Junos® OS for EX Series Ethernet Switches, Release 11.1: Class of Service

Published: 2011-07-07
Revision 4

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Junos® OS for EX Series Ethernet Switches, Release 11.1: Class of Service

Writing:
Editing:
Illustration:
Cover Design:

Revision History
July 2011—Revision 4
May 2011—Revision 3
April 2011—Revision 2
March 2011—Revision 1

The information in this document is current as of the date listed in the revision history.

END USER LICENSE AGREEMENT

READ THIS END USER LICENSE AGREEMENT ("AGREEMENT") BEFORE DOWNLOADING, INSTALLING, OR USING THE SOFTWARE. BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE OR OTHERWISE EXPRESSING YOUR AGREEMENT TO THE TERMS CONTAINED HEREIN, YOU (AS CUSTOMER OR IF YOU ARE NOT THE CUSTOMER, AS A REPRESENTATIVE/AGENT AUTHORIZED TO BIND THE CUSTOMER) CONSENT TO BE BOUND BY THIS AGREEMENT. IF YOU DO NOT OR CANNOT AGREE TO THE TERMS CONTAINED HEREIN, THEN (A) DO NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND (B) YOU MAY CONTACT JUNIPER NETWORKS REGARDING LICENSE TERMS.

1. **The Parties.** The parties to this Agreement are (i) Juniper Networks, Inc. (if the Customer's principal office is located in the Americas) or Juniper Networks (Cayman) Limited (if the Customer's principal office is located outside the Americas) (such applicable entity being referred to herein as "Juniper"), and (ii) the person or organization that originally purchased from Juniper or an authorized Juniper reseller the applicable license(s) for use of the Software ("Customer") (collectively, the "Parties").

2. **The Software.** In this Agreement, "Software" means the program modules and features of the Juniper or Juniper-supplied software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller, or which was embedded by Juniper in equipment which Customer purchased from Juniper or an authorized Juniper reseller. "Software" also includes updates, upgrades and new releases of such software. "Embedded Software" means Software which Juniper has embedded in or loaded onto the Juniper equipment and any updates, upgrades, additions or replacements which are subsequently embedded in or loaded onto the equipment.

3. **License Grant.** Subject to payment of the applicable fees and the limitations and restrictions set forth herein, Juniper grants to Customer a non-exclusive and non-transferable license, without right to sublicense, to use the Software, in executable form only, subject to the following use restrictions:

a. Customer shall use Embedded Software solely as embedded in, and for execution on, Juniper equipment originally purchased by Customer from Juniper or an authorized Juniper reseller.

b. Customer shall use the Software on a single hardware chassis having a single processing unit, or as many chassis or processing units for which Customer has paid the applicable license fees; provided, however, with respect to the Steel-Belted Radius or Odyssey Access Client software only, Customer shall use such Software on a single computer containing a single physical random access memory space and containing any number of processors. Use of the Steel-Belted Radius or IMS AAA software on multiple computers or virtual machines (e.g., Solaris zones) requires multiple licenses, regardless of whether such computers or virtualizations are physically contained on a single chassis.

c. Product purchase documents, paper or electronic user documentation, and/or the particular licenses purchased by Customer may specify limits to Customer's use of the Software. Such limits may restrict use to a maximum number of seats, registered endpoints, concurrent users, sessions, calls, connections, subscribers, clusters, nodes, realms, devices, links, ports or transactions, or require the purchase of separate licenses to use particular features, functionalities, services, applications, operations, or capabilities, or provide throughput, performance, configuration, bandwidth, interface, processing, temporal, or geographical limits. In addition, such limits may restrict the use of the Software to managing certain kinds of networks or require the Software to be used only in conjunction with other specific Software. Customer's use of the Software shall be subject to all such limitations and purchase of all applicable licenses.

d. For any trial copy of the Software, Customer's right to use the Software expires 30 days after download, installation or use of the Software. Customer may operate the Software after the 30-day trial period only if Customer pays for a license to do so. Customer may not extend or create an additional trial period by re-installing the Software after the 30-day trial period.

e. The Global Enterprise Edition of the Steel-Belted Radius software may be used by Customer only to manage access to Customer's enterprise network. Specifically, service provider customers are expressly prohibited from using the Global Enterprise Edition of the Steel-Belted Radius software to support any commercial network access services.

The foregoing license is not transferable or assignable by Customer. No license is granted herein to any user who did not originally purchase the applicable license(s) for the Software from Juniper or an authorized Juniper reseller.

4. **Use Prohibitions.** Notwithstanding the foregoing, the license provided herein does not permit the Customer to, and Customer agrees not to and shall not: (a) modify, unbundle, reverse engineer, or create derivative works based on the Software; (b) make unauthorized copies of the Software (except as necessary for backup purposes); (c) rent, sell, transfer, or grant any rights in and to any copy of the Software, in any form, to any third party; (d) remove any proprietary notices, labels, or marks on or in any copy of the Software or any product in which the Software is embedded; (e) distribute any copy of the Software to any third party, including as may be embedded in Juniper equipment sold in the secondhand market; (f) use any 'locked' or key-restricted feature, function, service, application, operation, or capability without first purchasing the applicable license(s) and obtaining a valid key from Juniper, even if such feature, function, service, application, operation, or capability is enabled without a key; (g) distribute any key for the Software provided by Juniper to any third party; (h) use the

Software in any manner that extends or is broader than the uses purchased by Customer from Juniper or an authorized Juniper reseller; (i) use Embedded Software on non-Juniper equipment; (j) use Embedded Software (or make it available for use) on Juniper equipment that the Customer did not originally purchase from Juniper or an authorized Juniper reseller; (k) disclose the results of testing or benchmarking of the Software to any third party without the prior written consent of Juniper; or (l) use the Software in any manner other than as expressly provided herein.

5. **Audit.** Customer shall maintain accurate records as necessary to verify compliance with this Agreement. Upon request by Juniper, Customer shall furnish such records to Juniper and certify its compliance with this Agreement.

6. **Confidentiality.** The Parties agree that aspects of the Software and associated documentation are the confidential property of Juniper. As such, Customer shall exercise all reasonable commercial efforts to maintain the Software and associated documentation in confidence, which at a minimum includes restricting access to the Software to Customer employees and contractors having a need to use the Software for Customer's internal business purposes.

7. **Ownership.** Juniper and Juniper's licensors, respectively, retain ownership of all right, title, and interest (including copyright) in and to the Software, associated documentation, and all copies of the Software. Nothing in this Agreement constitutes a transfer or conveyance of any right, title, or interest in the Software or associated documentation, or a sale of the Software, associated documentation, or copies of the Software.

8. **Warranty, Limitation of Liability, Disclaimer of Warranty.** The warranty applicable to the Software shall be as set forth in the warranty statement that accompanies the Software (the "Warranty Statement"). Nothing in this Agreement shall give rise to any obligation to support the Software. Support services may be purchased separately. Any such support shall be governed by a separate, written support services agreement. TO THE MAXIMUM EXTENT PERMITTED BY LAW, JUNIPER SHALL NOT BE LIABLE FOR ANY LOST PROFITS, LOSS OF DATA, OR COSTS OR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THIS AGREEMENT, THE SOFTWARE, OR ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. IN NO EVENT SHALL JUNIPER BE LIABLE FOR DAMAGES ARISING FROM UNAUTHORIZED OR IMPROPER USE OF ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. EXCEPT AS EXPRESSLY PROVIDED IN THE WARRANTY STATEMENT TO THE EXTENT PERMITTED BY LAW, JUNIPER DISCLAIMS ANY AND ALL WARRANTIES IN AND TO THE SOFTWARE (WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE), INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT DOES JUNIPER WARRANT THAT THE SOFTWARE, OR ANY EQUIPMENT OR NETWORK RUNNING THE SOFTWARE, WILL OPERATE WITHOUT ERROR OR INTERRUPTION, OR WILL BE FREE OF VULNERABILITY TO INTRUSION OR ATTACK. In no event shall Juniper's or its suppliers' or licensors' liability to Customer, whether in contract, tort (including negligence), breach of warranty, or otherwise, exceed the price paid by Customer for the Software that gave rise to the claim, or if the Software is embedded in another Juniper product, the price paid by Customer for such other product. Customer acknowledges and agrees that Juniper has set its prices and entered into this Agreement in reliance upon the disclaimers of warranty and the limitations of liability set forth herein, that the same reflect an allocation of risk between the Parties (including the risk that a contract remedy may fail of its essential purpose and cause consequential loss), and that the same form an essential basis of the bargain between the Parties.

9. **Termination.** Any breach of this Agreement or failure by Customer to pay any applicable fees due shall result in automatic termination of the license granted herein. Upon such termination, Customer shall destroy or return to Juniper all copies of the Software and related documentation in Customer's possession or control.

10. **Taxes.** All license fees payable under this agreement are exclusive of tax. Customer shall be responsible for paying Taxes arising from the purchase of the license, or importation or use of the Software. If applicable, valid exemption documentation for each taxing jurisdiction shall be provided to Juniper prior to invoicing, and Customer shall promptly notify Juniper if their exemption is revoked or modified. All payments made by Customer shall be net of any applicable withholding tax. Customer will provide reasonable assistance to Juniper in connection with such withholding taxes by promptly: providing Juniper with valid tax receipts and other required documentation showing Customer's payment of any withholding taxes; completing appropriate applications that would reduce the amount of withholding tax to be paid; and notifying and assisting Juniper in any audit or tax proceeding related to transactions hereunder. Customer shall comply with all applicable tax laws and regulations, and Customer will promptly pay or reimburse Juniper for all costs and damages related to any liability incurred by Juniper as a result of Customer's non-compliance or delay with its responsibilities herein. Customer's obligations under this Section shall survive termination or expiration of this Agreement.

11. **Export.** Customer agrees to comply with all applicable export laws and restrictions and regulations of any United States and any applicable foreign agency or authority, and not to export or re-export the Software or any direct product thereof in violation of any such restrictions, laws or regulations, or without all necessary approvals. Customer shall be liable for any such violations. The version of the Software supplied to Customer may contain encryption or other capabilities restricting Customer's ability to export the Software without an export license.

12. **Commercial Computer Software.** The Software is "commercial computer software" and is provided with restricted rights. Use, duplication, or disclosure by the United States government is subject to restrictions set forth in this Agreement and as provided in DFARS 227.7201 through 227.7202-4, FAR 12.212, FAR 27.405(b)(2), FAR 52.227-19, or FAR 52.227-14(ALT III) as applicable.

13. **Interface Information.** To the extent required by applicable law, and at Customer's written request, Juniper shall provide Customer with the interface information needed to achieve interoperability between the Software and another independently created program, on payment of applicable fee, if any. Customer shall observe strict obligations of confidentiality with respect to such information and shall use such information in compliance with any applicable terms and conditions upon which Juniper makes such information available.

14. **Third Party Software.** Any licensor of Juniper whose software is embedded in the Software and any supplier of Juniper whose products or technology are embedded in (or services are accessed by) the Software shall be a third party beneficiary with respect to this Agreement, and such licensor or vendor shall have the right to enforce this Agreement in its own name as if it were Juniper. In addition, certain third party software may be provided with the Software and is subject to the accompanying license(s), if any, of its respective owner(s). To the extent portions of the Software are distributed under and subject to open source licenses obligating Juniper to make the source code for such portions publicly available (such as the GNU General Public License ("GPL") or the GNU Library General Public License ("LGPL")), Juniper will make such source code portions (including Juniper modifications, as appropriate) available upon request for a period of up to three years from the date of distribution. Such request can be made in writing to Juniper Networks, Inc., 1194 N. Mathilda Ave., Sunnyvale, CA 94089, ATTN: General Counsel. You may obtain a copy of the GPL at http://www.gnu.org/licenses/gpl.html, and a copy of the LGPL at http://www.gnu.org/licenses/lgpl.html .

15. **Miscellaneous.** This Agreement shall be governed by the laws of the State of California without reference to its conflicts of laws principles. The provisions of the U.N. Convention for the International Sale of Goods shall not apply to this Agreement. For any disputes arising under this Agreement, the Parties hereby consent to the personal and exclusive jurisdiction of, and venue in, the state and federal courts within Santa Clara County, California. This Agreement constitutes the entire and sole agreement between Juniper and the Customer with respect to the Software, and supersedes all prior and contemporaneous agreements relating to the Software, whether oral or written (including any inconsistent terms contained in a purchase order), except that the terms of a separate written agreement executed by an authorized Juniper representative and Customer shall govern to the extent such terms are inconsistent or conflict with terms contained herein. No modification to this Agreement nor any waiver of any rights hereunder shall be effective unless expressly assented to in writing by the party to be charged. If any portion of this Agreement is held invalid, the Parties agree that such invalidity shall not affect the validity of the remainder of this Agreement. This Agreement and associated documentation has been written in the English language, and the Parties agree that the English version will govern. (For Canada: Les parties aux présentés confirment leur volonté que cette convention de même que tous les documents y compris tout avis qui s'y rattaché, soient redigés en langue anglaise. (Translation: The parties confirm that this Agreement and all related documentation is and will be in the English language)).

Table of Contents

About This Topic Collection

How to Use This Guide

Complete documentation for the EX Series product family is provided on webpages at http://www.juniper.net/techpubs/en_US/release-independent/information-products/pathway-pages/ex-series/product/index.html. We have selected content from these webpages and created a number of EX Series guides that collect related topics into a book-like format so that the information is easy to print and easy to download to your local computer.

The release notes are at http://www.juniper.net/techpubs/en_US/junos11.1/information-products/topic-collections/release-notes/11.1/junos-release-notes-11.1.pdf.

