



Modular QoS Configuration Guide for Cisco 8000 Series Routers, IOS XR Release 24.1.x, 24.2.x, 24.3.x, 24.4.x

First Published: 2024-03-13

Last Modified: 2024-12-16

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2024 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface ix

Changes to This Document ix

Communications, Services, and Additional Information ix

CHAPTER 1

New and Changed QoS Features 1

New and Changed QoS Features 1

CHAPTER 2

YANG Data Models for QoS Features 3

Using YANG Data Models 3

CHAPTER 3

Traffic Management Overview 5

Scope 5

Traditional Traffic Management 5

Traffic Management on Your Router 5

Limitations of the VoQ Model 6

Multicast Traffic Scheduling 7

QoS Policy Inheritance 8

Cisco Modular QoS CLI to Deploy QoS 9

Important Points about MQC Egress Queuing Policy 9

CHAPTER 4

Classify Packets to Identify Specific Traffic 11

Classify Packets to Identify Specific Traffic 11

Packet Classification Overview 11

QoS Support for VPWS 12

Specification of the CoS for a Packet with IP Precedence 16

IP Precedence Bits Used to Classify Packets 17

IP Precedence Value Settings	17
IP Precedence Compared to IP DSCP Marking	17
Packet Classification on Your Router	18
Classify and Remark Layer 3 Header on Layer 2 Interfaces	19
Traffic Class Elements	21
Default Traffic Class	22
Create a Traffic Class	22
Traffic Policy Elements	23
Create a Traffic Policy	24
Enhanced Running Configuration Display for Policy Maps and Class Maps	25
Attach a Traffic Policy to an Interface	27
QoS Policy Propagation via BGP (QPPB)	30
Configuration Workflow	33
Egress Interface Configuration	35
QoS on PWHE	36
Key Features and Functionalities of QoS on PWHE	37
General Features and Support Information for QoS on PWHE	37
Bandwidth Distribution for QoS on PWHE	39
Marking for QoS on PWHE	39
Policing and Queuing for QoS on PWHE	40
PWHE Aggregate Shaper	41
PW-Ether Subinterface Policy	41
Scale Information	42
PWHE without QoS policy	42
Limitations for QoS on PWHE	42
Configure QoS on PWHE	43
Configure Flat QoS on PWHE	43
Configure H-QoS on PWHE	47
Configure QoS with Aggregate Shaper on PWHE	52

CHAPTER 5
Mark Packets to Change Priority Settings 61

Packet Marking Overview	61
Default Marking	62
QoS Behavior for Generic Routing Encapsulation (GRE) Tunnels	63

Packet Marking	63
QoS Behavior for Generic Routing Encapsulation (GRE) Tunnels	65
Classification and Marking for L2 Traffic	65
Class-based Unconditional Packet Marking Feature and Benefits	69
Configure Class-based Unconditional Packet Marking	70
Mark IP DSCP and MPLS EXP Values for IP to MPLS Traffic	71
Mark IP DSCP and MPLS EXP Values for IP to MPLS Traffic: Benefits	72
Mark IP DSCP and MPLS EXP Values for IP to MPLS Traffic: Guidelines and Limitations	72
Configure IP DSCP and MPLS EXP Values for IP to MPLS Traffic	72
Class-based Unconditional Packet Marking: Examples	74
IP Precedence Marking Configuration: Example	74
IP DSCP Marking Configuration: Example	75
IP Encapsulation Marking Configuration: Example	75
QoS Group Marking Configuration: Example	76
CoS Marking Configuration: Example	76
MPLS Experimental Bit Imposition Marking Configuration: Example	76
MPLS Experimental Topmost Marking Configuration: Example	77
Queuing and Marking Policies: Examples	77
IP Precedence Compared to IP DSCP Marking	80
Configure DSCP CS7 (Precedence 7)	80
In-Place Policy Modification	80
Recommendations for Using In-Place Policy Modification	81

CHAPTER 6
Congestion Avoidance 83

Congestion Avoidance	83
Queuing Modes	85
Main Interface Queuing Policy	85
Subinterface Queueing Policy	86
Congestion Avoidance in VOQ	87
View VOQs Evicted to HBM	88
Sharing of VOQ Statistics Counters	89
Configuring Sharing of VOQ Statistics Counters	90
Dual Queue Limit	90
Restrictions	92

Virtual Output Queue Watchdog	93
Guidelines and Restrictions for Virtual Output Queue Watchdog	95
Configure Virtual Output Queue Watchdog	95
Equitable Traffic Flow Using Fair VOQ	96
Fair VOQ: Why	97
Fair VOQ: How	97
Fair VOQ Modes and Sharing of Counters	98
Fair VOQs and Slice (or Normal) VOQs: Key Differences	99
Guidelines and Limitations	99
Configure Fair VOQ	101
Modular QoS Congestion Avoidance	102
Tail Drop and the FIFO Queue	102
Configure Tail Drop	102
Random Early Detection and TCP	104
Configure Random Early Detection	104
Explicit Congestion Notification	106
Traffic Class Queue High Water Marks Monitoring	110

CHAPTER 7

Configure Priority Flow Control	117
Priority Flow Control Overview	117
buffer-internal mode	119
Restrictions and Guidelines	119
buffer-extended mode	119
Important Considerations	120
Hardware Support for Priority Flow Control	121
Configure Priority Flow Control	121
View Priority Flow Control TX Frames Per Traffic Class	127
Guidelines and Limitations for Viewing Priority Flow Control TX Frames Per Traffic Class	128
View Traffic Class Queue Pause Duration	128
Configurable ECN Threshold and Maximum Marking Probability Values	135
ECN Threshold and Maximum Marking Probability Values	135
Benefits of Configurable ECN Threshold and Maximum Marking Probability Values	136
ECN Threshold and Maximum Marking Probability Values: FAQs	136
Guidelines and Limitations	137

Configure ECN Threshold and Maximum Marking Probability Values	138
Priority Flow Control Watchdog Overview	140
Configure a Priority Flow Control Watchdog Interval	141
Global Statistics Counters for Priority Flow Control and Priority Flow Control Watchdog	144
PFC Watchdog Statistics to Monitor Traffic Drops	146
PFC Watchdog Statistics for Two Counters	147
PFC Watchdog Statistics Benefits	147
View PFC Watchdog Statistics	147
Detect High Bandwidth Memory Congestion	148
Benefits of Detecting High Bandwidth Memory Congestion	149
High Bandwidth Memory Congestion Detection: Guidelines and Limitations	149
Configure High Bandwidth Memory Congestion Detection	149
Available Shared Memory System and High Bandwidth Memory	154
Global Pause Frames for High Bandwidth Memory Congestion	155
Global Pause Frames for High Bandwidth Memory Congestion: Guidelines and Limitations	156
Configure Global Pause Frames for High Bandwidth Memory Congestion	156

CHAPTER 8

Congestion Management 159

Congestion Management Overview	159
Important Notes About Congestion Management	159
Ingress Traffic Management Model	160
Low Latency Queueing with Strict Priority Queueing	161
Configure Low Latency Queueing with Strict Priority Queueing	164
Traffic Shaping	167
Configure Traffic Shaping	167
Egress class-level traffic shaping for subinterfaces	169
Traffic Policing	169
Committed Bursts and Excess Bursts	170
Single-Rate Policer	170
Two-Rate Policer	173
Committed Bursts	175
Excess Bursts	176
Two-Rate Policer Details	176
References for Modular QoS Management	177

CHAPTER 9	Configure Modular QoS on Link Bundles	179
	QoS on Link Bundles	179
	Load Balancing	179
	Configure QoS on Link Bundles	180



Preface

This preface contains these sections:

- [Changes to This Document, on page ix](#)
- [Communications, Services, and Additional Information, on page ix](#)

Changes to This Document

This table lists the technical changes made to this document since it was first published.

Table 1: Changes to this Document

Date	Change Summary
March 2025	Republished for Release 25.1.1.
August 2024	Republished for Release 24.3.1.
June 2024	Republished for Release 24.2.11.
March 2024	Initial release of this document

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).
- To get the business results you’re looking for with the technologies that matter, visit [Cisco Services](#).
- To submit a service request, visit [Cisco Support](#).
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit [Cisco DevNet](#).
- To obtain general networking, training, and certification titles, visit [Cisco Press](#).
- To find warranty information for a specific product or product family, access [Cisco Warranty Finder](#).

Cisco Bug Search Tool

[Cisco Bug Search Tool](#) (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.



CHAPTER 1

New and Changed QoS Features

- [New and Changed QoS Features, on page 1](#)

New and Changed QoS Features

Table 2: QoS Features Added or Modified

Feature	Description	Changed in Release	Where Documented
QoS on PWHE	This feature was introduced.	Release 24.3.1	QoS on PWHE, on page 36
Virtual Output Queue Watchdog	This feature was introduced.	Release 24.2.11	Virtual Output Queue Watchdog
Set IP Marking for SRv6 Encapsulation	This feature was introduced.	Release 24.2.11	Packet Marking
Enhanced Running Configuration Display for Policy Maps and Class Maps	This feature was introduced.	Release 24.2.11	Enhanced Running Configuration Display for Policy Maps and Class Maps
Set VXLAN Outer IP Header DSCP Value to 0	This feature was introduced.	Release 24.2.11	Packet Marking Overview
Traffic Class Queue High Water Marks Monitoring	This feature was introduced.	Release 24.2.11	Traffic Class Queue High Water Marks Monitoring
View Traffic Class Queue Pause Duration	This feature was introduced.	Release 24.2.11	View Traffic Class Queue Pause Duration
View VOQs Evicted to HBM	This feature was introduced.	Release 24.2.11	Congestion Avoidance in VOQ
Available Shared Memory System and High Bandwidth Memory Buffers	This feature was introduced.	Release 24.1.1	Available Shared Memory System and High Bandwidth Memory, on page 154



CHAPTER 2

YANG Data Models for QoS Features

This chapter provides information about the YANG data models for QoS features.

- [Using YANG Data Models, on page 3](#)

Using YANG Data Models

Cisco IOS XR supports a programmatic way of configuring and collecting operational data of a network device using YANG data models. Although configurations using CLIs are easier and human-readable, automating the configuration using model-driven programmability results in scalability.

The data models are available in the release image, and are also published in the [Github](#) repository. Navigate to the release folder of interest to view the list of supported data models and their definitions. Each data model defines a complete and cohesive model, or augments an existing data model with additional XPaths. To view a comprehensive list of the data models supported in a release, navigate to the **Available-Content.md** file in the repository.

You can also view the data model definitions using the [YANG Data Models Navigator](#) tool. This GUI-based and easy-to-use tool helps you explore the nuances of the data model and view the dependencies between various containers in the model. You can view the list of models supported across Cisco IOS XR releases and platforms, locate a specific model, view the containers and their respective lists, leaves, and leaf lists presented visually in a tree structure. This visual tree form helps you get insights into nodes that can help you automate your network.

To get started with using the data models, see the *Programmability Configuration Guide*.



CHAPTER 3

Traffic Management Overview

- [Scope, on page 5](#)
- [Traditional Traffic Management, on page 5](#)
- [Traffic Management on Your Router, on page 5](#)
- [Limitations of the VoQ Model, on page 6](#)
- [Multicast Traffic Scheduling, on page 7](#)
- [QoS Policy Inheritance, on page 8](#)
- [Cisco Modular QoS CLI to Deploy QoS, on page 9](#)

Scope

Read this configuration guide to understand the overall architecture that powers the Cisco Quality of Service (QoS) technology, and also how to use its features to configure and manage the traffic bandwidth and packet loss parameters on your network.

Traditional Traffic Management


In traditional methods of traffic management, traffic packets are sent to the egress output queues without taking into consideration the egress interface availability to transmit.

Therein lies the problem as well. In case of a traffic congestion, the traffic packets may get dropped at the egress port. Which means that the network resources spent getting the packets from the ingress input queue across the switch fabric to the output queues at egress are wasted. That is not all—input queues buffer traffic meant for different egress ports, so congestion on one egress port could affect traffic on another port, an event referred to as head-of-line-blocking.

Traffic Management on Your Router

Your router's Network Processing Unit (NPU) uses a coupled ingress-egress virtual output queuing (VoQ)-based forwarding architecture to manage traffic.

Figure 1: Traffic flow from ingress port on slice 0 to an egress port on slot 3



Here, each ingress traffic class has a one-to-one VoQ mapping from each ingress slice (pipeline) to each egress port. Which means that every egress interface (#5 in the figure) has earmarked buffer space on every ingress pipeline (#1 in the figure) for each of its VoQs.

Here is how the packet travels in times of congestion on your router system:

#1: Packets A (colored green), B (colored pink), and C (colored brown) are at the ingress interface. This is where packet marking, classification, and policing takes place. (For details, see [Mark Packets to Change Priority Settings, on page 61](#), [Classify Packets to Identify Specific Traffic, on page 11](#), and [Congestion Management, on page 159](#).)

#2: These packets are stored in separate buffer storage spaces in dedicated VoQs. This is where queuing, VoQ transmit, and drop packet and byte counters come into play. (For details, see [Congestion Avoidance, on page 83](#).)

#3: Depending on the bandwidth available on the egress interface, these packets are subjected to egress scheduling, where egress credit and transmit schedulers are configured. In other words, the packets and the sequence in which they will now proceed towards the egress interface is determined here. This is where the fabric bandwidth is taken into consideration for egress scheduling.

#4: The packets are switched through the fabric.

#5: In the final phase, the egress marking and classification takes place, and the congestion is managed in a way that at this stage there is no packet dropped, and all the packets are transmitted to the next hop.

Limitations of the VoQ Model

While the VoQ model of traffic management offers distinct advantages (reducing memory bandwidth requirements, providing end-to-end QoS flow), it has this limitation:

The total egress queue scale is lower because each egress queue must be replicated as an ingress VoQ on each slice of each NPU/ASIC in the system. This means that with the addition of 1 NPU with 20 interfaces, the

number of VoQs used on each and every NPU in the system will increase by 20×8 (queue/interface) = 160. There is also an increase in the number of credit connectors from each scheduler for each egress port on pre-existing NPUs to each slice in the newly inserted NPU.

Multicast Traffic Scheduling

Table 3: Feature History Table

Feature Name	Release Information	Feature Description
Multicast Traffic Scheduling on Egress Queues	Release 24.3.1	<p>Introduced in this release on: Modular Systems (8800 [LC ASIC:ASIC: P100])(select variants only*)</p> <p>We have introduced multicast traffic scheduling on egress queues. Now you can have more granular control over multicast traffic by applying egress queuing policy map parameters such as traffic shaping, priority, and queuing, to each egress queue. This allows for specific management of multicast traffic for different receivers, ensuring efficient and prioritized handling of multicast data.</p> <p>*This feature is enabled by default and supported on:</p> <ul style="list-style-type: none">• 88-LC1-12TH24FH-E• 88-LC1-52Y8H-EM

What is this functionality all about?

The multicast traffic is now scheduled in the second pass, which means the egress queuing policy map parameters such as traffic shaping, priority, queuing, and so on, can be applied to the multicast traffic, and hence there's more control over the multicast traffic scheduling.

QoS Policy Inheritance

Table 4: Feature History Table

Feature Name	Release Information	Feature Description
QoS Policy Inheritance	Release 7.3.15	<p>To create QoS policies for subinterfaces, you had to apply the policy on each subinterface manually. From this release, all you do is create and apply a single QoS policy on the main interface and the subinterfaces automatically inherit the policy.</p> <p>The inheritance model provides an easily maintainable method for applying policies, enabling you to create targeted policies for a group of interfaces and their subinterfaces. This model saves your time and resources while creating QoS policies.</p>

- **What is this functionality all about?**—As the name suggests, the functionality is based on an inheritance model, where you create and apply a QoS policy to a main interface. The subinterfaces that are attached to the main interface automatically inherit the policy.

The inheritance model applies to all QoS operations including:

- Classification
 - Marking
 - Policing
 - Shaping
- **How does the inheritance model help?**—Previously, if you had, for example, eight subinterfaces, you created and applied policies to each of those subinterfaces separately. With the inheritance model, you save on time and resources, with just one policy automatically applied across a main interface and its subinterfaces.
 - **Do I need to do anything more to enable the inheritance model?**—No, you don't. The inheritance model is the default option.
 - **What if I want to override the inheritance option?**—Technically, you can't override this option. However, you can remove the policy from main interface and add policies to the subinterfaces except to the ones where you don't want the policy to be inherited.
 - **What about policy-map statistics?**—There's no change to this behavior. Running the [show policy-map interface](#) command displays the cumulative statistics for an interface, and these numbers include the subinterfaces as well.

- **Any limitations I need to be aware of?**—There's no support for ECN marking and egress marking policy on the same interface and subinterface combination. However, the QoS policy inheritance functionality accepts these multiple policies causing the ECN markings to fail. To prevent such failures:
 - Don't configure an egress marking policy on a subinterface and apply an ECN-enabled policy on the main interface.
 - Don't apply an ECN policy on a subinterface and configure an egress marking policy on the main interface.

Cisco Modular QoS CLI to Deploy QoS

Cisco Modular QoS CLI (MQC) framework is the Cisco IOS QoS user language that enables:

- a standard Command Line Interface (CLI) and semantics for QoS features.
- simple and accurate configurations.
- QoS provisioning within the context of an extensible language.

For your router, in the egress direction, two types of MQC policies are supported: queuing and marking. You use the queuing policy to configure credit scheduling hierarchy, rates, priority, buffering, and congestion avoidance. You use the marking policy to classify and mark packets that have been scheduled for transmission. Even when a queuing policy is not applied, there is an implicit queuing policy with TC7 - P1, TC6 - P2, TC5 - TC0 (6 x Pn), so packets marked with TC7 and control inject packets are always prioritized over other packets. In the ingress, only one policy is supported for classification and marking.

You can apply queuing and marking policy independent of each other or together in the egress direction. If you apply both policies together, the queuing policy actions are provisioned first, followed by marking policy actions.

Important Points about MQC Egress Queuing Policy

These are important points that you must know about the MQC egress queuing policy:

- The MQC queuing policy consists of a set of class maps, which are added to a policy map. You control queuing and scheduling parameters for that traffic class by applying actions to the policy.
- **class-default** matches **traffic-class 0**, and no other class can match **traffic-class 0**. When **set traffic-class** is used in the ingress policy-map or egress marking policy-map, then it should match the **traffic-class** in the egress queuing policy-map to have a predictable behaviour.
- Each unique combination of traffic classes that match **class-default** require a separate traffic class (TC) profile. The number of TC profiles are limited to 8 for main interfaces and 8 for sub-interfaces.
- You cannot configure multiple traffic classes with the same priority level.
- Each priority level, when configured, must be configured to the class that matches the corresponding TC as shown in the following table.

Priority Level	Traffic Class
P1	7

Priority Level	Traffic Class
P2	6
P3	5
P4	4
P5	3
P6	2
P7	1

- If all the priority levels configured in a policy-map are sorted, they must be contiguous. In other words, you cannot skip a priority level. For example, P1 P2 P4 (skipping P3), is not allowed.
- From IOS XR Release 7.3.1 onwards, you can create a single set of contiguous priority TCs. Ensure that you assign priority levels that increase for every TC and that you assign priority level 1 for traffic class 7. You need not configure unused traffic classes, so you can create only those many TCs that you require on the egress policy-map.
- MQC supports up to two levels (parent, child) of queuing policy. The parent level aggregates all the traffic classes and whereas the child level differentiates traffic classes using MQC classes.
- Only these actions are supported in the queuing policy:
 - priority
 - shape
 - bandwidth remaining ratio
 - queue-limit
 - Random Early Detection (RED)
 - Priority flow control
- You can have only one match traffic-class value in the class map.
- You cannot apply a queuing policy to a main interface and its sub-interfaces.



CHAPTER 4

Classify Packets to Identify Specific Traffic

- [Classify Packets to Identify Specific Traffic, on page 11](#)
- [Packet Classification Overview, on page 11](#)
- [Packet Classification on Your Router, on page 18](#)
- [Traffic Class Elements, on page 21](#)
- [Default Traffic Class, on page 22](#)
- [Create a Traffic Class, on page 22](#)
- [Traffic Policy Elements, on page 23](#)
- [Create a Traffic Policy, on page 24](#)
- [Enhanced Running Configuration Display for Policy Maps and Class Maps, on page 25](#)
- [Attach a Traffic Policy to an Interface, on page 27](#)
- [QoS Policy Propagation via BGP \(QPPB\), on page 30](#)
- [QoS on PWHE, on page 36](#)

Classify Packets to Identify Specific Traffic

Read this section to get an overview of packet classification and the different packet classification types for your router.

Packet Classification Overview

Packet classification involves categorizing a packet within a specific group (or class) and assigning it a traffic descriptor to make it accessible for QoS handling on the network. The traffic descriptor contains information about the forwarding treatment (quality of service) that the packet should receive. Using packet classification, you can partition network traffic into multiple priority levels or classes of service.

When traffic descriptors are used to classify traffic, the source agrees to adhere to the contracted terms and the network promises a quality of service. This is where traffic policers and traffic shapers come into the picture. Traffic policers and traffic shapers use the traffic descriptor of a packet—that is, its classification—to ensure adherence to the contract.

The Modular Quality of Service (QoS) command-line interface (MQC) is used to define the traffic flows that must be classified, where each traffic flow is called a class of service, or class. Later, a traffic policy is created and applied to a class. All traffic not identified by defined classes fall into the category of a default class.

QoS Support for VPWS

Table 5: Feature History Table

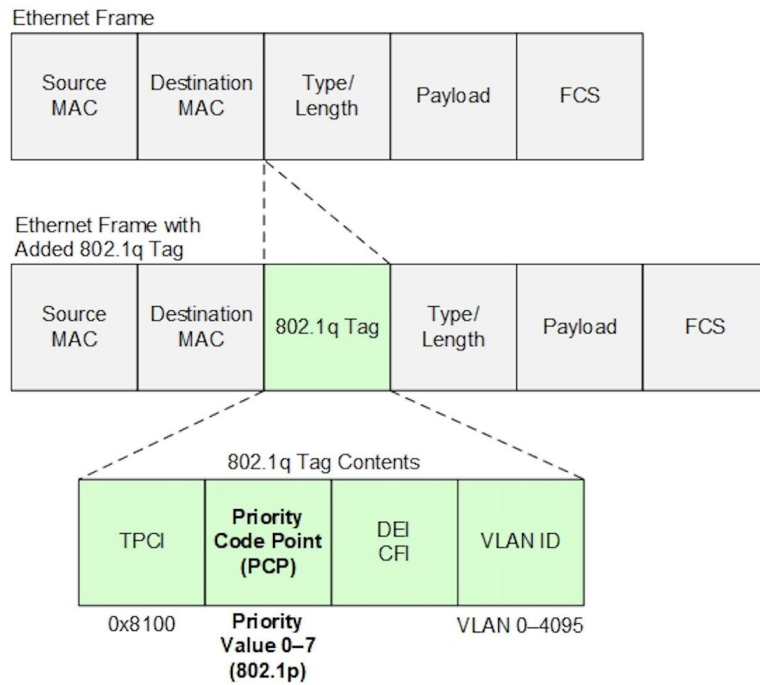
Feature Name	Release Information	Feature Description
QoS Support for VPWS	Release 24.4.1	Introduced in this release on: Fixed Systems (8700) (select variants only*) *QoS support for VPWS is now extended to Cisco 8712-MOD-M routers.
QoS Support for VPWS	Release 7.7.1	Now, you can apply QoS packet classification on Layer 2 subinterfaces, and for VPWS traffic. This feature provides ingress classification support for L2 attachment circuit (AC) traffic based on the 802.1p field priority value. With this feature, you can use QoS policies on VPWS traffic in your network on Cisco 8000 Series routers.

EVPN-VPWS is a BGP control plane solution for point-to-point services. It implements the signaling and encapsulation techniques for establishing an EVPN instance between a pair of PEs. It has the ability to forward traffic from one network to another without a MAC lookup.

By configuring QoS, you can provide preferential treatment to specific types of traffic at the expense of other traffic types. Without QoS, the device offers best-effort service for each packet, regardless of the packet contents or size. The device sends the packets without any assurance of reliability, delay bounds, or throughput.

The 802.1Q standard (and 802.1ad, a subset of 802.1Q) defines a system of VLAN tagging for Ethernet frames and also contains a provision for a QoS prioritization scheme known as 802.1p, which indicates the priority level of the frame, as shown in the figure. For more information on IEEE standards, browse the [IEEE website](#).

Figure 2: 802.1p



VLAN Tag Priority Field Based Classification

The QoS term class of service (CoS) is a 3-bit field called Priority Code Point (PCP) which specifies a priority value between 0 and 7 that is used by QoS to differentiate traffic. Drop Eligible Indicator (DEI) is a 1-bit field that is used to indicate frames eligible to be dropped during traffic congestion.

MQC allows configuration of class-map match condition based on the PCP and DEI fields. The classification is supported on 802.1Q and 802.1ad interfaces.



Note The ingress classification supports AC-to-AC traffic flow and AC-to-PWE traffic flow.

Configuration Example for QoS Support for VPWS

Follow these steps to enable this feature:

- Enable VPWS configuration. Refer the [L2VPN Configuration Guide](#) for details.
- Configure ingress traffic classification, based on the PCP and DEI fields in the VLAN header.
 - Create class maps for different traffic classes.
 - Associate them with a policy map.
- Configure ingress traffic remarking.

Configure Ingress Traffic Classification Based on PCP and DEI Fields in the VLAN Header

```
/* Create class maps for different traffic classes */
```

```

Router# configure terminal
Router(config)# class-map match-all CONTROL_PLANE
Router(config-cmap)# match cos 7
Router(config-cmap)# end-class-map

Router(config)# class-map match-all VOIP
Router(config-cmap)# match cos 6
Router(config-cmap)# end-class-map

Router(config)# class-map match-all VIDEO_STREAM
Router(config-cmap)# match cos 5
Router(config-cmap)# end-class-map

Router(config)# class-map match-all TRANSACTIONAL_DATA
Router(config-cmap)# match cos 4
Router(config-cmap)# end-class-map

Router(config)# class-map match-all DB_SYNC
Router(config-cmap)# match cos 3
Router(config-cmap)# match dei 1
Router(config-cmap)# end-class-map

Router(config)# class-map match-all BULK_DATA
Router(config-cmap)# match cos 2
Router(config-cmap)# match dei 1
Router(config-cmap)# end-class-map

Router(config)# class-map match-all SCAVENGER
Router(config-cmap)# match cos 1
Router(config-cmap)# match dei 1
Router(config-cmap)# end-class-map
Router(config)# commit

/* Create a policy map and associate the class maps to it */

Router# configure terminal
Router(config)# policy-map INGRESS_L2_AC
Router(config-pmap)# class CONTROL_PLANE
Router(config-pmap-c)# set traffic-class 7
Router(config-pmap-c)# exit

Router(config-pmap)# class VOIP
Router(config-pmap-c)# set traffic-class 6
Router(config-pmap-c)# exit

Router(config-pmap)# class VIDEO_STREAM
Router(config-pmap-c)# set traffic-class 5
Router(config-pmap-c)# exit

Router(config-pmap)# class TRANSACTIONAL_DATA
Router(config-pmap-c)# set traffic-class 4
Router(config-pmap-c)# exit

Router(config-pmap)# class DB_SYNC
Router(config-pmap-c)# set traffic-class 3
Router(config-pmap-c)# exit

Router(config-pmap)# class BULK_DATA
Router(config-pmap-c)# set traffic-class 2
Router(config-pmap-c)# exit

Router(config-pmap)# class SCAVENGER
Router(config-pmap-c)# set traffic-class 1

```



```

Router(config-pmap-c) # exit

Router(config-pmap) # class class-default
Router(config-pmap-c) # exit
Router(config-pmap) # end-policy-map
Router(config) # commit

```

Configure Ingress Traffic Remarking

/* Set CoS and DEI values for traffic classes, as needed */

```

Router# configure terminal
Router(config) # policy-map INGRESS_L2_AC
Router(config-pmap) # class CONTROL_PLANE
Router(config-pmap-c) # set traffic-class 7
Router(config-pmap-c) # exit

Router(config-pmap) # class VOIP
Router(config-pmap-c) # set traffic-class 6
Router(config-pmap-c) # set cos 7
Router(config-pmap-c) # exit

Router(config-pmap) # class VIDEO_STREAM
Router(config-pmap-c) # set traffic-class 5
Router(config-pmap-c) # set cos 5
Router(config-pmap-c) # exit

Router(config-pmap) # class TRANSACTIONAL_DATA
Router(config-pmap-c) # set traffic-class 4
Router(config-pmap-c) # set cos 5
Router(config-pmap-c) # exit

Router(config-pmap) # class DB_SYNC
Router(config-pmap-c) # set traffic-class 3
Router(config-pmap-c) # set dei 0
Router(config-pmap-c) # exit

Router(config-pmap) # class BULK_DATA
Router(config-pmap-c) # set traffic-class 2
Router(config-pmap-c) # set cos 3
Router(config-pmap-c) # set dei 0
Router(config-pmap-c) # exit

Router(config-pmap) # class SCAVENGER
Router(config-pmap-c) # set traffic-class 1
Router(config-pmap-c) # exit

Router(config-pmap) # class class-default
Router(config-pmap-c) # set dei 1
Router(config-pmap-c) # end-policy-map
Router(config) # commit

```

/* Associate Policy-Map INGRESS_L2_AC With the Designated Subinterface */

Before you enable the subinterface, ensure that the parent interface state is up.

```

Router(config) # interface hundredGigE 0/11/0/31.102
Router(config-subif) # service-policy input INGRESS_L2_AC
Router(config-subif) # commit

```

Verification

Verify ingress QoS policy configuration. In the output, you can see that traffic is classified and transmitted for some categories.


```
Router# show policy-map interface Hu0/11/0/31.102 input
```

```
HundredGigE0/11/0/31.102 input: INGRESS_L2_AC
```

```
Class CONTROL_PLANE
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :      9813350/13738690000      936847
    Transmitted                :      9813350/13738690000      936847
    Total Dropped              :              0/0              0
Class VOIP
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :     117760245/164864343000     11242157
    Transmitted                :     117760245/164864343000     11242157
    Total Dropped              :              0/0              0
Class VIDEO_STREAM
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :     49066792/68693508800      4684233
    Transmitted                :     49066792/68693508800      4684233
    Total Dropped              :              0/0              0
Class TRANSACTIONAL_DATA
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :     225707344/315990281600     21547467
    Transmitted                :     225707344/315990281600     21547467
    Total Dropped              :              0/0              0
Class DB_SYNC
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :              0/0              0
  Transmitted                  :              0/0              0
  Total Dropped                :              0/0              0
Class BULK_DATA
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :              0/0              0
  Transmitted                  :              0/0              0
  Total Dropped                :              0/0              0
Class SCAVENGER
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :              0/0              0
  Transmitted                  :              0/0              0
  Total Dropped                :              0/0              0
Class class-default
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :     500482413/700675378200     47779164
    Transmitted                :     500482413/700675378200     47779164
    Total Dropped              :              0/0              0
Policy Bag Stats time: 1657528790084 [Local Time: 07/11/22 08:39:50.084]
```

Specification of the CoS for a Packet with IP Precedence

Use of IP precedence allows you to specify the CoS for a packet. You can create differentiated service by setting precedence levels on incoming traffic and using them in combination with the QoS queuing features. So that, each subsequent network element can provide service based on the determined policy. IP precedence is usually deployed as close to the edge of the network or administrative domain as possible. This allows the rest of the core or backbone to implement QoS based on precedence.

Figure 3: IPv4 Packet Type of Service Field

You can use the three precedence bits in the type-of-service (ToS) field of the IPv4 header for this purpose. Using the ToS bits, you can define up to eight classes of service. Other features configured throughout the network can then use these bits to determine how to treat the packet in regard to the ToS to grant it. These other QoS features can assign appropriate traffic-handling policies, including congestion management strategy and bandwidth allocation. For example, queuing features such as LLQ can use the IP precedence setting of the packet to prioritize traffic.

IP Precedence Bits Used to Classify Packets

Use the three IP precedence bits in the ToS field of the IP header to specify the CoS assignment for each packet. You can partition traffic into a maximum of eight classes and then use policy maps to define network policies in terms of congestion handling and bandwidth allocation for each class.

Each precedence corresponds to a name. IP precedence bit settings 6 and 7 are reserved for network control information, such as routing updates. These names are defined in RFC 791.

IP Precedence Value Settings

By default, the routers leave the IP precedence value untouched. This preserves the precedence value set in the header and allows all internal network devices to provide service based on the IP precedence setting. This policy follows the standard approach stipulating that network traffic should be sorted into various types of service at the edge of the network and that those types of service should be implemented in the core of the network. Routers in the core of the network can then use the precedence bits to determine the order of transmission, the likelihood of packet drop, and so on.

