarding classes: network-control Queued: Transmitted: Packets 46714 : Bytes : 6901326 0 Tail-dropped packets : Packet Forwarding Engine Chassis Queues: Queues: 8 supported, 7 in use Queue: 0, Forwarding classes: best-effort Queued: Transmitted: Packets 739338141426 : 94635282101928 Bytes : Tail-dropped packets : 0 RED-dropped packets : 5606426444 5606426444 Low : High 0 : RED-dropped bytes 683262846464 : 683262846464 Low : 0 High : Queue: 1, Forwarding classes: assured-forwarding

Queued: Transmitted: Packets : Bytes : Tail-dropped packets : RED-dropped packets : Low : High : RED-dropped bytes : Low : High : Queue: 2, Forwarding classes: mcast-be Queued: Transmitted: Packets : Bytes : Tail-dropped packets : RED-dropped packets : Low : High : RED-dropped bytes : Low : High : Queue: 4, Forwarding classes: mcast-ef Queued: Transmitted: Packets : Bytes : Tail-dropped packets : RED-dropped packets : Low : High : RED-dropped bytes : Low : High : Queue: 5, Forwarding classes: expedited-forwarding Queued:

0

0

0

0

0

0

0

0

0

0

0

0

0

0

0

0

0

0

0

0

0

0

0

0

0

0

0

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 clas	sses: 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 clas	sses: network-contro	1
Queued:		
Transmitted:		
Packets	: 9	7990
Bytes	: 1498	7506
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

usonanouton chow interfo		A/2 clico clicol	
	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,		•	oh
Slice name: slice1, Sl		0	
Forwarding classes: 16 s			
Egress queues: 8 support		126	
CoS scheduler resource information: Maximum units supported per MIC/PIC: 84			
	•	54	
Configured units per M			
Maximum units allowed Configured units on th			
Queue: 0, Forwarding cla	·		
Queued: 0, Forwarding cia	SSES: DE		
Packets		4	0, ppg
	:	4 200	0 pps
Bytes Transmitted:	·	200	0 bps
Packets		4	0, ppg
	:	4 200	0 pps
Bytes Tail-dropped packets		200	0 bps
			0 pps
RL-dropped packets	:	0	0 pps
RL-dropped bytes	:	0	0 bps
RED-dropped packets	:	-	0 pps
Low Medium-low	:	0	0 pps
	:	0	0 pps
Medium-high	:	0	0 pps
High RED-drapped bytes		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.

12-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 TrafficMonitoring Interfaces That Have CoS Components | 58Defining CoS Schedulers and Scheduler Maps (CLI Procedure) | 138Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line
Cards (CLI Procedure) | 87ATM Interfaces User Guide for Routing Devices

Layer 2 Bridging, Address Learning, and Forwarding User Guide

show pfe statistics traffic cpu

IN THIS SECTION

- Syntax | 478
- Description | 478
- Options | **478**
- Required Privilege Level | 479
- Output Fields | 479
- Sample Output | 480
- Release Information | 482

Syntax

show pfe statistics traffic cpu <fpc fpc-slot>

Description

(On EX8200 switches only) Display count of multidestination packets ingressing from the physical interface to the CPU.

NOTE: Multidestination packets include unknown unicast, broadcast, and multicast packets.

Options

none	Displays the count of packets ingressing from all the physical interfaces (line cards) to the CPU.
fpc <i>fpc-</i> <i>slot</i>	(Optional) Displays the count of packets ingressing from the physical interface, referred to by the slot number, to the CPU.
	On an EX8200 switch, the FPC slot number is the slot number for the line card. Possible values are 0 through 7 on the EX8208 switch and 0 through 15 on the EX8216 switch.

Required Privilege Level

view

Output Fields

Table 57 on page 479 lists the output fields for the show pfe statistics traffic cpu command. Output fields are listed in the approximate order in which they appear.

Table 57: show pfe statistics	traffic cpu Output Fields
-------------------------------	---------------------------

Field Name	Field Description
Queue	CoS 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.
Packets	Number of packets transmitted by this queue.
Bytes	Number of bytes transmitted by this queue.
Tail-dropped packets	Count of packets dropped at the tail end of the queue because of lack of buffer space.
RED-dropped packets	Number of packets dropped because of Random Early Discard (RED):
	• Low–Number of low-loss priority packets dropped because of RED.
	• High —Number of high-loss priority packets dropped because of RED.
RED-dropped bytes	Number of bytes dropped because of Random Early Discard (RED):
	• Low-Number of low-loss priority bytes dropped because of RED.
	• High —Number of high-loss priority bytes dropped because of RED.