List of EX Series Guides for Junos OS Release 11.1

Title	Description
Complete Hardware Guide for EX2200 Ethernet Switches	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX2200 Ethernet switches
Complete Hardware Guide for EX3200 Ethernet Switches	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX3200 Ethernet switches
Complete Hardware Guide for EX4200 Ethernet Switches	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX4200 Ethernet switches
Complete Hardware Guide for EX4500 Ethernet Switches	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX4500 Ethernet switches

Title	Description
Complete Hardware Guide for EX8208 Ethernet Switches	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX8208 Ethernet switches
Complete Hardware Guide for EX8216 Ethernet Switches	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX8216 Ethernet switches
Complete Hardware Guide for the XRE200 External Routing Engine	Component descriptions, site preparation, installation, replacement, and safety and compliance information for the XRE200 External Routing Engine
Complete Software Guide for Junos® OS for EX Series Ethernet Switches, Release 11.1	Software feature descriptions, configuration examples, and tasks for Junos OS for EX Series switches
Software Topic Collections	Software feature descriptions, configuration examples and tasks, and reference pages for configuration statements and operational commands (This information also appears in the *Complete Software Guide for Junos® OS for EX Series Ethernet Switches, Release 11.1.*)
Junos® OS for EX Series Ethernet Switches, Release 11.1: Access Control	
Junos® OS for EX Series Ethernet Switches, Release 11.1: Configuration Management	
Junos® OS for EX Series Ethernet Switches, Release 11.1: Class of Service	
Junos® OS for EX Series Ethernet Switches, Release 11.1: Device Security	
Junos® OS for EX Series Ethernet Switches, Release 11.1: Ethernet Switching	
Junos® OS for EX Series Ethernet Switches, Release 11.1: EX4200 and EX4500 Virtual Chassis	
Junos® OS for EX Series Ethernet Switches, Release 11.1: EX8200 Virtual Chassis	
Junos® OS for EX Series Ethernet Switches, Release 11.1: Fibre Channel over Ethernet	
Junos® OS for EX Series Ethernet Switches, Release 11.1: High Availability	
Junos® OS for EX Series Ethernet Switches, Release 11.1: Interfaces	
Junos® OS for EX Series Ethernet Switches, Release 11.1: Layer 3 Protocols	
Junos® OS for EX Series Ethernet Switches, Release 11.1: MPLS	

Title	Description
Junos® OS for EX Series Ethernet Switches, Release 11.1: Multicast	
Junos® OS for EX Series Switches, Release 11.1: Network Management and Monitoring	
Junos® OS for EX Series Switches, Release 11.1: Port Security	
Junos® OS for EX Series Ethernet Switches, Release 11.1: Routing Policy and Packet Filtering	
Junos® OS for EX Series Ethernet Switches, Release 11.1: Software Installation	
Junos® OS for EX Series Ethernet Switches, Release 11.1: Spanning-Tree Protocols	
Junos® OS for EX Series Ethernet Switches, Release 11.1: System Monitoring	
Junos® OS for EX Series Ethernet Switches, Release 11.1: System Services	
Junos® OS for EX Series Ethernet Switches, Release 11.1: System Setup	
Junos® OS for EX Series Ethernet Switches, Release 11.1: User and Access Management	
Junos® OS for EX Series Ethernet Switches, Release 11.1: User Interfaces	

Downloading Software

You can download Junos OS for EX Series switches from the Download Software area at http://www.juniper.net/customers/support/ . To download the software, you must have a Juniper Networks user account. For information about obtaining an account, see http://www.juniper.net/entitlement/setupAccountInfo.do.

Documentation Symbols Key

Notice Icons		
Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.

Text and Syntax Conventions		
Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces important new terms. • Identifies book names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS System Basics Configuration Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Plain text like this	Represents names of configuration statements, commands, files, and directories; IP addresses; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Enclose optional keywords or variables.	**stub <default-metric** ***metric*****>;**

Text and Syntax Conventions		
Convention	Description	Examples
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** (***string1*** **|** ***string2*** **|** ***string3***)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Enclose a variable for which you can substitute one or more values.	**community name members [** ***community-ids*** **]**

Indention and braces ({ })	Identify a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  retain;
		  }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
J-Web GUI Conventions		
Bold text like this	Represents J-Web graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.
> (bold right angle bracket)	Separates levels in a hierarchy of J-Web selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. Send e-mail to techpubs-comments@juniper.net with the following:

- Document URL or title
- Page number if applicable
- Software version
- Your name and company

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://www.juniper.net/alerts/
- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

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

Opening a Case with JTAC

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

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

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

PART 1

Class of Service

CHAPTER 1

Class of Service (CoS)—Overview

Junos OS CoS for EX Series Switches Overview

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" on page 26 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" on page 26.**

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 5 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

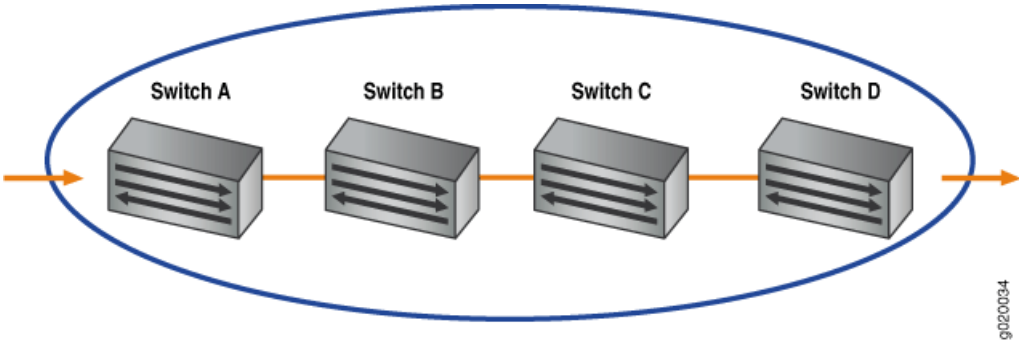

Default CoS Behavior on EX Series Switches

If you do not configure any CoS settings on the switch, the software performs some CoS functions to ensure 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 on page 6
- Understanding Junos OS EZQoS for CoS Configurations on EX Series Switches on page 24
- Example: Configuring CoS on EX Series Switches on page 33
- Example: Combining CoS with MPLS on EX Series Switches on page 48

Understanding Junos OS CoS Components for EX Series Switches

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: Rewrite rules are applied when the packets are routed. Rewrite rules are not applied when the packets are forwarded.

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

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.

Behavior aggregate classifiers use class-of-service (CoS) values such as Differentiated Services code points (DSCPs), IP precedence, and IEEE 802.1 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.

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

- dscp—Handles incoming IPv4 packets.
- ieee-802.1—Handles Layer 2 CoS.
- inet-precedence—Handles incoming IPv4 packets. IP precedence mapping requires only the upper three bits of the DSCP field.

This topic covers:

Default Code-Point Aliases

Table 1 on page 9 shows the default mappings between the bit values and standard aliases.

Table 1: Default Code-Point 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
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
IEEE 802.1p CoS Values	
be	000

Table 1: Default Code-Point Aliases *(continued)*

CoS Value Types	Mapping
be1	001
ef	100
ef1	101
af11	010
af12	011
nc1/cs6	110
nc2/cs7	111
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 on page 6
- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Code-Point Aliases (CLI Procedure) on page 64
- Defining CoS Code-Point Aliases (J-Web Procedure) on page 62

Understanding CoS Classifiers

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 assign packets to output queues based on the associated forwarding class. 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.

NOTE: When a source media access control (MAC) address is learned, the frame that contains the source MAC address is always sent out on queue 0 while egressing from the network interface, irrespective of the classifier applied to the ingress interface.

On Juniper Networks EX8200 Ethernet Switches, you can specify BA classifiers for bridged multidestination traffic and IP multidestination traffic. The 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 on page 11
- Multifield Classifiers on page 13

Behavior Aggregate Classifiers

The behavior aggregate classifier maps a CoS value to a forwarding class and loss priority. The forwarding class determines the output queue. The loss priority is used by a scheduler to control packet discard during periods of congestion.

There are three types of BA classifiers:

- Differentiated Services Code Point (DSCP) for IP DiffServ

- IP precedence bits
- IEEE 802.1p CoS bits

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

Default Behavior Aggregate Classification

Juniper Networks Junos operating system (Junos OS) automatically assigns implicit default classifiers to all logical interfaces based on the type of interface. Table 2 on page 12 lists different types of interfaces and the corresponding implicit default classifiers.

Table 2: Default BA Classification

Type of Interface	Default BA Classification
Trunk interface	**ieee8021p-default**
Layer 3 interface (IPv4)	**dscp-default**
Layer 3 interface (IPv6)	**dscp-ipv6-default**
Access interface	Untrusted
Routed VLAN interface (RVI)	No default classification

When you explicitly associate a classifier with a logical interface, you are in effect overriding the implicit default classifier with an explicit classifier.

On Juniper Networks EX3200 and EX4200 Ethernet Switches, you can apply classifier rules for each interface. Table 3 on page 12 describes the different classifier types you can configure on Layer 2 and Layer 3 interfaces.

Table 3: 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 can configure all the allowed classifier types on the same logical interface or on different logical interfaces. If you need to apply all classifier rules on the same logical interface, configure the classifier rules allowed for both IPv4 and IPv6 on the logical interface.

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

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

NOTE: **On EX8200 switches, only one classifier of type DSCP and of type IEEE 802.1p can be applied to an interface.**

You can configure routed VLAN interfaces (RVIs) to classify packets. 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 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 DSCP or IP precedence support 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 on page 6
- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Classifiers (CLI Procedure) on page 65
- Defining CoS Classifiers (J-Web Procedure) on page 66

Understanding CoS Forwarding Classes

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.

NOTE: **The forwarding classes multicast expedited-forwarding, multicast assured-forwarding, and multicast best-effort are applicable only to Juniper Networks EX8200 Ethernet 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. Therefore, if you configure more than eight forwarding classes, you must map multiple forwarding classes to single 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 on page 14

Default Forwarding Classes

Table 4 on page 15 shows the four default forwarding classes defined for unicast traffic, and Table 5 on page 15 shows the three default forwarding classes defined for multicast traffic.

NOTE: The default forwarding classes for multicast traffic are applicable only to EX8200 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 4: 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 5: Default Forwarding Classes for Multicast Traffic

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.
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.

Table 5: Default Forwarding Classes for Multicast Traffic *(continued)*

Forwarding Class Name	Comments
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.

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 on page 6
- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Forwarding Classes (CLI Procedure) on page 68
- Defining CoS Forwarding Classes (J-Web Procedure) on page 68

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.

On Juniper Networks EX Series Ethernet Switches, drop probability is implicitly set to **100 percent** and it cannot be modified.

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 on page 6
- Example: Configuring CoS on EX Series Switches on page 33
- Configuring CoS Tail Drop Profiles (CLI Procedure) on page 75

Understanding CoS Schedulers

You use schedulers to define the 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, 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. Only two forwarding classes, best-effort (queue0) and network-control (queue7) are used in the default configuration.

NOTE: On Juniper Networks EX8200 Ethernet Switches three forwarding classes—best-effort (queue0), multicast best-effort (queue2), and network-control (queue7)—are used in the default configuration.

By default, 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.

NOTE: On EX8200 switches, 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.

The expedited-forwarding (queue 5) and assured-forwarding (queue 1) classes have no scheduler because no resources are assigned to queue 5 and queue 1, by default. However, you can manually configure resources for the expedited-forwarding and assured-forwarding classes.

Also by default, each 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 need to carry traffic load that exceeds their allocated bandwidth.

Transmission Rate

The transmission-rate control determines the actual traffic bandwidth from 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 outgoing 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, configured amount of transmission rate is guaranteed for the queue. This property allows you to ensure that each queue receives the amount of bandwidth appropriate to its level of service.

Scheduler Buffer Size

To control congestion at the output stage, you can configure the delay-buffer bandwidth using the **buffer-size** command. The delay-buffer bandwidth 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.

The default scheduler transmission rate for queues 0 through 7 are 95, 0, 0, 0, 0, 0, 0, and 5 percent 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 of the total available buffer.

NOTE: **On EX8200 switches, the default scheduler transmission rates for queues 0 through 7 are 75, 0, 20, 0, 0, 0, 0, and 5 percent 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 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.

Priority Scheduling

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

Priority scheduling is accomplished through a procedure in which the scheduler examines the priority of the queue. The 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. This binary decision, which is re-evaluated on a regular time cycle, compares the amount of data transmitted by the queue against the bandwidth allocated to it by the scheduler. When 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 queue's allocated limit. Out-of-profile queue will be transmitted only if bandwidth is available. Otherwise, it will be buffered.

  A queue from the set is selected based on the shaped deficit weighted round robin (SDWRR) algorithm, which operates within the set.

- Strict-high—Strict-high priority queue receives preferential treatment over low priority queue. Unlimited bandwidth is assigned to strict-high priority queue. 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, in case of two high priority queues, the queue with higher queue number is processed first.

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.

Juniper Networks EX2200 Ethernet Switches allow you to associate up to five user-defined scheduler maps and Juniper Networks EX3200, EX4200, and EX4500 Ethernet Switches allow you to associate up to four user-defined scheduler maps with interfaces. EX8200 switches allow up to six user-defined scheduler maps with a port group, which is a set of network ports in an EX Series switch. A certain bandwidth is available for each port group and each port in a port group can utilize a portion of that bandwidth. However, the total utilization of bandwidth in a port group cannot exceed the bandwidth available for the port group.

EX8200 switches use four types of line cards:

- 8-port SFP+ line card
- 40-port SFP+ line card
- 48-port SFP line card
- 48-port RJ-45 line card

A 40-port SFP+ line card is an oversubscribed 10-Gigabit Ethernet line card for the EX8200 switch. The 40 ports in this line card are divided into eight port groups, with 5 ports in each port group. The eight port groups are ports 0-4, ports, 5-9, ports 10-14, ports 15-19, ports 20-24, ports 25-29, ports 30-34, and ports 35-39. Each port group shares 10 gigabits of bandwidth and you can configure a maximum of six scheduler maps for each port group.

In an 8-port 10-Gigabit Ethernet SFP+ line card, the eight ports in the line card are divided into four port groups with two network ports in each port group. The four port groups are ports 0-1, ports 2-3, ports 4-5, and ports 6-7. Each port in this line card supports 10 gigabits of bandwidth. In this line card, six scheduler maps are available for two ports in the port group, and hence, a separate scheduler maps is available for each of the ports.

In a 48-port line card, the 48 ports in the line card are divided into two port groups with 24 ports in each port group. The two ports groups are ports 0-23 and ports 24-47. The 24 ports in each port group share a bandwidth of 1 gigabit. In this case, six scheduler maps are available for the 24 ports in each port group.

If you configure more than the supported number of scheduler maps in a switch, an error is logged in the system log (syslog) and the default scheduler map is bound to all port groups. We recommend that you check the system log for errors (if any) after the commit operation to verify that you have not configured more than the maximum permitted number of scheduler maps.