Because traffic coming into your network can have the precedence set by outside devices, we recommend that you reset the precedence for all traffic entering your network. By controlling IP precedence settings, you prohibit users that have already set the IP precedence from acquiring better service for their traffic simply by setting a high precedence for all of their packets.

The class-based unconditional packet marking and LLQ features can use the IP precedence bits.

IP Precedence Compared to IP DSCP Marking

If you need to mark packets in your network and all your devices support IP DSCP marking, use the IP DSCP marking to mark your packets because the IP DSCP markings provide more unconditional packet marking options. If marking by IP DSCP is undesirable, however, or if you are unsure if the devices in your network support IP DSCP values, use the IP precedence value to mark your packets. The IP precedence value is likely to be supported by all devices in the network.

You can set up to 8 different IP precedence markings and 64 different IP DSCP markings.

Packet Classification on Your Router

Table 6: Feature History Table

Feature Name	Release Information	Feature Description
Ingress and Egress Packet Classification	Release 24.4.1	<p>Introduced in this release on: Fixed Systems (8700) (select variants only*)</p> <p>You can categorize packets into specific groups or classes and assign them traffic descriptors for QoS to classify and manage network traffic effectively.</p> <p>*This functionality is now supported on Cisco 8712-MOD-M routers.</p>
Ingress and Egress Packet Classification	Release 24.3.1	<p>Introduced in this release on: Modular Systems (8800 [LC ASIC: P100]) (select variants only*), Fixed Systems (8200) (select variants only*), Fixed Systems (8700 (P100, K100)) (select variants only*)</p> <p>You can categorize packets into specific groups or classes and assign them traffic descriptors for QoS to classify and manage network traffic effectively.</p> <p>*This feature is supported on:</p> <ul style="list-style-type: none"> • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM • 8212-48FH-M • 8711-32FH-M

Feature Name	Release Information	Feature Description
Ingress and Egress Packet Classification	Release 24.2.1	<p>Introduced in this release on: Modular Systems (8800 [LC ASIC: P100])(select variants only*)</p> <p>Categorizing packets into specific groups or classes and assigning them traffic descriptors for QoS helps classify and manage network traffic effectively. At the ingress, the QoS map and TCAM are used for classification. The QoS map is used for classification at the egress when policy matches only on DSCP, while TCAM is used for other criteria such as MPLS.</p> <p>*This feature is supported on 88-LC1-36EH.</p>

On your router, there are two types of packet classification systems:

- In the ingress direction, QoS map and Ternary Content Addressable Memory (TCAM).
- In the egress direction, QoS map.

The router selects a QoS map-based classification system when you specify a QoS policy for matching on:

- Differentiated Services Code Point (DSCP) or precedence value (also called DSCP or Precedence-based classification)
- Experimental bits (EXP)
- Priority Code Point (PCP)
- Drop Eligibility Indicator (DEI), or
- qos-group.

Otherwise, the router selects TCAM as the packet classification system.

The TCAM is an extension of the Content Addressable Memory (CAM) table concept. A CAM table takes in an index or key value (usually a MAC address) and looks up the resulting value (usually a switch port or VLAN ID). Table lookup is fast and always based on an exact key match consisting of two input values: 0 and 1 bits.

The QoS map is a table-based classification system for traffic packets.

Classify and Remark Layer 3 Header on Layer 2 Interfaces

When you need to mark packets for Layer 2 interface traffic that flows across bridge domains and bridge virtual interfaces (BVI), you can create a mixed QoS policy. This policy has both map-based and TCAM-based classification class-maps. The mixed policy ensures that both bridged (Layer 2) and Bridge Virtual Interface (BVI, or Layer 3) traffic flows are classified and remarked.

Guidelines

- A class-map with TCAM classification may not match bridged traffic. TCAM entries match only routed traffic while map entries match both bridged and BVI traffic.
- A class-map with map-based classification matches both bridged and BVI traffic.

Example

```

ipv4 access-list acl_v4
10 permit ipv4 host 100.1.1.2 any
20 permit ipv4 host 100.1.100.2 any
ipv6 access-list acl_v6
10 permit tcp host 50:1:1::2 any
20 permit tcp any host 50:1:200::2
class-map match-any c_match_acl
match access-group ipv4 acl_v4 ! This entry does not match bridged traffic
match access-group ipv6 acl_v6 ! This entry does not match bridged traffic
match dscp af11 This entry matches bridged and BVI traffic
class-map match-all c_match_all
match protocol udp ! This entry does not match bridged traffic
match prec 7
class-map match-any c_match_protocol
match protocol tcp ! This entry, and hence this class does not match bridged traffic
class-map match-any c_match_ef
match dscp ef ! This entry/class matches bridged and BVI traffic
class-map match-any c_qosgroup_1 This class matches bridged and BVI traffic
!
match qos-group 1
policy-map p_ingress
class c_match_acl
set traffic-class 1
set qos-group 1
!
class c_match_all
set traffic-class 2
set qos-group 2
!
class c_match_ef
set traffic-class 3
set qos-group 3
!
class c_match_protocol
set traffic-class 4
set qos-group 4
policy-map p_egress
class c_qosgroup_1
set dscp af23
interface FourHundredGigE0/0/0/0
l2transport
service-policy input p_ingress
service-policy output p_egress
!
!
interface FourHundredGigE0/0/0/1
ipv4 address 200.1.2.1 255.255.255.0
ipv6 address 2001:2:2::1/64
service-policy input p_ingress
service-policy output p_egress

```

Traffic Class Elements

The purpose of a traffic class is to classify traffic on your router. Use the **class-map** command to define a traffic class.

A traffic class contains three major elements:

- A name
- A series of **match** commands - to specify various criteria for classifying packets.
- An instruction on how to evaluate these **match** commands (if more than one **match** command exists in the traffic class)

Packets are checked to determine whether they match the criteria specified in the **match** commands. If a packet matches the specified criteria, that packet is considered a member of the class and is forwarded according to the QoS specifications set in the traffic policy. Packets that fail to meet any of the matching criteria are classified as members of the default traffic class.

This table shows the details of match types supported on the router.

Match Type Supported	Min, Max	Max Entries	Support for Match NOT	Support for Ranges	Direction Supported on Interfaces
IPv4 DSCP IPv6 DSCP	(0,63)	64	Yes	Yes	Ingress
DSCP					Egress
IPv4 Precedence IPv6 Precedence	(0,7)	8	Yes	No	Ingress
Precedence					Egress
MPLS Experimental Topmost	(0,7)	8	Yes	No	Ingress Egress
Access-group	Not applicable	8	No	Not applicable	Ingress
Match qos-group	(1-31)	7 + class-default	No	No	Egress
Protocol	(0, 255)	1	Yes	Not applicable	Ingress
CoS	(0,7)	8	Yes	No	Ingress and Egress
DEI	(0,1)	2	Yes	No	Ingress and Egress

Default Traffic Class

Unclassified traffic (traffic that doesn't meet the match criteria specified in the traffic classes) is treated as belonging to the default traffic class.

If the user doesn't configure a default class, packets are still treated as members of the default class. However, by default, the default class has no enabled features. Therefore, packets belonging to a default class with no configured features have no QoS functionality.

For egress classification, match on **qos-group** for seven groups with range (1-31) is supported. Match **qos-group 0** can't be configured. The class-default in the egress policy maps to **qos-group 0**.

This example shows how to configure a traffic policy for the default class:

```
configure
policy-map ingress_policy1
class class-default
  police rate percent 30
!
```

Create a Traffic Class

To create a traffic class containing match criteria, use the **class-map** command to specify the traffic class name, and then use the **match** commands in class-map configuration mode, as needed.

Guidelines

- Users can provide multiple values for a match type in a single line of configuration; that is, if the first value does not meet the match criteria, then the next value indicated in the match statement is considered for classification.
- Use the **not** keyword with the **match** command to perform a match based on the values of a field that are not specified.
- All **match** commands specified in this configuration task are considered optional, but you must configure at least one match criterion for a class.
- If you specify **match-any**, one of the match criteria must be met for traffic entering the traffic class to be classified as part of the traffic class. This is the default. If you specify **match-all**, the traffic must match all the match criteria.
- For the **match access-group** command, QoS classification based on the packet length or TTL (time to live) field in the IPv4 and IPv6 headers is not supported.
- For the **match access-group** command, when an ACL list is used within a class-map, the deny action of the ACL is ignored and the traffic is classified based on the specified ACL match parameters.
- The **match qos-group**, **traffic-class**, **DSCP/Prec**, and **MPLS EXP** are supported only in egress direction, and these are the only match criteria supported in egress direction.
- The egress default class implicitly matches **qos-group 0**.

- Multicast takes a system path that is different than unicast on router, and they meet later on the egress in a multicast-to-unicast ratio of 20:80 on a per interface basis. This ratio is maintained on the same priority level as that of the traffic.
- Egress QoS for multicast traffic treats traffic classes 0-5 as low-priority and traffic classes 6-7 as high priority. Currently, this is not user-configurable.
- Egress shaping does not take effect for multicast traffic in the high priority (HP) traffic classes. It only applies to unicast traffic.
- If you set a traffic class at the ingress policy and do not have a matching class at egress for the corresponding traffic class value, then the traffic at ingress with this class will not be accounted for in the default class at the egress policy map.
- Only traffic class 0 falls in the default class. A non-zero traffic class assigned on ingress but with no assigned egress queue, falls neither in the default class nor any other class.

Configuration Example

You have to accomplish the following to complete the traffic class configuration:

1. Creating a class map
2. Specifying the match criteria for classifying the packet as a member of that particular class

(For a list of supported match types, see [Traffic Class Elements, on page 21](#).)

```
Router# configure
Router(config)# class-map match-any qos-1
Router(config-cmap)# match qos-group 1
Router(config-cmap)# end-class-map
Router(config-cmap)# commit
```

Use this command to verify the class-map configuration:

```
Router#show class-map qos-1
1) ClassMap: qos-1      Type: qos
   Referenced by 2 Polycmaps
```

Also see, [Attach a Traffic Policy to an Interface, on page 27](#).

Related Topics

- [Traffic Class Elements, on page 21](#)
- [Traffic Policy Elements, on page 23](#)

Traffic Policy Elements

A traffic policy contains three elements:

- Name

- Traffic class
- QoS policies

After choosing the traffic class that is used to classify traffic to the traffic policy, the user can enter the QoS features to be applied to the classified traffic.

The MQC does not necessarily require that the users associate only one traffic class to one traffic policy.

The order in which classes are configured in a policy map is important. The match rules of the classes are programmed into the TCAM in the order in which the classes are specified in a policy map. Therefore, if a packet can possibly match multiple classes, only the first matching class is returned and the corresponding policy is applied.

The router supports 8 classes per policy-map in the ingress direction and 8 classes per policy-map in the egress direction.

This table shows the supported class-actions on the router.

Supported Action Types	Direction supported on Interfaces
bandwidth-remaining	egress
mark	See Packet Marking, on page 63
police	ingress
priority	egress (level 1 to level 7)
queue-limit	egress
shape	egress
red	egress

RED supports the **discard-class** option; the only values to be passed to the discard-class being 0 and 1.

Create a Traffic Policy

The purpose of a traffic policy is to configure the QoS features that should be associated with the traffic that has been classified in a user-specified traffic class or classes.

To configure a traffic class, see [Create a Traffic Class, on page 22](#).

After you define a traffic policy with the **policy-map** command, you can attach it to one or more interfaces to specify the traffic policy for those interfaces by using the **service-policy** command in interface configuration mode. With dual policy support, you can have two traffic policies, one marking and one queuing attached at the output. See, [Attach a Traffic Policy to an Interface, on page 27](#).

Configuration Example

You have to accomplish the following to complete the traffic policy configuration:

1. Creating a policy map that can be attached to one or more interfaces to specify a service policy
2. Associating the traffic class with the traffic policy

3. Specifying the class-action(s) (see [Traffic Policy Elements, on page 23](#))

```
Router# configure
Router(config)# policy-map test-shape-1
Router(config-pmap)# class qos-1

/* Configure class-action ('shape' in this example).
Repeat as required, to specify other class-actions */
Router(config-pmap-c)# shape average percent 40
Router(config-pmap-c)# exit

/* Repeat class configuration as required, to specify other classes */

Router(config-pmap)# end-policy-map
Router(config)# commit
```

Related Topics

- [Traffic Policy Elements, on page 23](#)
- [Traffic Class Elements, on page 21](#)

Enhanced Running Configuration Display for Policy Maps and Class Maps

Table 7: Feature History Table

Feature Name	Release Information	Feature Description
Enhanced Running Configuration Display for Policy Maps and Class Maps	Release 24.4.1	Introduced in this release on: Fixed Systems (8700) (select variants only*) *Enhanced running configuration display for policy maps and class maps is now supported on Cisco 8712-MOD-M routers.
Enhanced Running Configuration Display for Policy Maps and Class Maps	Release 24.2.11	You can view each class map or policy map running configuration instance on a separate line. The feature modifies the output display of this command: CLI: show run formal

Running Configuration Example

The class maps DSCP and MPLS are associated with the policy map p1.

After the configuration, its **show run formal** running configuration is displayed.

```
/* class map DSCP */
```

```

Router# config
Router(config)# class-map match-any DSCP
Router(config-cmap)# match dscp 1
Router(config-cmap)# end-class-map

/* class map MPLS */

Router(config)# class-map match-any MPLS
Router(config-cmap)# match mpls experimental topmost 2
Router(config-cmap)# end-class-map

/* DSCP-to-p1 association */

Router(config)# policy-map p1
Router(config-pmap)# class DSCP
Router(config-pmap-c)# bandwidth remaining ratio 80
Router(config-pmap-c)# root

/* MPLS-to-p1 association */

Router(config)# policy-map p1
Router(config-pmap)# class MPLS
Router(config-pmap-c)# bandwidth remaining ratio 60
Router(config-pmap-c)# exit
Router(config-pmap)# end-policy-map
Router(config)# commit

```

Verification

The **show run formal** running configuration displays each class map and policy map running configuration instance on a separate line:

```

Router# show run formal
..
class-map match-any DSCP match dscp 1
class-map match-any MPLS match mpls experimental topmost 2
policy-map p1 class DSCP bandwidth remaining ratio 80
policy-map p1 class MPLS bandwidth remaining ratio 60
..

```

show run and **show run formal** running configuration comparison:

show run displays indented configuration

```

class-map match-any DSCP
  match dscp 1
end-class-map
!
class-map match-any MPLS
  match mpls experimental topmost 2
end-class-map
!
policy-map p1
  class DSCP
    bandwidth remaining ratio 80
  !
  class MPLS
    bandwidth remaining ratio 60
  !
  class class-default
  !
end-policy-map

```

show run formal displays a running configuration instance on a separate line

```

class-map match-any DSCP match dscp 1
class-map match-any MPLS match mpls experimental topmost 2
policy-map p1 class DSCP bandwidth remaining ratio 80
policy-map p1 class MPLS bandwidth remaining ratio 60

```

Attach a Traffic Policy to an Interface

After the traffic class and the traffic policy are created, you must attach the traffic policy to interface, and specify the direction in which the policy should be applied.



Note Hierarchical policies are not supported.

When a policy-map is applied to an interface, the transmission rate counter of each class is not accurate. This is because the transmission rate counter is calculated based on the exponential decay filter.



Note When you enable BGP Flowspec on the ingress interface, if the Flowspec police rule is acting on the incoming data traffic and if the traffic matches the classification rule configured on the ingress policy-map, then the PBR ACL (policy-based routing access control list) police action configuration will override the metering pointer. As a result, the QoS input statistics (**show policy-map interface <interface> input**) for the ingress interface displays null values.

Configuration Example

You have to accomplish the following to attach a traffic policy to an interface:

1. Creating a traffic class and the associated rules that match packets to the class (see [Create a Traffic Class, on page 22](#))
2. Creating a traffic policy that can be attached to one or more interfaces to specify a service policy (see [Create a Traffic Policy, on page 24](#))
3. Associating the traffic class with the traffic policy
4. Attaching the traffic policy to an interface, in the ingress or egress direction

```
Router# configure
Router(config)# interface fourHundredGigE 0/0/0/2
Router(config-int)# service-policy output strict-priority
Router(config-int)# commit
```

Running Configuration

```
/* Class-map configuration */

class-map match-any traffic-class-7
match traffic-class 7
end-class-map

!class-map match-any traffic-class-6
match traffic-class 6
end-class-map
```

```

class-map match-any traffic-class-5
  match traffic-class 5
end-class-map

class-map match-any traffic-class-4
  match traffic-class 4
end-class-map

class-map match-any traffic-class-3
  match traffic-class 3

class-map match-any traffic-class-2
  match traffic-class 2
end-class-map

class-map match-any traffic-class-1
  match traffic-class 1
end-class-map

/* Traffic policy configuration */

policy-map test-shape-1
  class traffic-class-1
    shape average percent 40
  !

policy-map strict-priority
  class tc7
    priority level 1
    queue-limit 75 mbytes
  !
  class tc6
    priority level 2
    queue-limit 75 mbytes
  !
  class tc5
    priority level 3
    queue-limit 75 mbytes
  !
  class tc4
    priority level 4
    queue-limit 75 mbytes
  !
  class tc3
    priority level 5
    queue-limit 75 mbytes
  !
  class tc2
    priority level 6
    queue-limit 75 mbytes
  !
  class tc1
    priority level 7
    queue-limit 75 mbytes
  !
  class class-default
    queue-limit 75 mbytes
  !
end-policy-map

- - -
- - -

```

```

/* Attaching traffic policy to an interface in egress direction */
interface fourHundredGigE 0/0/0/2
 service-policy output strict-priority
 !

```

Verification

Router# **#show qos int fourHundredGigE 0/0/0/2 output**

NOTE:- Configured values are displayed within parentheses Interface FourHundredGigE0/0/0/2
ifh 0xf0001c0 -- output policy

```

NPU Id:                                0
Total number of classes:                8
Interface Bandwidth:                   4000000000 kbps
Policy Name:                           strict-priority
VOQ Base:                              2400
Accounting Type:                        Layer1 (Include Layer 1 encapsulation and above)
-----
Level1 Class (HP1)                     = tc7
Egressq Queue ID                       = 2407 (HP1 queue)
Queue Max. BW.                         = no max (default)
TailDrop Threshold                     = 74999808 bytes / 2 ms (75 megabytes)
WRED not configured for this class

Level1 Class (HP2)                     = tc6
Egressq Queue ID                       = 2406 (HP2 queue)
Queue Max. BW.                         = no max (default)
TailDrop Threshold                     = 74999808 bytes / 2 ms (75 megabytes)
WRED not configured for this class

Level1 Class (HP3)                     = tc5
Egressq Queue ID                       = 2405 (HP3 queue)
Queue Max. BW.                         = no max (default)
TailDrop Threshold                     = 74999808 bytes / 2 ms (75 megabytes)
WRED not configured for this class

Level1 Class (HP4)                     = tc4
Egressq Queue ID                       = 2404 (HP4 queue)
Queue Max. BW.                         = no max (default)
TailDrop Threshold                     = 74999808 bytes / 2 ms (75 megabytes)
WRED not configured for this class

Level1 Class (HP5)                     = tc3
Egressq Queue ID                       = 2403 (HP5 queue)
Queue Max. BW.                         = no max (default)
TailDrop Threshold                     = 74999808 bytes / 2 ms (75 megabytes)
WRED not configured for this class

Level1 Class (HP6)                     = tc2
Egressq Queue ID                       = 2402 (HP6 queue)
Queue Max. BW.                         = no max (default)
TailDrop Threshold                     = 74999808 bytes / 2 ms (75 megabytes)
WRED not configured for this class

Level1 Class (HP7)                     = tc1
Egressq Queue ID                       = 2401 (HP7 queue)
Queue Max. BW.                         = no max (default)
TailDrop Threshold                     = 74999808 bytes / 2 ms (75 megabytes)
WRED not configured for this class

Level1 Class                           = class-default

```

```

Egressq Queue ID           = 2400 (Default LP queue)
Queue Max. BW.             = no max (default)
Inverse Weight / Weight    = 1 / (BWR not configured)
TailDrop Threshold         = 74999808 bytes / 150 ms (75 megabytes)
WRED not configured for this class

```

!

Related Topics

- [Traffic Policy Elements, on page 23](#)
- [Traffic Class Elements, on page 21](#)

QoS Policy Propagation via BGP (QPPB)

Table 8: Feature History Table

Feature Name	Release Information	Feature Description
QoS Policy Propagation via BGP	Release 24.4.1	Introduced in this release on: Fixed Systems (8700) (select variants only*) *QoS Policy Propagation via BGP is now supported on Cisco 8712-MOD-M routers.
QoS Policy Propagation via BGP	Release 7.5.2	You now have the ability to install a BGP route in the routing table with a QoS Group so that IP packets that match the route receive the QoS policies associated with the QoS group. This functionality enables convenient classification and marking when BGP is deployed, overcoming the administrative challenges of classifying based on ACLs.

QoS Policy Propagation via Border Gateway Protocol (QPPB via BGP) is a mechanism that allows propagation of quality of service (QoS) policy and classification by the sending party that is based on the following:

- Access lists
- Community lists
- Autonomous system paths in the BGP

QPPB thus helps in classification that is based on the destination address instead of the source address.

QoS policies that differentiate between different types of traffic are defined for a single enterprise network. For instance, one enterprise may want to treat important web traffic, not-important web traffic, and all other data traffic as three different classes. And thereafter, use the different classes for the voice and video traffic.

Hence, QPPB overcomes:

- the administrative challenges of classifying that is based on ACLs.
- the administrative problems of just listing the networks that need premium services.

QPPB allows marking of packets that are based on QoS group value that is associated with a Border Gateway Protocol (BGP) route.

Benefits of QPPB

- QPPB provides an IP prefix-based QoS capability.
- Traffic to IP addresses that have specific IP prefixes can be prioritized above other IP addresses.
- IP prefixes of interest are tagged through the control plane that uses common BGP route-map techniques, including the community attribute.
- Traffic to the tagged BGP prefixes is then classified and prioritized via the data forwarding plane by using the IOS-XR MQC (Modular QoS CLI) mechanisms, such as remarking.
- QPPB provides the glue between the BGP control plane and the IP data forwarding plane in support of IP prefix-based QoS.
- BGP configuration within QPPB uses a table map to match specific prefixes learned through BGP neighbors, and then sets the router's local QoS Group variable that is maintained within the Forwarding Information Base (FIB) for those specific prefixes.

Guidelines and Limitations

- While configuring QPPB in the route policy for BGP, the range of the QoS group value that you can set using the **set qos-group** command is between 0 through 7.
- IP precedence-based QPPB isn't supported; only QoS group-based QPPB is supported.
- The order of precedence among ACL features (in ascending order):

With QoS-based group policy:

- Security ACL (0), BGP-FS (1), QPPB (2), Peering QoS (3)

• With ACL-based policy:


- Security ACL (0), ACL based QoS (1), BGP-FS (2), QPPB (3)

- You can configure QPPB on the following interface types:

- Physical
- Bundle
- Physical subinterface
- Bundle subinterface

- QPPB is not supported on Q100-based routers or line cards.
- QPPB overwrites the QoS Group value set by Security ACL.
- If you configure ACL-based Switched Port Analyzer (SPAN) on an interface, you can't enable QPPB on that interface. To enable QPPB, remove the ACL-based SPAN from the interface.
- If you configure QPPB on an interface, you can't enable ACL-based SPAN on that interface. To enable ACL-based SPAN, remove QPPB from the interface.
- Peering QoS overwrites the QoS Group value set by QPPB.
- Peering QoS and QPPB overwrite BGP FlowSpec actions except setting Drop and Policer. Currently, BGP FlowSpec sets DSCP and Meter.
- Remarking support is the same that QoS currently supports.
- Ingress policing support is the same that QoS currently supports.
- Egress policing is not supported.
- The router supports a subset of full QPPB options—only IP destination prefix mode on input policy is supported.

Figure 4: Sample Scenario



Router A learns routes from AS 200 and AS 100. QoS policy is applied to any ingress interface of Router A to match the defined route maps with destination prefixes of incoming packets. Matching packets on Router A to AS 200 or AS 100 are sent with the appropriate QoS policy from Router A.

BGP maintains a scalable database of destination prefixes, QPPB, by using BGP table maps. BGP adds the ability to map a qos-group value to desired IP destinations. These qos-group values are used in QoS policies applied locally on ingress interfaces. Whenever a packet that is bound for such destinations is encountered,

the qos-group value matching that destination route looks up with work inside the policy classmap, and marks that packet for any configured policy actions.

Configuration Workflow

Use the following configuration workflow for QPPB:

- Define routing policy.
- Put routing policy at table-policy attach point under BGP.
- Define classmaps and ingress policy to use the qos-groups that are used in table-policy.
- Enable IPv4 or IPv6 QPPB configuration under the desired interfaces.

Define Routing Policy

A routing policy instructs the router to inspect routes, filter them, and potentially modify their attributes as they are accepted from a peer, advertised to a peer, or redistributed from one routing protocol to another.

The routing policy language (RPL) provides a language to express routing policy. You must set up destination prefixes either to match inline values or one of a set of values in a prefix set.

Example:

```
prefix-set prefix-list-v4
  70.1.1.1,
  70.2.1.0/24,
  70.2.2.0/24 ge 28,
  70.2.3.0/24 le 28
end-set
prefix-set prefix-list-v6
  2001:300::2,
  2003:200::3
end-set

route-policy qppb1
  if destination in (60.60.0.2/24) then
    set qos-group 5
  elseif destination in prefix-list-v4 then
    set qos-group 4
  else
    set qos-group 1
  pass
endif
end-policy
route-policy qppb2
  if destination in prefix-list-v6 then
    set qos-group 5
  elseif destination in (2001:300::2) then
    set qos-group 4
  else
    set qos-group 1
  pass
endif
end-policy
```

Put Routing Policy at table-policy Attach Point Under BGP

The table-policy attach point permits the routing policy to perform actions on each route as they are installed into the RIB routing table. QPPB uses this attachment point to intercept all routes as they are received from peers. Ultimately the RIB will update the FIB in the hardware forwarding plane to store destination prefix routing entries, and in cases where table policy matches a destination prefix, the qos-group value is also stored with the destination prefix entry for use in the forwarding plane.

Example:

```
router bgp 900
  [vrf <name>]
  bgp router-id 22.22.22.22
  address-family ipv4 unicast
    table-policy qppb1
  address-family ipv6 unicast
    table-policy qppb2
  neighbor 30.2.2.1
    remote-as 500
    address-family ipv4 unicast
      route-policy pass in
      route-policy pass out
    address-family ipv6 unicast
      route-policy pass in
      route-policy pass out
```

Ingress QoS and IPv4 or IPv6 BGP Configuration

QPPB is enabled per interface and individually for IPv4 and IPv6. An ingress policy matches on the QoS groups marked by QPPB and takes the required action.

If a packet is destined for a destination prefix on which BGP route policy has stored a **qos-group**, but it ingresses on an interface on which QPPB is not enabled, it would not be remarked with **qos-group**.

Example:

```
class-map match-any qos-group5
  match qos-group 5
end-class-map

class-map match-any qos-group4
  match qos-group 4
end-class-map

policy-map ingress-marker-pol1
  class qos-group5
    set precedence 0
    set discard-class 0
    set traffic-class 1

  class qos-group4
    set precedence 1
    set discard-class 1
    set traffic-class 2
  class class-default

end-policy-map

Interface hun 0/0/0/0
  [vrf vrfA]
  ipv4 address 25.1.1.1/24
  ipv6 address 2001:db8:a0b:12f0::1/64
  ipv4 bgp policy propagation input qos-group destination
```

```
ipv6 bgp policy propagation input qos-group destination
service-policy input ingress-marker-pol
```

Egress Interface Configuration

The traffic-class set on ingress has no existence outside the device. Also, traffic-class is not a part of any packet header but is associated internal context data on relevant packets. It can be used as a match criteria in an egress policy to set up various fields on the outgoing packet or shape flows.

```
class-map match-any level1
    match traffic-class 1
end-class-map

class-map match-any level2
    match traffic-class 2
end-class-map

policy-map output-pol
    class level1
        bandwidth percent 50
    class level2
        bandwidth percent 20
        queue-limit 50 ms
end-policy-map

interface hun 0/5/0/0
    ipv4 address 30.1.1.1/24
    ipv6 address 2001:da8:b0a:12f0::1/64
    service-policy output output-pol
```

QoS on PWHE

Table 9: Feature History Table

Feature Name	Release Information	Feature Description
QoS on PWHE	Release 24.3.1	<p>Introduced in this release on: Modular Systems (8800 [LC ASIC: P100])(select variants only*)</p> <p>You can now prioritize critical services like voice, video, and real-time applications for network traffic entering a Pseudowire Headend (PWHE) interface, which then forwards the prioritized traffic into the pseudowire to be transported over a packet-switched network, such as an IP or MPLS network.</p> <p>The feature introduces the show qos-ea default-queue command.</p> <p>The feature modifies the <code>Cisco-IOS-XR-qos-ma-oper.yang</code> (see GitHub, YANG Data Models Navigator) data model.</p> <p>*This feature is supported on:</p> <ul style="list-style-type: none"> • 88-LC1-52Y8H-EM • 88-LC1-12TH24FH-E

QoS on Pseudowire Headend (PWHE) is a technology that

- applies QoS mechanisms to pseudowire headend interfaces in a network
- ensures that traffic transported over pseudowires is managed and prioritized effectively, and
- enhances the performance and reliability of Layer 2 and Layer 3 VPN services on an edge router.

Pseudowires—Pseudowires (PWs) are virtual point-to-point connections that emulate the properties of traditional physical wires. They allow for the transport of various types of data over a packet-switched network, such as an IP or MPLS network. Pseudowires enable the extension of Layer 2 services over a Layer 3 infrastructure, providing a seamless and transparent connection between two endpoints.

PWHE—PWHE refers to the technology and configuration used to extend Layer 2 services over a Layer 3 network by creating virtual point-to-point connections known as pseudowires. The headend is the starting point of the pseudowire, typically located at the service provider's edge router. PWHE enables the seamless transport of various data types, such as Ethernet, Frame Relay, or asynchronous transfer mode (ATM), over an IP or MPLS backbone.

For more information on PWHE, see the chapter *Configure Point-to-Point Layer 2 Services* in the *L2VPN Configuration Guide for Cisco 8000 Series Routers*.

Key Features and Functionalities of QoS on PWHE

Before configuring QoS on PWHE interfaces, it's essential to understand the features and functionalities this technology provides and the restrictions that accompany them. These include:

- **Supported Features:** Features supported on PWHE interfaces, including IPv4 and IPv6 support, policy maps for both ingress and egress traffic, and how policies are replicated across PWHE members.
- **Bandwidth Distribution:** How PWHE and non-PWHE traffic share scheduling resources, and configuring bandwidth to manage excess distribution effectively.
- **Classification and Marking:** Support for packet classification and marking based on Layer 2 and Layer 3 headers and specific configurations for PW-Ether main and PW-Ether subinterfaces.
- **Policing and Queuing:** The policing features supported on PWHE interfaces and subinterfaces, including the replication of policies on bundle members and the effect of these policies on traffic management.
- **PWHE Aggregate Shaper Policies:** PWHE aggregate shaper policies co-existing with subinterface policies, and various policy configurations and their implications on traffic flow.

General Features and Support Information for QoS on PWHE

- **IPv4 and IPv6 Support:** The feature supports both address families.
- **QoS Policies:**
 - Ingress: Support for policing and marking.
 - Egress: Support for marking and queuing.

Subinterface Policies:

- You can apply policy maps on PW-Ether subinterfaces.
- PW-Ether subinterfaces inherit policies from their main PW-Ether interface.

The table summarizes the policies you can configure on PWHE main and subinterfaces for egress.

On the egress PWHE main interface, when you configure...	On the egress PWHE subinterface, you can configure...
<ul style="list-style-type: none"> • two-level H-QoS queuing policy or flat QoS queuing policy, or • marking policy, or • both. 	no policies.

On the egress PWHE main interface, when you configure...	On the egress PWHE subinterface, you can configure...
no policies.	<ul style="list-style-type: none"> • two-level H-QoS queuing policy or flat QoS queuing policy, or • marking policy, or • both.
aggregate shaper, you can't configure a second policy.	marking policy, but you can't set traffic class. You can set bandwidth allocation in the queuing policy.

- **Simultaneous Policies:** Policies at the port for transit traffic can be applied simultaneously with policies for PWHE interfaces.
- **Traffic Mapping for PWHE Interfaces to Pin-Down Members:**

All the flows for a PWHE interface are mapped to single pin-down member.

Pinning down a member is a technique used to ensure that specific network resources remain consistently in use, providing stability, predictability, and control over network operations. This can be particularly useful in complex network environments where dynamic changes could lead to performance issues or disruptions.

For example, consider a case where the generic interface list has 10G and 100G interfaces, and 2 PWHE interfaces named PWHE-1 and PWHE-2 use this generic list.