Sample Output

show pfe statistics traffic cpu (EX8208 Switch)

user@switch> show pfe st	atistics traffic cpu		
Queue: 0, Forwarding cla Queued:	sses: Dest-errort		
-	: Not Available		
	: Not Available		
Bytes Packets	: NOU AVAILADIE	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets		0	0 003
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
High	•	0	0 bps
RED-dropped packets	· ·	0	0 pps
Low	· ·	ů 0	0 pps
High	:	0	0 pps
Queue: 1, Forwarding cla		-	, hb.
Queued:		0	
Packets	: Not Available		
Bytes	: Not Available		
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
High	:	0	0 bps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
High	:	0	0 pps
Queue: 2, Forwarding cla	sses: assured-forward	ding	
Queued:			
Packets	: Not Available		
Bytes	: Not Available		
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
High	:	0	0 bps

RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
High	:	0	0 pps
Queue: 3, Forwarding cla	asses: network-control		
Queued:			
Packets	: Not Available		
Bytes	: Not Available		
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	5 :	0	
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
High	:	0	0 bps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
High	:	0	0 pps
Queue: 4			
Packets	: Not Available		
Bytes	: Not Available		
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	5 :	0	
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
High	:	0	0 bps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
High	:	0	0 pps
Queue: 5			
Packets	: Not Available		
Bytes	: Not Available		
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	5 :	0	
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
High	:	0	0 bps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
High	:	0	0 pps
Queue: 6			
Packets	: Not Available		
Bytes	: Not Available		
-			

	Packets	:	0	0 pps
	Bytes	:	0	0 bps
	Tail-dropped packets	:	0	
	RED-dropped bytes	:	0	0 bps
	Low	:	0	0 bps
	High	:	0	0 bps
	RED-dropped packets	:	0	0 pps
	Low	:	0	0 pps
	High	:	0	0 pps
Que	ue: 7			
	Packets	: Not Available		
	Bytes	: Not Available		
	Packets	:	0	0 pps
	Bytes	:	0	0 bps
	Tail-dropped packets	:	0	
	RED-dropped bytes	:	0	0 bps
	Low	:	0	0 bps
	High	:	0	0 bps
	RED-dropped packets	:	0	0 pps
	Low	:	0	0 pps
	High	:	0	0 pps

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

show pfe statistics traffic multicast 485	
show pfe statistics traffic egress-queues 483	
show interfaces queue	
Monitoring Interface Status and Traffic	
Understanding Junos OS CoS Components for EX Series Switches 8	

show pfe statistics traffic egress-queues

IN THIS SECTION

- Syntax | 483
- Description | 483
- Options | **483**
- Required Privilege Level | 484
- Output Fields | 484
- Sample Output | 484
- Release Information | 484

Syntax

show pfe statistics traffic egress-queues <fpc fpc-slot>

Description

(On EX8200 switches only) Display count of multidestination packets dropped on egress ports when the egress queues are oversubscribed due to multidestination traffic.

NOTE: Multidestination packets include unknown unicast, broadcast, and multicast packets.

Options

none Displays count of packets dropped on egress ports of all physical interfaces (line cards) when egress queues are oversubscribed due to multidestination traffic.

fpc fpc-
slot(Optional) Displays count of packets dropped on egress ports of the physical interface (line
card) referred to by the slot number.

NOTE: On an EX8200 switch, the FPC slot number is the slot number for the line card. Possible values are **0** through **7** on the EX8208 switch and **0** through **15** on the EX8216 switch.

Required Privilege Level

view

Output Fields

Table 58 on page 484 lists the output fields for the show pfe statistics traffic egress-queues command. Output fields are listed in the approximate order in which they appear.

Table 58: show pfe statistics traffic egress-queues Output Fields

Field Name	Field Description
Tail-dropped packets	Number of arriving packets dropped because the output queue buffers are full.