Related Documentation

- Understanding Junos OS CoS Components for EX Series Switches on page 6
- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Schedulers (CLI Procedure) on page 70
- Defining CoS Schedulers (J-Web Procedure) on page 71

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 on page 6
- Understanding the Use of Policers in Firewall Filters
- Configuring Policers to Control Traffic Rates (CLI Procedure)

Understanding CoS Rewrite Rules

As packets enter or exit a network, edge switches might be required to alter the class-of-service (CoS) settings of the packets. 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. Only tagged Layer 3 interfaces and tagged routed VLAN interfaces (RVIs) automatically rewrite packets by using the default IEEE 802.1p rewrite rule. Multiple rewrite rules of different types can be assigned to a single interface.

NOTE: On Juniper Networks EX8200 Ethernet Switches, tagged Layer 3 interfaces and tagged RVIs do not automatically rewrite packets using the default IEEE 802.1p rewrite rule. You must explicitly assign the IEEE 802.1p rewrite rule to these interfaces for rewrites to occur.

Also, only one rewrite rule of each type can be assigned to any interface on an EX8200 switch.

In effect, the rewrite rule performs the opposite 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 76.

Table 6 on page 23 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 6: 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

Table 6: Default Packet Header Rewrite Mappings *(continued)*

Map from Forwarding Class	PLP Value	Map to DSCP/IEEE 802.1p/IP Precedence Value
network-control	low	nc1/cs6
network-control	high	nc2/cs7

Related Documentation

- Understanding Junos OS CoS Components for EX Series Switches on page 6
- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Rewrite Rules (CLI Procedure) on page 76
- Defining CoS Rewrite Rules (J-Web Procedure) on page 77

Understanding Port Shaping and Queue Shaping for CoS on EX Series Switches

If the amount of traffic on a switch's network interface is more than the maximum bandwidth allowed on the interface, it leads to congestion. Port shaping and queue shaping can be used to manage the excess traffic and avoid congestion. Port shaping defines the maximum bandwidth allocated to a port, while queue shaping defines a limit on excess-bandwidth usage per queue.

This topic covers:

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 or port rate.

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.

Related Documentation

- Understanding CoS Schedulers on page 18
- Defining CoS Schedulers (CLI Procedure) on page 70

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 on page 4
- Configuring Junos OS EZQoS for CoS (CLI Procedure) on page 81

Understanding Using CoS with MPLS Networks on EX Series Switches

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 prior to 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 DSCP, IP precedence, or IEEE 802.1 p 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. 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).

This topic includes:

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 *cannot* use more than one type of DSCP/IP precedence and not more than one type of IEEE 802.1p classifier on the CCC interfaces. Thus, if you configure one CCC interface to use DSCP1, you cannot configure another CCC interface to use DSCP2.

Likewise, if you configure one CCC interface to use IEEE1, you cannot configure another CCC interface on the same switch to use IEEE2. All the CCC interfaces on the switch must use the same DSCP classifier and the same type of IEEE 802.1p classifier.

- You *cannot* configure one CCC interface as DSCP and another CCC interface as IP precedence, because these classifier types overlap.
- You *can* configure one CCC interface as DSCP and another CCC interface as IEEE 802.1p.
- You *can* configure one CCC interface as both DSCP and IEEE 802.1p. If you configure a CCC interface with 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.

NOTE: **You can define multiple types of DSCP, IP precedence, and IEEE 802.1p on the switch and use the different classifier types for the non-CCC interfaces on the 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 that has been set up 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 regarding using multiple types of DSCP, IP precedence, and IEEE 802.1p on the same switch.

- You can modify the CoS classifier for a particular interface, but it is not required.
- You can configure one interface as DSCP1 and another as DSCP2 and another and IP precedence, and so forth.

Default Classifiers and Default Rewrite Rules

The default classifiers support only two forwarding classes, **best-effort** and **network-control**, and use only two queues, **0** and **7**. However, EX Series switches support up to sixteen forwarding classes and eight queues. To use the additional forwarding classes and queues, create a custom classifier. To modify the code point and loss priority for a specific forwarding class, configure a rewrite rule on the switch. The default rewrite rule for EXP is enabled in the default configuration. However, the default rewrite rules for the other classifiers are not enabled in the default configuration. You can display the default classifier mappings and default rewrite mappings by entering the **show class-of-service** command on the switch.

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 MPLS-enabled 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 should configure a policer on the ingress PE switch:

- If you are using MPLS with 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.

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 other custom forwarding classes, we recommend that you configure a scheduler to support that forwarding class. See “Understanding CoS Schedulers” on page 18.

Related Documentation

- Junos OS MPLS for EX Series Switches Overview
- Understanding CoS Classifiers on page 11
- Understanding CoS Schedulers on page 18
- Example: Configuring CoS on EX Series Switches on page 33
- Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect (CLI Procedure) on page 83
- Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS (CLI Procedure) on page 82
- Configuring Rewrite Rules for EXP Classifiers on MPLS Networks (CLI Procedure)
- Configuring CoS on Provider Switches of an MPLS Network (CLI Procedure) on page 86
- Defining CoS Rewrite Rules (CLI Procedure) on page 76
- Configuring Policers to Control Traffic Rates (CLI Procedure)

Understanding CoS Queues on the 40-port SFP+ Line Card on EX8200 Switches

The 40-port SFP+ line card is an oversubscribed 10-Gigabit Ethernet line card for the Juniper Networks EX8200 Ethernet Switches. The 40 ports on the line card are divided into 8 port groups, with each port group containing 5 ports. The ports in a port group share 10 gigabits of bandwidth. Because the port groups share bandwidth, class-of-service (CoS) ingress and egress queues are handled differently on the 40-port SFP+ line card than on other line cards for EX8200 switches.

This topic describes:

Ingress Queues on the 40-port SFP+ Line Card

Ingress packet classification and queuing occurs in two steps:

Preclassification of Packets and Port Ingress Queuing

Packets entering the ports on a port group are sent to one of two ingress queues. These ingress queues are used to schedule the traffic from the port group into the Packet Forwarding Engine.

- Low priority queue—Each interface in a port group 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 in the port group having an equal weight.
- High priority queue—The interfaces in a port group share a single high priority queue. Traffic on this queue is scheduled by strict-high priority.

For the purpose of port ingress queuing, packets are classified only by behavior aggregate (BA) classification. To control which ingress queue the packets get sent to, you 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 40-port SFP+ 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 are handled differently from other packets. Instead of using the BA classifier to classify them, the switch always sends critical network-control packets to the high priority queue. This handling ensures that these packets are not dropped because of congestion on the oversubscribed ports.

Full Classification of Packets and Fabric Ingress Queueing

When the packets from a port group 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 fabric ingress queues.

Egress Queues on the 40-port SFP+ Line Card

As with all EX Series switch interfaces, each interface on the 40-port SFP+ line card supports eight egress CoS queues. You can map up to 16 forwarding classes to these queues.

All interfaces in a port group also share a single set of eight egress chassis queues at the Packet Forwarding Engine. Egress traffic is fanned out from the Packet Forwarding Engine chassis 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 to the system log and the default scheduler map is used for the ports in the port group.

Related Documentation

- Understanding Junos OS CoS Components for EX Series Switches on page 6
- Understanding CoS Schedulers on page 18
- Understanding CoS Forwarding Classes on page 14
- Example: Configuring CoS on EX Series Switches on page 33
- Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure) on page 85

Understanding Priority-Based Flow Control

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 on page 30
- Calculations for Buffer Requirements When Using PFC PAUSE on page 31
- How PFC and Congestion Notification Profiles Work on page 31

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 in order to prevent buffer overflow. Upon receiving a PAUSE frame, the sender stops transmissions of any new packets until the receiver 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

Receivers must ensure that a PFC PAUSE frame is sent while there is sufficient receive buffer to absorb the data that might continue to be received while the system is responding to the PFC PAUSE.

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 reduced buffer space occurs 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.**

How PFC and Congestion Notification Profiles Work

PFC is triggered when the incoming frame has a User Priority (UP) field that matches the three-bit pattern specified for the PFC congestion notification profile, which you have configured. Table 7 on page 32 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 6 traffic classes and allow you to associate those classes with 6 different congestion notification profiles. (The switches support up to 16 forwarding classes.)

Table 7: 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	TC 0	queue 0
001	TC 1	queue 1
010	TC 2	queue 2
011	TC 3	queue 3
100	TC4	queue 4
101	TC 5	queue 5

Related Documentation

- Example: Configuring FIP Snooping and Priority-Based Flow Control on an FCoE Transit Switch
- Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure)
- **schedulers on page 128**
- congestion-notification-profile

CHAPTER 2

Examples: CoS Configuration

Example: Configuring CoS on EX Series Switches

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:

- Junos OS Release 9.0 or later for EX Series switches
- One Juniper Networks EX3200 switch

Overview and Topology

This example uses the topology shown in Figure 2 on page 34.

Figure 2: Topology for Configuring CoS

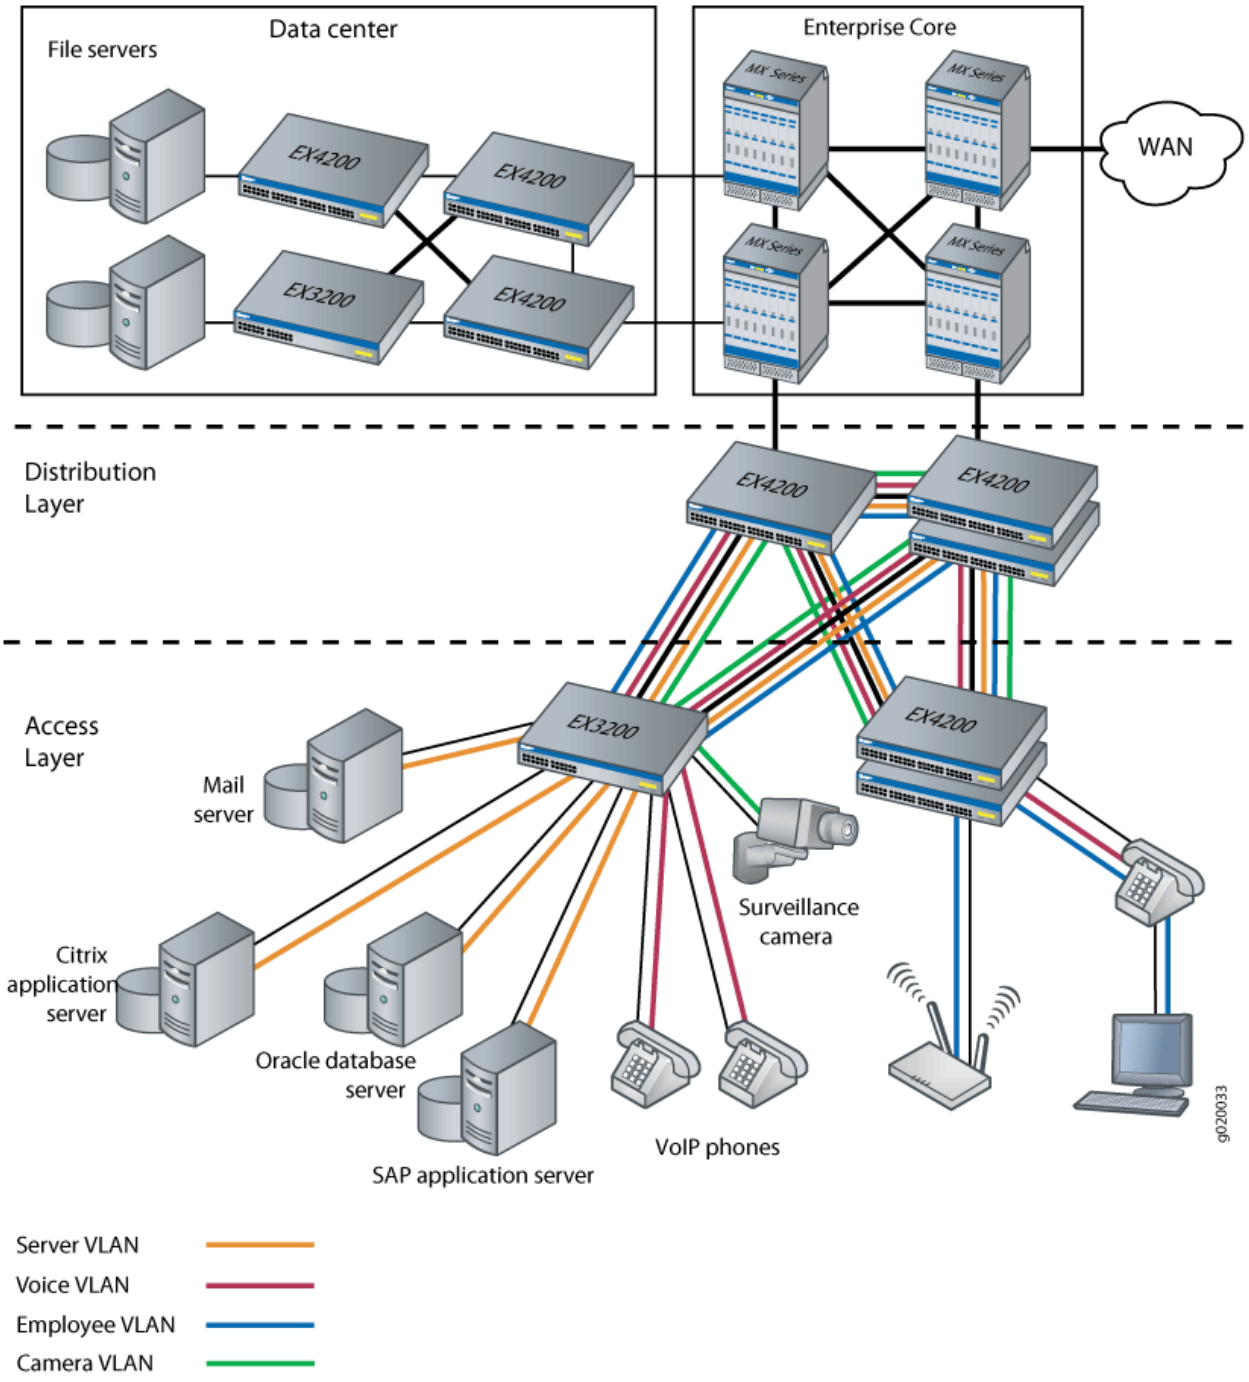

The topology for this configuration example consists of one EX Series switch at the access layer.

The EX Series access switch is configured to support VLAN membership. Switch ports **ge-0/0/0**and **ge-0/0/1** are assigned to the **voice-vlan** for two VoIP phones. Switch port **ge-0/0/2** is assigned to the **camera-vlan** for the surveillance camera. Switch ports

ge-0/0/3, **ge-0/0/4**, **ge-0/0/5**, and **ge-0/0/6** are assigned to the **server-vlan** for the servers hosting various applications such as those provided by Citrix, Microsoft, Oracle, and SAP.