If...	Then...
the traffic for PWHE-1 is hashed (or assigned) to the 10G interface	all the traffic for PWHE-1 will exclusively use the pinned-down 10G interface.
the traffic for PWHE-2 is hashed to the 100G interface	all the traffic for PWHE-2 will exclusively use the pinned-down 100G interface.

To summarize: Each PWHE interface's traffic is confined to a single physical interface from the generic interface list, ensuring that all data flows for that PWHE interface are routed through the same physical path.

- **Port Shaping Effect:** A port shaping policy on a member interface influences the amount of traffic that can pass through that port. This is especially important when traffic is routed through specific physical interfaces or pin-down members.

For example, if a pin-down member (a specific interface) has a shaping policy set to 10G on its default traffic class (**TC-0**), then the maximum bandwidth allowed on that interface is 10G. Consider a PWHE interface (named **PWHE-1**) that has a shaping policy that allows it to use up to 50% of the available bandwidth but its traffic is hashed to the pin-down member with 10G shaping policy. In such a case, the maximum bandwidth for PWHE-1 traffic on that pin-down member will be 5G.

To summarize: The port shaping policy on a pin-down member directly determines the amount of PWHE traffic that can pass through that port. Even if a PWHE interface has its own shaping policy, the actual bandwidth used will be constrained by the shaping policy of the pin-down member it is hashed to.

Bandwidth Distribution for QoS on PWHE

You can fine-tune how leftover bandwidth is distributed among various traffic classes after the highest-priority traffic has received its guaranteed bandwidth. This functionality helps maintain fairness and efficiency in bandwidth utilization across different traffic classes.

PWHE and non-PWHE traffic on the same pin-down member share scheduling resources. We recommend that you configure **bandwidth remaining** in the parent class-default of PWHE policies to control the distribution of excess bandwidth:

- between PWHE and non-PWHE traffic and
- between various PWHE interfaces and physical interfaces.

QoS Accounting

QoS accounting provides the data necessary to assess whether the bandwidth remaining ratio effectively distributes resources according to policy, allowing for informed adjustments to optimize network performance.

Some key points about QoS accounting for PWHE are:

- The packet length when performing QoS functions (such as policing, shaping, and gathering statistics) is based on the customer IP packet, customer L2 header and the configured additional overhead.
- QoS statistics include the customer IP packet, customer L2 header and configured additional overhead.
- Outer MPLS headers (such as VC labels and transport labels) and outer L2 header (Layer 2 encapsulation of the underlying physical interface) will not be included in the packet length when performing QoS on the PWHE virtual interface.

Marking for QoS on PWHE

Marking for PW-Ether in Ingress and Egress Directions

Proper packet marking on both ingress and egress interfaces is critical for maintaining QoS across the network for PW-Ether traffic. Ingress marking involves classifying, marking, and policing traffic as it enters the pseudowire, while egress marking focuses on preserving or adjusting those markings as traffic exits. To that end, knowing the supported QoS markings for PWHE traffic at the ingress and egress is essential.

Restrictions for Marking PW-Ether in Ingress and Egress Directions

- Marking of customer IP header, qos-group and discard-class are supported.
- Marking of CoS bits in transport L2 header are not supported.
- Depending on whether you want to apply conditional or unconditional markings and the traffic direction, you have different options.

Type of Marking	For Traffic Type	Supported Markings
Unconditional	Ingress and egress	<ul style="list-style-type: none">• DSCP or precedence value,• discard-class and,• qos-group.

Type of Marking	For Traffic Type	Supported Markings
Conditional policer	Ingress	At most two from <ul style="list-style-type: none"> • DSCP or precedence value, • qos-group, and • discard-class.
Unconditional policer	Egress	DSCP or precedence value and discard-class

Policing and Queuing for QoS on PWHE

Policing Features

All the policing features supported on normal L3 interfaces are also supported on the PWHE main interface and subinterface.

Egress Queuing Behavior

Depending on whether you attach a policy map to the PWHE interface or not, the queuing behavior changes for ingress and egress traffic.

Policy Map	Egress Queues	Ingress Policers
PWHE interface with no policy map	Each PWHE member has per port default queues. Egress traffic will use the member's port default queue. Queueing is supported only in the egress direction when you attach a service policy in the output direction.	Not applicable
PWHE interface with a policy map	Egress queues in the policy maps would be replicated on each PWHE member.	Policer bandwidth is calculated based on active members per slice, that is, members attached to the generic interface list on the PWHE interface.



Note If PWHE member is a bundle, policy maps will be replicated on bundle members.

Statistics

Show commands for a PWHE virtual interface and PWHE subinterface QoS policy provide egress statistics per pin-down member and per bundle member if the bundle is a pin-down member.

PWHE aggregate shaper statistics aggregate queuing statistics of all subinterfaces for every pin-down member.

PWHE Aggregate Shaper

You can configure a line card to allow PWHE aggregate shaper policy to co-exist with subinterface policies. The PWHE aggregate shaper policy will only have a class-default with shape and bandwidth remaining actions.

The PWHE aggregate shaper policy can have

- only shape action, or
- only bandwidth remaining action or,
- shape and bandwidth remaining actions (**recommended**).

An example of PWHE aggregate shaper policy with only shape action:

```
policy-map pwhe-aggregate-shaper
class class-default
shape average 1 gbps
!
end-policy-map
!
end
```

An example of PWHE aggregate shaper policy with only bandwidth remaining action:

```
policy-map pwhe-aggregate-shaper
class class-default
bandwidth remaining ratio 20
!
end-policy-map
!
end
```

Here is the example for PWHE aggregate shaper policy with shape and bandwidth remaining actions:

```
policy-map pwhe-aggregate-shaper
class class-default
shape average 1 gbps
bandwidth remaining ratio 20
!
end-policy-map
!
end
```



Note We recommend that you configure shape and bandwidth remaining actions for PWHE aggregate shaper policy.

PW-Ether Subinterface Policy

You can apply QoS policies on PW-Ether subinterfaces when there is no policy applied on the main PW-Ether interface.

You can set bandwidth limits using percentage values at the most granular level within the policy hierarchy applied to a PW-Ether subinterface. This allows for flexible and precise bandwidth management on subinterfaces.

For example, if you have a policy that manages traffic on a PW-Ether subinterface, you can specify that a particular class of traffic should use, say, 20% of the available bandwidth. This configuration is supported at the most granular level of the policy.

Scale Information

QoS on PWHE supports 256 PWHE interfaces with either ingress or egress policies per line card.



Note The scale numbers are applicable only if you apply configurations such that the queuing resource is not exhausted.

PWHE without QoS policy

These cases represent the default behavior of PWHE interfaces without QoS policies:

- PWHE ingress to core facing egress (access to core) - DSCP or precedence value from customer IP packet is copied to EXP of all imposed labels (VPN and transport) in the core-facing direction.
- PWHE egress (core to access) - DSCP or precedence value from customer IP packet is copied to EXP in the access-facing direction.

Limitations for QoS on PWHE

Supported Interfaces

You can configure PWHE in Layer 3 mode on PWHE-Ether main and subinterfaces.

QoS Accounting Scope

QoS accounting, which measures and records the packet length when performing QoS functions such as policing, shaping, and gathering statistics, doesn't include PW headers.

QoS Configuration Rules

- **match** commands are optional when configuring QoS on PWHE. However, you must configure at least one match criterion for a class.
- When using the **match access-group** command to configure the match criteria for a class map on the basis of the specified access control list (ACL), QoS classification based on the packet length or time to live (TTL) field in the IPv4 and IPv6 headers is not supported.

L2 Header Based Classification and Marking Scope

L2 header classification and marking are not supported on L3 PWHE interfaces.



Note The classification and marking applied on PW-Ether main interface are inherited by its subinterfaces without policy.

Configure QoS on PWHE

Depending on your network's specific requirements and goals, you have different options for configuring QoS on PWHE interfaces.

The table calls out some of the suitable scenarios for different configuration options.

Configuration option	Suitable for...
Flat QoS on PWHE	simple traffic management and basic prioritization without hierarchical structuring. Recommended for small to medium-sized networks with uniform traffic.
Hierarchical QoS (H-QoS) on PWHE	complex traffic management with multiple levels of QoS policies. Recommended for large networks requiring granular control and hierarchical structuring.
QoS with aggregate shaper on PWHE	controlling total bandwidth usage and smoothing traffic bursts. Recommended for networks needing to manage overall bandwidth and prevent congestion.

Configure Flat QoS on PWHE

This task configures QoS policies for a physical interface and a PWHE interface.

Before you begin

Determine which types of traffic are critical and need prioritization with finer control. Also, see [General Features and Support Information for QoS on PWHE, on page 37](#) for important information about the policies you can configure on PWHE main and subinterfaces for egress.

Procedure

- Step 1** Apply an input service policy named **MAIN_IN_PMAP5** to the **HundredGigE0/0/0/59** physical interface.

Example:

```
Router(config)#interface HundredGigE0/0/0/59
Router(config-if)#service-policy input MAIN_IN_PMAP5
Router(config-if)#commit
```

- Step 2** Apply two output service policies named **shape_flat_bwrr** and **MAIN_OUT_PMAP5_Q_TC_MARK** to the PWHE interface named **PW-Ether1**.

Example:

```
Router(config)#interface PW-Ether1
Router(config-if)#service-policy output shape_flat_bwrr
Router(config-if)#service-policy output MAIN_OUT_PMAP5_Q_TC_MARK
Router(config-if)#commit
```

Step 3 Configure bandwidth remaining ratio in the **shape_flat_bwrr** service policy on the **PW-Ether1** interface for different traffic classes.

Example:

```
Router(config)#policy-map type qos shape_flat_bwrr
Router(config-pmap)#class type qos MAIN_OUT_CMAP_1
Router(config-pmap-c)#bandwidth remaining ratio 20
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_CMAP_2
Router(config-pmap-c)#bandwidth remaining ratio 30
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_CMAP_3
Router(config-pmap-c)#bandwidth remaining ratio 40
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_CMAP_4
Router(config-pmap-c)#bandwidth remaining ratio 50
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_CMAP_5
Router(config-pmap-c)#bandwidth remaining ratio 60
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_CMAP_6
Router(config-pmap-c)#shape average percent 30
Router(config-pmap-c)#priority level 2
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_CMAP_7
Router(config-pmap-c)#shape average percent 40
Router(config-pmap-c)#priority level 1
Router(config-pmap)#end-policy-map
```

Step 4 Set traffic classes in the **MAIN_OUT_PMAP5_Q_TC_MARK** policy on the **PW-Ether1** interface.

Example:

```
Router(config)#policy-map type qos MAIN_OUT_PMAP5_Q_TC_MARK
Router(config-pmap)#class type qos MAIN_OUT_MARK_CMAP_1
Router(config-pmap-c)#set traffic-class 1
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_MARK_CMAP_2
Router(config-pmap-c)#set traffic-class 2

Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_MARK_CMAP_3
Router(config-pmap-c)#set traffic-class 3
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_MARK_CMAP_4
Router(config-pmap-c)#set traffic-class 4
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_MARK_CMAP_5
Router(config-pmap-c)#set traffic-class 5
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_MARK_CMAP_6
Router(config-pmap-c)#set traffic-class 6
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_MARK_CMAP_7
Router(config-pmap-c)#set traffic-class 7
Router(config-pmap-c)#exit
Router(config-pmap)#end-policy-map
```

Step 5 Set traffic classes and QoS groups in the **MAIN_IN_PMAP5** policy on the **HundredGigE0/0/0/59** physical interface.

Example:

```

Router(config)#policy-map type qos MAIN_IN_PMAP5
Router(config-pmap)#class type qos MAIN_IN_CMAP8_1
Router(config-pmap-c)#set traffic-class 1
Router(config-pmap-c)#set qos-group 1
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_IN_CMAP8_2
Router(config-pmap-c)#set traffic-class 2
Router(config-pmap-c)#set qos-group 2
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_IN_CMAP8_3
Router(config-pmap-c)#set traffic-class 3
Router(config-pmap-c)#set qos-group 3
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_IN_CMAP8_4
Router(config-pmap-c)#set traffic-class 4
Router(config-pmap-c)#set qos-group 4
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_IN_CMAP8_5
Router(config-pmap-c)#set traffic-class 5
Router(config-pmap-c)#set qos-group 5
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_IN_CMAP8_6
Router(config-pmap-c)#set traffic-class 6
Router(config-pmap-c)#set qos-group 6
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_IN_CMAP8_7
Router(config-pmap-c)#set traffic-class 7
Router(config-pmap-c)#set qos-group 7
Router(config-pmap-c)#exit
Router(config-pmap)#end-policy-map
Router(config-pmap)#commit

```

Step 6 Verify the flat QoS configuration for the PWHE interface named **PW-Ether1**.

Example:

Router#**show policy-map interface PW-Ether1 output**

PW-Ether1 output: shape_flat_bwrr

```

Class MAIN_OUT_CMAP_1
  Classification statistics          (packets/bytes)      (rate - kbps)
    Matched                        :      3713933519/885747042770      20297753
    Transmitted                    :      819154447/194810747182      4464282
    Total Dropped                  :      2894779072/690936295588      15833471
  Queueing statistics
    Queue ID                       : None (Pseudo-Wire)
    Taildropped(packets/bytes)     : 2894779072/690936295588
Class MAIN_OUT_CMAP_2
  Classification statistics          (packets/bytes)      (rate - kbps)
    Matched                        :      3713933859/885747126352      20297723
    Transmitted                    :      1226431948/292261487200      6697427
    Total Dropped                  :      2487501911/593485639152      13600296
  Queueing statistics
    Queue ID                       : None (Pseudo-Wire)
    Taildropped(packets/bytes)     : 2487501911/593485639152
Class MAIN_OUT_CMAP_3
  Classification statistics          (packets/bytes)      (rate - kbps)
    Matched                        :      3713934320/885747236808      20297726
    Transmitted                    :      1635701991/389672415000      8929683
    Total Dropped                  :      2078232329/496074821808      11368043
  Queueing statistics

```

```

Queue ID                               : None (Pseudo-Wire)
Taildropped(packets/bytes)             : 2078232329/496074821808
Class MAIN_OUT_CMAP_4
  Classification statistics              (packets/bytes)      (rate - kbps)
  Matched                               : 3713934620/885747307472      20297725
  Transmitted                            : 1962135193/467620808304      10715952
  Total Dropped                          : 1751799427/418126499168      9581773
Queueing statistics
Queue ID                               : None (Pseudo-Wire)
Taildropped(packets/bytes)             : 1751799427/418126499168
Class MAIN_OUT_CMAP_5
  Classification statistics              (packets/bytes)      (rate - kbps)
  Matched                               : 3713934983/885747394920      20297724
  Transmitted                            : 2453607357/584506995776      13394493
  Total Dropped                          : 1260327626/301240399144      6903231
Queueing statistics
Queue ID                               : None (Pseudo-Wire)
Taildropped(packets/bytes)             : 1260327626/301240399144
Class MAIN_OUT_CMAP_6
  Classification statistics              (packets/bytes)      (rate - kbps)
  Matched                               : 3713935277/885747464816      20297721
  Transmitted                            : 3713522895/885652168176      20295533
  Total Dropped                          : 412382/95296640              2188
Queueing statistics
Queue ID                               : None (Pseudo-Wire)
Taildropped(packets/bytes)             : 412382/95296640
Class MAIN_OUT_CMAP_7
  Classification statistics              (packets/bytes)      (rate - kbps)
  Matched                               : 3713935660/885747555696      20297721
  Transmitted                            : 3713935660/885747555696      20297721
  Total Dropped                          : 0/0                          0
Queueing statistics
Queue ID                               : None (Pseudo-Wire)
Taildropped(packets/bytes)             : 0/0
Class class-default
  Classification statistics              (packets/bytes)      (rate - kbps)
  Matched                               : 3713933237/885746977848      20297727
  Transmitted                            : 40921134/9742495832          223233
  Total Dropped                          : 3673012103/876004482016      20074494
Policy Bag Stats time: 1722242150998 [Local Time: 07/29/24 08:35:50.998]

```

PW-Ether1 output: MAIN_OUT_PMAP5_Q_TC_MARK

```

Class MAIN_OUT_MARK_CMAP_1
  Classification statistics              (packets/bytes)      (rate - kbps)
  Matched                               : 3714522947/885887620056      20317550
  Transmitted                            : 3714522947/885887620056      20317550
  Total Dropped                          : 0/0                          0
Class MAIN_OUT_MARK_CMAP_2
  Classification statistics              (packets/bytes)      (rate - kbps)
  Matched                               : 3714523360/885887719136      20317549
  Transmitted                            : 3714523360/885887719136      20317549
  Total Dropped                          : 0/0                          0
Class MAIN_OUT_MARK_CMAP_3
  Classification statistics              (packets/bytes)      (rate - kbps)
  Matched                               : 3714523737/885887808376      20317548
  Transmitted                            : 3714523737/885887808376      20317548
  Total Dropped                          : 0/0                          0
Class MAIN_OUT_MARK_CMAP_4
  Classification statistics              (packets/bytes)      (rate - kbps)
  Matched                               : 3714524115/885887898616      20317547
  Transmitted                            : 3714524115/885887898616      20317547
  Total Dropped                          : 0/0                          0
Class MAIN_OUT_MARK_CMAP_5

```



```

Classification statistics      (packets/bytes)      (rate - kbps)
Matched      :      3714524499/885887990496      20317545
Transmitted  :      3714524499/885887990496      20317545
Total Dropped :      0/0      0
Class MAIN_OUT_MARK_CMAP_6
Classification statistics      (packets/bytes)      (rate - kbps)
Matched      :      3714524873/885888079352      20317545
Transmitted  :      3714524873/885888079352      20317545
Total Dropped :      0/0      0
Class MAIN_OUT_MARK_CMAP_7
Classification statistics      (packets/bytes)      (rate - kbps)
Matched      :      3714525265/885888173128      20317546
Transmitted  :      3714525265/885888173128      20317546
Total Dropped :      0/0      0
Class class-default
Classification statistics      (packets/bytes)      (rate - kbps)
Matched      :      3714525658/885888266160      20317545
Transmitted  :      3714525658/885888266160      20317545
Total Dropped :      0/0      0
Policy Bag Stats time: 1722242151064 [Local Time: 07/29/24 08:35:51.064]

```

Configure H-QoS on PWHE

This task sets up H-QoS on a PW-Ether interface. The configuration involves applying a parent policy map to the PW-Ether interface, which includes shaping and bandwidth management rules. The parent policy map references a child policy map that further defines bandwidth allocation and queuing for different traffic classes.

Before you begin

Determine which types of traffic are critical and need prioritization with finer control. Also, see [General Features and Support Information for QoS on PWHE, on page 37](#) for important information about QoS policies on PWHE interfaces for egress.

Procedure

Step 1 Apply an output service policy named **parent_shape_hqos_bwrr** to the **PW-Ether1** interface.

Example:

```

Router(config)#interface PW-Ether1
Router(config-if)#service-policy output parent_shape_hqos_bwrr

```

This policy includes shaping and bandwidth management rules.

Step 2 Define a policy map named **parent_shape_hqos_bwrr** for QoS.

Example:

```

Router(config)#policy-map type qos parent_shape_hqos_bwrr
Router(config-pmap)#class type qos class-default
Router(config-pmap-c)#shape average percent 70
Router(config-pmap-c)#end-policy-map

```

Within the policy map, specify the class type as **class-default** and apply traffic shaping to it.

Step 3 Reference a child policy map named **shape_hqos_bwrr** within **class-default** of the parent policy map.

Example:

```
Router(config)#policy-map type qos parent_shape_hqos_bwrr
Router(config-pmap)#class type qos class-default
Router(config-pmap-c)#service-policy type qos shape_hqos_bwrr
```

Step 4 Define the child policy map **shape_hqos_bwrr** by specifying different classes within the map.

Example:

```
Router(config)#policy-map type qos shape_hqos_bwrr
Router(config-pmap)#class type qos MAIN_OUT_CMAP_1
Router(config-pmap-c)#bandwidth remaining ratio 20
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_CMAP_2
Router(config-pmap-c)#bandwidth remaining ratio 30
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_CMAP_3
Router(config-pmap-c)#bandwidth remaining ratio 40
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_CMAP_4
Router(config-pmap-c)#bandwidth remaining ratio 50
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_CMAP_5
Router(config-pmap-c)#bandwidth remaining ratio 60
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_CMAP_6
Router(config-pmap-c)#shape average percent 10
Router(config-pmap-c)#priority level 2
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_CMAP_7
Router(config-pmap-c)#shape average percent 20
Router(config-pmap-c)#priority level 1
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos class-default
Router(config-pmap-c)#bandwidth remaining ratio 10
Router(config-pmap-c)#end-policy-map
Router(config)#commit
```

Configure QoS parameters such as bandwidth remaining ratio and traffic shaping for these classes depending on the control you want to achieve.

Step 5 Verify the H-QoS configuration for the **PW-Ether1** interface.

Example:

```
Router# show policy-map interface PW-Ether1 output
```

```
PW-Ether1 output: MAIN_OUT_PMAP5_Q_TC_MARK
Class MAIN_OUT_MARK_CMAP_1
```

Classification statistics		(packets/bytes)	(rate - kbps)
Matched	:	567459747/135334786816	12797449
Transmitted	:	567459747/135334786816	12797449
Total Dropped	:	0/0	0

```
Class MAIN_OUT_MARK_CMAP_2
```

Classification statistics		(packets/bytes)	(rate - kbps)
---------------------------	--	-----------------	---------------

```

    Matched          :          567460159/135334884896      12797447
    Transmitted       :          567460159/135334884896      12797447
    Total Dropped     :              0/0                    0

Class MAIN_OUT_MARK_CMAP_3
    Classification statistics      (packets/bytes)      (rate - kbps)
    Matched          :          567460559/135334979696      12797443
    Transmitted       :          567460559/135334979696      12797443
    Total Dropped     :              0/0                    0

Class MAIN_OUT_MARK_CMAP_4
    Classification statistics      (packets/bytes)      (rate - kbps)
    Matched          :          567460923/135335066400      12797441
    Transmitted       :          567460923/135335066400      12797441
    Total Dropped     :              0/0                    0

Class MAIN_OUT_MARK_CMAP_5
    Classification statistics      (packets/bytes)      (rate - kbps)
    Matched          :          567461313/135335159920      12797441
    Transmitted       :          567461313/135335159920      12797441
    Total Dropped     :              0/0                    0

Class MAIN_OUT_MARK_CMAP_6
    Classification statistics      (packets/bytes)      (rate - kbps)
    Matched          :          567461672/135335245112      12797434
    Transmitted       :          567461672/135335245112      12797434
    Total Dropped     :              0/0                    0

Class MAIN_OUT_MARK_CMAP_7
    Classification statistics      (packets/bytes)      (rate - kbps)
    Matched          :          567462071/135335339784      12797433
    Transmitted       :          567462071/135335339784      12797433
    Total Dropped     :              0/0                    0

Class class-default
    Classification statistics      (packets/bytes)      (rate - kbps)
    Matched          :          567462460/135335433176      12797433
    Transmitted       :          567462460/135335433176      12797433

```

```

Total Dropped      :                0/0                0

Policy Bag Stats time: 1722242661147

PW-Ether1 output: parent_shape_hqos_bwrr

Class class-default

Classification statistics      (packets/bytes)      (rate - kbps)

Matched              :          4534918248/1081543126494      103380648

Transmitted          :          1754661178/418305616732      40446522

Total Dropped        :          2780257070/663237509762      62934126

Policy shape_hqos_bwrr Class MAIN_OUT_CMAP_1

Classification statistics      (packets/bytes)      (rate - kbps)

Matched              :          566863835/135192664022      13582722

Transmitted          :          116771460/27808308708      3393098

Total Dropped        :          450092375/107384355314      10189624

Queueing statistics

Queue ID              : None (Pseudo-Wire)

Taildropped(packets/bytes) : 450092375/107384355314

Policy shape_hqos_bwrr Class MAIN_OUT_CMAP_2

Classification statistics      (packets/bytes)      (rate - kbps)

Matched              :          566864236/135192760288      12828274

Transmitted          :          175109685/41713170712      3958057

Total Dropped        :          391754551/93479589576      8870217

Queueing statistics

Queue ID              : None (Pseudo-Wire)

Taildropped(packets/bytes) : 391754551/93479589576

Policy shape_hqos_bwrr Class MAIN_OUT_CMAP_3

Classification statistics      (packets/bytes)      (rate - kbps)

Matched              :          566864574/135192841920      12828272

Transmitted          :          233516796/55616629856      5277334

Total Dropped        :          333347778/79576212064      7550938

Queueing statistics

Queue ID              : None (Pseudo-Wire)

Taildropped(packets/bytes) : 333347778/79576212064

```

Policy shape_hqos_bwrr Class MAIN_OUT_CMAP_4

Classification statistics		(packets/bytes)	(rate - kbps)
Matched	:	566864964/135192934568	12828276
Transmitted	:	280180776/66740681128	6332876
Total Dropped	:	286684188/68452253440	6495400

Queueing statistics

Queue ID	:	None (Pseudo-Wire)
Taildropped(packets/bytes)	:	286684188/68452253440

Policy shape_hqos_bwrr Class MAIN_OUT_CMAP_5

Classification statistics		(packets/bytes)	(rate - kbps)
Matched	:	566865352/135193027832	12828278
Transmitted	:	350225162/83425674520	7916096
Total Dropped	:	216640190/51767353312	4912182

Queueing statistics

Queue ID	:	None (Pseudo-Wire)
Taildropped(packets/bytes)	:	216640190/51767353312

Policy shape_hqos_bwrr Class MAIN_OUT_CMAP_6

Classification statistics		(packets/bytes)	(rate - kbps)
Matched	:	566865746/135193120992	12828275
Transmitted	:	180140646/43035061400	4083496
Total Dropped	:	386725100/92158059592	8744779

Queueing statistics

Queue ID	:	None (Pseudo-Wire)
Taildropped(packets/bytes)	:	386725100/92158059592

Policy shape_hqos_bwrr Class MAIN_OUT_CMAP_7

Classification statistics		(packets/bytes)	(rate - kbps)
Matched	:	566866035/135193189376	12828277
Transmitted	:	360261283/86063059512	8166361
Total Dropped	:	206604752/49130129864	4661916

Queueing statistics

Queue ID	:	None (Pseudo-Wire)
Taildropped(packets/bytes)	:	206604752/49130129864

```

Policy shape_hqos_bwrr Class class-default

Classification statistics          (packets/bytes)      (rate - kbps)

Matched                          :      566863506/135192587496      12828274

Transmitted                      :           58455370/13903030896      1319204

Total Dropped                    :      508408136/121289556600      11509070

Queueing statistics

Queue ID                        : None (Pseudo-Wire)

Taildropped(packets/bytes)      :      508408136/121289556600

Policy Bag Stats time: 1722242661011

```

Configure QoS with Aggregate Shaper on PWHE

This task configures QoS policies for a PWHE interface and its subinterface. The main interface policy has a class-default (TC_0) with aggregate shaper with traffic shaping and bandwidth remaining rules.

Before you begin

Determine which types of traffic are critical and need prioritization with finer control. Also, see [General Features and Support Information for QoS on PWHE, on page 37](#) for important information about the policies you can configure on PWHE main and subinterfaces for egress. Also, see [PWHE Aggregate Shaper, on page 41](#) for details about aggregate shaper policies.

Procedure

- Step 1** Apply an output service policy with aggregate shaper named **port_shaper_combination_hqos_bwrr** to the **PW-Ether1** main interface.

Example:

```

Router(config)#interface PW-Ether1
Router(config-if)#service-policy output port_shaper_combination_hqos_bwrr
Router(config-if)#commit

```

- Step 2** Apply an output service policy named **MAIN_OUT_DSCP_MARK** to the **PW-Ether1.1** subinterface.

Example:

```

Router(config)#interface PW-Ether1.1
Router(config-if)#service-policy output MAIN_OUT_DSCP_MARK
Router(config-if)#commit

```

Use this policy to mark DSCP values on outgoing traffic.

- Step 3** Define a policy map named **MAIN_OUT_DSCP_MARK** for QoS. Define a class map named **MAIN_OUT_MARK_DSCP_CMAP_1** within the policy map.

Example:

```

Router(config)#policy-map type qos MAIN_OUT_DSCP_MARK
Router(config-pmap)#class type qos MAIN_OUT_MARK_DSCP_CMAP_1
Router(config-pmap-c)#set dscp af11
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_MARK_DSCP_CMAP_2
Router(config-pmap-c)#set dscp af12
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_MARK_DSCP_CMAP_3
Router(config-pmap-c)#set dscp af13
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_MARK_DSCP_CMAP_4
Router(config-pmap-c)#set dscp af41
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_MARK_DSCP_CMAP_5
Router(config-pmap-c)#set dscp af42
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_MARK_DSCP_CMAP_6
Router(config-pmap-c)#set dscp af43
Router(config-pmap-c)#exit
Router(config-pmap)#class type qos MAIN_OUT_MARK_DSCP_CMAP_7
Router(config-pmap-c)#set dscp ef
Router(config-pmap-c)#end-policy-map

Router(config-pmap)#commit

```

Set the DSCP value to **af11** for traffic matching the **MAIN_OUT_MARK_DSCP_CMAP_1** class, and so on for class maps **MAIN_OUT_MARK_DSCP_CMAP_2** through **MAIN_OUT_MARK_DSCP_CMAP_6**. For the class map named **MAIN_OUT_MARK_DSCP_CMAP_7**, set the DSCP value to **ef** for traffic matching the **MAIN_OUT_MARK_DSCP_CMAP_7** class. This step sets specific DSCP values for different traffic classes. This policy is applied to the subinterface to ensure that outgoing traffic is marked with the appropriate DSCP values for QoS.

- Step 4** Verify that the aggregate shaper is enforced on the default traffic class **TC_0** for the egress traffic from **PW-Ether1** main interface.

Example:

```
Router#show qos-ea default-queue PW-Ether 1 member interface HundredGigE0/7/0/2 output
```

Interface Name	=	PE1		
Interface Handle	=	1c000180		
Location	=	0/7/CPU0		
Asic Instance	=	0		
VOQ Base	=	53312		
Port Speed(kbps)	=	100000000		
Local Port	=	etm_local		
VOQ Mode	=	8		
	ReceivedPkts	ReceivedBytes	DroppedPkts	DroppedBytes
-----	-----	-----	-----	-----
TC_0 = 150295879	150596322940	762151909	763676346056	
TC_1 = 0	0	0	0	
TC_2 = 0	0	0	0	
TC_3 = 0	0	0	0	
TC_4 = 0	0	0	0	
TC_5 = 0	0	0	0	
TC_6 = 0	0	0	0	
TC_7 = 0	0	0	0	

- Step 5** View detailed QoS statistics and configuration information for output traffic from the **PW-Ether1** main interface.