Sample Output

show pfe statistics traffic egress-queues fpc 4 (EX8208 Switch)

user@switch> show pfe statistics traffic egress-queues fpc 4
Tail-dropped packets : 0

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

show pfe statistics traffic cpu | 478

show pfe statistics traffic multicast | 485

show interfaces queue

Monitoring Interface Status and Traffic

Understanding Junos OS CoS Components for EX Series Switches | 8

show pfe statistics traffic multicast

IN THIS SECTION

- Syntax | 485
- Description | 485
- Options | 486
- Required Privilege Level | 486
- Output Fields | 486
- Sample Output | 487
- Release Information | 490

Syntax

show pfe statistics traffic multicast <fpc fpc-slot dev-number>

Description

(On EX8200 switches only) Display class-of-service (CoS) queue information for multidestination traffic on a physical interface (line card).

NOTE: Multidestination packets include unknown unicast, broadcast, and multicast packets.

NOTE: To view statistical information for unicast traffic, use the show interfaces queue command.

Options

fpc *fpcslot devnumber* (Optional) Displays class-of-service (CoS) queue information for multidestination traffic on the physical interface (line card) referred to by the slot number and device number.

NOTE: On an EX8200 switch, the FPC slot number is the slot number for the line card. Possible values for the FPC slot number are **0** through **7** on the EX8208 switch and **0** through **15** on the EX8216 switch. The value for the device number ranges from 0–5, where 0–4 values correspond to the statistics only from that specific device and the value 5 corresponds to the combined statistics from all the devices in the FPC.

Required Privilege Level

view

Output Fields

Table 59 on page 486 lists the output fields for the show pfe statistics traffic multicast command. Output fields are listed in the approximate order in which they appear.

Field Name	Field Description
Queue	CoS 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.
Packets	Number of packets transmitted by this queue.
Bytes	Number of bytes transmitted by this queue.

Field Name	Field Description
Tail-dropped packets	Count of packets dropped at the tail end of the queue because of lack of buffer space.
RED-dropped packets	 Number of packets dropped because of Random Early Discard (RED): Low-Number of low-loss priority packets dropped because of RED. High-Number of high-loss priority packets dropped because of RED.
RED-dropped bytes	 Number of bytes dropped because of Random Early Discard (RED): Low-Number of low-loss priority bytes dropped because of RED. High-Number of high-loss priority bytes dropped because of RED.

Table 59: show pfe statistics traffic multicast Output Fields (Continued)

Sample Output

show pfe statistics traffic multicast fpc 0 2(EX8208 Switch)

user@switch> show pfe st Queue: 0, Forwarding cla			
Queued:			
Packets	: Not Available		
Bytes	: Not Available		
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	; :	0	
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
High	:	0	0 bps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
High	:	0	0 pps
Queue: 1, Forwarding cla	asses: expedited-forwa	irding	
Queued:			
Packets	: Not Available		

Bytes	: Not Available		
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
High	:	0	0 bps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
High	:	0	0 pps
Queue: 2, Forwarding clas	sses: assured-forwardi	ng	
Queued:			
Packets	: Not Available		
Bytes	: Not Available		
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
High	:	0	0 bps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
High	:	0	0 pps
Queue: 3, Forwarding cla	sses: network-control		
Queued:			
Packets	: Not Available		
Bytes	: Not Available		
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
High	:	0	0 bps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
High	:	0	0 pps
Queue: 4			
Packets	: Not Available		
Bytes	: Not Available		
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	
RED-dropped bytes	:	0	0 bps

Low	:	0	0 bps
High	:	0	0 bps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
High	:	0	0 pps
Queue: 5			
Packets	: Not Available		
Bytes	: Not Available		
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
High	:	0	0 bps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
High	:	0	0 pps
Queue: 6			
Packets	: Not Available		
Bytes	: Not Available		
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
High	:	0	0 bps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
High	:	0	0 pps
Queue: 7			
Packets	: Not Available		
Bytes	: Not Available		
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
High	:	0	0 bps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
High	:	0	0 pps

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

show pfe statistics traffic cpu | 478

show pfe statistics traffic egress-queues | 483

show interfaces queue

Monitoring Interface Status and Traffic

Understanding Junos OS CoS Components for EX Series Switches | 8