Table 8 on page 35 shows the VLAN configuration components.

Table 8: Configuration Components: VLANs

VLAN Name	VLAN ID	VLAN Subnet and Available IP Addresses	VLAN Description
voice-vlan	10	**192.168.1.0/32** **192.168.1.1** through **192.168.1.11** **192.168.1.12** is the subnet's broadcast address.	Voice VLAN used for employee VoIP communication.
camera-vlan	20	**192.168.1.13/32** **192.168.1.14** through **192.168.1.20** **192.168.1.21** is the subnet's broadcast address.	VLAN for the surveillance cameras.
server-vlan	30	**192.168.1.22/32** **192.168.1.23** through **192.168.1.35** **192.168.1.36** is the subnet's broadcast address.	VLAN for the servers hosting enterprise applications.

Ports on the EX Series switches support Power over Ethernet (PoE) to provide both network connectivity and power for VoIP telephones connecting to the ports. Table 9 on page 35 shows the switch interfaces that are assigned to the VLANs and the IP addresses for devices connected to the switch ports:

Table 9: Configuration Components: Switch Ports 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** through **192.168.1.2**	Two VoIP telephones.
ge-0/0/2	**camera-vlan**	**192.168.1.14**	Surveillance camera.
ge-0/0/3, **ge-0/0/4**, **ge-0/0/5**, **ge-0/0/6**	**sevrer-vlan**	**192.168.1.23** through **192.168.1.26**	Four servers hosting applications such as those provided by Citrix, Microsoft, Oracle, and SAP.

NOTE: This example shows how to configure CoS on a single 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.

Configuration

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 192.168.1.1/32
set firewall family ethernet-switching filter voip_class term voip from source-address 192.168.1.2/32
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 forwarding-class
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 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
192.168.1.14/32
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 forwarding-class
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
192.168.1.23/32
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
loss-priority low
set firewall family ethernet-switching filter app_class term mail from source-address
192.168.1.24/32
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 192.168.1.25/32
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 loss-priority
low
set firewall family ethernet-switching filter app_class term erp from source-address 192.168.1.26/32
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
loss-priority 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 forwarding-class
network-control loss-priority low
set firewall family ethernet-switching filter app_class term best_effort_traffic then forwarding-class
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 buffer-size percent 10
set class-of-service schedulers voice-sched priority strict-high
set class-of-service schedulers voice-sched transmit-rate percent 10
set class-of-service schedulers video-sched buffer-size percent 15
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 buffer-size percent 5
set class-of-service schedulers nc-sched priority strict-high
set class-of-service schedulers nc-sched transmit-rate percent 5
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
```

Step-by-Step Procedure

To configure and apply CoS:

1. Configure one-to-one mapping 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 192.168.1.1/32
user@switch# set family ethernet-switching filter voip_class term voip from
source-address 192.168.1.2/32
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**:

```
[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:

```
[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. Define the firewall filter **video_class** to classify the video traffic:

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

8. Define the term **video**:

```
[edit firewall]
user@switch# set family ethernet-switching filter video_class term video from
source-address 192.168.1.14/32
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
```

9. 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
```

10. Define the term **best_effort_traffic** (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
```

11. 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
```

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

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

13. Define the term **app**:

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term app from
source-address 192.168.1.23/32
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
```

14. Define the term **mail**:

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term mail from
source-address 192.168.1.24/32
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
```

15. Define the term **db**:

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term db from source-address
192.168.1.25/32
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
```

16. Define the term **erp**:

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term erp from
source-address 192.168.1.26/32
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
```

17. 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
```

18. 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
```

19. 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
```

20. Configure schedulers:

```
[edit class-of-service]
user@switch# set schedulers voice-sched buffer-size percent 10
user@switch# set schedulers voice-sched priority strict-high
user@switch# set schedulers voice-sched transmit-rate percent 10
user@switch# set schedulers video-sched buffer-size percent 15
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 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
```

21. 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
voice-sched
user@switch# set scheduler-maps ethernet-cos-map forwarding-class video scheduler
video-sched
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-sched
```

22. Associate the scheduler map with the outgoing interface:

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

Results Display the results of the configuration:

```
user@switch# show firewall

firewall family ethernet-switching {
  filter voip_class {
    term voip {
      from {
        source-address {
          192.168.1.1/32;
          192.168.1.2/32;
        }
        protocol udp;
        source-port 2698;
      }
      then {
        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 {
          192.168.1.14/32;
        }
        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 {
          192.168.1.23/32;
        }
        protocol tcp;
        source-port [1491 2512 2513 2598 2897];
      }
```

```
    then {
      forwarding-class app;
      loss-priority low;
    }
  }
  term mail {
    from {
      source-address {
        192.168.1.24/32;
      }
      protocol tcp;
      source-port [25 143 389 691 993 3268 3269];
    }
    then {
      forwarding-class mail;
      loss-priority low;
    }
  }
  term db {
    from {
      source-address {
        192.168.1.25/32;
      }
      protocol tcp;
      source-port [1521 1525 1527 1571 1810 2481];
    }
    then {
      forwarding-class db;
      loss-priority low;
    }
  }
  term erp {
    from {
      source-address {
        192.168.1.26/32;
      }
      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;
}
schedulers {
    voice-sched {
        buffer-size percent 10;
        priority strict-high;
        transmit-rate percent 10;
    }
    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 {
        buffer-size percent 5;
        priority strict-high;
        transmit-rate percent 5;
    }
    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;
    }
  }
```

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

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

- Verifying That the Defined Forwarding Classes Exist and Are Mapped to Queues on page 46
- Verifying That the Forwarding Classes Have Been Assigned to Schedulers on page 47
- Verifying That the Scheduler Map Has Been Applied to the Interface on page 48

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

Purpose Verify that the following forwarding classes **app**, **db**, **erp**, **mail**, **video**, and **voice** have been defined and mapped to queues.

Action

```
user@switch> show class-of-service forwarding-class
Forwarding class                          ID    Queue
  app                                      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
    Transmit rate: 5 percent, Rate Limit: none, Buffer size: 15 percent,
    Priority: Strict-high
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             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>
      High            TCP             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         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  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>
      High            TCP             1    <default-drop-profile>

  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
      High            non-TCP         1    <default-drop-profile>
      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
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  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>
```

```
    High            TCP             1     <default-drop-profile>

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

Purpose Verify that the scheduler map has been applied to the interface.

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
...
```

Meaning This output shows that the scheduler map (**ethernet-cos-map**) has been applied to the interface (**ge-0/0/20**).

Related Documentation

- Defining CoS Code-Point Aliases (CLI Procedure) on page 64
- Defining CoS Classifiers (CLI Procedure) on page 65
- Defining CoS Forwarding Classes (CLI Procedure) on page 68
- Defining CoS Schedulers (CLI Procedure) on page 70
- Configuring CoS Tail Drop Profiles (CLI Procedure) on page 75
- Assigning CoS Components to Interfaces (CLI Procedure) on page 79
- Configuring Firewall Filters (CLI Procedure)

Example: Combining CoS with MPLS on EX Series Switches

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 EX Series Switches. This example assumes that an MPLS network has been configured using a cross circuit-connect (CCC).

Overview and Topology

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 **assured-forwarding**, **expedited-forwarding**, or other custom forwarding classes, we recommend that you configure a scheduler to support that forwarding class. See “Defining CoS Schedulers (CLI Procedure)” on page 70.

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 EX Series 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 loss-priority 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 10 on page 50 shows the CoS configuration components added to the ingress PE switch.

Table 10: 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 11 on page 50 shows the CoS configuration components added to the egress PE switch in this example.

Table 11: 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

Table 11: CoS Configuration Components of the Egress PE Switch *(continued)*

Property	Settings	Description
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 12 on page 51 shows the MPLS configuration components used for the provider switch in this example.

Table 12: 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.
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.

Configuring the Local PE Switch

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 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 ]
   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

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

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

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

- Verifying That the Policer Firewall Filter Is Operational on page 55
- Verifying That the CoS Classifiers Are Going to the Right Queue on page 55
- Verifying the CoS Forwarding Table Mapping on page 58
- Verifying the Rewrite Rules on page 59

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@switch> show class-of-service forwarding-table classifier

Classifier table index: 7, # entries: 64, Table type: DSCP
```

Entry #	Code point	Forwarding-class #	PLP
0	000000	0	0
1	000001	0	0
2	000010	0	0
3	000011	0	0
4	000100	0	0
5	000101	0	0
6	000110	0	0
7	000111	0	0
8	001000	0	0
9	001001	0	0
10	001010	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, 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

Classifier table index: 12, # entries: 8, Table type: IPv4 precedence
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

Classifier table index: 16, # entries: 8, Table type: Untrust
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        0       0
   7           111        0       0

Classifier table index: 9346, # entries: 64, Table type: DSCP
Entry #   Code point   Forwarding-class #   PLP
   0        000000        0       0
   1        000001        0       0
   2        000010        0       0
   3        000011        0       0
   4        000100        0       0
   5        000101        0       0
   6        000110        0       0
   7        000111        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@switch>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.

Related Documentation

- Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure)
- Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure)
- Understanding Using CoS with MPLS Networks on EX Series Switches on page 26
- Monitoring CoS Forwarding Classes on page 88

CHAPTER 3

Configuring CoS

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) on page 66
- Defining CoS Code-Point Aliases (J-Web Procedure) on page 62
- Defining CoS Forwarding Classes (J-Web Procedure) on page 68
- Defining CoS Rewrite Rules (J-Web Procedure) on page 77
- Defining CoS Schedulers (J-Web Procedure) on page 71
- Assigning CoS Components to Interfaces (J-Web Procedure) on page 79

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

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 in 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 for details about all commit options.

2. Click one:
   - **Add**—Adds a code-point alias. Enter information into the code point alias page as described in Table 13 on page 63.
   - **Edit**—Modifies an existing code-point alias. Enter information into the code point alias page as described in Table 13 on page 63.
   - **Delete**—Deletes an existing code-point alias.

Table 13 on page 63 describes the related fields.

Table 13: 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.
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 xxxxxx, 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**.

Related Documentation

- Defining CoS Code-Point Aliases (CLI Procedure) on page 64
- Monitoring CoS Value Aliases on page 93
- Example: Configuring CoS on EX Series Switches on page 33

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—Handles incoming IPv4 packets.
- IEEE 802.1p—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 code-point (**110001**):

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

Related Documentation

- Defining CoS Code-Point Aliases (J-Web Procedure) on page 62
- Example: Configuring CoS on EX Series Switches on page 33
- Monitoring CoS Value Aliases on page 93
- Understanding CoS Code-Point Aliases on page 8

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.
- 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 and apply it 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 14 on page 65, 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).

Table 14: BA-classifier Loss Priority Assignments

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**

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
```

Related Documentation

- Defining CoS Classifiers (J-Web Procedure) on page 66
- Example: Configuring CoS on EX Series Switches on page 33
- Assigning CoS Components to Interfaces (CLI Procedure) on page 79
- Monitoring CoS Classifiers on page 87
- Understanding CoS Classifiers on page 11
- Troubleshooting a CoS Classifier Configuration for a TCAM Space Error on page 96

Defining CoS Classifiers (J-Web Procedure)

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 in 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 for details about all commit options.

2. Click one:
   - **Add**—Adds a classifier. Enter information into the classifier page as described in Table 15 on page 67.
   - **Edit**—Modifies an existing classifier. Enter information into the classifier page as described in Table 15 on page 67.
   - **Delete**—Deletes an existing classifier.

Table 15: 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.
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. 4. 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.

Related Documentation

- Defining CoS Classifiers (CLI Procedure) on page 65
- Example: Configuring CoS on EX Series Switches on page 33
- Monitoring CoS Classifiers on page 87
- Understanding CoS Classifiers on page 11

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) on page 68
- Example: Configuring CoS on EX Series Switches on page 33
- Assigning CoS Components to Interfaces (CLI Procedure) on page 79
- Monitoring CoS Forwarding Classes on page 88
- Understanding CoS Forwarding Classes on page 14

Defining CoS Forwarding Classes (J-Web Procedure)

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 in 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 for details about all commit options.

2. Click one:

   - **Add**—Adds a forwarding class. Enter information into the forwarding class page as described in Table 16 on page 69.
   - **Edit**—Modifies an existing forwarding class. Enter information into the forwarding class page as described in Table 16 on page 69.
   - **Delete**—Deletes an existing forwarding class.

Table 16: Forwarding Classes Configuration Fields

Field	Function	Your Action
Forwarding Class Summary		
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 7, appropriate for your platform.
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**.

Related Documentation

- Defining CoS Forwarding Classes (CLI Procedure) on page 68
- Example: Configuring CoS on EX Series Switches on page 33
- Monitoring CoS Forwarding Classes on page 88
- Assigning CoS Components to Interfaces (J-Web Procedure) on page 79
- Understanding CoS Forwarding Classes on page 14

Defining CoS Schedulers (CLI Procedure)

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 tail 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.

You can associate up to four user-defined scheduler maps with the interfaces.

This topic describes:

- Configuring CoS Schedulers on page 70
- Assigning Scheduler Maps to Interfaces on a 40-port SFP+ Line Card on page 70

Configuring CoS Schedulers

To configure CoS schedulers:

1. Create a scheduler (**be-sched**) and assign it a priority:

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

2. Configure a scheduler map (**be-map**) that associates the scheduler (**be-sched**) with the forwarding class (**best-effort**):

   ```
   [edit class-of-service scheduler-maps]
   user@switch# set be-map forwarding-class best-effort scheduler be-sched
   ```

3. Assign the scheduler map (**be-map**) to one or more Ethernet interfaces:

   - To assign the scheduler map to one interface (**ge-0/0/1**):

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

   - To assign the scheduler map to more than one interface by using a wildcard (all Gigabit Ethernet interfaces):

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

Assigning Scheduler Maps to Interfaces on a 40-port SFP+ Line Card

For interfaces on a 40-port SFP+ line card, 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. The switch automatically uses that scheduler map for the interfaces in the port group when you bring the interfaces up.

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 to the system log. When you bring an interface in the port group up, the default scheduler map is used. If you assign a scheduler map to an interface that is down that is different from the scheduler map being used by the currently operating interfaces in a port group, the default scheduler map is used by all interfaces in a port group, even the currently operating ones, when you bring the interface up.