Example:

```

Router#show qos interface PW-Ether1 output member HundredGigE0/7/0/2
NOTE:- Configured values are displayed within parentheses
Node 0/7/CPU0, Interface PW-Ether1 Ifh 0x7800104c (PW-Ether) -- output policy
NPU Id: 0
Total number of classes: 9
Interface Bandwidth: 100000000 kbps
Policy Name: parent_shape_hqos_bwrr
Accounting Type: Layer1 (Include Layer 1 encapsulation and above)
-----
Level1 Class = class-default

Level2 Class = MAIN_OUT_CMAP_1

Level2 Class = MAIN_OUT_CMAP_2

Level2 Class = MAIN_OUT_CMAP_3

Level2 Class = MAIN_OUT_CMAP_4

Level2 Class = MAIN_OUT_CMAP_5

Level2 Class = MAIN_OUT_CMAP_6

Level2 Class = MAIN_OUT_CMAP_7

Level2 Class = class-default
Node 0/7/CPU0, Interface PW-Ether1 Ifh 0x7800104c (PW-Ether) -- output policy
NPU Id: 0
Total number of classes: 8
Interface Bandwidth: 100000000 kbps
Policy Name: MAIN_OUT_PMAP5_Q_TC_MARK
Accounting Type: Layer1 (Include Layer 1 encapsulation and above)
-----
Level1 Class = MAIN_OUT_MARK_CMAP_1
New traffic class = 1

Level1 Class = MAIN_OUT_MARK_CMAP_2
New traffic class = 2

Level1 Class = MAIN_OUT_MARK_CMAP_3
New traffic class = 3

Level1 Class = MAIN_OUT_MARK_CMAP_4
New traffic class = 4

Level1 Class = MAIN_OUT_MARK_CMAP_5
New traffic class = 5

Level1 Class = MAIN_OUT_MARK_CMAP_6
New traffic class = 6

Level1 Class = MAIN_OUT_MARK_CMAP_7
New traffic class = 7

Level1 Class = class-default
Node 0/7/CPU0, Interface PW-Ether1 Ifh 0x7800104c (PW-Ether) -- output policy
NPU Id: 3
Total number of classes: 9
Interface Bandwidth: 100000000 kbps
Policy Name: parent_shape_hqos_bwrr
Accounting Type: Layer1 (Include Layer 1 encapsulation and above)
-----
Level1 Class = class-default

```



```

Level2 Class                               =   MAIN_OUT_CMAP_1

Level2 Class                               =   MAIN_OUT_CMAP_2

Level2 Class                               =   MAIN_OUT_CMAP_3

Level2 Class                               =   MAIN_OUT_CMAP_4

Level2 Class                               =   MAIN_OUT_CMAP_5

Level2 Class                               =   MAIN_OUT_CMAP_6

Level2 Class                               =   MAIN_OUT_CMAP_7

Level2 Class                               =   class-default
Node 0/7/CPU0, Interface PW-Ether1 Ifh 0x7800104c (PW-Ether) -- output policy
NPU Id:                                    3
Total number of classes:                   8
Interface Bandwidth:                       100000000 kbps
Policy Name:                              MAIN_OUT_PMAP5_Q_TC_MARK
Accounting Type:                           Layer1 (Include Layer 1 encapsulation and above)
-----
Level1 Class                               =   MAIN_OUT_MARK_CMAP_1
New traffic class                          =   1
Level1 Class                               =   MAIN_OUT_MARK_CMAP_2
New traffic class                          =   2

Level1 Class                               =   MAIN_OUT_MARK_CMAP_3
New traffic class                          =   3

Level1 Class                               =   MAIN_OUT_MARK_CMAP_4
New traffic class                          =   4

Level1 Class                               =   MAIN_OUT_MARK_CMAP_5
New traffic class                          =   5

Level1 Class                               =   MAIN_OUT_MARK_CMAP_6
New traffic class                          =   6

Level1 Class                               =   MAIN_OUT_MARK_CMAP_7
New traffic class                          =   7

Level1 Class                               =   class-default
Interface HundredGigE0/7/0/2 Ifh 0x1c000180 (PinDown Member) -- output policy
NPU Id:                                    0
Total number of classes:                   9
Interface Bandwidth:                       100000000 kbps
Policy Name:                              parent_shape_hqos_bwrr
VOQ Base:                                  53312
Accounting Type:                           Layer1 (Include Layer 1 encapsulation and above)
VOQ Mode:                                  8
Shared Counter Mode:                       1
-----
Level1 Class                               =   class-default
Queue Max. BW.                             =   70000000 kbps (70 %)
Inverse Weight / Weight                    =   255 / (BWR not configured)

Level2 Class                               =   MAIN_OUT_CMAP_1
Egressq Queue ID                           =   53313 (LP queue)
Queue Max. BW.                             =   no max (default)
Inverse Weight / Weight                    =   12 / (20)
Guaranteed service rate                    =   4711538 kbps
TailDrop Threshold                         =   614400 bytes / 1 ms (614400 bytes)
WRED not configured for this class

```

```

Level2 Class                               = MAIN_OUT_CMAP_2
Egressq Queue ID                           = 53314 (LP queue)
Queue Max. BW.                             = no max (default)
Inverse Weight / Weight                    = 8 / (30)
Guaranteed service rate                     = 7067307 kbps
TailDrop Threshold                         = 614400 bytes / 702 us (614400 bytes)
WRED not configured for this class

Level2 Class                               = MAIN_OUT_CMAP_3
Egressq Queue ID                           = 53315 (LP queue)
Queue Max. BW.                             = no max (default)
Inverse Weight / Weight                    = 6 / (40)
Guaranteed service rate                     = 9423076 kbps
TailDrop Threshold                         = 614400 bytes / 527 us (614400 bytes)
WRED not configured for this class

Level2 Class                               = MAIN_OUT_CMAP_4
Egressq Queue ID                           = 53316 (LP queue)
Queue Max. BW.                             = no max (default)
Inverse Weight / Weight                    = 5 / (50)
Guaranteed service rate                     = 11307692 kbps
TailDrop Threshold                         = 614400 bytes / 421 us (614400 bytes)
WRED not configured for this class

Level2 Class                               = MAIN_OUT_CMAP_5
Egressq Queue ID                           = 53317 (LP queue)
Queue Max. BW.                             = no max (default)
Inverse Weight / Weight                    = 4 / (60)
Guaranteed service rate                     = 14134615 kbps
TailDrop Threshold                         = 614400 bytes / 351 us (614400 bytes)
WRED not configured for this class

Level2 Class (HP2)                         = MAIN_OUT_CMAP_6
Egressq Queue ID                           = 53318 (HP2 queue)
Queue Max. BW.                             = 7000000 kbps (10 %)
Guaranteed service rate                     = 7000000 kbps
TailDrop Threshold                         = 614400 bytes / 702 us (614400 bytes)
WRED not configured for this class

Level2 Class (HP1)                         = MAIN_OUT_CMAP_7
Egressq Queue ID                           = 53319 (HP1 queue)
Queue Max. BW.                             = 14000000 kbps (20 %)
Guaranteed service rate                     = 14000000 kbps
TailDrop Threshold                         = 614400 bytes / 351 us (614400 bytes)
WRED not configured for this class

Level2 Class                               = class-default
Egressq Queue ID                           = 53312 (Default LP queue)
Queue Max. BW.                             = no max (default)
Inverse Weight / Weight                    = 24 / (10)
Guaranteed service rate                     = 2355769 kbps
TailDrop Threshold                         = 614400 bytes / 2 ms (614400 bytes)
WRED not configured for this class

Interface HundredGigE0/7/0/2 Ifh 0x1c000180 (PinDown Member) -- output policy
NPU Id:                                    0
Total number of classes:                   8
Interface Bandwidth:                       100000000 kbps
Policy Name:                               MAIN_OUT_PMAP5_Q_TC_MARK
VOQ Base:                                  0
Accounting Type:                            Layer1 (Include Layer 1 encapsulation and above)
VOQ Mode:                                   8
Shared Counter Mode:                       1

```

```

-----
Level1 Class                      =  MAIN_OUT_MARK_CMAP_1
New traffic class                  =  1

Level1 Class                      =  MAIN_OUT_MARK_CMAP_2
New traffic class                  =  2

Level1 Class                      =  MAIN_OUT_MARK_CMAP_3
New traffic class                  =  3

Level1 Class                      =  MAIN_OUT_MARK_CMAP_4
New traffic class                  =  4

Level1 Class                      =  MAIN_OUT_MARK_CMAP_5
New traffic class                  =  5

Level1 Class                      =  MAIN_OUT_MARK_CMAP_6
New traffic class                  =  6

Level1 Class                      =  MAIN_OUT_MARK_CMAP_7
New traffic class                  =  7

Level1 Class                      =  class-default

Node 0/0/CPU0, Interface PW-Ether1 Ifh 0x7800104c (PW-Ether) -- output policy
NPU Id:                          0
Total number of classes:          9
Interface Bandwidth:              25000000 kbps
Policy Name:                     parent_shape_hqos_bwrr
Accounting Type:                  Layer1 (Include Layer 1 encapsulation and above)
-----
Level1 Class                      =  class-default

    Level12 Class                  =  MAIN_OUT_CMAP_1

    Level12 Class                  =  MAIN_OUT_CMAP_2

    Level12 Class                  =  MAIN_OUT_CMAP_3

    Level12 Class                  =  MAIN_OUT_CMAP_4

    Level12 Class                  =  MAIN_OUT_CMAP_5

    Level12 Class                  =  MAIN_OUT_CMAP_6

    Level12 Class                  =  MAIN_OUT_CMAP_7

    Level12 Class                  =  class-default
Node 0/0/CPU0, Interface PW-Ether1 Ifh 0x7800104c (PW-Ether) -- output policy
NPU Id:                          0
Total number of classes:          8
Interface Bandwidth:              25000000 kbps
Policy Name:                     MAIN_OUT_PMAP5_Q_TC_MARK
Accounting Type:                  Layer1 (Include Layer 1 encapsulation and above)
-----
Level1 Class                      =  MAIN_OUT_MARK_CMAP_1
New traffic class                  =  1

Level1 Class                      =  MAIN_OUT_MARK_CMAP_2
New traffic class                  =  2

Level1 Class                      =  MAIN_OUT_MARK_CMAP_3
New traffic class                  =  3

```

```

Level1 Class                =   MAIN_OUT_MARK_CMAP_4
New traffic class           =   4

Level1 Class                =   MAIN_OUT_MARK_CMAP_5
New traffic class           =   5

Level1 Class                =   MAIN_OUT_MARK_CMAP_6
New traffic class           =   6

Level1 Class                =   MAIN_OUT_MARK_CMAP_7
New traffic class           =   7

Level1 Class                =   class-default

```

Step 6 View detailed information about the QoS policy maps applied to the egress traffic on a specific member interface (**HundredGigE0/7/0/2**) on the **PW-Ether1** main interface.

Example:

```

Router#show policy-map interface PW-Ether1 output member HundredGigE0/7/0/2
Interface:PW-Ether1 Member:HundredGigE0/7/0/2 output: parent_shape_hqos_bwrr

```

```

Class class-default
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :      1048830183/250133071040      0
  Transmitted                  :      405970112/96745121320      0
  Total Dropped                :      642860071/153387949720      0

Policy shape_hqos_bwrr Class MAIN_OUT_CMAP_1
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :      131102864/31266417264      0
  Transmitted                  :      27020068/6431611720      0
  Total Dropped                :      104082796/24834805544      0
  Queueing statistics
    Queue ID                   :      53313
    Taildropped(packets/bytes) :      104082796/24834805544

Policy shape_hqos_bwrr Class MAIN_OUT_CMAP_2
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :      131103199/31266497640      0
  Transmitted                  :      40521792/9647311792      0
  Total Dropped                :      90581407/21619185848      0
  Queueing statistics
    Queue ID                   :      53314
    Taildropped(packets/bytes) :      90581407/21619185848

Policy shape_hqos_bwrr Class MAIN_OUT_CMAP_3
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :      131103636/31266600664      0
  Transmitted                  :      54004138/12863485816      0
  Total Dropped                :      77099498/18403114848      0
  Queueing statistics
    Queue ID                   :      53315
    Taildropped(packets/bytes) :      77099498/18403114848

Policy shape_hqos_bwrr Class MAIN_OUT_CMAP_4
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :      131103920/31266669280      0
  Transmitted                  :      64802769/15436026336      0
  Total Dropped                :      66301151/15830642944      0
  Queueing statistics
    Queue ID                   :      53316
    Taildropped(packets/bytes) :      66301151/15830642944

```

```

Policy shape_hqos_bwrr Class MAIN_OUT_CMAP_5
  Classification statistics      (packets/bytes)      (rate - kbps)
    Matched      :      131104311/31266762056      0
    Transmitted   :      80969302/19295430376      0
    Total Dropped :      50135009/11971331680      0
  Queueing statistics
    Queue ID      :      53317
  Taildropped(packets/bytes)    : 50135009/11971331680

Policy shape_hqos_bwrr Class MAIN_OUT_CMAP_6
  Classification statistics      (packets/bytes)      (rate - kbps)
    Matched      :      131104735/31266863416      0
    Transmitted   :      41698712/9952700928      0
    Total Dropped :      89406023/21314162488      0
  Queueing statistics
    Queue ID      :      53318
  Taildropped(packets/bytes)    : 89406023/21314162488

Policy shape_hqos_bwrr Class MAIN_OUT_CMAP_7
  Classification statistics      (packets/bytes)      (rate - kbps)
    Matched      :      131105067/31266942536      0
    Transmitted   :      83425746/19902839360      0
    Total Dropped :      47679321/11364103176      0
  Queueing statistics
    Queue ID      :      53319
  Taildropped(packets/bytes)    : 47679321/11364103176

Policy shape_hqos_bwrr Class class-default
  Classification statistics      (packets/bytes)      (rate - kbps)
    Matched      :      131102451/31266318184      0
    Transmitted   :      13527585/3215714992      0
    Total Dropped :      117574866/28050603192      0
  Queueing statistics
    Queue ID      :      53312
  Taildropped(packets/bytes)    : 117574866/28050603192
Policy Bag Stats time: 1722238640916 [Local Time: 07/29/24 07:37:20.916]

Interface:PW-Ether1 Member:HundredGigE0/7/0/2 output: MAIN_OUT_PMAP5_Q_TC_MARK

Class MAIN_OUT_MARK_CMAP_1
  Classification statistics      (packets/bytes)      (rate - kbps)
    Matched      :      0/0      0
    Transmitted   :      0/0      0
    Total Dropped :      0/0      0
Class MAIN_OUT_MARK_CMAP_2
  Classification statistics      (packets/bytes)      (rate - kbps)
    Matched      :      0/0      0
    Transmitted   :      0/0      0
    Total Dropped :      0/0      0
Class MAIN_OUT_MARK_CMAP_3
  Classification statistics      (packets/bytes)      (rate - kbps)
    Matched      :      0/0      0
    Transmitted   :      0/0      0
    Total Dropped :      0/0      0
Class MAIN_OUT_MARK_CMAP_4
  Classification statistics      (packets/bytes)      (rate - kbps)
    Matched      :      0/0      0
    Transmitted   :      0/0      0
    Total Dropped :      0/0      0
Class MAIN_OUT_MARK_CMAP_5
  Classification statistics      (packets/bytes)      (rate - kbps)
    Matched      :      0/0      0
    Transmitted   :      0/0      0

```

```

Total Dropped      :                0/0                0
Class MAIN_OUT_MARK_CMAP_6
  Classification statistics
    Matched        :                0/0                0
    Transmitted     :                0/0                0
    Total Dropped   :                0/0                0
Class MAIN_OUT_MARK_CMAP_7
  Classification statistics
    Matched        :                0/0                0
    Transmitted     :                0/0                0
    Total Dropped   :                0/0                0
Class class-default
  Classification statistics
    Matched        :                0/0                0
    Transmitted     :                0/0                0
    Total Dropped   :                0/0                0
Policy Bag Stats time: 1722238674599 [Local Time: 07/29/24 07:37:54.599]

```



CHAPTER 5

Mark Packets to Change Priority Settings

- [Packet Marking Overview](#), on page 61
- [Class-based Unconditional Packet Marking Feature and Benefits](#), on page 69
- [Configure Class-based Unconditional Packet Marking](#), on page 70
- [Mark IP DSCP and MPLS EXP Values for IP to MPLS Traffic](#), on page 71
- [Class-based Unconditional Packet Marking: Examples](#), on page 74
- [IP Precedence Compared to IP DSCP Marking](#), on page 80
- [In-Place Policy Modification](#), on page 80

Packet Marking Overview

Table 10: Feature History Table

Feature Name	Release Information	Feature Description
Packet Marking for IPv4 and IPv6 Traffic	Release 24.4.1	<p>Introduced in this release on: Fixed Systems (8700) (select variants only*)</p> <p>You can use uniform or pipe modes to implement QoS for IPv4 and IPv6 traffic.</p> <p>*This functionality is now extended to the Cisco 8712-MOD-M routers.</p>

Feature Name	Release Information	Feature Description
Packet Marking for IPv4 and IPv6 Traffic	Release 24.3.1	<p>Introduced in this release on: Modular Systems (8800 [LC ASIC: P100]) (select variants only*), Fixed Systems (8200) (select variants only*), Fixed Systems (8700 (P100, K100)) (select variants only*)</p> <p>You can use uniform or pipe modes to implement QoS for IPv4 and IPv6 traffic.</p> <p>*This feature is supported on:</p> <ul style="list-style-type: none"> • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM • 8212-48FH-M • 8711-32FH-M
Packet Marking for IPv4 and IPv6 Traffic	Release 24.2.1	<p>Introduced in this release on: Modular Systems (8800 [LC ASIC: P100])(select variants only*)</p> <p>Using uniform or pipe modes, you can implement QoS for IPv4 and IPv6 traffic. Uniform mode maintains DSCP marking for packets across different network segments and is useful where the network policy requires consistent treatment of packets across multiple domains. Pipe mode separates QoS policies between network segments. In this mode, ingress DSCP markings classify and queue the packets, but the packets can remarked before they egress. Such treatment is beneficial when you need to apply one set of QoS policies within your network and a different set when packets exit the network towards your customer.</p> <p>*This feature is supported on 88-LC1-36EH.</p>

You can use packet marking in input policy maps to set or modify the attributes for traffic belonging to a specific class. For example, you can change the CoS value in a class or set IP DSCP or IP precedence values for a specific type of traffic. These new values are then used to determine how the traffic should be treated.

Default Marking

When an ingress or egress interface adds VLAN tags or MPLS labels, it requires a default value for the class of service and EXP values that go into those tags and labels.

On the router, one ingress default QoS mapping profile and one egress default QoS mapping profile are created and configured per device during initialization.

QoS Behavior for Generic Routing Encapsulation (GRE) Tunnels

Table 11: Feature History Table

Feature Name	Release Information	Feature Description
QoS Behavior for Generic Routing Encapsulation (GRE) Tunnels: Default Marking	Release 7.3.1	With the support for GRE encapsulation and decapsulation tunnel interfaces, there are some important updates to QoS behavior for GRE tunnels. These updates are applicable for default packet marking and involve Type of Service (ToS) and MPLS experimental bits.

GRE Encapsulation

If you do not configure Type of Service (ToS), the outer IP precedence value or the differentiated services code point (DSCP) value is copied from the inner IP header. If you configure ToS, the outer IP precedence value or DSCP value is as per the ToS configuration.

GRE Decapsulation

During decapsulation, the MPLS experimental bits (EXP) are derived from the outer IP packet.

For more information on GRE tunnels, see the *Interfaces Configuration Guide for Cisco 8000 Series Routers, IOS XR Release 7.3.x*.

Packet Marking

Table 12: Feature History Table

Feature Name	Release Information	Feature Description
Set IP Marking for SRv6 Encapsulation	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100]) (select variants only*); Modular Systems (8800 [LC ASIC: P100]) * Set IP Marking for SRv6 Encapsulation feature is now supported on Cisco 8712-MOD-M routers.
Set IP Marking for SRv6 Encapsulation	Release 24.2.11	With this feature support for IP marking for SRv6 packets that are encapsulated, there are some important updates to the QoS behavior. This is an explicit packet marking feature that applies only to ingress QoS policies. CLI: This feature introduces the set ip encapsulation command.

The packet marking feature, also called explicit marking, provides users with a means to differentiate packets based on the designated markings. The router supports ingress and egress packet marking.

The following list provides the usage guidelines for the IP marking for SRv6 encapsulation:

- The **set ip encapsulation class-of-service** command is not allowed in egress QoS policies.
- The described behavior is applicable only in pipe mode. When uniform mode or **propagate** mode is active, regardless of QoS policy encapsulation settings, the class-of-service (CoS) from the inner header is always copied to the outer header.
- The **set qos-group** action cannot be used in conjunction with the set ip encapsulation class-of-service command.
- The IP encapsulation marking (**set ip encapsulation class-of-service**) and the MPLS experimental imposition marking (**set mpls experimental**) features are mutually exclusive. Both the configurations are not allowed at the same time.

Table 13: Feature History Table

Feature Name	Release Information	Feature Description
Set VXLAN Outer IP Header DSCP Value to 0	Release 24.2.11	When a PE device transports IP traffic over a VXLAN tunnel that originates on the device, it automatically sets the DSCP value in the VXLAN outer IP header to 0 (CS0).

Supported Packet Marking Operations

This table shows the supported packet marking operations.

Table 14: Supported Packet Marking Operations for Egress and Ingress

Supported Mark Types	Range	Support for Unconditional Marking	Support for Conditional Marking
set ip encapsulation	0-63	ingress	No
set discard-class	0-1	ingress	No
set dscp	0-63	ingress, egress	No
set mpls experimental topmost	0-7	ingress, egress	No
set precedence	0-7	ingress, egress	No
set qos-group	0-31	ingress	No
set cos	0-7	ingress, egress	No
set dei	0-1	ingress, egress	No

QoS Behavior for Generic Routing Encapsulation (GRE) Tunnels

Table 15: Feature History Table

Feature Name	Release Information	Feature Description
QoS Behavior for Generic Routing Encapsulation (GRE) Tunnels: Explicit Marking	Release 7.3.1	With the support for GRE encapsulation and decapsulation tunnel interfaces, there are some important updates to QoS behavior for GRE tunnels. These updates are applicable for explicit packet marking and involve QoS behavior during ingress and egress.

GRE Encapsulation

During encapsulation of IPv4/IPv6 payload inside the GRE header, QoS behavior is as follows:

- Ingress: QoS supports classification on the payload Layer 3 fields or EXP and remarking payload IP header DSCP.
- Egress: QoS supports setting outer GRE IP header DSCP. It doesn't overwrite the Tunnel Type of Service (ToS) configuration and doesn't remark GRE IP header DSCP.

GRE Decapsulation

During decapsulation of the outer GRE header (during which the inner IPv4/IPv6/MPLS payload is forwarded to the next-hop router), QoS behavior is as follows:

- Ingress: QoS supports classification on Layer 3 fields of outer GRE using the **set qos-group** command. Setting DSCP on the ingress interface sets DSCP for the inner headers.
- Egress: QoS supports classification using **qos-group** to set DSCP or EXP for egress packets.

For more information on GRE tunnels, see the *Interfaces Configuration Guide for Cisco 8000 Series Routers, IOS XR Release 7.3.x*.

Classification and Marking for L2 Traffic

Table 16: Feature History Table

Feature Name	Release Information	Feature Description
Classification and Marking for L2 Traffic	Release 24.4.1	Introduced in this release on: Fixed Systems (8700) (select variants only*) *The ability to perform packet remarking is now supported on Cisco 8712-MOD-M routers.

Classification and Marking for L2 Traffic	Release 7.7.1	You can now perform packet remarking on Layer 2 subinterfaces. Traffic marking allows you to fine tune the traffic attributes on your network. The increased granularity helps single out traffic that requires special handling, and achieves optimal application performance.
---	---------------	---

Traffic marking is a method used to identify certain traffic types for unique handling, effectively partitioning network traffic into different categories. Cisco 8000 Series routers support Layer 2 marking of Ethernet packets in the egress and ingress directions.

Configuration Example for Classification and Marking for L2 Traffic

Follow these steps to enable this feature:

For guidelines on creating a traffic class and egress traffic classification, refer the *Create a Traffic Class* section in the *Classify Packets to Identify Specific Traffic* chapter.

- Enable Layer 2 configuration. Refer the [L2VPN Configuration Guide](#) for details.
- Enable egress remarking with the **qos-group** match type.
- Verify the egress remark QoS policy configuration.
- Enable egress traffic remarking.

Egress Remarking With *qos-group* Match Type

/ Associate a specific qos-group with a traffic class using the **set qos-group** command */*

```
Router# configure terminal
Router(config)# policy-map INGRESS_L2_AC
Router(config-pmap)# class CONTROL_PLANE
Router(config-pmap-c)# set traffic-class 7
Router(config-pmap-c)# set qos-group 27
Router(config-pmap-c)# exit

Router(config-pmap)# class VOIP
Router(config-pmap-c)# set traffic-class 6
Router(config-pmap-c)# set qos-group 26
Router(config-pmap-c)# exit

Router(config-pmap)# class VIDEO_STREAM
Router(config-pmap-c)# set traffic-class 5
Router(config-pmap-c)# set qos-group 25
Router(config-pmap-c)# exit

Router(config-pmap)# class TRANSACTIONAL_DATA
Router(config-pmap-c)# set traffic-class 4
Router(config-pmap-c)# set qos-group 24
Router(config-pmap-c)# exit

Router(config-pmap)# class DB_SYNC
Router(config-pmap-c)# set traffic-class 3
Router(config-pmap-c)# set qos-group 23
Router(config-pmap-c)# exit

Router(config-pmap)# class BULK_DATA
Router(config-pmap-c)# set traffic-class 2
Router(config-pmap-c)# set qos-group 22
Router(config-pmap-c)# exit
```

```

Router(config-pmap)# class SCAVENGER
Router(config-pmap-c)# set traffic-class 1
Router(config-pmap-c)# set qos-group 21
Router(config-pmap-c)# exit

Router(config-pmap)# class class-default
Router(config-pmap-c)# end-policy-map
Router(config)# commit

/* Associate egress policy map with the qos-groups */

Router# configure terminal
Router(config)# policy-map REMARK_EGRESS_QG_L2_AC
Router(config-pmap)# class QoS_GROUP_27
Router(config-pmap-c)# set cos 7
Router(config-pmap-c)# exit

Router(config-pmap)# class QoS_GROUP_26
Router(config-pmap-c)# set cos 7
Router(config-pmap-c)# exit

Router(config-pmap)# class QoS_GROUP_25
Router(config-pmap-c)# set cos 5
Router(config-pmap-c)# exit

Router(config-pmap)# class QoS_GROUP_24
Router(config-pmap-c)# set cos 5
Router(config-pmap-c)# exit

Router(config-pmap)# class QoS_GROUP_23
Router(config-pmap-c)# set cos 3
Router(config-pmap-c)# set dei 0
Router(config-pmap-c)# exit

Router(config-pmap)# class QoS_GROUP_22
Router(config-pmap-c)# set cos 3
Router(config-pmap-c)# set dei 1
Router(config-pmap-c)# exit

Router(config-pmap)# class QoS_GROUP_21
Router(config-pmap-c)# set cos 1
Router(config-pmap-c)# set dei 1
Router(config-pmap-c)# exit

Router(config-pmap)# class class-default
Router(config-pmap-c)# set dei 1
Router(config-pmap-c)# exit
Router(config)# commit

/* Associate Policy-Map REMARK_EGRESS_QG_L2_AC With the Designated Subinterface */

Router(config)# interface bundle-Ether 1.104
Router(config-subif)# service-policy output REMARK_EGRESS_QG_L2_AC
Router(config-subif)# commit

```

Verification

Verify the egress remark QoS policy configuration:

```

Router# show policy-map interface Bundle-Ether 1.104 output

Bundle-Ether1.104 output: REMARK_EGRESS_QG_L2_AC

```

```

Class QoS_GROUP_7
  Classification statistics      (packets/bytes)      (rate - kbps)
    Matched      :      21319554/29847375600      985424
    Transmitted   :      21319554/29847375600      985424
    Total Dropped :      0/0      0
Class QoS_GROUP_6
  Classification statistics      (packets/bytes)      (rate - kbps)
    Matched      :      21581741/30214437400      997540
    Transmitted   :      21581741/30214437400      997540
    Total Dropped :      0/0      0
Class QoS_GROUP_5
  Classification statistics      (packets/bytes)      (rate - kbps)
    Matched      :      106597824/149236953600      4927114
    Transmitted   :      106597824/149236953600      4927114
    Total Dropped :      0/0      0
Class QoS_GROUP_4
  Classification statistics      (packets/bytes)      (rate - kbps)
    Matched      :      490350115/686490161000      22664721
    Transmitted   :      490350115/686490161000      22664721
    Total Dropped :      0/0      0
Class QoS_GROUP_3
  Classification statistics      (packets/bytes)      (rate - kbps)
    Matched      :      0/0      0
    Transmitted   :      0/0      0
    Total Dropped :      0/0      0
Class QoS_GROUP_2
  Classification statistics      (packets/bytes)      (rate - kbps)
    Matched      :      0/0      0
    Transmitted   :      0/0      0
    Total Dropped :      0/0      0
Class QoS_GROUP_1
  Classification statistics      (packets/bytes)      (rate - kbps)
    Matched      :      0/0      0
    Transmitted   :      0/0      0
    Total Dropped :      0/0      0
Class class-default
  Classification statistics      (packets/bytes)      (rate - kbps)
    Matched      :      1087299229/1522218920600      50256530
    Transmitted   :      1087299229/1522218920600      50256530
    Total Dropped :      0/0      0
Policy Bag Stats time: 1657531125148 [Local Time: 07/11/22 09:18:45.148]

```

Remarking of Egress Traffic

```

Router# configure terminal
Router(config)# policy-map REMARK_EGRESS_L2_AC
Router(config-pmap)# class CONTROL_PLANE
Router(config-pmap-c)# set traffic-class 7
Router(config-pmap-c)# exit

Router(config-pmap)# class VOIP
Router(config-pmap-c)# set traffic-class 6
Router(config-pmap-c)# set cos 7
Router(config-pmap-c)# exit

Router(config-pmap)# class VIDEO_STREAM
Router(config-pmap-c)# set traffic-class 5
Router(config-pmap-c)# set cos 5
Router(config-pmap-c)# exit

Router(config-pmap)# class TRANSACTIONAL_DATA
Router(config-pmap-c)# set traffic-class 4

```

```

Router(config-pmap-c) # set cos 5
Router(config-pmap-c) # exit

Router(config-pmap) # class DB_SYNC
Router(config-pmap-c) # set traffic-class 3
Router(config-pmap-c) # set dei 0
Router(config-pmap-c) # exit

Router(config-pmap) # class BULK_DATA
Router(config-pmap-c) # set traffic-class 2
Router(config-pmap-c) # set cos 3
Router(config-pmap-c) # set dei 0
Router(config-pmap-c) # exit

Router(config-pmap) # class SCAVENGER
Router(config-pmap-c) # set traffic-class 1
Router(config-pmap-c) # exit

Router(config-pmap) # class class-default
Router(config-pmap-c) # set dei 1
Router(config-pmap-c) # end-policy-map
Router(config) # commit

```

Class-based Unconditional Packet Marking Feature and Benefits

The packet marking feature allows you to partition your network into multiple priority levels or classes of service, as follows:

- Use QoS unconditional packet marking to set the IP precedence or IP DSCP values for packets entering the network. Routers within your network can then use the newly marked IP precedence values to determine how the traffic should be treated.

On ingress direction, after matching the traffic based on either the IP Precedence or DSCP value, you can set it to a particular discard-class. Weighted random early detection (WRED), a congestion avoidance technique, thereby uses discard-class values to determine the probability that a packet is dropped.

- Use QoS unconditional packet marking to assign MPLS packets to a QoS group. The router uses the QoS group to determine how to prioritize packets for transmission. To set the QoS group identifier on MPLS packets, use the **set qos-group** command in policy map class configuration mode.



Note Setting the QoS group identifier does not automatically prioritize the packets for transmission. You must first configure an egress policy that uses the QoS group.

- Mark Multiprotocol Label Switching (MPLS) packets by setting the EXP bits within the imposed or topmost label.
- Mark packets by setting the value of the *qos-group* argument.
- Mark packets by setting the value of the *discard-class* argument.



Note *qos-group* and *discard-class* are variables internal to the router, and are not transmitted.

The configuration task is described in the [Configure Class-based Unconditional Packet Marking, on page 70](#).

Configure Class-based Unconditional Packet Marking

This configuration task explains how to configure the following class-based, unconditional packet marking features on your router:

- IP precedence value
- IP DSCP value
- IP Encapsulation value
- QoS group value (ingress only)
- CoS value
- MPLS experimental value
- Discard class



Note IPv4 and IPv6 QoS actions applied to MPLS tagged packets are not supported. The configuration is accepted, but no action is taken.

Configuration Example

Follow these steps to configure unconditional packet marking features on your router.

1. Create or modify a policy map that can be attached to one or more interfaces to specify a service policy and enter the policy map configuration mode.
2. Configure an interface and enter the interface configuration mode.
3. Attach a policy map to an input or output interface to be used as the service policy for that interface.

Configuration Example

```
router# configure
router(config)# interface hundredGigE 0/0/0/24
router(config-pmap)# policy-map policy1
Router(config-int)# commit
```

Running Configuration

```
router(config)# policy-map policy1

class-map match-any class1
  match protocol ipv4
end-class-map
!
!
```



```

policy-map policy1
  class class1
    set precedence 1
  !
  class class-default
  !
end-policy-map
!
interface HundredGigE0/0/0/24
  service-policy input policy1
!

```

Verification

Run this command to display policy configuration information for all classes configured for all service policies on the specified interface.

```
router# show run interface hundredGigE 0/0/0/24
```

Mark IP DSCP and MPLS EXP Values for IP to MPLS Traffic

Table 17: Feature History Table

Feature Name	Release Information	Feature Description
Mark IP DSCP and MPLS EXP Values for IP to MPLS Traffic	Release 7.5.4	<p>We've provided you with more granular control for IP to MPLS traffic, so you can collect billing information for your customers upstream at the MPLS core and correlate them with subsequent next-hop IP nodes. This helps you accurately bill your customers as they pay for specific services based on their network requirements.</p> <p>Such control is possible because we've enabled the remarking of IP DSCP values in the inner IP header—that remain even after the POP tag removes the outer header—along with the existing ability to remark MPLS EXP values on the outer MPLS header.</p> <p>There's no action required for you to enable this functionality, and there are no changes to command options.</p>

In earlier releases, when you remarked the DSCP value in the ingress policy for IP to MPLS traffic, here's what happened post the MPLS imposition:

- The inner IP header didn't receive the remarked DSCP value, meaning that the QoS action you required for the IP packet across the network didn't take place.
- The outer header received the remarked value for EXP bits based on your remarking action.