To change the scheduler map assigned to a port group:

1. Delete the current scheduler map from the interfaces (**xe-0/0/1** and **xe-0/0/2**) it is currently assigned to:

   ```
   [edit class-of-service interfaces]
   user@switch# delete xe-0/0/1 scheduler-map

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

2. Assign the new scheduler map (**ef-map**) to at least one interface in the port group:

   ```
   [edit class-of-service interfaces]
   user@switch# set xe-0/0/1 scheduler-map ef-map
   ```

Related Documentation

- Defining CoS Schedulers (J-Web Procedure) on page 71
- Example: Configuring CoS on EX Series Switches on page 33
- Assigning CoS Components to Interfaces (CLI Procedure) on page 79
- Monitoring CoS Scheduler Maps on page 91
- Understanding CoS Schedulers on page 18

Defining CoS Schedulers (J-Web Procedure)

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 in 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 for details about all commit options.

2. Click one:

- **Add**—Adds a scheduler. Enter information into the schedulers page as described in Table 17 on page 72.
- **Edit**—Modifies an existing scheduler. Enter information into the schedulers page as described in Table 17 on page 72.
- **Delete**—Deletes an existing scheduler.

Table 17: 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 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.
Buffer Size	Defines the size of the delay buffer. By default, queues 0 through 7 are allotted the following percentage of the total available buffer space: • Queue 0—95 percent • Queue 1—0 percent • Queue 2—0 percent • Queue 3—0 percent • Queue 4—0 percent • Queue 5—0 percent • Queue 6—0 percent • Queue 7—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.	To define a delay buffer size for a scheduler, select the appropriate option: • To specify no buffer size, select the blank. • To specify buffer size as a percentage of the total buffer, select **Percent** and type an integer from 1 through 100. • To specify buffer size as the remaining available buffer, select **Remainder**. NOTE: On EX8200 switches, you can specify the buffer size as a temporal value. The queuing algorithm will then drop packets once it has queued a computed number of bytes. This number is the product of the logical interface speed and the configured temporal value.
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.

Table 17: Schedulers Configuration Page *(continued)*

Field	Function	Your Action
Transmit Rate	Defines the transmission rate of a scheduler. The transmit rate determines the traffic bandwidth from each forwarding class you configure. By default, queues 0 through 7 are allotted the following percentage of the transmission capacity: • Queue 0—95 percent • Queue 1—0 percent • Queue 2—0 percent • Queue 3—5 percent • Queue 4—0 percent • Queue 6—0 percent • Queue 7—5 percent	To define a transmit rate, select the appropriate option: • To enforce the exact transmission rate, select **Rate** and enter a value. • To specify the remaining transmission capacity, select **Remainder Available**. • To specify a percentage of transmission capacity, select **Percent** and type an integer from 1 through 100. • To specify no transmit rate, select the blank.

Related Documentation

- Defining CoS Schedulers (CLI Procedure) on page 70
- Example: Configuring CoS on EX Series Switches on page 33
- Monitoring CoS Scheduler Maps on page 91

Defining CoS Scheduler Maps (J-Web Procedure)

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

To configure scheduler maps:

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

NOTE: After you make changes to the configuration in 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 for details about all commit options.

2. Click one:
   - **Add**—Adds a scheduler map. Enter information into the scheduler map page as described in Table 18 on page 74.
   - **Edit**—Modifies an existing scheduler map. Enter information into the scheduler map page as described in Table 18 on page 74.
   - **Delete**—Deletes an existing scheduler map.

Table 18: 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**.
Scheduler Mapping	Allows you to associate a preconfigured scheduler with a forwarding class. After scheduler maps have been applied to an interface, they affect the hardware queues and packet schedulers.	To associate a scheduler with a forwarding class, locate the forwarding class and select the scheduler in the box next to it. For example, for the **best-effort** forwarding class, select the configured scheduler from the list.

Related Documentation

- Defining CoS Schedulers (J-Web Procedure) on page 71
- Defining CoS Schedulers (CLI Procedure) on page 70
- Example: Configuring CoS on EX Series Switches on page 33
- Monitoring CoS Scheduler Maps on page 91

Defining CoS Drop Profiles (J-Web Procedure)

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 in 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 for details about all commit options.

2. Click one:
   - **Add**—Adds a drop profile. Enter information into the drop profiles page as described in Table 19 on page 74.
   - **Edit**—Modifies an existing drop file. Enter information into the drop profiles page as described in Table 19 on page 74.
   - **Delete**—Deletes an existing drop profile.

Table 19: Drop Profiles Configuration parameters

Field	Function	Your Action
Drop Profile Name	Specifies the name for a drop profile.	Type the name.

Table 19: Drop Profiles Configuration parameters *(continued)*

Field	Function	Your Action
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 on page 93
- Example: Configuring CoS on EX Series Switches on page 33

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 on page 33
- Understanding CoS Tail Drop Profiles on page 17

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. After the rewrite rule is created, enable it on an interface. You can also apply an existing rewrite rule on an interface.

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.

NOTE: Custom rewrite-rule bindings are implemented through filters. And custom rewrite rules cannot be bound to routed VLAN interfaces (RVIs).

To create rewrite rules and enable them on interfaces:

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

```
[edit class-of-service rewrite-rules]
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority low
code-point 000
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority high
code-point 001
user@switch# set ieee-802.1 customup-rw forwarding-class 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 101
user@switch# set ieee-802.1 customup-rw forwarding-class nc loss-priority low
code-point 110
user@switch# set ieee-802.1 customup-rw forwarding-class nc loss-priority high
code-point 111
```

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

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

- To enable an 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
```

Related Documentation

- Defining CoS Rewrite Rules (J-Web Procedure) on page 77
- Example: Configuring CoS on EX Series Switches on page 33
- Monitoring CoS Rewrite Rules on page 90
- Understanding CoS Rewrite Rules on page 22

Defining CoS Rewrite Rules (J-Web Procedure)

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 in 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 for details about all commit options.

2. Click one:
   - **Add**—Adds a rewrite rule. Enter information into the rewrite rule page as described in Table 20 on page 77.
   - **Edit**—Modifies an existing rewrite rule. Enter information into the rewrite rule page as described in Table 20 on page 77.
   - **Delete**—Deletes an existing rewrite rule.

Table 20: 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.

Table 20: Rewrite Rules Configuration Page Summary *(continued)*

Field	Function	Your Action
Code Point Mapping	Rewrites outgoing CoS values of a packet based on the forwarding class and loss priority. Allows you to remove a code point mapping entry.	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 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. 4. 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. To edit an existing code point mapping, select it and click **Edit**. To remove a code point mapping entry, select it and click **Remove**.

Related Documentation

- Defining CoS Rewrite Rules (CLI Procedure) on page 76
- Understanding CoS Rewrite Rules on page 22
- Monitoring CoS Rewrite Rules on page 90
- Example: Configuring CoS on EX Series Switches on page 33

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) on page 79
- Example: Configuring CoS on EX Series Switches on page 33
- Monitoring Interfaces That Have CoS Components on page 89
- Understanding Junos OS CoS Components for EX Series Switches on page 6

Assigning CoS Components to Interfaces (J-Web Procedure)

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 in 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 for details about all commit options.

2. To configure interface association, select an interface from the list and click **Edit**.
3. Select one:
   - **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, the J-Web interface does not allow you to commit your changes 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 service assignment on a logical interface, click one:
   - **Add**—Adds a CoS service to a logical interface on a specified physical interface. Enter information as described in Table 21 on page 80.
   - **Edit**—Modifies a CoS service assignment to a logical interface. Enter information as described in Table 21 on page 80.
   - **Delete**—Deletes the CoS service assignment to a logical interface.

Table 21: Assigning CoS Components to Logical Interfaces

Field	Function	Your Action
Unit	Specifies the name of a logical interface. Allows you to assign CoS components while configuring a logical interface on a physical interface at the same time.	Type the interface name. To assign CoS services 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	Allows 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.

Table 21: Assigning CoS Components to Logical Interfaces *(continued)*

Field	Function	Your Action
Rewrite Rules	Allows 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.

Related Documentation

- Assigning CoS Components to Interfaces (CLI Procedure) on page 79
- Example: Configuring CoS on EX Series Switches on page 33
- Monitoring Interfaces That Have CoS Components on page 89

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 on page 33
- Understanding Junos OS EZQoS for CoS Configurations on EX Series Switches on page 24

Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS (CLI Procedure)

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 now 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.

For this procedure, we assume that the switch has already been configured for MPLS. See Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure).

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. 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]
   set policer mypolicer if-exceeding bandwidth-limit 500m
   ```

5. 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
   ```

6. Discard traffic that exceeds the rate limits for this policer:

   ```
   [edit firewall policer]
   set mypolicer then discard
   ```

7. 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
```

8. Apply the filter to the customer-edge interface:

```
[edit interfaces]
user@switch# set ge-2/0/3 unit 0 family inet address 121.121.121.1/16 policing filter
myfilter
```

NOTE: You can also configure schedulers and shapers as needed. See “Defining CoS Schedulers (CLI Procedure)” on page 70.

Related Documentation

- Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure)
- Assigning CoS Components to Interfaces (CLI Procedure) on page 79
- Configuring Policers to Control Traffic Rates (CLI Procedure)
- Understanding the Use of Policers in Firewall Filters

Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect (CLI Procedure)

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 provider edge (PE) switch that is using MPLS over circuit-cross connect (CCC).

NOTE: If you are using MPLS with CCC, you can use only one type of DSCP/IP precedence and only one type of IEEE 802.1p on the CCC interfaces.

This procedure creates a custom DSCP classifier and a custom EXP rewrite rule on the ingress PE. It also enables 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.

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 ]
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
```

NOTE: You can also configure schedulers and shapers as needed. See “Defining CoS Schedulers (CLI Procedure)” on page 70.

Related Documentation

- Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure)
- Assigning CoS Components to Interfaces (CLI Procedure) on page 79
- Configuring Policers to Control Traffic Rates (CLI Procedure)
- Understanding the Use of Policers in Firewall Filters

Configuring CoS Traffic Classification for Ingress Queuing on 40-port SFP+ Line Cards (CLI Procedure)

Packets arriving on a port in the 40-port SFP+ line card in an EX8200 switch are directed to either a high priority or a low priority class-of-service (CoS) ingress queue. These queues are used for scheduling traffic from the port group 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, are sent to the low priority ingress queue by default. This procedure describes how you can direct high priority 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 on page 29

Configuring CoS on Provider Switches of an MPLS Network (CLI Procedure)

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
```

NOTE: You can also configure schedulers and shapers as needed. See “Defining CoS Schedulers (CLI Procedure)” on page 70.

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33

CHAPTER 4

Verifying CoS Configuration

Monitoring CoS Classifiers

Purpose Use the monitoring functionality to display the mapping of incoming CoS values to 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 22 on page 87 summarizes key output fields for CoS classifiers.

Table 22: 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 (+).
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.	

Table 22: Summary of Key CoS Classifier Output Fields *(continued)*

Field	Values	Additional Information
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.	

Related Documentation

- Defining CoS Classifiers (CLI Procedure) on page 65
- Defining CoS Classifiers (J-Web Procedure) on page 66
- Example: Configuring CoS on EX Series Switches on page 33

Monitoring CoS Forwarding Classes

Purpose View the current assignment of class-of-service (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 23 on page 89 summarizes key output fields for CoS forwarding classes.

Table 23: 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 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.
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
Fabric Priority	(EX8200 switches only) Fabric priority for the forwarding class, either **high** or **low**. The fabric priority determines the priority of packets ingressing the switch fabric.

Related Documentation

- Defining CoS Forwarding Classes (CLI Procedure) on page 68
- Defining CoS Forwarding Classes (J-Web Procedure) on page 68
- Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure) on page 85
- Example: Configuring CoS on EX Series Switches on page 33

Monitoring Interfaces That Have CoS Components

Purpose Use the monitoring functionality to display details about the physical and logical interfaces and the CoS components assigned to them.

Action To monitor interfaces that have CoS components in the 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 24 on page 90 summarizes key output fields for CoS interfaces.

Table 24: 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**.	
Type	Type of an object—for example, **dscp** for a classifier.	
Index	Index of this interface or the internal index of a specific object.	

Related Documentation

- Assigning CoS Components to Interfaces (CLI Procedure) on page 79
- Assigning CoS Components to Interfaces (J-Web Procedure) on page 79
- Example: Configuring CoS on EX Series Switches on page 33

Monitoring CoS Rewrite Rules

Purpose Use the monitoring functionality to display information about CoS value rewrite rules, which are based on the forwarding class and loss priority.

Action To monitor CoS rewrite rules in the 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 25 on page 91 summarizes key output fields for CoS rewrite rules.

Table 25: 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. • **inet-precedence**—For IPv4 traffic.	To display forwarding classes, loss priorities, and rewritten CoS values, click the plus sign (+).
Index	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.	
Rewrite CoS Value To	Value that the CoS value is rewritten to.	

Related Documentation

- Defining CoS Rewrite Rules (CLI Procedure) on page 76
- Defining CoS Rewrite Rules (J-Web Procedure) on page 77
- Example: Configuring CoS on EX Series Switches on page 33

Monitoring CoS Scheduler Maps

Purpose 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 26 on page 92 summarizes key output fields for CoS scheduler maps.

Table 26: 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.	
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.	
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.	
Index	Index of a specific object—scheduler maps, schedulers, or drop profiles.	

Related Documentation

- Defining CoS Schedulers (CLI Procedure) on page 70
- Defining CoS Schedulers (J-Web Procedure) on page 71
- Example: Configuring CoS on EX Series Switches on page 33

Monitoring CoS Value Aliases

Purpose 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 27 on page 93 summarizes key output fields for CoS value aliases.

Table 27: 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) on page 64
- Defining CoS Code-Point Aliases (J-Web Procedure) on page 62
- Example: Configuring CoS on EX Series Switches on page 33

Monitoring CoS Drop Profiles

Purpose 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 28 on page 94 summarizes the key output fields for CoS RED drop profiles.

Table 28: Summary of 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.
Type	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.	

Related Documentation

- Defining CoS Drop Profiles (J-Web Procedure) on page 74
- Example: Configuring CoS on EX Series Switches on page 33

CHAPTER 5

Troubleshooting CoS Configuration

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

Problem 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 (CLI Procedure) on page 70
- Understanding CoS Queues on EX8200 Line Cards That Include Oversubscribed Ports on page 29

Troubleshooting a CoS Classifier Configuration for a TCAM Space Error

Problem 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-0/0/1
     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-0/0/6
     user@switch# delete class-of-service interfaces ge-0/0/7
     user@switch# delete class-of-service interfaces ge-0/0/8
     ```

     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 loss-priority
low code-points 000001
user@switch# set class-of-service classifiers dscp d2 forwarding-class fc2 loss-priority
low code-points 000010
user@switch# set class-of-service classifiers dscp d2 forwarding-class fc3 loss-priority
low code-points 000011
user@switch# set class-of-service classifiers dscp d2 forwarding-class fc4 loss-priority
low code-points 000100
user@switch# set class-of-service classifiers dscp d2 forwarding-class fc5 loss-priority
low code-points 000101
user@switch# set class-of-service classifiers dscp d2 forwarding-class fc6 loss-priority
low code-points 000110
user@switch# set class-of-service classifiers dscp d2 forwarding-class fc7 loss-priority
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 best-effort
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 forwarding-class
assured-forwarding
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:

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

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 on page 11
- Defining CoS Classifiers (CLI Procedure) on page 65

CHAPTER 6

Configuration Statements for CoS

[edit class-of-service] Configuration Statement Hierarchy

```
class-of-service {
  classifiers {
    (dscp | ieee-802.1 | inet-precedence) classifier-name {
      import (classifier-name | default);
      forwarding-class class-name {
        loss-priority loss-priority {
          code-points [ aliases ] [ 6 bit-patterns ];
        }
      }
    }
  }
  code-point-aliases {
    (dscp | ieee-802.1 | inet-precedence) {
      alias-name bits;
    }
  }
  congestion-notification-profile profile-name {
    input{
      ieee-802.1 {
        code-point up-bits pfc;
    }
  }
  forwarding-classes {
    class class-name queue-num queue-number priority ( high | low );
  }
  interfaces {
    interface-name  {
      congestion-notification-profile profile-name {
        input {
          ieee-802.1 {
            code-point up-bits pfc;
            }
          }
        }
      }
      scheduler-map map-name;
      unit logical-unit-number {
```

```
            forwarding-class class-name;
            classifiers {
              (dscp | ieee-802.1 | inet-precedence) (classifier-name | default);
            }
          }
        }
      }
      multi-destination {
        family {
          ethernet {
            broadcast forwarding-class-name;
          }
          inet {
            classifiers {
              (dscp | inet-precedence) classifier-name;
            }
          }
        }
        scheduler-map  map-name;
      }
      rewrite-rules {
        (dscp | ieee-802.1 | inet-precedence) rewrite-name {
          import (rewrite-name | default);
          forwarding-class class-name {
            loss-priority loss-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 (rate | percent percentage | remainder);
        }
      }
    }
```

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Code-Point Aliases (CLI Procedure) on page 64 or Defining CoS Code-Point Aliases (J-Web Procedure) on page 62
- Defining CoS Classifiers (CLI Procedure) on page 65 or Defining CoS Classifiers (J-Web Procedure) on page 66
- Defining CoS Forwarding Classes (CLI Procedure) on page 68 or Defining CoS Forwarding Classes (J-Web Procedure) on page 68

- Configuring CoS Tail Drop Profiles (CLI Procedure) on page 75
- Defining CoS Schedulers (CLI Procedure) on page 70 or Defining CoS Schedulers (J-Web Procedure) on page 71
- Defining CoS Rewrite Rules (CLI Procedure) on page 76 or Defining CoS Rewrite Rules (J-Web Procedure) on page 77
- Assigning CoS Components to Interfaces (CLI Procedure) on page 79 or Assigning CoS Components to Interfaces (J-Web Procedure) on page 79
- Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure) on page 85

broadcast

Syntax broadcast *forwarding-class-name*;

Hierarchy Level [edit **class-of-service multi-destination** family **ethernet**]

Release Information Statement introduced in Junos OS Release 9.5 for EX Series switches.

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.

Related Documentation

- Understanding CoS Schedulers on page 18
- Understanding CoS Forwarding Classes on page 14
- Understanding CoS Classifiers on page 11

buffer-size

Syntax buffer-size (exact | percent *percentage* | remainder);

Hierarchy Level [edit **class-of-service schedulers** *scheduler-name*]

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

Description Specify buffer size.

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

exact—Enforce the exact buffer size. When this option is configured, sharing is disabled on the queue, restricting the usage to guaranteed buffers only.

percent*percentage*—Buffer size as a percentage of total buffer.

remainder—Remaining buffer available.

Required Privilege Level

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Schedulers (CLI Procedure) on page 70 or Defining CoS Schedulers (J-Web Procedure) on page 71
- Understanding CoS Schedulers on page 18

class

Syntax class *class-name* queue-num *queue-number* priority (high | low);

Hierarchy Level [edit **class-of-service forwarding-classes**]

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

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.
Values: **high** or **low**
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.

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Forwarding Classes (CLI Procedure) on page 68 or Defining CoS Forwarding Classes (J-Web Procedure) on page 68
- Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure) on page 85

class-of-service

Syntax

```
class-of-service {
  classifiers {
    (dscp | 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 | 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 | ieee-802.1 | inet-precedence) (classifier-name | default);
        }
      }
    }
  }
  multi-destination {
    family {
      ethernet {
        broadcast forwarding-class-name;
      }
      inet {
        classifiers {
          (dscp | inet-precedence) classifier-name;
        }
      }
    }
    scheduler-map  map-name;
  }
  rewrite-rules {
    (dscp | 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 (rate | percent percentage | remainder);
    }
  }
}
```

Hierarchy Level [edit]

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

Description Configure class-of-service (CoS) parameters on EX Series switches.

The remaining statements are explained separately.

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.

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Code-Point Aliases (CLI Procedure) on page 64 or Defining CoS Code-Point Aliases (J-Web Procedure) on page 62
- Defining CoS Classifiers (CLI Procedure) on page 65 or Defining CoS Classifiers (J-Web Procedure) on page 66
- Defining CoS Forwarding Classes (CLI Procedure) on page 68 or Defining CoS Forwarding Classes (J-Web Procedure) on page 68
- Configuring CoS Tail Drop Profiles (CLI Procedure) on page 75
- Defining CoS Schedulers (CLI Procedure) on page 70 or Defining CoS Schedulers (J-Web Procedure) on page 71
- Defining CoS Rewrite Rules (CLI Procedure) on page 76 or Defining CoS Rewrite Rules (J-Web Procedure) on page 77
- Assigning CoS Components to Interfaces (CLI Procedure) on page 79 or Assigning CoS Components to Interfaces (J-Web Procedure) on page 79
- Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure) on page 85

classifiers

Syntax

```
classifiers {
  (dscp | 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*]

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Expanded to include EXP classifiers in Junos OS Release 10.1 for EX Series switches.

Description

Apply a CoS aggregate behavior classifier to a logical interface. You can apply a default classifier or a custom classifier.

The statements are explained separately.

Required Privilege Level

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Example: Combining CoS with MPLS on EX Series Switches on page 48
- Defining CoS Classifiers (CLI Procedure) on page 65 or Defining CoS Classifiers (J-Web Procedure) on page 66
- Assigning CoS Components to Interfaces (CLI Procedure) on page 79 or Assigning CoS Components to Interfaces (J-Web Procedure) on page 79
- Understanding CoS Classifiers on page 11

code-point-aliases

Syntax

```
code-point-aliases {
  (dscp | ieee-802.1 | inet-precedence) |{
    alias-name  bits;
  }
}
```

Hierarchy Level [edit class-of-service]

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

Description Define an alias for a CoS marker.

The statements are explained separately.

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Code-Point Aliases (CLI Procedure) on page 64 or Defining CoS Code-Point Aliases (J-Web Procedure) on page 62
- Understanding CoS Code-Point Aliases on page 8

code-points

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*]

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

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.

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Classifiers (CLI Procedure) on page 65 or Defining CoS Classifiers (J-Web Procedure) on page 66
- Understanding CoS Classifiers on page 11

drop-profile-map

Syntax drop-profile-map **loss-priority** *loss-priority* **protocol** *protocol* drop-profile *profile-name*;

Hierarchy Level [edit **class-of-service schedulers** *scheduler-name*]

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

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.

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Schedulers (CLI Procedure) on page 70 or Defining CoS Schedulers (J-Web Procedure) on page 71
- Understanding CoS Schedulers on page 18

dscp

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]

Release Information

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

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.

Required Privilege Level

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Code-Point Aliases (CLI Procedure) on page 64 or Defining CoS Code-Point Aliases (J-Web Procedure) on page 62
- Defining CoS Classifiers (CLI Procedure) on page 65 or Defining CoS Classifiers (J-Web Procedure) on page 66
- Defining CoS Rewrite Rules (CLI Procedure) on page 76 or Defining CoS Rewrite Rules (J-Web Procedure) on page 77
- Assigning CoS Components to Interfaces (CLI Procedure) on page 79 or Assigning CoS Components to Interfaces (J-Web Procedure) on page 79
- Understanding CoS Classifiers on page 11

dscp-ipv6

Syntax

```
dscp-ipv6 classifier-name {
  import (classifier-name | default);
  forwarding-classclass-name {
    loss-priority level {
      code-points [ aliases ] [ 6–bit-patterns ];
    }
  }
}
```

Hierarchy Level

```
[edit class-of-serviceclassifiers],
  [edit class-of-servicecode-point-aliases],
  [editclass-of-serviceinterfacesinterface-nameunitlogical-unit-number classifiers]
  [editclass-of-serviceinterfacesinterface-nameunitlogical-unit-number rewrite-rules]
  [edit class-of-servicerewrite-rules]
```

Release Information Statement introduced in Junos OS Release 10.2 for EX Series switches.

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.

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Code-Point Aliases (CLI Procedure) on page 64 or Defining CoS Code-Point Aliases (J-Web Procedure) on page 62
- Defining CoS Classifiers (CLI Procedure) on page 65 or Defining CoS Classifiers (J-Web Procedure) on page 66
- Defining CoS Rewrite Rules (CLI Procedure) on page 76 or Defining CoS Rewrite Rules (J-Web Procedure) on page 77
- Assigning CoS Components to Interfaces (CLI Procedure) on page 79 or Assigning CoS Components to Interfaces (J-Web Procedure) on page 79
- Understanding CoS Classifiers on page 11

ethernet

Syntax

```
ethernet {
  broadcast forwarding-class-name;
}
```

Hierarchy Level [edit **class-of-service multi-destination family**]

Release Information Statement introduced in Junos OS Release 9.5 for EX Series switches.

Description Specify the Ethernet broadcast traffic family.

The remaining statement is explained separately.

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

Related Documentation

- Understanding CoS Schedulers on page 18
- Understanding CoS Forwarding Classes on page 14
- Understanding CoS Classifiers on page 11

exp

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]

Release Information Statement introduced in Junos OS Release 10.1 for EX Series switches.

Description Define the experimental bits (EXP) code point mapping that is applied to the MPLS packets.

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 to an individual interface and you cannot disable it.

Options ***classifier-name***—Name of the classifier.

The remaining statements are explained separately.

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

Related Documentation

- Understanding Using CoS with MPLS Networks on EX Series Switches on page 26
- Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure)
- Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure)
- Configuring CoS on Provider Switches of an MPLS Network (CLI Procedure) on page 86

family

Syntax

```
family {
  ethernet {
    broadcast forwarding-class-name;
  }
  inet {
    classifiers{
      (dscp | ieee-802.1 | inet-precedence) classifier-name;
    }
  }
}
```

Hierarchy Level [edit class-of-service multi-destination]

Release Information Statement introduced in Junos OS Release 9.5 for EX Series switches.

Description Specify the multidestination traffic family.

The remaining statements are explained separately.

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

Related Documentation

- Understanding CoS Schedulers on page 18
- Understanding CoS Forwarding Classes on page 14
- Understanding CoS Classifiers on page 11

forwarding-class

Syntax

```
forwarding-class class-name {
  loss-priority level {
    code-points [ aliases ] [ 6–bit-patterns ];
  }
}
```

Hierarchy Level

[edit **class-of-service classifiers** (**dscp** | **ieee-802.1** | **inet-precedence**) *classifier-name*],
[edit**class-of-service interfaces** *interface-name* **unit** *logical-unit-number*],
[edit **class-of-service rewrite-rules**] (**dscp** | **ieee-802.1** | **inet-precedence**) *rewrite-name*],
[edit **class-of-service scheduler-maps** *map-name*]

Release Information

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

Description

Define forwarding class name and option values.

Options

class-name —Name of the forwarding class.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Forwarding Classes (CLI Procedure) on page 68 or Defining CoS Forwarding Classes (J-Web Procedure) on page 68
- Understanding CoS Forwarding Classes on page 14

forwarding-classes

Syntax

```
forwarding-classes {
  class class-name  queue-num queue-number;
}
```

Hierarchy Level [edit class-of-service]

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

Description Associate the forwarding class with a queue name and number.

The statement is explained separately.