From Cisco IOS XR Release 7.5.4, for IP traffic flowing across an MPLS core, here's the QoS behavior when you remark or set the DSCP value in the ingress policy:

- At the ingress, when the MPLS header is imposed on an IP packet, the inner IP header receives the DSCP value you remarked in the ingress policy. In other words, the IP DSCP value in the inner header is remarked.
- As in earlier releases, the top label in the MPLS header receives the remarked value for EXP bits based on your remarking action. This value for the EXP bits is derived from the three most significant bits of the remarked DSCP value.

Mark IP DSCP and MPLS EXP Values for IP to MPLS Traffic: Benefits

With the ability to remark IP DSCP and MPLS EXP values at the ingress, you have granular control of QoS actions across hops in your network. You can now collect customer-specific billing information upstream at the MPLS core and collate it with the traffic service provided at various IP locations.

Mark IP DSCP and MPLS EXP Values for IP to MPLS Traffic: Guidelines and Limitations

- This functionality is available only on Cisco Silicon One Q200-based fixed and distributed systems.
- You can use a Cisco Silicon One Q100-based line card for distributed systems, but the egress line card must be Cisco Silicon One Q200-based. Else, the feature won't work.

Configure IP DSCP and MPLS EXP Values for IP to MPLS Traffic

Configuration Example

In this example, you set the DSCP value (or assured forwarding per-hop behavior [PHB]) as AF41 for the policy map named **INGRESS_REMARK** at the ingress. At the ingress PE router, the configuration would be:

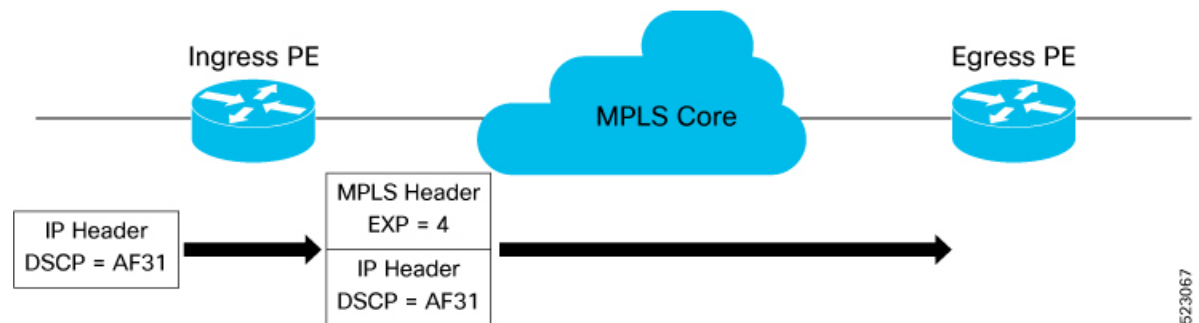
```
Router(config)#class-map AF31
Router(config-cmap)#match dscp AF31
Router(config-cmap)#end-class-map
Router(config)#
Router(config)#policy-map INGRESS_REMARK
Router(config-pmap)#class AF31
Router(config-pmap-c)#set dscp AF41
Router(config-pmap-c)#end-policy-map
Router(config)#
Router(config)#interface FourHundredGigE0/0/0/0
Router(config-if)#service-policy input INGRESS_REMARK
Router(config-if)#commit
```

At the egress PE router, your configuration would be:

```
Router(config)#class-map AF31
Router(config-cmap)#match dscp AF31
Router(config-cmap)#end-class-map
Router(config)#
Router(config)#class-map AF41
Router(config-cmap)#match dscp AF41
Router#end-class-map
Router(config)#
Router(config)#policy-map INGRESS_CLASSIFY
Router(config-pmap)#class AF31
Router(config-pmap-c)#set traffic-class 3
Router(config-pmap-c)#exit
Router(config-pmap)#
Router(config-pmap)#class AF41
Router(config-pmap-c)#set traffic-class 4
Router(config-pmap-c)#exit
Router(config-pmap)#end-policy-map
Router(config)#
Router(config)#int FourHundredGigE0/0/0/1
Router(config-if)#service-policy input INGRESS_CLASSIFY
Router(config-if)#commit
```

Verification

Figure 5: MPLS QoS Behavior Before Cisco IOS XR Release 7.5.4



Before this release, even though you configured the new DSCP value as AF41 for the policy map named **INGRESS_REMARK** at the ingress, the output for the **show policy-map interface** showed that the assured forwarding marking wasn't applied for the PHB. Thus, the statistics counter displayed zero for Class AF41 at the egress PE router.

```
Router#show policy-map interface FourHundredGigE0/0/0/1 input
```

```
FourHundredGigE0/0/0/1 input: INGRESS_CLASSIFY
```

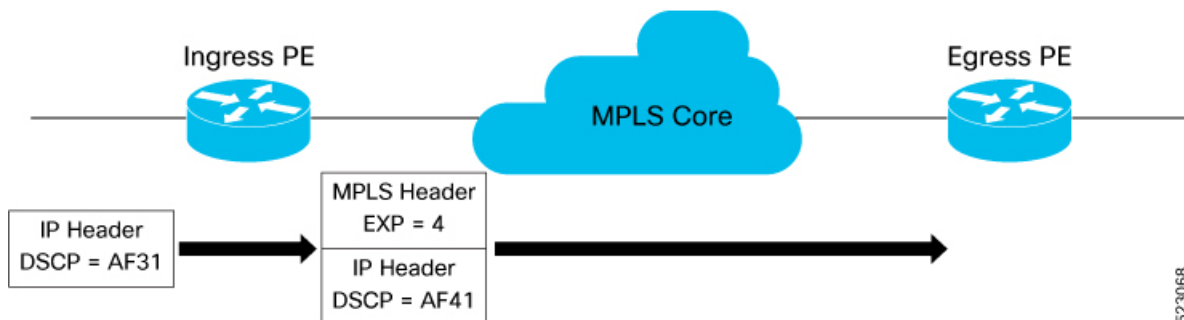
```
Class AF31
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :      9813350/13738690000      936847
  Transmitted                   :      9813350/13738690000      936847
  Total Dropped                 :              0/0              0
Class AF41
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :              0/0              0
  Transmitted                   :              0/0              0
  Total Dropped                 :              0/0              0
Class class-default
```

```

Classification statistics      (packets/bytes)      (rate - kbps)
Matched                      :                    0/0              0
Transmitted                  :                    0/0              0
Total Dropped                :                    0/0              0

```

Figure 6: MPLS QoS Behavior From Cisco IOS XR Release 7.5.4



With the ability to remark IP DSCP values in the inner IP header, the **show policy-map interface** command now displays traffic statistics for Class AF41 at the egress PE router. These statistics confirm that the requisite QoS action—which is, DSCP remark—has taken place at the ingress PE router.

```
Router#show policy-map interface FourHundredGigE0/0/0/1 input
```

```
FourHundredGigE0/0/0/1 input: INGRESS_CLASSIFY
```

```

Class AF31
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :                    0/0              0
  Transmitted                  :                    0/0              0
  Total Dropped                :                    0/0              0
Class AF41
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :          9813350/13738690000      936847
  Transmitted                  :          9813350/13738690000      936847
  Total Dropped                :                    0/0              0
Class class-default
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      :                    0/0              0
  Transmitted                  :                    0/0              0
  Total Dropped                :                    0/0              0

```

Class-based Unconditional Packet Marking: Examples

These are typical examples for class-based unconditional packet marking.

IP Precedence Marking Configuration: Example

In this example, a service policy called *policy1* is created. This service policy is associated to a previously defined class map called *class1* through the use of the **class** command, and then the service policy is attached to the output HundredGigE interface 0/7/0/1. The IP precedence bit in the ToS byte is set to 1:

```

policy-map policy1
  class class1
    set precedence 1

```

```
!
interface HundredGigE 0/7/0/1
  service-policy output policy1
```

IP DSCP Marking Configuration: Example

In this example, a service policy called `policy1` is created. This service policy is associated to a previously defined class map through the use of the **class** command. In this example, it is assumed that a class map called `class1` was previously configured and new class map called `class2` is created.

In this example, the IP DSCP value in the ToS byte is set to 5:

```
policy-map policy1
  class class1
    set dscp 5

  class class2
    set dscp ef
```

After you configure the settings shown for voice packets at the edge, all intermediate routers are configured to provide low-latency treatment to the voice packets, as follows:

```
class-map voice
  match dscp ef
policy-map qos-policy
  class voice
    priority level 1
    police rate percent 10
```

IP Encapsulation Marking Configuration: Example

In this example, a service policy called *policy1* is created. This service policy is associated to a class map called *class1* by using the **class** command. The service policy is attached in the input direction on the interface HundredGigE 0/0/0/24.

```
class-map match-any class1
  match protocol ipv4
  match access-group ipv4 10
  class class-default
end-class-map

policy-map policy1
  class class1
    set ip encapsulation class-of-service 55
  class class-default
end-policy-map
!
interface interface HundredGigE 0/0/0/24
  service-policy input policy1
```

Guidelines and Limitations:

- The **set ip encapsulation class-of-service** command is not allowed in egress QoS policies.
- The described behavior is applicable only in pipe mode. When in uniform mode or when propagate mode is active, regardless of QoS policy encapsulation settings, the class-of-service (CoS) from the inner header is always copied to the outer header.

- The **set qos-group** action cannot be used in conjunction with the **set ip encapsulation class-of-service** command.
- The IP encapsulation marking (**set ip encapsulation class-of-service**) and the MPLS experimental imposition marking (**set mpls experimental**) features are mutually exclusive. Both the configurations are not allowed at the same time.

QoS Group Marking Configuration: Example

In this example, a service policy called *policy1* is created. This service policy is associated to a class map called *class1* through the use of the **class** command, and then the service policy is attached in the input direction on a HundredGigE 0/7/0/1. The qos-group value is set to 1.

```
class-map match-any class1
  match protocol ipv4
  match access-group ipv4 101

policy-map policy1
  class class1
    set qos-group 1
  !
interface HundredGigE 0/7/0/1
  service-policy input policy1
```



Note The **set qos-group** command is supported only on an ingress policy.

CoS Marking Configuration: Example

In this example, a service policy called *policy1* is created. This service policy is associated to a class map called *class1* through the use of the **class** command, and then the service policy is attached in the output direction on a HundredGigE 0/7/0/1.100. The IEEE 802.1p (CoS) bits in the Layer 2 header are set to 1.

```
class-map match-any class1
  match protocol ipv4
  match access-group ipv4 101

policy-map policy1
  class class1
    set cos 1
  !
interface HundredGigE 0/7/0/1.100
  service-policy output policy1
```

MPLS Experimental Bit Imposition Marking Configuration: Example

In this example, a service policy called *policy1* is created. This service policy is associated to a class map called *class1* through the use of the **class** command, and then the service policy is attached in the input direction on a HundredGigE 0/7/0/1. The MPLS EXP bits of all imposed labels are set to 1.

```

class-map match-any class1
  match protocol ipv4
  match access-group ipv4 101

policy-map policy1
  class class1
    set mpls exp imposition 1
  !
interface HundredGigE 0/7/0/1
  service-policy input policy1

```



Note The `set mpls exp imposition` command is supported only on an ingress policy.

MPLS Experimental Topmost Marking Configuration: Example

In this example, a service policy called *policy1* is created. This service policy is associated to a class map called *class1* through the use of the **class** command, and then the service policy is attached in the output direction on a HundredGigE 0/7/0/1. The MPLS EXP bits on the TOPMOST label are set to 1:

```

class-map match-any class1
  match mpls exp topmost 2

policy-map policy1
  class class1
    set mpls exp topmost 1
  !
interface HundredGigE 0/7/0/1
  service-policy output policy1

```

Queuing and Marking Policies: Examples

These examples explain how queuing and marking can be applied to an interface.

Ingress Traffic Classification and QoS Mapping Configuration for Input Interfaces

In this example, a policy map named "INGRESS_L3" is configured, specifically for classifying and marking incoming traffic.

```
RP/0/RP0/CPU0:Router#show policy-map pmap-name INGRESS_L3 detail
```

```

class-map match-all CONTROL_PLANE
match dscp cs7
end-class-map
!
class-map match-all VOIP
match dscp cs6
end-class-map
!
class-map match-all VIDEO_STREAM
match dscp cs5
end-class-map
!
class-map match-all TRANSACTIONAL_DATA
match dscp cs4
end-class-map

```

```

!
class-map match-all DB_SYNC
match dscp cs3
end-class-map
!
class-map match-all BULK_DATA
match dscp cs2
end-class-map
!
class-map match-all SCAVENGER
match dscp cs1
end-class-map
!
policy-map INGRESS_L3
class CONTROL_PLANE
    set traffic-class 7
    set qos-group 27
!
class VOIP
    set traffic-class 6
    set qos-group 26
!
class VIDEO_STREAM
    set traffic-class 5
    set qos-group 25
!
class TRANSACTIONAL_DATA
    set traffic-class 4
    set qos-group 24
!
class DB_SYNC
    set traffic-class 3
    set qos-group 23
!
class BULK_DATA
    set traffic-class 2
    set qos-group 22
!
class SCAVENGER
    set traffic-class 1
    set qos-group 21
!
class class-default
!
end-policy-map
!

```

Egress Queueing Policy Configuration

In this example, a policy map named "EGRESS_QOS_QUEUEING" is configured for the egress (outgoing) direction.

```
RP/0/RP0/CPU0:Router#show policy-map pmap-name EGRESS_QOS_QUEUEING detail
```

```

class-map match-any TC7_CONTROL_PLANE
match traffic-class 7
end-class-map
!
class-map match-any TC6_VOIP
match traffic-class 6
end-class-map
!
class-map match-any TC5_VIDEO_STREAM
match traffic-class 5

```



```

    end-class-map
    !
    class-map match-any TC4_TRANSACTIONAL_DATA
    match traffic-class 4
    end-class-map
    !
    class-map match-any TC3_DB_SYNC
    match traffic-class 3
    end-class-map
    !
    class-map match-any TC2_BULK_DATA
    match traffic-class 2
    end-class-map
    !
    class-map match-any TC1_SCAVENGER
    match traffic-class 1
    end-class-map
    !
    policy-map EGRESS_QOS_QUEUEING
    class TC7_CONTROL_PLANE
    priority level 1
    !
    class TC6_VOIP
    priority level 2
    !
    class TC5_VIDEO_STREAM
    bandwidth remaining ratio 20
    random-detect 50 mbytes 100 mbytes
    !
    class TC4_TRANSACTIONAL_DATA
    bandwidth remaining ratio 20
    queue-limit 75 mbytes
    !
    class TC3_DB_SYNC
    bandwidth remaining ratio 10
    !
    class TC2_BULK_DATA
    bandwidth remaining ratio 10
    !
    class TC1_SCAVENGER
    bandwidth remaining ratio 5
    !
    class class-default
    !
    end-policy-map
    !

```

Egress Marking Policy Configuration

In this example, a policy map named "EGRESS_REMARK" is configured for the egress (outgoing) direction.

```
RP/0/RP0/CPU0:Router#show policy-map pmap-name EGRESS_REMARK detail
```

```

class-map match-any QG4
match qos-group 4
end-class-map
!
class-map match-any QG3
match qos-group 3
end-class-map
!
class-map match-any QG2
match qos-group 2
end-class-map

```

```

!
policy-map EGRESS_REMARK
class QG4
  set dscp af42
!
class QG3
  set dscp af33
!
class QG2
  set dscp af21
!
class class-default
!
end-policy-map
!

```

Output Interface Configuration with Queueing Policy and Remark Policy Applied

```

interface FourHundredGigE0/2/0/29
mtu 10000
service-policy output EGRESS_REMARK
service-policy output EGRESS_QOS_QUEUEING
ipv4 address 18.29.0.1 255.255.255.0
ipv6 address 18:1d::1/96
!

```

IP Precedence Compared to IP DSCP Marking

If you need to mark packets in your network and all your devices support IP DSCP marking, use the IP DSCP marking to mark your packets because the IP DSCP markings provide more unconditional packet marking options. If marking by IP DSCP is undesirable, however, or if you are unsure if the devices in your network support IP DSCP values, use the IP precedence value to mark your packets. The IP precedence value is likely to be supported by all devices in the network.

You can set up to 8 different IP precedence markings and 64 different IP DSCP markings.

Configure DSCP CS7 (Precedence 7)

See the following example to configure options in DSCP for a particular source address in IPv4 packets.

Configuration Example

```

policy-map policy1
class class1
  set dscp cs7
!

```

In-Place Policy Modification

The In-Place policy modification feature allows you to modify a QoS policy even when the QoS policy is attached to one or more interfaces. A modified policy is subjected to the same checks that a new policy is subject to when it is bound to an interface. If the policy-modification is successful, the modified policy takes

effect on all the interfaces to which the policy is attached. However, if the policy modification fails on any one of the interfaces, an automatic rollback is initiated to ensure that the pre-modification policy is in effect on all the interfaces.

You can also modify any class map used in the policy map. The changes made to the class map take effect on all the interfaces to which the policy is attached.

**Note**

- The QoS statistics for the policy that is attached to an interface are lost (reset to 0) when the policy is modified.
- When a QoS policy attached to an interface is modified, there might not be any policy in effect on the interfaces in which the modified policy is used for a short period of time.
- An in-place modification of an ACL does not reset the policy-map statistics counter.

Verification

If unrecoverable errors occur during in-place policy modification, the policy is put into an inconsistent state on target interfaces. No new configuration is possible until the configuration session is unblocked. It is recommended to remove the policy from the interface, check the modified policy and then re-apply accordingly.

Recommendations for Using In-Place Policy Modification

For a short period of time while a QoS policy is being modified, there might not be any policy in effect on the interfaces in which the modified policy is used. For this reason, modify QoS policies that affect the fewest number of interfaces at a time. Use the **show policy-map targets** command to identify the number of interfaces that will be affected during policy map modification.



CHAPTER 6

Congestion Avoidance

- [Congestion Avoidance](#), on page 83
- [Queuing Modes](#), on page 85
- [Congestion Avoidance in VOQ](#), on page 87
- [Virtual Output Queue Watchdog](#), on page 93
- [Equitable Traffic Flow Using Fair VOQ](#), on page 96
- [Modular QoS Congestion Avoidance](#) , on page 102
- [Tail Drop and the FIFO Queue](#), on page 102
- [Random Early Detection and TCP](#), on page 104
- [Explicit Congestion Notification](#) , on page 106
- [Traffic Class Queue High Water Marks Monitoring](#), on page 110

Congestion Avoidance

Table 18: Feature History Table

Feature Name	Release Information	Feature Description
Queueing for Congestion Avoidance	Release 24.4.1	<p>Introduced in this release on: Fixed Systems (8700) (select variants only*)</p> <p>You can shape traffic to control the traffic flow from queues and also configure queues to ensure certain traffic classes get a guaranteed amount of bandwidth.</p> <p>*This functionality is now supported on Cisco 8712-MOD-M routers.</p>

Feature Name	Release Information	Feature Description
Queueing for Congestion Avoidance	Release 24.3.1	<p>Introduced in this release on: Modular Systems (8800 [LC ASIC: P100]) (select variants only*), Fixed Systems (8200) (select variants only*), Fixed Systems (8700 (P100, K100)) (select variants only*)</p> <p>You can shape traffic to control the traffic flow from queues and also configure queues to ensure certain traffic classes get a guaranteed amount of bandwidth.</p> <p>*This feature is supported on:</p> <ul style="list-style-type: none"> • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM • 8212-48FH-M • 8711-32FH-M
Queueing for Congestion Avoidance	Release 24.2.1	<p>Introduced in this release on: Modular Systems (8800 [LC ASIC: P100])(select variants only*)</p> <p>By placing packets in different queues based on priority, queueing helps prevent traffic congestion and ensures that high-priority traffic is transmitted with minimal delay. You can shape traffic to control the traffic flow from queues and also configure queues to ensure certain traffic classes get a guaranteed amount of bandwidth. Queueing provides buffers to temporarily store packets during bursts of traffic and also supports strategies that enable dropping lower-priority packets when congestion builds up.</p> <p>*This feature is supported on 88-LC1-36EH.</p>

Queueing provides a way to temporarily store data when the received rate of data is larger than what can be sent. Managing the queues and buffers is the primary goal of congestion avoidance. As a queue starts to fill up with data, it is important to try to make sure that the available memory in the ASIC/NPU does not fill up completely. If this happens, subsequent packets coming into the port are dropped, irrespective of the priority that they received. This could have a detrimental effect on the performance of critical applications. For this reason, congestion avoidance techniques are used to reduce the risk of a queue from filling up the memory

completely and starving non-congested queues for memory. Queue thresholds are used to trigger a drop when certain levels of occupancy are exceeded.

Scheduling is the QoS mechanism that is used to empty the queues of data and send the data onward to its destination.

Shaping is the act of buffering traffic within a port or queue until it is able to be scheduled. Shaping smoothens traffic, making traffic flows much more predictable. It helps ensure that each transmit queue is limited to a maximum rate of traffic.

Queuing Modes

Two network queuing modes are supported for network interface queuing: the default mode of 8xVOQ (virtual output queuing) and 4xVOQ. To change the mode from one to another, configure the **hw-module profile qos voq-mode** command and reload all line cards in the system.

In the 8xVOQ mode, eight VoQs and their associated resources are allocated for each interface. These queues are allocated regardless of the exact policy configuration on that interface. This mode supports a separate VOQ for each of the eight internal traffic classes.

In the 4xVOQ mode, four VoQs and their associated resources are allocated to each interface, and these queues are allocated regardless of the exact policy applied. In this mode the system supports twice the number of logical interfaces, but the eight traffic classes must be mapped down by configuration to four VoQs, not eight VoQs.



Note From Cisco IOS XR Release 7.2.12 onwards, all the queuing features that are supported on Layer 3 interfaces are also supported on Layer 2 interfaces. However, these features apply only to the main interface (physical and bundle interfaces), and not on the sub-interfaces.

Main Interface Queuing Policy

The main interface default queues are created as part of the main interface creation.

When you apply a queuing policy to the main interface, it will override the default queuing and scheduling parameters for the traffic classes you have configured.

In the 8xVOQ mode, a P1+P2+6PN hierarchy is used for the main interface queues (default queuing and scheduling). The default queues are used for all traffic to the main interface and traffic to any sub-interface without a queuing policy applied. The control/protocol traffic uses traffic class 7 (TC7), priority 1 (P1) to avoid drops during congestion.

Subinterface Queueing Policy

Table 19: Feature History Table

Feature Name	Release Information	Feature Description
Subinterface Queueing Policy	Release 24.4.1	<p>Introduced in this release on: Fixed Systems (8700) (select variants only*)</p> <p>You can manage traffic flows with fine granularity by configuring QoS policies at the subinterface level.</p> <p>*This functionality is now supported on Cisco 8712-MOD-M routers.</p>
Subinterface Queueing Policy	Release 24.3.1	<p>Introduced in this release on: Modular Systems (8800 [LC ASIC: P100]) (select variants only*), Fixed Systems (8200) (select variants only*), Fixed Systems (8700 (P100, K100)) (select variants only*)</p> <p>You can manage traffic flows with fine granularity by configuring QoS policies at the subinterface level.</p> <p>*This feature is supported on:</p> <ul style="list-style-type: none"> • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM • 8212-48FH-M • 8711-32FH-M

Feature Name	Release Information	Feature Description
Subinterface Queueing Policy	Release 24.2.1	<p>Introduced in this release on: Modular Systems (8800 [LC ASIC: P100])(select variants only*)</p> <p>To manage traffic flows with fine granularity, you can configure QoS policies at the subinterface level. Also, with QoS support on mixed Layer 2 and Layer 3 subinterfaces under the same main interface, you can ensure that multiple traffic types with varying QoS requirements coexist on a single main physical interface.</p> <p>*This feature is supported on 88-LC1-36EH.</p>

Each subinterface supports up to three policies: an ingress policy, an egress marking policy, and an egress queuing policy. To create and configure a separate set of VoQs for a subinterface, apply a queuing policy on that subinterface. When you remove the subinterface queuing policy, the associated VoQs are freed and the subinterface traffic reverts to using the main interface VoQs.

Congestion Avoidance in VOQ

Table 20: Feature History Table

Feature Name	Release Information	Feature Description
View VOQs Evicted to HBM	Release 24.2.11	<p>The newly introduced command displays the virtual output queues (VOQs) that are evicted to the High Bandwidth Memory (HBM) and the VOQs' HBM buffer usage details. You can use this information whilst monitoring and debugging congestion scenarios.</p> <p>This feature introduces the show controllers npu voq in-extended-memory instance command.</p> <p>This feature modifies the Cisco-IOS-XR-8000-platforms-npu-evict-voq-buff-oper.yang (see GitHub, YANG Data Models Navigator) data model.</p>

Congestion avoidance within a VOQ block is done by applying a congestion management profile to a VOQ. This profile defines the admission criteria and checks performed at the enqueue time. Under normal traffic conditions the packet is enqueued into the Shared Memory System (SMS) buffers. (The shared memory system is the primary packet storage area.) If the SMS VOQ is congested beyond a set threshold, the VOQ is moved

to the external High Band Memory (HBM) block. When the HBM queue drains, it is returned to the on-chip SMS. The queue size in HBM is adaptive and decreases when the total HBM usage is high.



Note Random Early Detect (RED) is available only for VOQs in HBM. The hardware does not support Weighted Random Early Detect (WRED).

View VOQs Evicted to HBM

Important Caveat

Do not use the **show controllers npu voq in-extended-memory instance** command in an automation script.

Guidelines and Limitations to View VOQs Evicted to HBM

- If the PFC buffer-extended mode is enabled on a device, the associated VOQs are evicted to the HBM on priority. The remaining VOQs are retained in the Shared Memory Switch (SMS).
- If PFC buffer-internal mode is enabled on a device, the associated VOQs are retained in the SMS. The remaining VOQs are evicted to the HBM.
- If PFC isn't enabled on a device, VOQs are evicted to the HBM based on the VOQs' age and buffer usage.
- With this feature, you can view up to 4000 records for these interfaces:
 - 400G and 100G (PFC buffer-internal and PFC buffer-extended modes)
 - 40G (PFC buffer-internal mode)

Verification

Output pointers:

- **Egress Interface**—The egress interface of the virtual output queue.
- **VOQ_Base**—Base VOQ ID
- **TC**—Traffic Class number
- **Slice**—Source slice number
- **Buff_Usage and In_Bytes**—Buffer usage in blocks and in bytes

Example 1: The following output displays VOQs that are evicted to the HBM for node 0/6/cpu0 and all instances. In this case, VOQs from device instance **Device 1** are evicted to the HBM.

```
Router# show controllers npu voq in-extended-memory instance all location 0/6/CPU0
```

* Use this CLI with caution.

* This should not be integrated with any automation scripts.

```
Total Entries  0  | Slot 6  | Device  0  | Percent in Evict Voq Buff(s) 0.000000
```

```
-----
Egress Interface | VOQ_Base | TC  | Slice | Buff_Usage | In_Bytes |
```

```
Total Entries  2  | Slot 6  | Device  1  | Percent in Evict Voq Buff(s) 0.004883
```

```
-----
Egress Interface | VOQ_Base | TC  | Slice | Buff_Usage | In_Bytes |
```

```

-----
FH0/6/0/0      28672      2      0      14722      87837459
FH0/6/0/3      28728      2      0      1049       6335369

```

```
Total Entries  0 | Slot 6 | Device 2 | Percent in Evict Voq Buff(s) 0.000000
```

```
-----
Egress Interface | VOQ_Base | TC | Slice | Buff_Usage | In_Bytes |
```

Example 2: The following output displays VOQs that are evicted to the HBM for node 0/6/cpu0 and all instances. In this case, the virtual output queue from device instance **Device 0** is evicted to the HBM.

```
Router# show controllers npu voq in-extended-memory instance all location 0/6/CPU0
```

```
* Use this CLI with caution.
```

```
* This should not be integrated with any automation scripts.
```

```
Total Entries  1 | Slot 6 | Device 0 | Percent in Evict Voq Buff(s) 0.002441
```

```
-----
Egress Interface | VOQ_Base | TC | Slice | Buff_Usage | In_Bytes |
```

```
-----
FH0/6/0/22      28776      6      1      7140      46969050

```

```
Total Entries  0 | Slot 6 | Device 1 | Percent in Evict Voq Buff(s) 0.000000
```

```
-----
Egress Interface | VOQ_Base | TC | Slice | Buff_Usage | In_Bytes |
```

```
Total Entries  0 | Slot 6 | Device 2 | Percent in Evict Voq Buff(s) 0.000000
```

```
-----
Egress Interface | VOQ_Base | TC | Slice | Buff_Usage | In_Bytes |
```

Sharing of VOQ Statistics Counters

Every network processor on the router has multiple slices (or pipelines), and every slice has a set of VOQs associated with every interface on the router. To maintain counters at high packet rates, two sets of counters are associated with each interface on each network slice. As an example, consider a device with six slices (12 interfaces), each with 24,000 VOQs, where you want both transmitted and dropped events counted. In this scenario, you would require $12 \times 24,000 \times 2 = 5,76,000$ counters, which alone exceeds the counter capacity of the device.

It is to mitigate such a scenario that the router supports configurable sharing of VOQ counters. You can configure the sharing such that a counter is shared by {1,2,4,8} VOQs.

Each set of VoQs sharing counters has two counters that measure:

- Enqueued packets count in packets and bytes units.
- Dropped packets count in packets and bytes units.

For the feature to take effect:

- Delete egress queuing policy-map configuration from all interfaces.
- Run the command **# reload location all** to reload all the nodes on your router.

Configuring Sharing of VOQ Statistics Counters

To configure VOQs sharing counters, use the `#hw-module profile stats voqs-sharing-counters` and specify the number of VOQ counters for each queue.

```
RP/0/RP0/CPU0:ios(config)#hw-module profile stats ?
  voqs-sharing-counters  Configure number of voqs (1, 2, 4) sharing counters
RP/0/RP0/CPU0:ios(config)#hw-module profile stats voqs-sharing-counters ?
  1  Counter for each queue
  2  2 Queues share counters
  4  4 Queues share counters
RP/0/RP0/CPU0:ios(config)#hw-module profile stats voqs-sharing-counters 1
RP/0/RP0/CPU0:ios(config)#hw-module profile stats voqs-sharing-counters 2
RP/0/RP0/CPU0:ios(config)#commit
RP/0/RP0/CPU0:ios#reload location all
```

Running Configuration

```
RP/0/RP0/CPU0:ios#show run | in hw-mod
Mon Feb 10 13:57:35.296 UTC
Building configuration...
hw-module profile stats voqs-sharing-counters 2
RP/0/RP0/CPU0:ios#
```

Verification

```
RP/0/RP0/CPU0:ios#show controllers npu stats voq ingress interface hundredGigE 0/0/0/16
instance all location 0/RP0/CPU0
Mon Feb 10 13:58:26.661 UTC
```

```
Interface Name      = Hu0/0/0/16
Interface Handle    = f0001b0
Location            = 0/RP0/CPU0
Asic Instance       = 0
VOQ Base            = 10288
Port Speed(kbps)    = 100000000
Local Port          = local
VOQ Mode            = 8
Shared Counter Mode = 2
  ReceivedPkts    ReceivedBytes    DroppedPkts    DroppedBytes
-----
TC_{0,1} = 114023724      39908275541      113945980      39881093000
TC_{2,3} = 194969733      68239406550      196612981      68814543350
TC_{4,5} = 139949276      69388697075      139811376      67907466750
TC_{6,7} = 194988538      68242491778      196612926      68814524100
```

Related Commands `hw-module profile stats voqs-sharing-counters`

Dual Queue Limit

The dual queue limit option is added to **queue-limit** command on the CLI of your router and displays as **discard-class**. What the **discard-class** option does is give you the flexibility to configure two queue limits on a single policy map—one for high-priority traffic and the other for low-priority traffic. This option ensures that the high priority traffic flow continues unaffected (up to the derived threshold from the **discard-class 0** queue-limit) while the low-priority traffic continues up to the lower threshold (per **discard-class 1** queue-limit).

Tell Me More

You can configure the two queue limits per these details:

- One for flow that you mark as **discard-class 0** (higher priority) on ingress via ingress-policy.
- second, for flow that you mark as **discard-class 1** (lower priority) on ingress via ingress policy.

The **discard-class 1** flow (for low-priority traffic) begins to drop when the queue length hits the size limit that you configured for discard-class 1. Conversely, the flow for **discard-class 1** stops dropping when queue-length falls below its configured value.