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Forwarding Classes (CLI Procedure) on page 68 or Defining CoS Forwarding Classes (J-Web Procedure) on page 68
- Understanding CoS Forwarding Classes on page 14

ieee-802.1

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]
```

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

Description Apply an IEEE-802.1 rewrite rule.

Options ***classifier-name*** —Name of the classifier.

The remaining statements are explained separately.

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Classifiers (CLI Procedure) on page 65 or Defining CoS Classifiers (J-Web Procedure) on page 66
- Defining CoS Code-Point Aliases (CLI Procedure) on page 64 or Defining CoS Code-Point Aliases (J-Web Procedure) on page 62
- Defining CoS Rewrite Rules (CLI Procedure) on page 76 or Defining CoS Rewrite Rules (J-Web Procedure) on page 77
- Understanding CoS Classifiers on page 11
- Understanding CoS Rewrite Rules on page 22

import

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*]

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

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.

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Classifiers (CLI Procedure) on page 65 or Defining CoS Classifiers (J-Web Procedure) on page 66
- Defining CoS Rewrite Rules (CLI Procedure) on page 76 or Defining CoS Rewrite Rules (J-Web Procedure) on page 77
- Understanding CoS Classifiers on page 11
- Understanding CoS Rewrite Rules on page 22

inet

Syntax

```
inet {
  classifiers {
    (dscp | ieee-802.1 | inet-precedence) classifier-name ;
  }
}
```

Hierarchy Level [edit class-of-service multi-destination family]

Release Information Option **inet** introduced in Junos OS Release 9.5 for EX Series switches.
The remaining statements are explained separately.

Description Specify the IP multicast family.

The remaining statements are explained separately.

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

Related Documentation

- Understanding CoS Schedulers on page 18
- Understanding CoS Forwarding Classes on page 14
- Understanding CoS Classifiers on page 11

inet-precedence

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]

Release Information

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

Description

Apply an IPv4 precedence rewrite rule.

Options

classifier-name—Name of the classifier.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Classifiers (CLI Procedure) on page 65 or Defining CoS Classifiers (J-Web Procedure) on page 66
- Defining CoS Code-Point Aliases (CLI Procedure) on page 64 or Defining CoS Code-Point Aliases (J-Web Procedure) on page 62
- Defining CoS Rewrite Rules (CLI Procedure) on page 76 or Defining CoS Rewrite Rules (J-Web Procedure) on page 77
- Understanding CoS Classifiers on page 11
- Understanding CoS Rewrite Rules on page 22

interfaces

Syntax

```
interfaces {
  interface-name  {
    congestion-notification-profile profile-name {
      input {
        ieee-802.1 {
          code-point up-bits pfc;
          }
        }
      }
    }
    scheduler-map map-name;
    unit logical-unit-number {
      forwarding-class class-name;
      classifiers {
        (dscp | ieee-802.1 | inet-precedence) (classifier-name | default);
      }
    }
  }
}
```

Hierarchy Level [edit class-of-service]

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

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.

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Classifiers (CLI Procedure) on page 65 or Defining CoS Classifiers (J-Web Procedure) on page 66
- Defining CoS Forwarding Classes (CLI Procedure) on page 68 or Defining CoS Forwarding Classes (J-Web Procedure) on page 68
- Defining CoS Schedulers (CLI Procedure) on page 70 or Defining CoS Schedulers (J-Web Procedure) on page 71
- Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure)

loss-priority

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-name* **forwarding-class** *class-name*]

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement expanded to apply to EXP classifiers in Junos OS Release 10.1 for EX Series switches.

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.
- **low**—Packet has low loss priority.

The remaining statement is explained separately.

Required Privilege Level

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Classifiers (CLI Procedure) on page 65 or Defining CoS Classifiers (J-Web Procedure) on page 66
- Defining CoS Rewrite Rules (CLI Procedure) on page 76 or Defining CoS Rewrite Rules (J-Web Procedure) on page 77
- Understanding CoS Classifiers on page 11
- Understanding CoS Rewrite Rules on page 22

multi-destination

Syntax

```
multi-destination {
  family {
    ethernet {
      broadcast forwarding-class-name;
    }
    inet {
      classifiers {
        (dscp | ieee-802.1 | inet-precedence) classifier-name;
      }
    }
  }
  scheduler-map map-name;
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 9.5 for EX Series switches.

Description Define the CoS configuration for multidestination traffic.

The remaining statements are explained separately.

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

Related Documentation

- Understanding CoS Schedulers on page 18
- Understanding CoS Forwarding Classes on page 14
- Understanding CoS Classifiers on page 11

policing

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*]

Release Information Statement introduced in Junos OS Release 10.1 for EX Series switches.

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.

Related Documentation

- policer
- Configuring Policers to Control Traffic Rates (CLI Procedure)
- Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect (CLI Procedure) on page 83
- Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS (CLI Procedure) on page 82

priority

Syntax priority *priority*;

Hierarchy Level [edit **class-of-service schedulers** *scheduler-name*]

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

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.

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Schedulers (CLI Procedure) on page 70 or Defining CoS Schedulers (J-Web Procedure) on page 71
- Understanding CoS Schedulers on page 18

protocol

Syntax protocol *protocol* drop-profile *profile-name*;

Hierarchy Level [edit **class-of-service schedulers** *scheduler-name*]

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

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.

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Configuring CoS Tail Drop Profiles (CLI Procedure) on page 75
- Understanding CoS Tail Drop Profiles on page 17

rewrite-rules

Syntax

```
rewrite-rules {
  (dscp | 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]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement expanded for use with global EXP classifiers in Junos OS Release 10.1 for EX Series switches.

Description Specify a rewrite-rules mapping for the traffic that passes through all queues on the interface.

The remaining statements are explained separately.

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Rewrite Rules (CLI Procedure) on page 76 or Defining CoS Rewrite Rules (J-Web Procedure) on page 77
- Understanding CoS Rewrite Rules on page 22
- Understanding Using CoS with MPLS Networks on EX Series Switches on page 26

scheduler-map

Syntax scheduler-map *map-name*;

Hierarchy Level [edit **class-of-service interfaces**],
[edit **class-of-service multi-destination**]

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

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.

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Assigning CoS Components to Interfaces (CLI Procedure) on page 79 or Assigning CoS Components to Interfaces (J-Web Procedure) on page 79
- Understanding CoS Schedulers on page 18
- Understanding CoS Classifiers on page 11

scheduler-maps

Syntax

```
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
```

Hierarchy Level [edit class-of-service]

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

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.

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Forwarding Classes (CLI Procedure) on page 68 or Defining CoS Forwarding Classes (J-Web Procedure) on page 68
- Understanding CoS Schedulers on page 18
- Understanding CoS Forwarding Classes on page 14

schedulers

Syntax

```
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 (rate | percent percentage | remainder);
  }
}
```

Hierarchy Level [edit class-of-service]

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

Description Specify scheduler name and parameter values.

Options *scheduler-name* —Name of the scheduler.

The remaining statements are explained separately.

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Schedulers (CLI Procedure) on page 70 or Defining CoS Schedulers (J-Web Procedure) on page 71
- Understanding CoS Schedulers on page 18

shaping-rate

Syntax shaping-rate (percent *percentage* | rate);

Hierarchy Level [edit **class-of-service schedulers** *scheduler-name*]

Release Information Statement introduced in Junos OS Release 9.3 for EX Series switches.

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

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Understanding Junos OS CoS Components for EX Series Switches on page 6

shared-buffer

Syntax shared-buffer percent *percentage*

Hierarchy Level [edit **class-of-service**],

Release Information Statement introduced in Junos OS Release 10.1 for EX Series switches.

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.

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Understanding Junos OS CoS Components for EX Series Switches on page 6

transmit-rate

Syntax transmit-rate (*rate* | percent *percentage* | remainder);

Hierarchy Level [edit **class-of-service schedulers** *scheduler-name*]

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

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 160,000,000,000 bps

percent *percentage* —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.

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Defining CoS Schedulers (CLI Procedure) on page 70 or Defining CoS Schedulers (J-Web Procedure) on page 71
- Understanding CoS Schedulers on page 18

unit

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*]

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

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.

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

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Assigning CoS Components to Interfaces (CLI Procedure) on page 79 or Assigning CoS Components to Interfaces (J-Web Procedure) on page 79

CHAPTER 7

Operational Commands for CoS

show class-of-service

Syntax **show class-of-service**

Release Information Command introduced in Junos OS Release 9.0 for EX Series switches.
EXP classifiers added in Junos OS Release 10.1 for EX Series switches.

Description Display the class-of-service (CoS) information.

Required Privilege Level view

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Monitoring CoS Value Aliases on page 93
- Monitoring CoS Classifiers on page 87
- Monitoring CoS Forwarding Classes on page 88
- Monitoring CoS Scheduler Maps on page 91
- Monitoring CoS Rewrite Rules on page 90

List of Sample Output **show class-of- service on page 135**
show class-of-service rewrite-rule on page 138

Output Fields Table 29 on page 134 lists the output fields for the **show class-of-service** command. Output fields are listed in the approximate order in which they appear.

Table 29: 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

Table 29: show class-of-service Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
Type	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 high packets beyond which high packets are dropped.	All levels
Scheduler	Name of the scheduler.	All levels
Transmit rate	Transmission rate of the scheduler.	All levels
Buffer size	Delay buffer size in the queue.	All levels
Drop profiles	Drop profiles configured for the specified scheduler.	All levels
Protocol	Transport protocol corresponding to the drop profile.	All levels
Name	Name of the drop profile.	All levels
Queues supported	Number of queues that can be configured on the interface.	All levels
Queues in use	Number of queues currently configured.	All levels
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

Sample Output

show class-of- service

```
user@switch> show class-of-service
Forwarding class                          ID      Queue
  best-effort                              0       0
  expedited-forwarding                     1       5
  assured-forwarding                       2       1
  network-control                          3       7

Code point type: dscp
  Alias              Bit pattern
  af11               001010
```

```
  af12              001100
  ...               ...

Code point type: ieee-802.1
  Alias             Bit pattern
  af11              010
  ...               ...

Code point type: inet-precedence
  Alias             Bit pattern
  af11              001
  ...               ...

Classifier: dscp-default, Code point type: dscp, Index: 7
  Code point        Forwarding class                  Loss priority
  000000            best-effort                       low
  000001            best-effort                       low
  ...               ...                               ...

Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11
  Code point        Forwarding class                  Loss priority
  000               best-effort                       low
  001               best-effort                       low
  010               best-effort                       low
  011               best-effort                       low
  100               best-effort                       low
  101               best-effort                       low
  110               network-control                   low
  111               network-control                   low

Classifier: ipprec-default, Code point type: inet-precedence, 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

Classifier: ieee8021p-untrust, Code point type: ieee-802.1, Index: 16
  Code point        Forwarding class                  Loss priority
  000               best-effort                       low
  001               best-effort                       low
  010               best-effort                       low
  011               best-effort                       low
  100               best-effort                       low
  101               best-effort                       low
  110               best-effort                       low
  111               best-effort                       low

Rewrite rule: dscp-default, Code point type: dscp, Index: 27
  Forwarding class                  Loss priority     Code point
  best-effort                       low               000000
  best-effort                       high              000000
  expedited-forwarding              low               101110
  expedited-forwarding              high              101110
  assured-forwarding                low               001010
  assured-forwarding                high              001100
  network-control                   low               110000
```

```
  network-control                     high                111000

Rewrite rule: ieee8021p-default, Code point type: ieee-802.1, Index: 30
  Forwarding class                    Loss priority       Code point
  best-effort                         low                 000
  best-effort                         high                001
  expedited-forwarding                low                 100
  expedited-forwarding                high                101
  assured-forwarding                  low                 010
  assured-forwarding                  high                011
  network-control                     low                 110
  network-control                     high                111

Rewrite rule: ipprec-default, Code point type: inet-precedence, Index: 31
  Forwarding class                    Loss priority       Code point
  best-effort                         low                 000
  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>
      High            TCP             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: <default-fabric>, Index: 23
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
```

```
      High            TCP             1    <default-drop-profile>

Fabric priority: high
  Scheduler: <default-fabric>, Index: 23
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>
```

show class-of-service rewrite-rule

```
user@switch> show class-of-service rewrite-rule
Rewrite rule: dscp-default, Code point type: dscp, Index: 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 point 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, Code 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-precedence, 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
```

show class-of-service classifier

Syntax

show class-of-service classifier
<name *name*>
<type dscp | type dscp-ipv6 | type exp | type ieee-802.1 | type inet-precedence>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.

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 *name*—(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

List of Sample Output

show class-of-service classifier type ieee-802.1 on page 140
show class-of-service classifier type ieee-802.1 (QFX Series) on page 140

Output Fields

Table 30 on page 139 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 30: 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.

Table 30: show class-of-service classifier Output Fields *(continued)*

Field Name	Field Description
Loss priority	Loss priority value used for classification. For most platforms, the value is **high** or **low**. For some platforms, the value is **high**, **medium-high**, **medium-low**, or **low**.

Sample Output

show class-of-service classifier type ieee-802.1

```
user@host>  show class-of-service classifier type ieee-802.1
Classifier: ieee802.1-default, Code point type: ieee-802.1, Index: 3
Code Point         Forwarding Class                  Loss priority
  000              best-effort                         low
  001              best-effort                         high
  010              expedited-forwarding                low
  011              expedited-forwarding                high
  100              assured-forwarding                  low
  101              assured-forwarding                  medium-high
  110              network-control                     low
  111              network-control                     high

Classifier: users-ieee802.1, Code point type: ieee-802.1
  Code point         Forwarding class                  Loss priority
  100                expedited-forwarding              low
```

show class-of-service classifier type ieee-802.1 (QFX Series)

```
user@switch>  show class-of-service classifier type ieee-802.1
Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11
  Code point         Forwarding class                  Loss priority
  000                best-effort                       low
  001                best-effort                       low
  010                best-effort                       low
  011                best-effort                       low
  100                best-effort                       low
  101                best-effort                       low
  110                network-control                   low
  111                network-control                   low

Classifier: ieee-mcast, Code point type: ieee-802.1, Index: 46
  Code point         Forwarding class                  Loss priority
  000                mcast-be                          low
  001                mcast-be                          low
  010                mcast-be                          low
  011                mcast-be                          low
  100                mcast-be                          low
  101                mcast-be                          low
  110                mcast-nc                          low
  111                mcast-nc                          low
```

show class-of-service code-point-aliases

Syntax show class-of-service code-point-aliases
<dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence>

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display the mapping of class-of-service (CoS) code point aliases to corresponding bit patterns.

Options none—Display code point aliases of all code point types.

dscp—(Optional) Display Differentiated Services code point (DSCP) aliases.

dscp-ipv6—(Optional) Display IPv6 DSCP aliases.

exp—(Optional) Display MPLS EXP code point aliases.

ieee-802.1—(Optional) Display IEEE-802.1 code point aliases.

inet-precedence—(Optional) Display IPv4 precedence code point aliases.

Required Privilege Level view

List of Sample Output **show class-of-service code-point-aliases exp on page 142**

Output Fields Table 31 on page 141 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 31: 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 or the QFX Series), **exp** (not on EX Series switch or the QFX Series), **ieee-802.1**, or **inet-precedence** (not on the QFX Series).
Alias	Alias for a bit pattern.
Bit pattern	Bit pattern for which the alias is displayed.

Sample Output

show class-of-service code-point-aliases exp

```
user@host> show class-of-service code-point-aliases exp
Code point type: exp
  Alias               Bit pattern
  af11                100
  af12                101
  be                  000
  be1                 001
  cs6                 110
  cs7                 111
  ef                  010
  ef1                 011
  nc1                 110
  nc2                 111
```

show class-of-service drop-profile

Syntax show class-of-service drop-profile
<profile-name *profile-name*>

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.

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

List of Sample Output **show class-of-service drop-profile on page 144**

Output Fields Table 32 on page 143 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 32: show class-of-service drop-profile Output Fields

Field Name	Field Description
Drop profile	Name of a drop profile.
Type	Type of this drop profile: **discrete** or **interpolated**.
Index	Internal index of this drop profile.
Fill Level	Percentage fullness of a queue.
Drop probability	Drop probability at this fill level.

Sample Output

show class-of-service drop-profile

```
user@host> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: user-drop-profile, Type: interpolated, Index: 2989
  Fill level    Drop probability
           0                   0
           1                   1
           2                   2
           4                   4
           5                   5
           6                   6
           8                   8
          10                  10
          12                  15
          14                  20
          15                  23
        ... 64 entries total
          90                  96
          92                  96
          94                  97
          95                  98
          96                  98
          98                  99
          99                  99
         100                 100
```

show class-of-service forwarding-class

Syntax show class-of-service forwarding-class

Release Information Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display information about forwarding classes, including the mapping of forwarding classes to queue numbers.

Required Privilege Level view

Related Documentation

- Example: Configuring CoS on EX Series Switches on page 33
- Monitoring CoS Forwarding Classes on page 88
- Defining CoS Forwarding Classes (CLI Procedure) on page 68
- Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure) on page 85

List of Sample Output

show class-of-service forwarding-class on page 145
show class-of-service forwarding-class (EX8200 Switch) on page 146
show class-of-service forwarding-class (QFX Series) on page 146

Output Fields Table 33 on page 145 describes the output fields for the **show class-of-service forwarding-class** command. Output fields are listed in the approximate order in which they appear.

Table 33: show class-of-service forwarding-class Output Fields

Field Name	Field Description
Forwarding class	Name of forwarding class.
ID	Forwarding class identifier.
Queue	CoS queue mapped to the forwarding class.
Policing priority	Not supported on EX Series switches or the QFX Series and can be ignored.
Fabric priority	(EX8200 switches only) Fabric priority for the forwarding class, either **high** or **low**. Determines the priority of packets ingressing the switch fabric.

Sample Output

show class-of-service forwarding-class

```
user@switch> show class-of-service forwarding-class
Forwarding class                ID                    Queue  Policing priority
  best-effort                    0                        0           normal
```

```
  expedited-forwarding            1                    5          normal
  assured-forwarding              2                    1          normal
  network-control                 3                    7          normal
```

Sample Output

show class-of-service forwarding-class (EX8200 Switch)

```
user@switch> show class-of-service forwarding-class
Forwarding class                  ID      Queue  Fabric priority
  best-effort                      0          0  low
  expedited-forwarding             1          5  low
  assured-forwarding               2          1  low
  network-control                  3          7  low
  mcast-be                         4          2  low
  mcast-ef                         5          4  low
  mcast-af                         6          6  low
```

Sample Output

show class-of-service forwarding-class (QFX Series)

```
user@switch> show class-of-service forwarding-class
Forwarding class                  ID                    Queue  Policing priority

  best-effort                      0                        0          normal
  fcoe                             1                        1          normal
  no-loss                          2                        2          normal
  network-control                  3                        3          normal
  mcast-be                         8                        8          normal
  mcast-ef                         9                        9          normal
  mcast-af                        10                       10          normal
  mcast-nc                        11                       11          normal
```

show class-of-service interface

Syntax show class-of-service interface
<*interface-name*>

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Forwarding class map information added in Junos OS Release 9.4.
Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display the logical and physical interface associations for the classifier, rewrite rules, and scheduler map objects.

Options none—Display class-of-service (CoS) associations for all physical and logical interfaces.

interface-name—(Optional) Display CoS associations for the specified interface.

Required Privilege Level view

List of Sample Output
show class-of-service interface (Physical) on page 148
show class-of-service interface (Logical) on page 148
show class-of-service interface (Gigabit Ethernet) on page 149

Output Fields Table 34 on page 147 describes the output fields for the **show class-of-service interface** command. Output fields are listed in the approximate order in which they appear.

Table 34: show class-of-service interface Output Fields

Field Name	Field Description
Physical interface	Name of a physical interface.
Index	Index of this interface or the internal index of this object.
Dedicated Queues	Status of dedicated queues configured on an interface. Supported on Trio MPC/MIC interfaces on MX Series routers only.
Queues supported	Number of queues you can configure on the interface.
Queues in use	Number of queues currently configured.
Total non-default queues created	Number of queues created in addition to the default queues. Supported on Trio MPC/MIC interfaces on MX Series routers.
Shaping rate	Maximum transmission rate on the physical interface. You can configure the shaping rate on the physical interface, or on the logical interface, but not both. Therefore, the **Shaping rate** field is displayed for the physical interface or the logical interface, but not both.
Scheduler map	Name of the output scheduler map associated with this interface.
Input shaping rate	For Gigabit Ethernet IQ2 PICs, maximum transmission rate on the input interface.

Table 34: show class-of-service interface Output Fields *(continued)*

Field Name	Field Description
Input scheduler map	For Gigabit Ethernet IQ2 PICs, name of the input scheduler map associated with this interface.
Chassis scheduler map	Name of the scheduler map associated with the packet forwarding component queues.
Rewrite	Name and type of the rewrite rules associated with this interface.
Classifier	Name and type of classifiers associated with this interface.
Forwarding-class-map	Name of the forwarding map associated with this interface.
Congestion-notification	Congestion notification state, **enabled** or **disabled** (QFX Series only).
Logical interface	Name of a logical interface.
Shaping rate	Maximum transmission rate on the logical interface. You can configure the shaping rate on the physical interface, or on the logical interface, but not both. Therefore, the **Shaping rate** field is displayed for the physical interface or the logical interface, but not both.
Object	Category of an object: **Classifier**, **Fragmentation-map** (for LSQ interfaces only), **Scheduler-map**, **Rewrite**, or **Translation Table** (for IQE PICs only).
Name	Name of an object.
Type	Type of an object: **dscp**, **dscp-ipv6**, **exp**, **ieee-802.1**, **ip**, or **inet-precedence**.

Sample Output

show class-of-service interface (Physical)

```
user@host> show class-of-service interface so-0/2/3
Physical interface: so-0/2/3, Index: 135
Queues supported: 8, Queues in use: 4
Total non-default queues created: 4
  Scheduler map: <default>, Index: 2032638653

  Logical interface: fe-0/0/1.0, Index: 68, Dedicated Queues: no
    Shaping rate: 32000
    Object                  Name                    Type
Index
    Scheduler-map           <default>
  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 (Logical)

```
user@host> show class-of-service interface so-0/2/3.0
Logical interface: so-0/2/3.0, Index: 68, Dedicated Queues: no
    Shaping rate: 32000
```

```
    Object                  Name                    Type
Index
    Scheduler-map           <default>
  27
    Rewrite                 exp-default             exp
  21
    Classifier              exp-default             exp
   5
    Classifier              ipprec-compatibility    ip
   8
    Forwarding-class-map    exp-default             exp
   5
```

show class-of-service interface (Gigabit Ethernet)

```
user@host> show class-of-service interface ge-6/2/0
Physical interface: ge-6/2/0, Index: 175
Queues supported: 4, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Input scheduler map: <default>, Index: 3
  Chassis scheduler map: <default-chassis>, Index: 4
```

show pfe statistics traffic

Syntax show pfe statistics traffic

Release Information Command introduced in Junos OS Release 9.5 for EX Series switches.

Description Display the packet forwarding engine traffic statistics.

Options none—Display statistics about all the traffic handled by the packet forwarding engine.

Required Privilege Level admin

List of Sample Output **show pfe statistics traffic on page 151**

Output Fields Table 35 on page 150 lists the output fields for the **show pfe statistics traffic** command. Output fields are listed in the approximate order in which they appear.

Table 35: show pfe statistics traffic Output Fields

Field Name	Field Description
Packet Forwarding Engine Traffic statistics	Information about Packet Forwarding Engine traffic: • **Input Packets**—Number and rate of input packets. • **Output Packets**—Number and rate of output packets.
Packet Forwarding Engine Local Traffic statistics	Information about Packet Forwarding Engine local traffic: • **Local packets input**—Number of local input packets. • **Local packets output**—Number of local output packets. • **Software input high drops**—Number of software input high-priority drops. • **Software input medium drops**—Number of software input medium-priority drops. • **Software input low drops**—Number of software input low-priority drops. • **Software output drops**—Number of software output drops. • **Hardware input drops**—Number of hardware input drops.

Table 35: show pfe statistics traffic Output Fields *(continued)*

Field Name	Field Description
Packet Forwarding Engine Local Protocol statistics	Information about the Packet Forwarding Engine Local Protocol: • **HDLC keepalives**—Number of HDLC keepalive packets. • **ATM OAM**—Number of Asynchronous Transfer Mode (ATM) Operation, Administration, and Maintenance (OAM) packets. • **Frame Relay LMI**—Number of Frame Relay Local Management Interface (LMI) packets. • **PPP LCP/NCP**—Number of Point-to-Point Protocol (PPP) Link Control Protocol (LCP) or Network Control Protocol (NCP) packets. • **OSPF hello**—Number of Open Shortest Path First (OSPF) hello packets. • **OSPF3 hello**—Number of Open Shortest Path First version 3 (OSPFv3) hello packets. • **RSVP hello**—Number of Reservation Setup Protocol (RSVP) hello packets. • **LDP hello**—Number of Label Distribution Protocol (LDP) hello packets. • **BFD**—Number of Bidirectional Forwarding Detection Protocol (BFD) hello packets. • **IS-IS IIH**—Number of Intermediate System-to-Intermediate System Hello (IIH) packets. • **LACP**—Number of Link Aggregation Control Protocol (LACP) packets. • **ARP**—Number of Address Resolution Protocol (ARP) packets. • **ETHER OAM**—Number of Ethernet Operations, Administration, and Management (OAM) packets. • **Unknown**—Number of unknown packets not matching any of the packet types listed above.
Packet Forwarding Engine Hardware Discard statistics	Information about Packet Forwarding Engine hardware discards: • **Timeout**—Number of packets discarded because of timeouts. • **Truncated key**—Number of packets discarded because of truncated keys. • **Bits to test**—Number of bits to test. • **Data error**—Number of packets discarded because of data errors. • **Stack underflow**—Number of packets discarded because of stack underflows. • **Stack overflow**—Number of packets discarded because of stack overflows. • **Normal discard**—Number of packets discarded because of discard routes. • **Extended discard**—Number of packets discarded because of illegal next hops. • **Invalid interface**—Number of packets discarded because of invalid incoming interfaces. • **Info cell drops**—Number of information cell drops. • **Fabric drops**—Number of fabric drops.

Sample Output

show pfe statistics traffic

```
user@host> show pfe statistics traffic
Packet Forwarding Engine traffic statistics:
    Input  packets:              102682                    5 pps
    Output packets:               58033                    4 pps
Packet Forwarding Engine local traffic statistics:
    Local packets input                :               44628
    Local packets output               :               46146
    Software input control plane drops  :                   0
    Software input high drops          :                   0
    Software input medium drops        :                   0
    Software input low drops           :                   0
    Software output drops              :                   0
    Hardware input drops               :                   0
```

```
Packet Forwarding Engine local protocol statistics:
    HDLC keepalives            :                    0
    ATM OAM                    :                    0
    Frame Relay LMI            :                    0
    PPP LCP/NCP                :                 5597
    OSPF hello                 :                 3195
    OSPF3 hello                :                    0
    RSVP hello                 :                    0
    LDP hello                  :                 7478
    BFD                        :                    0
    IS-IS IIH                  :                    0
    LACP                       :                    0
    ARP                        :                    0
    ETHER OAM                  :                    0
    Unknown                    :                    8
Packet Forwarding Engine hardware discard statistics:
    Timeout                    :                    0
    Truncated key              :                    0
    Bits to test               :                    0
    Data error                 :                    0
    Stack underflow            :                    0
    Stack overflow             :                    0
    Normal discard             :                    0
    Extended discard           :                    0
    Invalid interface          :                    0
    Info cell drops            :                    0
    Fabric drops               :                    0
Packet Forwarding Engine Input IPv4 Header Checksum Error and Output MTU Error
statistics:
    Input Checksum             :                    0
    Output MTU                 :                    0
```

show pfe statistics traffic cpu

Syntax show pfe statistics traffic cpu <fpc *fpc-slot*>

Release Information Command introduced in Junos OS Release 9.5 for EX Series switches.

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 *fpc-slot*—(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

Related Documentation

- **show pfe statistics traffic multicast on page 159**
- **show pfe statistics traffic egress-queues on page 157**
- show interfaces queue
- Monitoring Interface Status and Traffic
- Understanding Junos OS CoS Components for EX Series Switches on page 6

List of Sample Output **show pfe statistics traffic cpu (EX8208 Switch) on page 154**

Output Fields Table 36 on page 153 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 36: 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.

Table 36: show pfe statistics traffic cpu Output Fields *(continued)*

Field Name	Field Description
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 statistics traffic cpu

Queue: 0, Forwarding classes: best-effort
  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 classes: expedited-forwarding
  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 classes: assured-forwarding
  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 classes: 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
```

show pfe statistics traffic egress-queues

Syntax show pfe statistics traffic egress-queues <fpc *fpc-slot*>

Release Information Command introduced in Junos OS Release 9.5 for EX Series switches.

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

Related Documentation

- **show pfe statistics traffic cpu on page 153**
- **show pfe statistics traffic multicast on page 159**
- show interfaces queue
- Monitoring Interface Status and Traffic
- Understanding Junos OS CoS Components for EX Series Switches on page 6

List of Sample Output **show pfe statistics traffic egress-queues fpc 4 (EX8208 Switch) on page 158**

Output Fields Table 37 on page 157 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 37: 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
```

show pfe statistics traffic multicast

Syntax show pfe statistics traffic multicast <fpc *fpc-slot dev-number*>

Release Information Command introduced in Junos OS Release 9.5 for EX Series switches.

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 *fpc-slot dev-number*—(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

Related Documentation

- **show pfe statistics traffic cpu on page 153**
- **show pfe statistics traffic egress-queues on page 157**
- show interfaces queue
- Monitoring Interface Status and Traffic
- Understanding Junos OS CoS Components for EX Series Switches on page 6

List of Sample Output **show pfe statistics traffic multicast fpc 0 2 (EX8208 Switch) on page 160**

Output Fields Table 38 on page 160 lists the output fields for the **show pfe statistics traffic multicast** command. Output fields are listed in the approximate order in which they appear.

Table 38: show pfe statistics traffic multicast 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 multicast fpc 0 2 (EX8208 Switch)

```
user@switch> show pfe statistics traffic multicast fpc 0 2

Queue: 0, Forwarding classes: best-effort
  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 classes: expedited-forwarding
  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 classes: assured-forwarding
  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 classes: 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
```