As an example, consider this configuration:

```
policy-map egress_pol_dql
class tc7
  queue-limit discard-class 0 100 mbytes
  queue-limit discard-class 1 50 mbytes
  priority level 1
!
class class-default
  bandwidth remaining ratio 1
!
end-policy-map
!
```

Also consider the verification:

```
RP/0/RP0/CPU0:ios#
RP/0/RP0/CPU0:ios#show qos interface hundredGigE 0/0/0/30 output
NOTE:- Configured values are displayed within parentheses
Interface HundredGigE0/0/0/30 ifh 0xf000210 -- output policy
NPU Id: 0
Total number of classes: 2
Interface Bandwidth: 100000000 kbps
Policy Name: egress_pol_dql
VOQ Base: 464
Accounting Type: Layer1 (Include Layer 1 encapsulation and above)
VOQ Mode: 8
Shared Counter Mode: 1
-----
Level1 Class (HP1) = tc7
Egressq Queue ID = 471 (HP1 queue)
Queue Max. BW. = no max (default)
Discard Class 1 Threshold = 25165824 bytes / 2 ms (50 mbytes)
Discard Class 0 Threshold = 75497472 bytes / 5 ms (100 mbytes)
WRED not configured for this class

Level1 Class = class-default
Egressq Queue ID = 464 (Default LP queue)
Queue Max. BW. = no max (default)
Inverse Weight / Weight = 1 / (1)
TailDrop Threshold = 749568 bytes / 6 ms (default)
WRED not configured for this class
```

In the preceding example, there are two traffic flows that are marked as **discard-class 0** (higher priority) and **discard-class 1** (lower priority).

As long as the queue length of the two flows remains below 25165824 bytes (the threshold for **discard-class 1**), packets from both flows continue without any drops. When the queue length reaches 25165824 bytes, **discard-class 1** packets are not enqueued, ensuring all remaining bandwidth is used for the higher priority flow (**discard-class 0**).

The higher priority flow drops only when the queue length reaches 75497472 bytes.

**Note**

- This option protects the high-priority traffic from loss due to congestion, but not necessarily from latency due to congestion.
- These thresholds are derived from hardware-specific queue regions.

Restrictions

Ensure that you read these restrictions about the dual queue limit option.

- Both the queue-limits must use the same unit of measurement.
- The queue limit for **discard-class 0** must always be greater than that for **discard-class 1**.
- When the discard-class option is not used to configure the queue-limit, packets marked with **discard-class 0** and **discard-class 1** have the same queue-limit; in other words, they receive identical treatment.
- A queue-limit that is configured with only **discard-class 0** or **discard-class 1** is rejected.

Virtual Output Queue Watchdog

Table 21: Feature History Table

Feature Name	Release Information	Feature Description
Virtual Output Queue Watchdog	Release 24.2.11	<p>We ensure the continuous movement of traffic queues, which is crucial for enforcing QoS policies, even when hardware issues disrupt the Virtual Output Queue (VOQ) and impede the flow of traffic. With this feature, if the router detects a stuck queue on a line card, it shuts down the line card, and if it detects a stuck queue on a fabric card, the router triggers a hard reset on the NPU. A queue is considered stuck only when there is no transmission for one minute.</p> <p>The feature is disabled by default and can be enabled using the command hw-module voq-watchdog feature enable.</p> <p>The feature is supported only on Cisco 8000 Series Routers (Modular) with Cisco Silicon One Q100 or Q200 ASICs.</p> <p>The feature introduces these changes:</p> <p>CLI:</p> <ul style="list-style-type: none"> hw-module voq-watchdog feature enable hw-module voq-watchdog cardshut disable

Virtual Output Queue (VOQ) is a mechanism to manage a situation where a packet at the front of the queue is unable to move forward due to a busy destination, causing a delay for all packets behind it, even those with available output ports. VOQ prevents packet delay or loss by providing each output port with a dedicated queue for every input port. This mechanism ensures that packets are transmitted in the correct order and maintains quality of service (QoS) by managing congestion.

A VOQ can become stuck when it fails to transmit traffic for a certain period despite the non-empty queue. There are a few scenarios in which this can happen. For example, when a port or traffic class (Output Queue, OQ) is PFC-paused (Priority Flow Control) by the peer, all the VOQs sending traffic to this OQ are also paused. Another scenario is when a port's higher priority traffic class (TC) consumes all the port bandwidth,

causing lower priority TCs to run out of credits and their corresponding VOQs to become stuck. This can cause VOQs to stop functioning, potentially leading to traffic loss.

VOQ watchdog actively identifies and resolves issues in VOQs that have not sent packets for one minute. It detects stuck VOQs in both line cards (LCs) and fabric cards (FCs). The feature is disabled by default and can be enabled using the command **hw-module voq-watchdog feature enable**.

VOQ Watchdog Behavior on Line Cards

By default, the VOQ watchdog feature is disabled on the line cards; as a result, stuck VOQs are not detected, and there are no interrupts or syslog notifications.

However, after you enable this feature using the **hw-module voq-watchdog feature enable** command, the router regularly checks the line cards for packets stuck in VOQs. If the router detects any such packets, it will raise a notification and shut down the line card.

The router displays the following messages when it detects a stuck VOQ.

Syslog Notification - Stuck VOQ; Action: Line card Shutdown

```
LC/0/0/CPU0:Feb 22 09:16:56.090 UTC: npu_drvr[203]: %FABRIC-NPU_DRV-3-VOQ_HARDWARE_WATCHDOG
:
[7127] : npu[1]: hardware_watchdog.voq_error: VOQ slc:2 voqnum:19955 isinhbm:1 smscntxtnum:3
hbmcntxtnum:14
isstuck:1 nochangesec:64 rdptr:1728 wrptr:1735 avblcrdts:-16668 is_fabric:0
```

Stuck VOQ; Action: Line card Shutdown

```
LC/0/0/CPU0:Jan 30 15:10:57.299 UTC: npu_drvr[241]: %PKT_INFRA-FM-2-FAULT_CRITICAL :
ALARM_CRITICAL :VOQ WATCHDOG Alarm :DECLARE :: Shutdown card due to voq watchdog error on
ASIC 1.
```

Line cards automatically shut down if stuck VOQs are detected, making it impossible to identify the root cause of the problem. However, you can prevent the line cards from shutting down by using the **hw-module voq-watchdog cardshut disable** command to disable the shutdown function. This way, the router will send syslog notifications when it detects stuck VOQs without shutting down the line card.

After disabling the shutdown action on the line card, the router displays the following messages when it detects a stuck VOQ.

Syslog Notification - Stuck VOQ; Action: None

```
LC/0/0/CPU0:Feb 22 09:16:56.090 UTC: npu_drvr[203]: %FABRIC-NPU_DRV-3-VOQ_HARDWARE_WATCHDOG
:
[7127] : npu[1]: hardware_watchdog.voq_error: VOQ slc:2 voqnum:19955 isinhbm:1 smscntxtnum:3
hbmcntxtnum:14
isstuck:1 nochangesec:64 rdptr:1728 wrptr:1735 avblcrdts:-16668 is_fabric:0
```

Stuck VOQ; Action: None

```
LC/0/0/CPU0:Feb 22 09:16:56.090 UTC: npu_drvr[203]: %FABRIC-NPU_DRV-3-VOQ_HARDWARE_WATCHDOG
:
[7127] : npu[1]: hardware_watchdog.voq_error: VOQ Watchdog Action Handling Skipped Due to
User Configuration
```

VOQ Watchdog Behavior on Fabric Cards

By default, the VOQ watchdog feature is disabled on the fabric cards; as a result, stuck VOQs are not detected, and there are no interrupts or syslog notifications.

After you enable the feature using the **hw-module voq-watchdog feature enable** command, the router regularly checks the fabric cards for any packets stuck in VOQs. If such packets are detected, the router raises a syslog notification and hard resets the fabric element (FE) device. After five hard resets, the fabric card undergoes a graceful reload.

The router logs the following syslog notification upon detecting a stuck VOQ on an FC.

```
Syslog: RP/0/RP0/CPU0:Feb 22 09:16:47.721 UTC: npu_drvr[335]:
%FABRIC-NPU_DRV-3-ASIC_ERROR_ACTION :
[7912] : npu[6]: HARD_RESET needed for hardware_watchdog.voq_error
```

Guidelines and Restrictions for Virtual Output Queue Watchdog

- You can enable or disable the feature on both line and fabric cards. But, you can't enable or disable the feature on just one card, either line or fabric card.
- Disabling the shutdown action for line cards has no effect on fabric cards.
- You can't disable the shutdown action on fabric cards (FCs).
- The feature is supported only in Cisco 8000 Series Routers with Cisco Silicon One Q200 ASICs.

Configure Virtual Output Queue Watchdog

This section describes how to configure VOQ Watchdog.

Configuration Example

To enable the feature on your router:

```
Router# config
Router(config)# hw-module voq-watchdog feature enable
Router(config)# commit
```

To disable the shutdown action on the line cards:

```
Router# config
Router(config)# hw-module voq-watchdog cardshut disable
Router(config)# commit
```

Running Configuration

```
Router# show running-config
...
hw-module voq-watchdog feature enable
hw-module voq-watchdog cardshut disable
end
```

Verification

After the feature is enabled, the following syslog message is logged on detecting a stuck VOQ on the line card:

```
LC/0/0/CPU0:Feb 22 09:16:56.090 UTC: npu_drvr[203]:
%FABRIC-NPU_DRV-3-VOQ_HARDWARE_WATCHDOG: [7127] : npu[1]: hardware_watchdog.voq_error:
VOQ slc:2 voqnum:19955 isinhbm:1 smscntxtnum:3 hbmcntxtnum:14 isstuck:1 nochangesec:64
rdptr:1728 wrptr:1735 avblcrdts:-16668 is_fabric:0
```

On disabling the shutdown action of line cards, the following syslog message is logged on detecting a stuck VOQ on a line card:

```
LC/0/0/CPU0:Feb 22 09:16:56.090 UTC: npu_drvr[203]:
%FABRIC-NPU_DRV-3-VOQ_HARDWARE_WATCHDOG : [7127] : npu[1]:
hardware_watchdog.voq_error: VOQ Watchdog Action Handling Skipped Due to User Configuration
```

After the feature is enabled, the following syslog message is logged on detecting a stuck VOQ on the fabric card:

```
Syslog: RP/0/RP0/CPU0:Feb 22 09:16:47.721 UTC: npu_drvr[335]:
%FABRIC-NPU_DRV-3-ASIC_ERROR_ACTION :
[7912] : npu[6]: HARD_RESET needed for hardware_watchdog.voq_error
```

Equitable Traffic Flow Using Fair VOQ

Table 22: Feature History Table


Feature Name	Release Information	Feature Description
Equitable Traffic Flow Using Fair VOQ	Release 7.3.3	<p>Configuring this feature ensures that ingress traffic from various source ports on every network slice of an NPU is assigned a unique virtual output queue (VOQ) for every source port and destination port pair. This action ensures that the bandwidth available at the destination port for a given traffic class is distributed equally to all source ports requesting bandwidth.</p> <p>In earlier releases, the traffic wasn't distributed equitably because each slice wasn't given its fair share of the output queue bandwidth.</p> <p>This feature introduces the fair-4 and fair-8 keywords in the hw-module profile qos voq-mode command.</p>
Equitable Traffic Flow Using Fair VOQ on Cisco 8201-32FH Routers	Release 7.5.2	<p>You can now ensure that the bandwidth available at the destination port for a given traffic class is distributed equally to all source ports requesting bandwidth on Cisco 8201-32FH routers.</p>

Fair VOQ: Why

Per default behavior, every network slice of an NPU is assigned a set of 4 or 8 Virtual Output Queues (VOQ) per destination port. With such an assignment, it's challenging to ensure that the right amount of buffering is available via VOQs. With this configuration, the ingress traffic from various source ports on a slice (or pipeline) on an NPU destined to a destination port is assigned to a VOQ per slice. In other words, multiple source ports sending traffic to the same destination port use the same VOQ. However, when sending traffic to different destination ports, traffic is enqueued to different VOQs. This means that the traffic isn't distributed equitably because each slice doesn't get its fair share of the output queue bandwidth. In a scenario where one slice has two ports and another slice has only one port, the bandwidth drops for the ports sharing a slice, even though the two ports handle more traffic than the single port.

Consider the following example where two 100G ports—port-0 and port-1—that belong to the same slice (slice-0) are sending traffic to port-3 on the output queue (OQ). You have a 100G port on another slice (slice-1) on the same NPU that is also scheduled to send traffic to port-3. The ingress VOQ is shared between the two ports in slice-0, whereas the ingress VOQ in slice-1 is available exclusively for port-3. This arrangement results in port-0 and port-1 getting 25% of the buffer traffic, while port-3 receives 50% of the buffer traffic.

Figure 7: Existing behavior : Source ports on slice share one VOQ per destination port



The fair VOQ feature solves this disparity in traffic distribution.

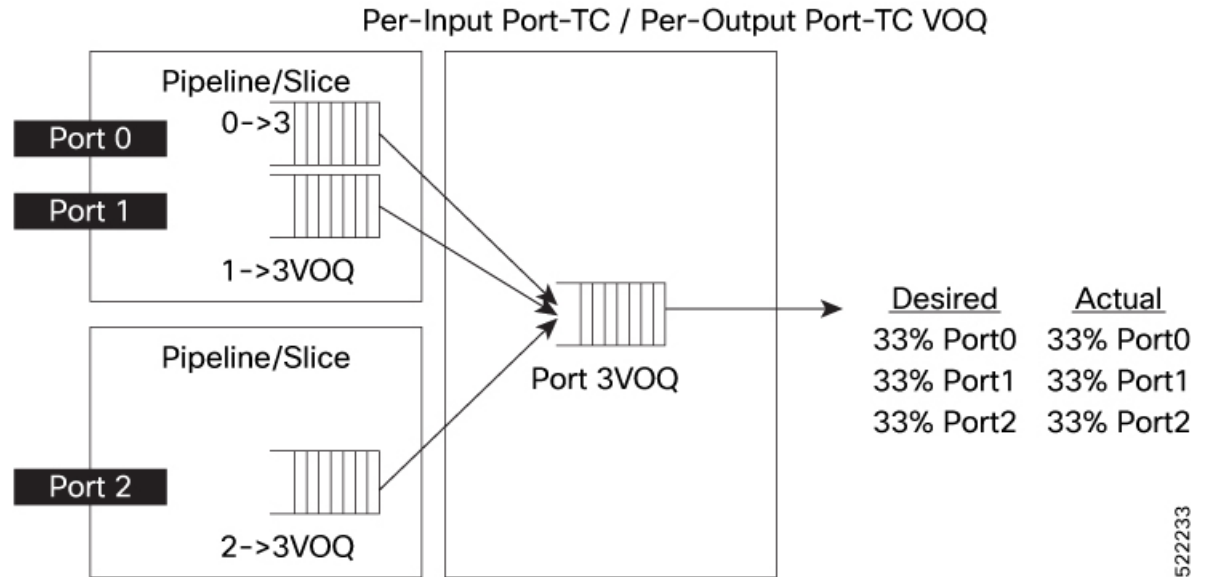
Fair VOQ: How

The fair VOQ feature tackles the default behavior that treats source ports on each NPU slice equally, regardless of the number of active source ports. It does so by redesigning the way bandwidth is allocated from the output queue. Rather than distributing bandwidth at a slice level, fair VOQ distributes bandwidth directly to the source ports. When you configure the command **hw-module profile qos voq-mode** and reload your router, the functionality creates a dedicated VOQ for every source port and destination port pair. This arrangement ensures that bandwidth available at the destination port for a given traffic class is distributed equally to all source ports requesting bandwidth.

Extending the preceding example to understand the fair VOQ functionality, there are now dedicated VOQs for each ingress port that connect to the port on the output queue. Thus, port-0 and port-1 now don't share a

VOQ, and port-3 has its VOQ as before, as shown in the following figure. This fair VOQ arrangement results in traffic queued on dedicated queues, thus improving the traffic performance.

Figure 8: Fair VOQ behavior: each source port on slice has one dedicated VOQ per destination port



Fair VOQ Modes and Sharing of Counters

You can configure fair VOQ for 8xVOQ mode (fair-8) and 4xVOQ mode (fair-4) using the following options in the **hw-module profile qos voq-mode** command:

- **hw-module profile qos voq-mode fair-8**
- **hw-module profile qos voq-mode fair-4**

You can also share VOQ statistics counters in both fair VOQ modes, as shown in the following table. (For details on why sharing counters is essential and how to configure counters sharing, see [Sharing of VOQ Statistics Counters](#), on page 89.)

Table 23: Fair VOQ Modes and Sharing Counters

Fair VOQ Mode	Sharing Counters Mode	Important Notes
fair-8	2, 4	<ul style="list-style-type: none"> Eight VOQs configured per source port and destination pair Counters are shared by {2, 4} VOQs. fair-8 mode doesn't support dedicated counter mode (counter mode1, where there's a counter for every queue)

Fair VOQ Mode	Sharing Counters Mode	Important Notes
fair-4	1, 2, 4	<ul style="list-style-type: none"> Four VOQs configured per source port and destination pair Counters are shared by {1, 2, 4} VOQs.

Fair VOQs and Slice (or Normal) VOQs: Key Differences

The following table is a snapshot to outline the key differences between fair VOQs and slice or regular VOQs.

Table 24: Fair VOQs and Normal VOQs

Fair VOQ	Normal VOQ
fair-8 mode: eight VOQs configured per source port and destination pair	8: <ul style="list-style-type: none"> Eight VOQs per destination port per slice These VOQs are shared by all source ports within an NPU slice.
fair-4 mode: four VOQs configured per source port and destination pair	4: <ul style="list-style-type: none"> Four VOQs per destination port per slice These VOQs are shared by all source ports within an NPU slice.

Guidelines and Limitations

- The fair VOQ feature is supported on the Cisco 8202 router (12 QSFP56-DD 400G and 60 QSFP28 100G ports) and Cisco 8201-32FH router (32 QSFP-DD 400G ports).
- The following table details the maximum interfaces (with basic IPv4 configurations and no other scale configuration such as QoS policy, ACL, and subinterface configuration) allowed based on VOQ mode and sharing counter mode.

Table 25: Maximum Interfaces Based on Fair VOQ Mode and Sharing Counter Mode

VOQ Mode	Sharing Counter Mode	Maximum Interfaces for Cisco 8202	Maximum Interfaces for Cisco 8201-32FH
fair-8	1	The router doesn't support this combination.	The router doesn't support this combination.

VOQ Mode	Sharing Counter Mode	Maximum Interfaces for Cisco 8202	Maximum Interfaces for Cisco 8201-32FH
fair-8	2	96 = 60 (100G) + 8x4 + 4 (400G) ==> you can configure only eight 400G interfaces in 4x10G or 4x25G breakout mode.	128 = 32x4 (4x10G or 4x25G - breakout on all 32 ports - 400G)
fair-8	4	108 = 60 + 12 x 4 (breakout on all 12 ports - 400G)	
fair-4	1	96 = 60(100G) + 8x4 + 4 (400G) ==> you can configure only eight 400G interfaces in 4x10G or 4x25G breakout mode.	
fair-4	2	108 = 60 + 12 x4 (breakout on all 12 ports - 400G)	
fair-4	4	108 = 60 + 12 x4 (breakout on all 12 ports - 400G)	



Note We recommend using sharing counter mode 4 in breakout modes and sharing counter mode 2 for nonbreakout modes.

If the counter usage of applications is high, sharing counter mode 2 is not supported.



Note Breakout mode isn't supported on 100G interfaces.

- Ensure that you reload the router for the configuration to take effect.
- Layer 2 traffic isn't supported in **fair-voq** mode (**fair-4** and **fair-8**).
- Subinterface queueing isn't supported. (This applies to bundle sub-interfaces as well). This means that you can't attach egress service-policies that require dedicated VOQs. However, egress marking is supported for subinterfaces.
- ERSPAN mirroring traffic is not supported in **fair-voq** mode.
- **hw-module profile stats voqs-sharing-counters 1** isn't supported in **fair-8** mode. Ensure that you configure **hw-module profile voq sharing-counters 2** or **hw-module profile voq sharing-counters 4**

along with **hw-module profile qos voq-mode fair-4** or **hw-module profile qos voq-mode fair-8** before reloading the router.

- Breakout is supported only on 400G interfaces in **fair-voq** mode (both **fair-4** and **fair-8**) on the Cisco 8202 router.
- **src-interface** and **src-slice** keywords in **show controller npu stats** are visible only when you configure the VOQ mode to either **fair-8** or **fair-4**.

Configure Fair VOQ

To configure fair VOQ:

1. Configure sharing of VOQ statistics counters. This example configures 2 counters.



Note Configuring **fair-8** mode without counter sharing may cause configuration failure or other unexpected behavior.

2. Configure fair VOQ mode. This example shows how to configure **fair-8** mode.
3. Restart the router for the configuration to take effect.
4. You have successfully enabled the fair VOQ feature to ensure equitable traffic distribution between every source port and destination port pair.

```
/*Configure sharing of VOQ statistics counters; we're configuring 2 counters per queue*/
Router(config)#hw-module profile stats ?
    voqs-sharing-counters  Configure number of voqs (1, 2, 4) sharing counters
Router(config)#hw-module profile stats voqs-sharing-counters ?
    1  Counter for each queue
    2  2 Queues share counters
    4  4 Queues share counters
Router(config)#hw-module profile stats voqs-sharing-counters 2

/*Configure fair-voq mode; we're configuring fair-8 VOQ mode here*/
Router#config
Router(config)#hw-module profile qos voq-mode fair-8
Router(config)#commit
Router#reload location all
```

Running Configuration

```
hw-module profile stats voqs-sharing-counters 2
!
hw-module profile qos voq-mode fair-8
!
```

Verification

Run the **show controller npu stats voq ingress interface <> instance <> location <>** command to verify the fair VOQ configuration.

```
Router#show controllers npu stats voq ingress interface hundredGigE 0/0/0/20 instance 0
location 0/RP0/CPU0
```

```
Interface Name      =  Hu0/0/0/20
```

```

Interface Handle      =      f000118
Location              =      0/RP0/CPU0
Asic Instance         =      0
Port Speed(kbps)      =      100000000
Local Port            =      local
Src Interface Name    =      ALL
VOQ Mode             =      Fair-8
Shared Counter Mode   =      2

```

	ReceivedPkts	ReceivedBytes	DroppedPkts	DroppedBytes
TC_{0,1} = 11110		1422080	0	0
TC_{2,3} = 0		0	0	0
TC_{4,5} = 0		0	0	0
TC_{6,7} = 0		0	0	0

```

RP/0/RP0/CPU0:ios#

```

Associated Commands

[hw-module profile qos voq-mode](#)

Modular QoS Congestion Avoidance

Congestion avoidance techniques monitor traffic flow in an effort to anticipate and avoid congestion at common network bottlenecks. Avoidance techniques are implemented before congestion occurs as compared with congestion management techniques that control congestion after it has occurred.

Congestion avoidance is achieved through packet dropping. The router supports these QoS congestion avoidance techniques:

- [Tail Drop and the FIFO Queue, on page 102](#)
- [Random Early Detection and TCP, on page 104](#)

Tail Drop and the FIFO Queue

Tail drop is a congestion avoidance technique that drops packets when an output queue is full until congestion is eliminated. Tail drop treats all traffic flow equally and does not differentiate between classes of service. It manages the packets placed into a first-in, first-out (FIFO) queue, and forwarded at a rate determined by the available underlying link bandwidth.

Configure Tail Drop

Packets satisfying the match criteria for a class accumulate in the queue reserved for the class until they are serviced. The **queue-limit** command is used to define the maximum threshold for a class. When the maximum threshold is reached, the enqueued packets to the class queue result in tail drop (packet drop).

Restrictions

- When configuring the **queue-limit** command, you must configure one of the following commands: **priority**, **shape average**, or **bandwidth remaining**, except for the default class.

Configuration Example

You have to accomplish the following to complete the tail drop configuration:

1. Creating (or modifying) a policy map that can be attached to one or more interfaces to specify a service policy
2. Associating the traffic class with the traffic policy
3. Specifying the maximum limit the queue can hold for a class policy configured in a policy map.
4. Specifying priority to a class of traffic belonging to a policy map.
5. (Optional) Specifying the bandwidth allocated for a class belonging to a policy map or specifying how to allocate leftover bandwidth to various classes.
6. Attaching a policy map to an output interface to be used as the service policy for that interface.

```
Router# configure
Router(config)# policy-map test-qlimit-1
Router(config-pmap)# class qos-1
Router(config-pmap-c)# queue-limit 100 us
Router(config-pmap-c)# priority level 7
Router(config-pmap-c)# exit
Router(config-pmap)# exit

Router(config)# interface HundredGigE 0/6/0/18
Router(config-if)# service-policy output test-qlimit-1
Router(config-if)# commit
```

Running Configuration

```
policy-map test-qlimit-1
  class qos-1
    queue-limit 100 us
    priority level 7
  !
  class class-default
  !
end-policy-map
!
```

Verification

```
Router# show qos int hundredGigE 0/6/0/18 output
```

```
NOTE:- Configured values are displayed within parentheses
Interface HundredGigE0/6/0/18 ifh 0x3000220 -- output policy
NPU Id: 3
Total number of classes: 2
Interface Bandwidth: 100000000 kbps
VOQ Base: 11176
VOQ Stats Handle: 0x88550ea0
Accounting Type: Layer1 (Include Layer 1 encapsulation and above)
-----
Level1 Class (HP7) = qos-1
Egress Queue ID = 11177 (HP7 queue)
```

```

TailDrop Threshold                = 1253376 bytes / 100 us (100 us)
WRED not configured for this class

Levell Class                      = class-default
Egressq Queue ID                 = 11176 (Default LP queue)
Queue Max. BW.                  = 101803495 kbps (default)
Queue Min. BW.                  = 0 kbps (default)
Inverse Weight / Weight          = 1 (BWR not configured)
TailDrop Threshold               = 1253376 bytes / 10 ms (default)
WRED not configured for this class

```

Related Topics

- [Tail Drop and the FIFO Queue, on page 102](#)

Random Early Detection and TCP

The Random Early Detection (RED) congestion avoidance technique takes advantage of the congestion control mechanism of TCP. By randomly dropping packets prior to periods of high congestion, RED tells the packet source to decrease its transmission rate. Assuming the packet source is using TCP, it decreases its transmission rate until all packets reach their destination, indicating that the congestion is cleared. You can use RED as a way to cause TCP to slow transmission of packets. TCP not only pauses, but it also restarts quickly and adapts its transmission rate to the rate that the network can support.

RED distributes losses in time and maintains normally low queue depth while absorbing traffic bursts. It achieves this by taking action on the average queue size, and not the instantaneous queue size. When enabled on an interface, RED begins dropping packets when congestion occurs at a rate you select during configuration.

Configure Random Early Detection

The **random-detect** command with the minimum threshold and maximum threshold keywords must be used to enable random early detection (RED).

Guidelines

- If you configure the **random-detect** **<min threshold>** **<max threshold>** command on any class including class-default, configure one of the following commands: **shape average** or **bandwidth remaining**.
- If you configure a queue-limit that is lesser than the minimum supported value, the configured value automatically adjusts to the supported minimum value.

While configuring **random-detect**, if you set the **<min threshold>** and **<max-threshold>** values lesser than the minimum supported threshold value:

- The **<min threshold>** value automatically adjusts to the minimum supported value.
- The **<max-threshold>** value doesn't autoadjust to a value above the minimum supported threshold value. This results in a failed **random-detect** configuration. To prevent this error, configure the **<max-threshold>** value such that it exceeds the **<min threshold>** value that your system supports.

Configuration Example

Accomplish the following to complete the random early detection configuration:

1. Creating (or modifying) a policy map that can be attached to one or more interfaces to specify a service policy
2. Associating the traffic class with the traffic policy
3. Enabling RED with minimum and maximum thresholds.
4. Configure **one** of the following:
 - Specifying how to allocate leftover bandwidth to various classes. **OR**
 - Shaping traffic to the specified bit rate or a percentage of the available bandwidth.
5. Attaching a policy map to an output interface to be used as the service policy for that interface.

```
Router# configure
Router(config)# policy-map red-abs-policy
Router(config-pmap)# class qos-1
Router(config-pmap-c)# random-detect <min threshold> <max threshold>
Router(config-pmap-c)# shape average percent 10
Router(config-pmap-c)# end-policy-map
Router(config)# commit
Router(config)# interface HundredGigE0/0/0/12
Router(config-if)# service-policy output red-abs-policy
Router(config-if)# commit
```

Running Configuration

```
policy-map red-abs-policy
class tc7
  priority level 1
  queue-limit 75 mbytes
!
class tc6
  priority level 2
  queue-limit 75 mbytes
!
class tc5
  shape average 10 gbps
  queue-limit 75 mbytes
!
class tc4
  shape average 10 gbps
  queue-limit 75 mbytes
!
class tc3
  shape average 10 gbps
  queue-limit 75 mbytes
!
class tc2
  shape average 10 gbps
  queue-limit 75 mbytes
!
class tc1
  shape average 10 gbps
```

```

random-detect ecn
random-detect 100 mbytes 200 mbytes
!
class class-default
  shape average 10 gbps
  random-detect 100 mbytes 200 mbytes
!
end-policy-map
!

interface HundredGigE0/0/0/12
service-policy output red-abs-policy
shutdown
!
```

Verification

Router# **show qos int hundredGigE 0/6/0/18 output**

```

NOTE:- Configured values are displayed within parentheses
Interface HundredGigE0/0/0/12 ifh 0x3000220  -- output policy
NPU Id:                                     3
Total number of classes:                   2
Interface Bandwidth:                       100000000 kbps
VOQ Base:                                  11176
VOQ Stats Handle:                          0x88550ea0
Accounting Type:                           Layer1 (Include Layer 1 encapsulation and above)
-----
Level1 Class                               = qos-1
Egressq Queue ID                           = 11177 (LP queue)
Queue Max. BW.                             = 10082461 kbps (10 %)
Queue Min. BW.                             = 0 kbps (default)
Inverse Weight / Weight                    = 1 (BWR not configured)
Guaranteed service rate                    = 10000000 kbps
TailDrop Threshold                         = 12517376 bytes / 10 ms (default)

Default RED profile
RED Min. Threshold                         = 12517376 bytes (10 ms)
RED Max. Threshold                         = 12517376 bytes (10 ms)

Level1 Class                               = class-default
Egressq Queue ID                           = 11176 (Default LP queue)
Queue Max. BW.                             = 101803495 kbps (default)
Queue Min. BW.                             = 0 kbps (default)
Inverse Weight / Weight                    = 1 (BWR not configured)
Guaranteed service rate                    = 50000000 kbps
TailDrop Threshold                         = 62652416 bytes / 10 ms (default)
WRED not configured for this class
```

Related Topics

- [Random Early Detection and TCP, on page 104](#)

Explicit Congestion Notification

Random Early Detection (RED) is implemented at the core routers of a network. Edge routers assign IP precedences to packets, as the packets enter the network. With RED, core routers then use these precedences

to determine how to treat different types of traffic. RED provides a single threshold and weights per traffic class or queue for different IP precedences.

ECN is an extension to RED. ECN marks packets instead of dropping them when the average queue length exceeds a specific threshold value. When configured, ECN helps routers and end hosts to understand that the network is congested and slow down sending packets. However, if the queue length is above the maximum threshold for the extended memory, packets are dropped. This is the identical treatment that a packet receives when RED is enabled without ECN configured on the router.

RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*, states that with the addition of active queue management (for example, RED) to the Internet infrastructure, routers are no longer limited to packet loss as an indication of congestion.



Note You cannot use this feature when you have set qos-group or mpls experimental along with a traffic class in the ingress policy.

Implementing ECN

Implementing ECN requires an ECN-specific field that has two bits—the ECN-capable Transport (ECT) bit and the CE (Congestion Experienced) bit—in the IP header. The ECT bit and the CE bit can be used to make four code points of 00 to 11. The first number is the ECT bit and the second number is the CE bit.

Table 26: ECN Bit Setting

ECT Bit	CE Bit	Combination Indicates
0	0	Not-ECN-capable.
0	1	Endpoints of the transport protocol are ECN-capable.
1	0	Endpoints of the transport protocol are ECN-capable.
1	1	Congestion experienced.

The ECN field combination 00 indicates that a packet is not using ECN. The code points 01 and 10—Called ECT(1) and ECT(0), respectively—are set by the data sender to indicate that the endpoints of the transport protocol are ECN-capable. Routers treat these two code points identically. Data senders can use either one or both of these two combinations. The ECN field combination 11 indicates congestion to the endpoints. Packets arriving a full queue of a router will be dropped.

Packet Handling When ECN Is Enabled

When ECN is enabled, all packets between <min_threshold> and <max tail drop threshold> are marked with ECN. Three different scenarios arise if the queue length is between the minimum threshold and the maximum threshold:

- If the ECN field on the packet indicates that the endpoints are ECN-capable (that is, the ECT bit is set to 1 and the CE bit is set to 0, or the ECT bit is set to 0 and the CE bit is set to 1)—and the RED algorithm determines that the packet should have been dropped based on the drop probability—the ECT and CE

bits for the packet are changed to 1, and the packet is transmitted. This happens because ECN is enabled and the packet gets marked instead of dropped.

- If the ECN field on the packet indicates that neither endpoint is ECN-capable (that is, the ECT bit is set to 0 and the CE bit is set to 0), the packet is transmitted. If, however, the max tail drop threshold is exceeded, the packet is dropped. This is the identical treatment that a packet receives when RED is enabled without ECN configured on the router.
- If the ECN field on the packet indicates that the network is experiencing congestion (that is, both the ECT bit and the CE bit are set to 1), the packet is transmitted. No further marking is required.

Configuration Example

```
Router# configure
Router(config)# policy-map policy1
Router(config-pmap)# class class1
Router(config-pmap-c)# bandwidth percent 50
Router(config-pmap-c)# random-detect 1000 packets 2000 packets
Router(config-pmap-c)# random-detect ecn
Router(config-pmap-c)# exit
Router(config-pmap)# exit
Router(config)# commit
```

Verification

Use the **show policy-map interface** to verify the configuration.

```
Router# show policy-map int hu 0/0/0/35 output
TenGigE0/0/0/6 output: pm-out-queue

HundredGigE0/0/0/35 output: egress_qosgrp_ecn

Class tc7
  Classification statistics          (packets/bytes)      (rate - kbps)
    Matched                        :      195987503/200691203072      0
    Transmitted                    :      188830570/193362503680      0
    Total Dropped                  :       7156933/7328699392      0
  Queueing statistics
    Queue ID                       :      18183
    Taildropped(packets/bytes)     :      7156933/7328699392

    WRED profile for
      RED Transmitted (packets/bytes) : N/A
      RED random drops(packets/bytes) : N/A
      RED maxthreshold drops(packets/bytes) : N/A
      RED ecn marked & transmitted(packets/bytes): 188696802/193225525248

Class tc6
  Classification statistics          (packets/bytes)      (rate - kbps)
    Matched                        :      666803815/133360763000      0
    Transmitted                    :      642172362/128434472400      0
    Total Dropped                  :      24631453/4926290600      0
  Queueing statistics
    Queue ID                       :      18182
    Taildropped(packets/bytes)     :      24631453/4926290600

    WRED profile for
      RED Transmitted (packets/bytes) : N/A
      RED random drops(packets/bytes) : N/A
      RED maxthreshold drops(packets/bytes) : N/A
      RED ecn marked & transmitted(packets/bytes): 641807908/128361581600
```

```

Class tc5
Classification statistics          (packets/bytes)      (rate - kbps)
  Matched                        :      413636363/82727272600      6138
  Transmitted                    :      398742312/79748462400      5903
  Total Dropped                  :      14894051/2978810200       235
Queueing statistics
  Queue ID                      :      18181
  Taildropped (packets/bytes)    :      14894051/2978810200

WRED profile for
RED Transmitted (packets/bytes)   : N/A
RED random drops (packets/bytes) : N/A
RED maxthreshold drops (packets/bytes) : N/A
RED ecn marked & transmitted(packets/bytes): 398377929/79675585800

```

**Note**

The **RED ecn marked & transmitted(packets/bytes)** row displays the statistics for ECN marked packets. To begin with, it displays *0/0*.

Traffic Class Queue High Water Marks Monitoring

Table 27: Feature History Table

Feature Name	Release Information	Feature Description
Traffic Class Queue High Water Marks Monitoring	Release 24.2.11	<p><i>Introduced in this release on Cisco 8000 Series Routers with Cisco Silicon One Q200 network processors. The Cisco 8608 router is not currently supported.</i></p> <p>This feature monitors egress interface traffic class queues and records the queue occupancy and queue delay high water marks information for each traffic class. This information includes the virtual output queue that experienced the high water mark and a timestamp indicating when the high water mark was recorded.</p> <p>You can use this data to identify network bottlenecks and prevent traffic congestion.</p> <p>This feature introduces these changes:</p> <p>Configuration CLI:</p> <ul style="list-style-type: none"> • hw-module profile qos high-water-marks <p>EXEC commands:</p> <ul style="list-style-type: none"> • show controllers npu qos high-water-marks • clear controller npu qos high-water-marks <p>YANG Data Models:</p> <ul style="list-style-type: none"> • cisco-IOS-XR-ofa-npu-qos-oper.yang • cisco-IOS-XR-ofa-npu-qos-act.yang • cisco-IOS-XR-um-8000-hw-module-profile-cfg.yang • cisco-IOS-XR-npu-hw-profile-cfg.yang

Queue occupancy and queue delay are used to measure the amount of congestion experienced by an egress interface's traffic class queue. The high water mark values for a traffic class's queue occupancy and delay can be monitored to help identify potential network bottlenecks.

You can view the following high water marks data for each traffic class on an egress interface after enabling the feature:

Queue Occupancy High Water Marks

- Maximum queue occupancy observed on the traffic class, in kilobytes and as a percentage of the queue size
- Queue delay at the time the occupancy high water mark was detected, in nanoseconds

- Information indicating the virtual output queue that experienced the occupancy high water mark
- Timestamp indicating when the queue occupancy high water mark was recorded

Queue Delay High Water Marks

- Maximum queue delay observed on the traffic class, in nanoseconds
- Queue occupancy at the time the delay high water mark was detected, in kilobytes and as a percentage of the queue size
- Information indicating the virtual output queue that experienced the delay high water mark
- Timestamp indicating when the queue delay high water mark was recorded

Limitations and Guidelines for Traffic Class Queue High Water Marks Monitoring

- The Traffic Class Queue High Water Marks Monitoring and Priority Flow Control (PFC) features can co-exist. In cases where PFC buffer-internal mode is configured, the optional **ecn** and **maximum-threshold** parameters must be specified.

Example:

```
Router# hw-module profile priority-flow-control location 0/0/CPU0
Router(config-pfc-loc)# buffer-internal traffic-class 3 pause-threshold 1574400 bytes
headroom 1651200 bytes ecn 629760 bytes max-threshold 1416960 bytes probability-percentage
50
Router(config-pfc-loc)# commit
```

- When a bundle interface is configured, high water marks are monitored on each bundle member interface rather than the bundle interface.
- Packets that egress a sub-interface or bundle sub-interface contribute to the high water marks on the corresponding parent interface
- Guidelines for a router that includes a mix of line cards with Cisco Silicon One Q200 network processors and other network processors:
 - Queue occupancy monitoring isn't supported for packets with a source interface on a non-Q200 network processor. So, these packets do not contribute to the egress interface's queue occupancy high water marks
 - Queue delay monitoring is supported for packets with a source interface on a non-Q200 network processor. So, these packets contribute to the egress interface's queue delay high water marks
- Several virtual output queues (VOQs) are mapped to each egress traffic class queue. The specific number of VOQs depends on the number of slices present on the router and the configured Virtual Output Queueing (VOQ) mode. The high water marks indicate the maximum queue occupancy and queue delay experienced on any of the VOQs mapped to the egress traffic class queue
- Perform a manual reload of the chassis or all line cards after enabling or disabling this feature
- The queue occupancy high water mark percentage is calculated based on the maximum supported queue-limit and does not consider any customized queue-limit configured by the end user

Configure Traffic Class Queue High Water Marks Monitoring

```
Router(config)# hw-module profile qos high-water-marks
```

A manual reload of the chassis or all line cards is required to enable/disable Traffic Class High Water Marks Monitoring

```
Router(config)# commit
```

Verification

Use the **show controllers npu qos high-water-marks** command after reloading to verify that the feature is enabled.

The command output will be displayed when the feature is enabled. A message similar to the following will be displayed if the feature is not enabled:

```
RP/0/RP0/CPU0:ios#show controllers npu qos high-water-marks interface all
Mon Jun  3 06:02:50.138 UTC
Not supported or not enabled on location 0/0/CPU0
RP/0/RP0/CPU0:ios#
```

Field/Output Description

Table 28: Common fields

Field	Description
Interval Start and End (periodic only)	The periodic collection interval number, and the start and end time of the interval.
TC_Number = (<i>Number</i> range is 0-7)	Indicates the traffic class for the high water marks data displayed on that line. For periodic output, TC_ <i>Number</i> is only displayed for the traffic class's first periodic interval.

Table 29: Queue Occupancy fields

Field	Description
Max Occupancy %	The maximum queue occupancy experienced by this traffic class as a percentage of the total queue size. Due to limited queue quantization thresholds provided by the NPU, the max occupancy percentage and max occupancy kilobytes value below are an estimate of the actual maximum queue occupancy.
Max Occupancy kilobytes	The maximum queue occupancy experienced by this traffic class in kilobytes. The kilobytes value is calculated with the assumption that all buffers are fully packed (i.e., all 384 bytes utilized in an SMS buffer). As a result, the displayed kilobytes value will be higher than the actual number of kilobytes queued in many cases.
Queue Delay ns	The delay in nanoseconds at the time the maximum queue occupancy high water mark occurred.

Field	Description
Src Sys Port Slot/NPU/Slc/Gid	<p>The Slot, NPU, Slice and GID identify the virtual output queue on which the queue occupancy high water mark occurred. The GID is the global identifier of the source system port whose packet was dequeued when the maximum queue occupancy high water mark occurred.</p> <p>In most cases, all ports on each slice share a virtual output queue. Although the identified source system port sent the packet that was detected as the high water mark, other ports using the same virtual output queue may have also contributed to the burst of packets that caused the queue occupancy high water mark.</p> <p>Note When the fair-4 or fair-8 VOQ mode is configured, each source port has its own virtual output queue. When PFC buffer-internal mode is configured, each port shares a virtual output queue with other ports on the same slice interface group (IFG). There are two IFGs per slice.</p> <p>Use the show controllers npu voq-usage command to see which other ports share a virtual output queue with the identified source system port.</p>
Timestamp (monotonic only)	<p>The timestamp when the maximum queue occupancy high water mark was recorded. The timestamp corresponds to the time the high water mark information was read from the NPU, and not the timestamp when the high water mark was detected by the NPU.</p> <p>The NPU is queried for high water mark info every 30 seconds, so the timestamp indicates the end time of a 30-second timeframe that the high water mark occurred within.</p> <p>For example, a timestamp of 16:56:44 indicates that the high water mark was observed sometime between 16:56:14 and 16:56:44.</p>

Table 30: Queue Delay fields

Field	Description
Max Queue Delay ns	The maximum delay experienced by this traffic class in nanoseconds.
Queue Occupancy %	<p>The queue occupancy as a percentage of the total queue size at the time the maximum queue delay high water mark occurred.</p> <p>Due to limited queue quantization thresholds provided by the NPU, the occupancy percentage and the occupancy kilobytes value below are an estimate of the actual maximum queue occupancy.</p>
Queue Occupancy kilobytes	<p>The queue occupancy in kilobytes at the time the maximum queue delay high water mark occurred.</p> <p>The kilobytes value is calculated with the assumption that all buffers are fully packed (i.e., all 384 bytes utilized in an SMS buffer). As a result, the displayed kilobytes value will be higher than the actual number of kilobytes queued in many cases.</p>

Field	Description
Src Sys Port Slot/NPU/Slc/Gid	<p>The Slot, NPU, Slice and GID identify the virtual output queue on which the queue delay high water mark occurred. The GID is the global identifier of the source system port whose packet was dequeued when the maximum delay high water mark occurred.</p> <p>In most cases, all ports on each slice share a virtual output queue. Although the identified source system port sent the packet that was detected as the high water mark, other ports using the same virtual output queue may have also contributed to the burst of packets that caused the maximum delay high water mark.</p> <p>Note When the fair-4 or fair-8 VOQ mode is configured, each source port has its own virtual output queue. When PFC buffer-internal mode is configured, each port shares a virtual output queue with other ports on the same slice interface group (IFG). There are two IFGs per slice.</p> <p>Use the show controllers npu voq-usage command to see which other ports share a virtual output queue with the identified source system port.</p>
Timestamp (monotonic only)	<p>The timestamp when the maximum delay high water mark was recorded. The timestamp corresponds to the time the high water mark information was read from the NPU, and not the timestamp when the high water mark was detected by the NPU.</p> <p>The NPU is queried for high water mark info every 30 seconds, so the timestamp indicates the end time of a 30-second timeframe that the high water mark occurred within.</p> <p>For example, a timestamp of 16:56:44 indicates that the high water mark was observed sometime between 16:56:14 and 16:56:44.</p>

Example 1: Monotonic High Water Marks for All Traffic Classes

The following output displays monotonic high water marks data (since bootup or the last clear operation) for all traffic classes on interface fourHundredGigE 0/0/0/11:



Note Monotonic high water marks are displayed if neither the **monotonic** or **periodic** keyword is used.

```
Router# show controllers npu qos high-water-marks interface fourHundredGigE 0/0/0/11
```

```

Interface Name      =   FH0/0/0/11
Interface Handle    =           0x1F8
System Port Gid     =           96
Asic Instance       =           0

      Queue Occupancy High Water Marks                                Queue Delay High
Water Marks

      Max Occupancy   Queue      Src Sys Port                        Max Queue   Occupancy
      Src Sys Port
      %   kilobytes   Delay ns   Slot/NPU/Slc/Gid Timestamp        Delay ns   %
kilobytes Slot/NPU/Slc/Gid Timestamp

```

```

TC_0 = 6.00 30965 73728 0/0/2/40 04/08/23 08:39:35 102400 3.00
15482 0/0/1/44 04/05/23 12:22:05
TC_1 = 0.00 0 0 0/0/0/0 - 0 0.00 0
0/0/0/0 -
TC_2 = 25.00 129024 1114112 0/0/0/48 04/07/23 01:10:23 1179648 15.00
77414 0/0/0/48 04/07/23 21:40:53
TC_3 = 70.00 361267 8912896 0/1/1/56 04/02/23 08:41:44 8912896 70.00
361267 0/1/1/58 04/02/23 08:41:44
TC_4 = 40.00 206438 2228224 3/0/2/4 04/09/23 06:38:35 2359296 25.00
129024 3/0/2/5 04/04/23 18:30:56
TC_5 = 0.00 0 0 0/0/0/0 - 0 0.00 0
0/0/0/0 -
TC_6 = 78.00 599 6437184 3/1/0/24 04/10/23 16:35:00 8628192 64.00 492
7/0/2/76 04/10/23 16:35:00
TC_7 = 25.00 129024 139264 3/0/0/14 04/06/23 08:39:41 155648 15.00
77414 0/2/2/66 04/08/23 08:39:41

[ ----- Occupancy High Water Marks ----- ] [ -----
Delay High Water Marks ----- ]

```

Example 2: Monotonic High Water Marks for a Single Traffic Class

The following output displays monotonically increasing high water marks data (since bootup or the last clear operation) for traffic class 5 on interface fourHundredGigE 0/0/0/2.

```
Router# show controllers npu qos high-water-marks monotonic interface fourHundredGigE 0/0/0/2
traffic-class 5
```

```

Interface Name      = FH0/0/0/2
Interface Handle    = 0xF000120
System Port Gid     = 6
Asic Instance       = 0

      Queue Occupancy High Water Marks      Queue Delay High
Water Marks

      Max Occupancy   Queue   Src Sys Port      Max Queue   Occupancy
      Src Sys Port
      %      kilobytes Delay ns Slot/NPU/Slc/Gid Timestamp      Delay ns      %
kilobytes Slot/NPU/Slc/Gid Timestamp
-----
TC_5 = 40.00 206438 6815744 3/0/0/15 11/11/23 17:43:30 1811939328 25.00
129024 7/1/2/89 11/27/23 11:21:26
[ ----- Occupancy High Water Marks ----- ] [ -----
Delay High Water Marks ----- ]

```

Example 3: Periodic High Water Marks for a Single Traffic Class

The following output displays high water marks data for the last three periodic collection intervals for traffic class 7 on interface fourHundredGigE 0/0/0/5.

```
Router# show controllers npu qos high-water-marks periodic last 3 interface fourHundredGigE
0/0/0/5 traffic-class 7
```

```

Interface Name      = FH0/0/0/5
Interface Handle    = 0xF000138
System Port Gid     = 9
Asic Instance       = 0

```

```

Queue Delay High Water Marks

Queue Occupancy High Water Marks

Queue  Occupancy      Src Sys Port      Max Occupancy  Queue  Src Sys Port      Max
      Interval Start      End      %      kilobytes Delay ns Slot/NPU/Slc/Gid
Delay ns      %      kilobytes Slot/NPU/Slc/Gid
-----
TC_7 = 1 12/01/23 17:46:30 12/01/23 17:46:59 50.00 258048      34680274 7/1/2/91
34680274      50.00 258048      7/1/2/91
      2 12/01/23 17:45:58 12/01/23 17:46:30 60.00 309657      52296260 0/2/1/68
61348106      50.00 258048      7/1/2/91
      3 12/01/23 17:45:30 12/01/23 17:45:58 40.00 206438      15290430 0/2/1/68
15290430      40.00 206438      0/2/1/68

----- Delay High Water Marks ----- ] [----- Occupancy High Water Marks ----- ] [

```



CHAPTER 7

Configure Priority Flow Control

- [Priority Flow Control Overview, on page 117](#)
- [View Priority Flow Control TX Frames Per Traffic Class , on page 127](#)
- [View Traffic Class Queue Pause Duration, on page 128](#)
- [Configurable ECN Threshold and Maximum Marking Probability Values, on page 135](#)
- [Priority Flow Control Watchdog Overview, on page 140](#)
- [Global Statistics Counters for Priority Flow Control and Priority Flow Control Watchdog, on page 144](#)
- [PFC Watchdog Statistics to Monitor Traffic Drops , on page 146](#)
- [Detect High Bandwidth Memory Congestion, on page 148](#)

Priority Flow Control Overview

Table 31: Feature History Table

Feature Name	Release Information	Feature Description
Priority Flow Control on Cisco 8808 and Cisco 8812 Modular Chassis Line Cards	Release 7.5.3	<p>Priority Flow Control is now supported on the following line card in the buffer-internal mode:</p> <ul style="list-style-type: none">• 88-LC0-34H14FH <p>The feature is supported in the buffer-internal and buffer-extended modes on:</p> <ul style="list-style-type: none">• 88-LC0-36FH <p>Apart from the buffer-external mode, support for this feature now extends to the buffer-internal mode on the following line cards:</p> <ul style="list-style-type: none">• 88-LC0-36FH-M• 8800-LC-48H

Feature Name	Release Information	Feature Description
Shortlink Priority Flow Control	Release 7.3.3	This feature and the hw-module profile priority-flow-control command are supported on 88-LC0-36FH line card.
Priority Flow Control Support on Cisco 8800 36x400 GbE QSFP56-DD Line Cards (88-LC0-36FH-M)	Release 7.3.15	<p>This feature and the hw-module profile priority-flow-control command are supported on 88-LC0-36FH-M and 8800-LC-48H line cards.</p> <p>All previous functionalities and benefits of this feature are available on these line cards. However, the buffer-internal mode is not supported.</p> <p>In addition, to use the buffer-extended mode on these line cards, you are required to configure the performance capacity or headroom values. This configuration requirement ensures that you can better provision and balance workloads to achieve lossless behavior, which in turn ensures efficient use of bandwidth and resources.</p>
Priority Flow Control	Release 7.3.1	This feature and the hw-module profile priority-flow-control command are not supported.

Priority-based Flow Control (IEEE 802.1Qbb), which is also referred to as Class-based Flow Control (CBFC) or Per Priority Pause (PPP), is a mechanism that prevents frame loss that is due to congestion. PFC is similar to 802.x Flow Control (pause frames) or link-level flow control (LFC). However, PFC functions on a per class-of-service (CoS) basis.

During congestion, PFC sends a pause frame to indicate the CoS value to pause. A PFC pause frame contains a 2-octet timer value for each CoS that indicates the length of time to pause the traffic. The unit of time for the timer is specified in pause quanta. A quanta is the time required for transmitting 512 bits at the speed of the port. The range is from 0 through 65535 quanta.

PFC asks the peer to stop sending frames of a particular CoS value by sending a pause frame to a well-known multicast address. This pause frame is a one-hop frame and isn't forwarded when received by the peer. When the congestion mitigates, the router stops sending the PFC frames to the upstream node.

You can configure PFC for each line card using the **hw-module profile priority-flow-control** command in one of two modes:

- buffer-internal
- buffer-extended



Note PFC threshold configurations are deprecated in `pause` command. Use the `hw-module profile priority-flow-control` command to configure PFC threshold configurations.



Tip You can programmatically retrieve the operational state of the PFC configuration using `Cisco-IOS-XR-ofa-npu-pfc-oper.yang` Cisco IOS XR native data model. To get started with using data models, see the *Programmability Configuration Guide for Cisco 8000 Series Routers*.

Related Topics

- [Configure Priority Flow Control, on page 121](#)
- [Priority Flow Control Watchdog Overview, on page 140](#)

buffer-internal mode

Use this mode if PFC-enabled devices aren't more than 1 km apart.

You can set values for pause-threshold, headroom (both related to PFC), and ECN for the traffic class using the `hw-module profile priority-flow-control` command in this mode. The buffer-internal configuration applies to all ports that the line card hosts, which mean that you can configure a set of these values per line card.

The existing queue limit and ECN configuration in the queueing policy attached to the interface has no impact in this mode.

The effective queue limit for this mode = pause-threshold + headroom (in bytes)

Restrictions and Guidelines

The following restrictions and guidelines apply while configuring the PFC threshold values using the buffer-internal mode.

- The PFC feature isn't supported on fixed chassis systems.
- Ensure that there's no breakout configured on a chassis that has the PFC configured. Configuring PFC and breakout on the same chassis may lead to unexpected behavior, including traffic loss.
- The feature isn't supported on bundle and non-bundle sub-interface queues.
- The feature is supported on 40GbE, 100 GbE, and 400 GbE interfaces.
- The feature isn't supported in the 4xVOQ queueing mode.
- The feature isn't supported when sharing of VOQ counters is configured.

buffer-extended mode

Use this mode for PFC-enabled devices with long-haul connections.

You can set the value for pause-threshold using the **hw-module profile priority-flow-control** command in this mode. You must, however, configure the queuing policy attached to the interface to set the ECN and queuing limits. The buffer-extended configuration applies to all ports that the line card hosts, which mean that you can configure a set of these values per line card.

Configuration Guidelines

- Important points while configuring the buffer-extended mode on 88-LC0-36FH-M line cards:
 - Apart from pause-threshold, you must also configure values for headroom.
 - The headroom value range is from 4 through 75000.
 - Specify pause-threshold and headroom values in units of kilobytes (KB) or megabytes (MB).
- Important points while configuring the buffer-extended mode on 8800-LC-48H line cards:
 - Configure values only for **pause-threshold**. Don't configure headroom values.
 - Configure **pause-threshold** in units of milliseconds (ms) or microseconds.
 - Don't use units of kilobytes (KB) or megabytes (MB) units, even though the CLI displays them as options. Only use units of milliseconds (ms) or microseconds.

(Also see [Configure Priority Flow Control, on page 121](#))

Important Considerations

- If you configure PFC values in the buffer-internal mode, then the ECN value for the line card is derived from the buffer-internal configuration. If you configure PFC values in the buffer-extended mode, then the ECN value is derived from the policy map. (For details on the ECN feature, see [Explicit Congestion Notification , on page 106](#).)
- The buffer-internal and buffer-extended modes can't coexist on the same line card.
- For Cisco 8808 and Cisco 8812 chassis, configure PFC on all line cards in the chassis, regardless of whether you're configuring the buffer-internal or buffer-extended mode. Otherwise, your network may experience traffic loss.
- If you add or remove traffic-class actions on a line card, you must reload the line card.
- When using the buffer-internal mode, you can change values of the following parameters without having to reload the line card. However, if you add a new traffic class and configure these values for the first time on that traffic class, you must reload the line card for the values to come into effect.
 - pause-threshold
 - headroom
 - ECN
- If you add or remove ECN configuration using the **hw-module profile priority-flow-control** command, you must reload the line card for the ECN changes to take effect.
- The PFC threshold value ranges for the buffer-internal mode are as follows.

Threshold	Configured (bytes)
pause (min)	307200
pause (max)	422400
headroom (min)	345600
headroom (max)	537600
ecn (min)	153600
ecn (max)	403200

- For a traffic-class, the ECN value must always be lesser than the configured pause-threshold value.
- The combined configured values for pause-threshold and headroom must not exceed 844800 bytes. Else, the configuration is rejected.
- The pause-threshold value range for buffer-extended mode is from 2 milliseconds (ms) through 25 ms and from 2000 microseconds through 25000 microseconds.

Hardware Support for Priority Flow Control

The table lists the PIDs that support PFC per release and the PFC mode in which the support is available.

Table 32: PFC Hardware Support Matrix

Release	PID	PFC Mode
Release 7.5.3	88-LC0-36FH	buffer-extended and buffer-internal
	88-LC0-36FH-M	buffer-extended and buffer-internal
	8800-LC-48H	buffer-extended and buffer-internal
	88-LC0-34H14FH	buffer-internal
Release 7.3.15	<ul style="list-style-type: none"> • 88-LC0-36FH-M • 88-LC0-36FH 	buffer-extended
Release 7.0.11	8800-LC-48H	buffer-internal

Configure Priority Flow Control

You can configure PFC to enable the no-drop behavior for the CoS as defined by the active network QoS policy.



Note

The system enables shortlink PFC by default when you enable PFC.

**YANG Data Model**

You can programmatically manage the flow control parameters using the `openconfig-qos.yang` data model. To get started with using data models, see the Programmability Configuration Guide for Cisco 8000 Series Routers.

Configuration Example

You must accomplish the following to complete the PFC configuration:

1. Enable PFC at the interface level.
2. Configure ingress classification policy.
3. Attach the PFC policy to the interface.
4. Configure PFC threshold values using either the buffer-internal or buffer-extended mode.

```
Router# configure
Router(config)# priority-flow-control mode on
/*Configure ingress classification policy*/
Router(config)# class-map match-any prec7
Router(config-cmap)# match precedence
Router(config)# class-map match-any tc7
/*Ingress policy attach*/
Router(config-if)# service-policy input QOS_marking
/*Egress policy attach*/
Router(config-if)# service-policy output qos_queuing
Router(config-pmap-c)# exit
Router(config-pmap)# exit
Router(config)#show controllers npu priority-flow-control location <loc>
```

Running Configuration

```
*Interface Level*
interface HundredGigE0/0/0/0
    priority-flow-control mode on

*Ingress:*
class-map match-any prec7

    match precedence 7

end-class-map

!

class-map match-any prec6

    match precedence 6

end-class-map

!

class-map match-any prec5

    match precedence 5
```

```
end-class-map

!

class-map match-any prec4

match precedence 4

end-class-map

!

class-map match-any prec3
match precedence 3
end-class-map
!
class-map match-any prec2
match precedence 2
end-class-map
!
class-map match-any prec1
match precedence 1
end-class-map
!
!
policy-map QOS_MARKING
class prec7
    set traffic-class 7
    set qos-group 7
!
class prec6
    set traffic-class 6
    set qos-group 6
!
class prec5
    set traffic-class 5
    set qos-group 5
!
class prec4
    set traffic-class 4
    set qos-group 4
!
class prec3
    set traffic-class 3
    set qos-group 3
!
class prec2
    set traffic-class 2
    set qos-group 2
!
class prec1
    set traffic-class 1
    set qos-group 1
!
class class-default
    set traffic-class 0
    set qos-group 0
!

*Egress:*
class-map match-any tc7
match traffic-class 7
end-class-map
!
```

```
class-map match-any tc6
  match traffic-class 6
end-class-map
!
class-map match-any tc5
  match traffic-class 5
end-class-map
!

class-map match-any tc4

  match traffic-class 4

end-class-map
!

class-map match-any tc3

  match traffic-class 3

end-class-map
!

class-map match-any tc2
match traffic-class 2
end-class-map
!
class-map match-any tc1
match traffic-class 1
end-class-map
!
policy-map QOS_QUEUEING
class tc7
  priority level 1
  shape average percent 10
!
class tc6
  bandwidth remaining ratio 1
  queue-limit 100 ms
!
class tc5
  bandwidth remaining ratio 20
  queue-limit 100 ms
!
class tc4
  bandwidth remaining ratio 20
  random-detect ecn
  random-detect 6144 bytes 100 mbytes
!
class tc3
  bandwidth remaining ratio 20
  random-detect ecn
  random-detect 6144 bytes 100 mbytes
!
class tc2
  bandwidth remaining ratio 5
  queue-limit 100 ms
!
class tc1
  bandwidth remaining ratio 5
  queue-limit 100 ms
```

```

!
class class-default
  bandwidth remaining ratio 20
  queue-limit 100 ms
!
[buffer-extended]

hw-module profile priority-flow-control location 0/0/CPU0
  buffer-extended traffic-class 3 pause-threshold 10 ms
  buffer-extended traffic-class 4 pause-threshold 10 ms
!

[buffer-internal]

hw-module profile priority-flow-control location 0/1/CPU0
  buffer-internal traffic-class 3 pause-threshold 403200 bytes headroom 441600 bytes ecn
  224640 bytes
  buffer-internal traffic-class 4 pause-threshold 403200 bytes headroom 441600 bytes ecn
  224640 bytes

```

Verification

```

Router#sh controllers hundredGigE0/0/0/22 priority-flow-control
Priority flow control information for interface HundredGigE0/0/0/22:
Priority Flow Control:
Total Rx PFC Frames : 0
Total Tx PFC Frames : 313866
Rx Data Frames Dropped: 0
CoS Status Rx Frames
---
0   on   0
1   on   0
2   on   0
3   on   0
4   on   0
5   on   0
6   on   0
7   on   0

/*[buffer-internal]*/
Router#show controllers hundredGigE 0/9/0/24 priority-flow-control

Priority flow control information for interface HundredGigE0/9/0/24:

Priority Flow Control:
Total Rx PFC Frames : 0
Total Tx PFC Frames : 313866
Rx Data Frames Dropped: 0
CoS Status Rx Frames
---
0   on   0
1   on   0
2   on   0
3   on   0
4   on   0
5   on   0
6   on   0
7   on   0
...

/*[buffer-internal, tc3 & tc4 configured. TC4 doesn't have ECN]*/

```

```

Router#show controllers npu priority-flow-control location <loc>
Location Id:                                0/1/CPU0
PFC:                                         Enabled
PFC-Mode:                                   buffer-internal
TC   Pause                                Headroom    ECN
-----
3    86800 bytes                          120000 bytes  76800 bytes
4    86800 bytes                          120000 bytes  Not-configured

/*[buffer-extended PFC, tc3 & tc4 configured]*/

Router#show controllers npu priority-flow-control location <loc>
Location Id:                                0/1/CPU0
PFC:                                         Enabled
PFC-Mode:                                   buffer-extended
TC   Pause                                -----
3    5000 us
4    10000 us

/*[No PFC]*/

Router#show controllers npu priority-flow-control location <loc>
Location Id:                                0/1/CPU0
PFC:                                         Disabled

```

Related Topics

- [Priority Flow Control Overview, on page 117](#)

Related Commands hw-module profile priority-flow-control location

View Priority Flow Control TX Frames Per Traffic Class

Table 33: Feature History Table

Feature Name	Release Information	Feature Description
View Priority Flow Control TX Frames Per Traffic Class	Release 7.5.4	<p>You can now view an estimation of the transmission (Tx) of Priority Flow Control (PFC) frames per traffic class, which informs you that the traffic flow has crossed configured PFC pause thresholds. This information allows you to rebalance traffic flows to ensure that network resources are efficiently used. Such verification is possible because we've added counters for PFC Tx frames.</p> <p>This functionality modifies the following to add the PFC Tx frame counters:</p> <ul style="list-style-type: none"> • YANG data model (at Github under the 754 folder): Cisco-IOS-XR-ofa-npu-pfc-oper • show controllers

Beginning Cisco IOS XR Release 7.5.4, we have added an additional counter for PFC Tx pause frames per traffic class in the **show controllers priority-flow-control statistics** command form. With this counter, you receive timely information that the traffic flow has exceeded the configured PFC pause thresholds and hence that PFC Tx frames are sent out of the PFC-enabled interfaces sourcing such traffic on that network processing core. You can use this information to investigate and analyze why such traffic flows are causing congestion and improve traffic distribution across interfaces to utilize network resources efficiently.

Run the **show controllers priority-flow-control statistics** to view an estimation of the statistics for PFC Tx frames per traffic class (**Tx Frames**):

```
Router#show controllers hundredGigE 0/0/0/4 priority-flow-control statistics
```

Priority flow control information for interface HundredGigE0/0/0/4:

Priority Flow Control:

```

Total Rx PFC Frames: 0
Total Tx PFC Frames: 4832680
Rx Data Frames Dropped: 1442056 (possible overflow)

```

CoS	Status	Rx Frames	Tx Frames
0	on	0	0
1	on	0	0
2	on	0	0
3	on	0	2416374
4	on	0	2416306
5	on	0	0
6	on	0	0
7	on	0	0

Guidelines and Limitations for Viewing Priority Flow Control TX Frames Per Traffic Class

This functionality is supported on:

- 88-LC0-36FH-M
- 88-LC0-34H14FH

View Traffic Class Queue Pause Duration

Table 34: Feature History Table

Feature Name	Release Information	Feature Description
View Traffic Class Queue Pause Duration	Release 24.2.11	<p><i>Introduced in this release on Cisco 8000 Series Routers with Cisco Silicon One Q200 network processors that support the PFC buffer-extended mode function.</i></p> <p>For traffic flows between routers, you can view the pause duration of output and input queues in the transmitting and receiving routers, respectively.</p> <p>The pause duration values of the impacted traffic class queues are displayed for regular intervals within a specified time duration.</p> <p>With the information, you can view the extent of congestion on PFC-enabled interfaces over a period of time and identify whether traffic congestion is due to small bursts of traffic or other causes.</p> <p>The feature introduces these changes:</p> <p>CLI:</p> <ul style="list-style-type: none"> • show controllers npu packet-memory interface <p>YANG Data Models:</p> <ul style="list-style-type: none"> • Cisco-IOS-XR-platforms-ofa-oper (see GitHub, YANG Data Models Navigator)

Pre-requisite to View Traffic Class Queue Pause Duration


- Enable PFC buffer-extended mode on the transmitting and receiving router nodes. For PFC-related information and PFC buffer-extended mode configuration guidelines, refer the [Priority Flow Control Overview](#) and [buffer-extended mode](#) > **Configuration Guidelines** sections.

Traffic Class Queue Pause Duration Scenario

The feature is explained through a sample scenario. In this network topology, PFC buffer-extended mode is enabled on routers R2 and R3. R2 is the transmitting router and R3 is the receiving router. After the explanation of the scenario, pause duration values of the output queues on R2 and input queues on R3 are provided. For ease-of-use, virtual output queues are referred as *output queues* in this feature.

PFC buffer-extended mode configuration on R2 and R3:

```
R2,R3(config)# hw-module profile priority-flow-control location 0/0/CPU0
R2,R3(config-pfc-loc)# buffer-extended traffic-class 2 pause-threshold 10000 kbytes headroom
10000 kbytes
R2,R3(config-pfc-loc)# buffer-extended traffic-class 3 pause-threshold 10000 kbytes headroom
10000 kbytes
R2,R3(config-pfc-loc)# commit
```



Traffic Flow Details:

1. R1 sends packets to R2 at a line rate of 400G.
2. R2 forwards the packets to R3 over two ECMP links at a line rate of 200G each, through its PFC-enabled interfaces. The two links have a 400G capacity per link.
3. R3 receives the packets over its PFC-enabled interfaces. A small part of R3's bandwidth, say 1G, is used for sending NO-PAUSE (or X-ON) frames. So, as R2 continuously forwards 400G packets to R3, R3 forwards 399G of it to R4 and stores 1G packets in its input queue buffer. This leads to mild congestion on R3 and its interfaces.
4. Slowly, the input queue buffers of R3 increase, and over time, it reaches the pause-threshold value. R3 sends PAUSE frames to R2.
5. R2 starts buffering the packets in its HBM. In turn, R2 sends PAUSE frames to R1.
6. During the PAUSE state, R3 drains the packets in its input queue by forwarding it to R4. When R3's input queue buffer size dips below the pause-threshold value, R3 resumes sending NO-PAUSE frames to R2.
7. Since R2 has a sizeable output queue buffer to clear, and it can send packets at a line rate of 400G, it sends out a short burst of traffic at its maximum possible line rate, 400G.
8. When R3 starts receiving 400G line rate packets over its two interfaces, its input queue buffer size increases (again) and, upon crossing the threshold value, R3 starts sending PAUSE frames to R2 (again).
9. R3 can only receive and forward nearly 400G out of the 800G traffic burst it receives from R2, or 50% of incoming packets. In other words, the pause duration of the queues is 50% of the time. The 50% value is called *pause-percentage*.
10. Since R3 has been sending PAUSE frames to R2 corresponding to its receiving capability, and R2 has to meet R3's requirement, R2's output queues are also PAUSED for 50% of the time.
11. The traffic burst instances occur for short periods of time, and they aren't a cause of persistent congestion. The high pause duration values (such as pause-percentage = 50%) are only seen during such occurrences.

Using the **show controllers npu packet-memory interface** command, you can view traffic class queue pause duration for regular intervals within time durations ranging from 30 seconds to five minutes, and identify whether traffic congestion is due to short-lived traffic bursts or other causes.

Verification

Pause duration parameters:

Choose one option: {**rx-pause-percent**|**tx-pause-percent**}

- **rx-pause-percent**—Specifies the input queues on the receiving router R3.
- **tx-pause-percent**—Specifies the output queues on the transmitting router R2.

(Optional) Choose one option: [**detail**|**five-minute**|**one-minute**|**verbose**]

- **one-minute**—Specifies the average pause duration of the queues in percentage for the last minute.
- **five-minute**—Specifies the average pause duration of the queues in percentage for the last five minutes.
- **detail**—Displays a maximum of 120 records at a frequency of one record per 250-millisecond interval.
- If you don't use **detail**, **five-minute** or **one-minute**, 30 records are displayed for each traffic class at a frequency of one record per second. If you want to view the time stamp in a raw mode, choose **verbose**.

In the following examples, input queue and output queue samples are provided for each example:

Example 1: Average Pause Duration in the Last Minute

Input Queue on R3: Displays the average pause duration of input queues on node 0/6/CPU0 interfaces for the last one-minute duration. Traffic for the traffic class 2 queue on interfaces FH0/6/0/21 and FH0/6/0/23 is paused for more than 50% of the time.

```
R3# show controllers npu packet-memory interface all tx-pause-percent one-minute location 0/6/CPU0
```

```
-----
Node ID: 0/6/CPU0
Source Queue Pause Percentage Info for interface(s) all
Intf          TC          Pause-Percentage
name
-----
FH0/6/0/10      2          0.00000
FH0/6/0/11      2          0.00000
FH0/6/0/13      2          0.00000
FH0/6/0/14      2          0.00000
FH0/6/0/15      2          0.00000
FH0/6/0/16      2          0.00000
FH0/6/0/18      2          0.00000
FH0/6/0/21      2          53.01604
FH0/6/0/22      2          0.00000
FH0/6/0/23      2          53.13991
```

Output Queue on R2: Displays the average pause duration of output queues on node 0/6/CPU0 interfaces for the last one-minute duration. Traffic for the traffic class 2 queue on interfaces FH0/6/0/21 and FH0/6/0/23 is paused for more than 50% of the time.

```
R2# show controllers npu packet-memory interface all rx-pause-percent one-minute location 0/6/CPU0
```

```
-----
Node ID: 0/6/CPU0
Out Queue Pause Percentage Info for interface(s) all
Intf          TC          Pause-Percentage
name
-----
FH0/6/0/10      2          0.00000
FH0/6/0/11      2          0.00000
FH0/6/0/13      2          0.00000
FH0/6/0/14      2          0.00000
FH0/6/0/15      2          0.00000
FH0/6/0/16      2          0.00000
```

FH0/6/0/18	2	0.00000
FH0/6/0/21	2	53.01604
FH0/6/0/22	2	0.00000
FH0/6/0/23	2	53.13991

Example 2: Average Pause Duration for the Last Minute on One Interface

Input Queue on R3: Displays the average pause duration of input queues on node 0/6/CPU0 interface FourHundredGigE0_6_0_21 for the last one-minute duration. Traffic for the traffic class 2 queue on interface FourHundredGigE0_6_0_21 is paused for more than 50% of the time.

```
R3# show controllers npu packet-memory interface FH0/6/0/21 tx-pause-percent one-minute
location 0/6/CPU0
```

```
-----
Node ID: 0/6/CPU0
Source Queue Pause Percentage Info for interface(s) FourHundredGigE0_6_0_21
Time                                         TC      Pause-Percentage
-----
Thu 2024-03-28 18:53:30.264 UTC           2      53.49027
```

Output Queue on R2: Displays the average pause duration of output queues on node 0/6/CPU0 interface FourHundredGigE0_6_0_21 for the last one-minute duration. Traffic for the traffic class 2 queue on interface FH0/6/0/21 is paused for more than 50% of the time.

```
R2# show controllers npu packet-memory interface FH0/6/0/21 rx-pause-percent one-minute
location 0/6/CPU0
```

```
-----
Node ID: 0/6/CPU0
Out Queue Pause Percentage Info for interface(s) FourHundredGigE0_6_0_21
Time                                         TC      Pause-Percentage
-----
Thu 2024-03-28 18:53:30.264 UTC           2      53.49027
```

Example 3: Average Pause Duration for the Last 5 Minutes

Input Queue on R3: Displays the average pause duration of input queues on node 0/6/CPU0 interfaces for the last five minutes. Traffic for the traffic class 2 queue on interfaces FH0/6/0/21 and FH0/6/0/23 is paused for more than 50% of the time.

```
R3# show controllers npu packet-memory interface all tx-pause-percent five-minute location
0/6/CPU0
```

```
-----
Node ID: 0/6/CPU0
Source Queue Pause Percentage Info for interface(s) all
Intf          TC      Pause-Percentage
name
-----
FH0/6/0/18    2      0.00000
FH0/6/0/21    2      53.24746
FH0/6/0/22    2      0.00000
FH0/6/0/23    2      53.23203
FH0/6/0/24    2      0.00000
FH0/6/0/30    2      0.00000
Hu0/6/0/17    2      0.00000
Hu0/6/0/20    2      0.00000
```

Output Queue on R2: Displays the average pause duration of output queues on node 0/6/CPU0 interfaces for the last five minutes. Traffic for the traffic class 2 queue on interfaces FH0/6/0/21 and FH0/6/0/23 is paused for more than 50% of the time.

```
R2# show controllers npu packet-memory interface all rx-pause-percent five-minute location
0/6/CPU0
```

```

Node ID: 0/6/CPU0
Out Queue Pause Percentage Info for interface(s) all
Intf          TC          Pause-Percentage
name
-----
FH0/6/0/10    2          0.00000
FH0/6/0/11    2          0.00000
FH0/6/0/13    2          0.00000
FH0/6/0/14    2          0.00000
FH0/6/0/15    2          0.00000
FH0/6/0/16    2          0.00000
FH0/6/0/18    2          0.00000
FH0/6/0/21    2          53.01604
FH0/6/0/22    2          0.00000
FH0/6/0/23    2          53.13991

```

Example 4: Average Pause Duration for the Last 5 Minutes on One Interface

Input Queue on R3: Displays the average pause duration of input queues on interface FH0/6/0/21 on node 0/6/CPU0 for the last five minutes. Traffic for the traffic class 2 queue on interface FH0/6/0/21 is paused for more than 50% of the time.

```

R3# show controllers npu packet-memory interface FH0/6/0/21 tx-pause-percent five-minute
location 0/6/CPU0

```

```

Node ID: 0/6/CPU0
Source Queue Pause Percentage Info for interface(s) FourHundredGigE0_6_0_21
Time          TC          Pause-Percentage
-----
Thu 2024-03-28 18:53:00.242 UTC          2          53.31213

```

Output Queue on R2: Displays the average pause duration of output queues on interface FH0/6/0/21 on node 0/6/CPU0 for the last five minutes. Traffic for the traffic class 2 queue on interface FH0/6/0/21 is paused for more than 50% of the time.

```

R2# show controllers npu packet-memory interface FH0/6/0/21 rx-pause-percent five-minute
location 0/6/CPU0

```

```

Node ID: 0/6/CPU0
Out Queue Pause Percentage Info for interface(s) FourHundredGigE0_6_0_21
Time          TC          Pause-Percentage
-----
                2          53.31213

```

Example 5: Pause Duration for the Last 30 Seconds With a One-Second Interval

Input Queue on R3: Displays the pause duration of input queues on interface FH0/6/0/21. The traffic class 2 queue is paused for more than 50% of the time for the entire time duration.

```

R3# show controllers npu packet-memory interface FH0/6/0/21 tx-pause-percent location
0/6/CPU0

```

```

Node ID: 0/6/CPU0
Source Queue Pause Percentage Info for interface(s) FourHundredGigE0_6_0_21
Time          TC          Pause-Percentage
-----
Thu 2024-03-28 18:53:31.265 UTC          2          53.57590
Thu 2024-03-28 18:53:32.266 UTC          2          54.04276
Thu 2024-03-28 18:53:33.267 UTC          2          53.99954
Thu 2024-03-28 18:53:34.268 UTC          2          53.38932
Thu 2024-03-28 18:53:35.269 UTC          2          54.52591
Thu 2024-03-28 18:53:36.270 UTC          2          54.27677

```

```

Thu 2024-03-28 18:53:37.271 UTC      2      52.78747
Thu 2024-03-28 18:53:38.271 UTC      2      53.42725
Thu 2024-03-28 18:53:39.271 UTC      2      53.34946
Thu 2024-03-28 18:53:40.272 UTC      2      53.06505
Thu 2024-03-28 18:53:41.272 UTC      2      53.50647
Thu 2024-03-28 18:53:42.272 UTC      2      53.49165
Thu 2024-03-28 18:53:43.272 UTC      2      52.52429
Thu 2024-03-28 18:53:44.272 UTC      2      53.44903
..
..

```

Output Queue on R2: Displays the pause duration of output queues on interface FH0/6/0/21. The traffic class 2 queue is paused for more than 50% of the time for the entire time duration.

R2# **show controllers npu packet-memory interface FH0/6/0/21 rx-pause-percent location 0/6/CPU0**

```

-----
Node ID: 0/6/CPU0
Out Queue Pause Percentage Info for interface(s) FourHundredGigE0_6_0_21
Time                                         TC      Pause-Percentage
-----
Thu 2024-03-28 18:53:31.265 UTC      2      53.57590
Thu 2024-03-28 18:53:32.266 UTC      2      54.04276
Thu 2024-03-28 18:53:33.267 UTC      2      53.99954
Thu 2024-03-28 18:53:34.268 UTC      2      53.38932
Thu 2024-03-28 18:53:35.269 UTC      2      54.52591
Thu 2024-03-28 18:53:36.270 UTC      2      54.27677
Thu 2024-03-28 18:53:37.271 UTC      2      52.78747
Thu 2024-03-28 18:53:38.271 UTC      2      53.42725
Thu 2024-03-28 18:53:39.271 UTC      2      53.34946
Thu 2024-03-28 18:53:40.272 UTC      2      53.06505
Thu 2024-03-28 18:53:41.272 UTC      2      53.50647
Thu 2024-03-28 18:53:42.272 UTC      2      53.49165
Thu 2024-03-28 18:53:43.272 UTC      2      52.52429
Thu 2024-03-28 18:53:44.272 UTC      2      53.44903
..
..

```

Example 6: Pause Duration for the Last 30 Seconds With a 250 ms Interval

Input Queue on R3: Displays the pause duration of the input queues on interface FH0/6/0/21. The traffic class 2 queue is paused for more than 50% of the time.

R3# **show controllers npu packet-memory interface FH0/6/0/21 tx-pause-percent detail location 0/6/CPU0**

```

-----
This CLI will fetch large number of records.
Integration with automation scripts is not recommended
-----
Node ID: 0/6/CPU0
Source Queue Pause Percentage Info for interface(s) FourHundredGigE0_6_0_21
Time                                         TC      Pause-Percentage
-----
Thu 2024-03-28 18:53:30.514 UTC      2      52.22349
Thu 2024-03-28 18:53:30.764 UTC      2      53.67103
Thu 2024-03-28 18:53:31.014 UTC      2      53.63637
Thu 2024-03-28 18:53:31.265 UTC      2      54.77273
Thu 2024-03-28 18:53:31.515 UTC      2      53.72371
Thu 2024-03-28 18:53:31.765 UTC      2      54.29221
Thu 2024-03-28 18:53:32.015 UTC      2      53.04150
Thu 2024-03-28 18:53:32.266 UTC      2      55.11364
..
..

```

Output Queue on R2: Displays the pause duration of the output queues on interface FH0/6/0/21. The traffic class 2 queue is paused for more than 50% of the time.

R2# **show controllers npu packet-memory interface FH0/6/0/21 rx-pause-percent detail location 0/6/CPU0**

 This CLI will fetch large number of records.
 Integration with automation scripts is not recommended

Node ID: 0/6/CPU0

Out Queue Pause Percentage Info for interface(s) FourHundredGigE0_6_0_21

Time	TC	Pause-Percentage
Thu 2024-03-28 18:53:30.514 UTC	2	52.22349
Thu 2024-03-28 18:53:30.764 UTC	2	53.67103
Thu 2024-03-28 18:53:31.014 UTC	2	53.63637
Thu 2024-03-28 18:53:31.265 UTC	2	54.77273
Thu 2024-03-28 18:53:31.515 UTC	2	53.72371
Thu 2024-03-28 18:53:31.765 UTC	2	54.29221
Thu 2024-03-28 18:53:32.015 UTC	2	53.04150
Thu 2024-03-28 18:53:32.266 UTC	2	55.11364

Configurable ECN Threshold and Maximum Marking Probability Values

Table 35: Feature History Table

Feature Name	Release Information	Feature Description
Configurable ECN Threshold and Maximum Marking Probability Values	Release 7.5.4Release 7.3.6	<p>While configuring PFC in the buffer-internal mode, you can now optimize congestion notification from the end router to the transmitting router, thus preventing the aggressive throttle of the source traffic. This optimization is possible because we've provided the flexibility to configure the minimum and maximum values for the ECN threshold and the maximum value for marking probability. With these values configured, the probability percentage marking is applied linearly, starting at the ECN minimum threshold till the ECN Max threshold.</p> <p>Earlier releases fixed the maximum ECN marking probability at 100% at the maximum ECN threshold.</p> <p>This functionality adds the following options to the hw-module profile priority-flow-control command:</p> <ul style="list-style-type: none"> • max-threshold • probability-percentage

ECN Threshold and Maximum Marking Probability Values


So far, the maximum ECN marking probability was not configurable and was fixed at 100%. You couldn't configure the ECN maximum threshold value either. Such an arrangement of preset marking probabilities and fixed maximum threshold values meant traffic rates began to drop as a function of the queue length. Because of the linear increase in the ECN marking probability—and the consequent congestion signaling from the end host to the transmitting host—traffic rates could begin slowing down even if your link had the necessary bandwidth.

With this feature, you can configure the minimum and maximum ECN threshold values and, depending on your link capability and traffic requirements, configure a marking probability such that the number of packets being marked between the minimum and maximum thresholds is reduced.

The figure depicts the feature implementation visually, where:

- Pmax denotes the configurable maximum marking probability.
- The ECN mark probability is 1 from the maximum threshold to the tail drop threshold.
- Mark probability increases linearly from 0 at the minimum threshold to Pmax at the maximum threshold.

Figure 9: ECN Mark Probability vs. Queue Length (VoQ Fill Levels)



You configure the **max-threshold** and **probability-percentage** options for this feature in the **hw-module profile priority-flow-control** command. We have provided you the flexibility to choose one of these configuration options:

- Default **max-threshold** and **probability-percentage** values.
- Configurable **max-threshold** and **probability-percentage** values.
- PFC in the buffer-internal mode without the new options, as you did before Release 7.5.4/Release 7.3.6.

See [Configure ECN Threshold and Maximum Marking Probability Values, on page 138](#) for details.

Benefits of Configurable ECN Threshold and Maximum Marking Probability Values

A lower mark probability allows the end host to better identify the aggressor flow in the queue and only throttle it without impacting other well-behaved flows. This in turn leads to optimized and efficient deployment and usage of bandwidth.

ECN Threshold and Maximum Marking Probability Values: FAQs

- What if I set the ECN marking probability percentage to 100%?

ECN Marking is linear once queue length exceeds the ECN minimum value. The tail drop comes into play when the average queue size reaches the marking probability percentage of 100%.

- **How does traffic behave beyond my configured value of maximum ECN marking probability?**

Let us say you've set the maximum ECN marking probability to 5%. Any more increase in the average queue length beyond 5% shifts the marking probability to 100%, and [Tail Drop and the FIFO Queue, on page 102](#) comes into effect.

Guidelines and Limitations

- This functionality is available only when you configure PFC in the buffer-internal mode.
- If you configure PFC values in the buffer-internal mode, the ECN value for the line card is derived from the buffer-internal configuration, unlike in the buffer-extended mode where the ECN value is derived from the policy map.
- Line Cards that support this functionality:
 - 88-LC0-36FH
 - 88-LC0-36FH-M
- The following interface types support this functionality:
 - Physical interfaces
 - Bundle interfaces
 - Subinterfaces
 - Bundle subinterfaces
- This functionality is supported for all interface speeds.
- If you have a policy map with one or multiple classes with maximum ECN marking probability enabled, you can:
 - Apply the map to any of the supported interface types.
 - Remove the map from any of the supported interface types.
 - Modify the map while you're attaching it to multiple interfaces.
- The probability percentage option is supported only with **random-detect ecn** configured in the same class. Else, the policy is rejected when applied on an interface.
- Ensure that the probability percentage configuration is the same for all traffic classes because this is a device-level configuration.

Table 36: Useful Tips

If you...	...you must...
<p>Want to switch from the default configuration mode to the custom configuration mode</p> <p>OR</p> <p>Want to switch from the custom configuration mode to the default configuration</p>	<ol style="list-style-type: none"> 1. Use the no form of the hw-module profile priority-flow-control command to remove the existing configuration. 2. Configure the new mode and settings using the hw-module profile priority-flow-control command. 3. Reload the line card
Configured PFC in buffer-internal mode without configuring the max-threshold and probability-percentage parameters, but now want to configure them	<ol style="list-style-type: none"> 1. Configure the max-threshold and probability-percentage parameters using the hw-module profile priority-flow-control command. 2. Reload the line card.
Want to change the buffer-internal parameters in the custom mode	Configure the buffer-internal parameters using the hw-module profile priority-flow-control command. You needn't reload the line card.
Want to continue configuring PFC in buffer-internal mode the way you did in releases prior to Release 7.5.4.	<p>Configure the buffer-internal parameters using the hw-module profile priority-flow-control command, but ensure that you only configure values for pause-threshold, headroom, and ecn. If you don't configure values for max-threshold, the router takes the ecn value as the ECN maximum threshold value.</p> <p>See Option 3 under Configure ECN Threshold and Maximum Marking Probability Values, on page 138 for details.</p>

Configure ECN Threshold and Maximum Marking Probability Values

To configure the ECN maximum threshold and maximum probability values:

1. Enable PFC at the interface level.

```
Router(config)#int fourHundredGigE 0/6/0/1
Router(config-if)#priority-flow-control mode on
```

2. [Option 1: Default configuration mode] Configure PFC in buffer-internal mode with pre-defined buffer values enabled along with default **max-threshold** and **probability-percentage** values.

```
Router(config)#hw-module profile priority-flow-control location 0/6/0/1
Router(config-pfc-loc)#buffer-extended traffic-class 3
Router(config-pfc-loc)#buffer-extended traffic-class 4
Router(config-pfc-loc)#commit
```

[Option 2: Custom configuration mode] Configure PFC in buffer-internal mode with custom values for all parameters including **max-threshold** and **probability-percentage**.

```
Router(config)#hw-module profile priority-flow-control location 0/6/0/1
Router(config-pfc-loc)#buffer-internal traffic-class 3 pause-threshold 1574400 bytes
headroom 1651200 bytes ecn 629760 bytes max-threshold 1416960 bytes
probability-percentage 50
Router(config-pfc-loc)#buffer-internal traffic-class 4 pause-threshold 1574400 bytes
headroom 1651200 bytes ecn 629760 bytes max-threshold 1416960 bytes
probability-percentage 50
Router(config-pfc-loc)#commit
```

[Option 3: Without **max-threshold** and **probability-percentage** parameters] Based on your requirements, you can also configure PFC in buffer-internal mode without these parameters.

```
Router(config)#hw-module profile priority-flow-control location 0/6/0/1
Router(config-pfc-loc)#buffer-internal traffic-class 3 pause-threshold 1574400 bytes
headroom 1651200 bytes ecn 629760 bytes
Router(config-pfc-loc)#buffer-internal traffic-class 4 pause-threshold 1574400 bytes
headroom 1651200 bytes ecn 629760 bytes
Router(config-pfc-loc)#commit
```



Note See **Useful Tips** at [ECN Threshold and Maximum Marking Probability Values: FAQs, on page 136](#) to know more about the nuances of switching between configuration modes.

Running Configuration

[Option 1: Default configuration mode]

```
/*Interface Level*/
priority-flow-control mode on
!
/*Line card*/
hw-module profile priority-flow-control location 0/6/0/1
buffer-extended traffic-class 3
buffer-extended traffic-class 4
!
```

[Option 2: Custom configuration mode]

```
/*Interface Level*/
priority-flow-control mode on
!
/*Line card*/

hw-module profile priority-flow-control location 0/6/0/1
buffer-internal traffic-class 3 pause-threshold 1574400 bytes headroom 1651200 bytes ecn
629760 bytes max-threshold 1416960 bytes probability-percentage 50
buffer-internal traffic-class 4 pause-threshold 1574400 bytes headroom 1651200 bytes ecn
629760 bytes max-threshold 1416960 bytes probability-percentage 50
!
```

[Option 3: Without **max-threshold** and **probability-percentage** parameters]

```
/*Interface Level*/
priority-flow-control mode on
!
/*Line card*/
buffer-internal traffic-class 3 pause-threshold 1574400 bytes headroom 1651200 bytes ecn
629760 bytes
```

```
buffer-internal traffic-class 4 pause-threshold 1574400 bytes headroom 1651200 bytes ecn
629760 bytes
!
```

Verification

[Option 1: Default configuration mode]

```
Router#show controllers npu priority-flow-control location all
```

```
Location:          0/6/CPU0
PFC:               Enabled
PFC Mode:          buffer-internal
TC    Pause-threshold  Headroom      ECN          ECN-MAX      Prob-per
-----
3     1574400 bytes    1651200 bytes  629760 bytes  1416960 bytes  5
4     1574400 bytes    1651200 bytes  629760 bytes  1416960 bytes  5
```

[Option 2: Custom configuration mode]

```
Router#show controllers npu priority-flow-control location all
```

```
Location:          0/6/CPU0
PFC:               Enabled
PFC Mode:          buffer-internal
TC    Pause-threshold  Headroom      ECN          ECN-MAX      Prob-per
-----
3     1574400 bytes    1651200 bytes  629760 bytes  1416960 bytes  50
4     1574400 bytes    1651200 bytes  629760 bytes  1416960 bytes  50
```

[Option 3: Without **max-threshold** and **probability-percentage** parameters]

```
Router#show controllers npu priority-flow-control location all
```

```
Location:          0/6/CPU0
PFC:               Enabled
PFC Mode:          buffer-internal
TC    Pause-threshold  Headroom      ECN          ECN-MAX      Prob-per
-----
3     1574400 bytes    1651200 bytes  629760 bytes  not-configured not-configured
4     1574400 bytes    1651200 bytes  629760 bytes  not-configured not-configured
```

Priority Flow Control Watchdog Overview

PFC Watchdog is a mechanism to identify any PFC storms (queue-stuck condition) in the network. It also prevents the PFC from propagating on the network and running in a loop. You can configure a PFC watchdog interval to detect whether packets in a no-drop queue are drained within a specified time period. When the time period is exceeded, all outgoing packets are dropped on interfaces that match the PFC queue that is not being drained.

This requires monitoring PFC receiving on each port and detecting ports seeing an unusual number of sustained pause frames. Once detected, the watchdog module can enforce several actions on such ports, which include generating a syslog message for network management systems, shutting down the queue, and autorestoring the queue (after the PFC storm stops).

Here's how the PFC Watchdog works:

1. The Watchdog module monitors the PFC-enabled queues to determine the reception of an unusual amount of PFC pause frames in a given interval (Watchdog interval.)
2. Your hardware notifies the Watchdog module when too many PFC frames are received and traffic on the corresponding queues is halted for a time interval.

3. On receiving such notifications, the Watchdog module starts the shutdown timer and moves the queue state to wait-to-shutdown state.
4. At regular intervals during the shutdown interval, the queue is checked for PFC frames and if the traffic in the queue is stuck. If the traffic isn't stuck because the queue didn't receive any PFC frames, the queue moves back to the monitored state.
5. If the traffic is stuck for a longer time and the shutdown-timer expires, the queue switches to a drop state and the PFC Watchdog begins to drop all packets.
6. At regular intervals, the Watchdog checks the queue for PFC frames and whether the traffic in the queue is still stuck. If traffic is stuck in the queue as PFC packets keep arriving, the queue remains in the drop or shutdown state.
7. When the traffic's no longer stuck, the autorestore timer starts. At regular intervals, the module checks if the queue is stuck because of PFC frames.
8. If the queue receives PFC frames during the last autorestore interval, the auto-restore timer is reset upon expiry.
9. If the queue receives no PFC frames during the last autorestore interval, the Watchdog module restores the queue, and traffic resumes.

Related Topics

- [Priority Flow Control Overview, on page 117](#)

Configure a Priority Flow Control Watchdog Interval

You can configure PFC Watchdog parameters (Watchdog interval, shutdown multiplier, auto-restore multiplier) at the global or interface levels. Note that:

- When global Watchdog mode is disabled or off, Watchdog is disabled on all interfaces. This condition is regardless of the interface level Watchdog mode settings.
- When global Watchdog mode is enabled or on, the interface level Watchdog mode configuration settings override the global Watchdog mode values.
- When you configure interface level Watchdog attributes such as interval, shutdown multiplier, and auto-restore multiplier, they override the global Watchdog attributes.



Note Configuring the PFC mode and its policies is a prerequisite for PFC Watchdog.

Configuration Example

You can configure the Watchdog at the global or at the interface level.



Note Watchdog is enabled by default, with system default values of:

Watchdog interval = 100 ms

Shutdown multiplier = 1

Auto-restore multiplier = 10

```
P/0/RP0/CPU0:ios#show controllers hundredGigE 0/2/0/0 priority-flow-control
```

Priority flow control information for interface HundredGigE0/2/0/0:

Priority flow control watchdog configuration:

(D) : Default value

U : Unconfigured

Configuration Item	Global	Interface	Effective
PFC watchdog state :	U	U	Enabled (D)
Poll interval :	U	U	100 (D)
Shutdown multiplier :	U	U	1 (D)
Auto-restore multiplier :	U	U	10 (D)

```
RP/0/RP0/CPU0:ios#config
```

```
RP/0/RP0/CPU0:ios(config)#priority-flow-control watchdog mode off
```

```
RP/0/RP0/CPU0:ios(config)#commit
```

```
RP/0/RP0/CPU0:ios(config)#do show controllers hundredGigE 0/2/0/0 priority-flo$
```

Priority flow control information for interface HundredGigE0/2/0/0:

Priority flow control watchdog configuration:

(D) : Default value

U : Unconfigured

Configuration Item	Global	Interface	Effective
PFC watchdog state :	Disabled	U	Disabled
Poll interval :	U	U	100 (D)
Shutdown multiplier :	U	U	1 (D)
Auto-restore multiplier :	U	U	10 (D)

```
RP/0/RP0/CPU0:ios(config)#interface hundredGigE 0/2/0/0 priority-flow-control $
```

```
RP/0/RP0/CPU0:ios(config)#commit
```

```
RP/0/RP0/CPU0:ios(config)#do show controllers hundredGigE 0/2/0/0 priority-flo$
```

Priority flow control information for interface HundredGigE0/2/0/0:

Priority flow control watchdog configuration:

(D) : Default value

U : Unconfigured

Configuration Item	Global	Interface	Effective
PFC watchdog state :	Disabled	Enabled	Disabled
Poll interval :	U	U	100 (D)
Shutdown multiplier :	U	U	1 (D)
Auto-restore multiplier :	U	U	10 (D)


```
RP/0/RP0/CPU0:ios(config)#interface hundredGigE 0/2/0/1 priority-flow-control $
RP/0/RP0/CPU0:ios(config)#commit
```

```
RP/0/RP0/CPU0:ios(config)#do show controllers hundredGigE 0/2/0/1 priority-flo$
```

Priority flow control information for interface HundredGigE0/2/0/1:

Priority flow control watchdog configuration:

(D) : Default value

U: Unconfigured

Configuration Item	Global	Interface	Effective
PFC watchdog state :	Enabled	Disabled	Disabled
Poll interval :	U	U	100 (D)
Shutdown multiplier :	U	U	1 (D)
Auto-restore multiplier :	U	U	10 (D)

Verification

To verify that the PFC Watchdog is enabled globally, run the **sh run priority-flow-control watchdog mode** command.

```
Router#sh run priority-flow-control watchdog mode
priority-flow-control watchdog mode on
```

To verify your PFC Watchdog global configuration, run the **priority-flow-control watchdog** command.

```
Router#sh run priority-flow-control watchdog
priority-flow-control watchdog interval 100
priority-flow-control watchdog auto-restore-multiplier 2
priority-flow-control watchdog mode on
priority-flow-control watchdog shutdown-multiplier 2
```

Related Topics

- [Priority Flow Control Watchdog Overview, on page 140](#)

Global Statistics Counters for Priority Flow Control and Priority Flow Control Watchdog

Table 37: Feature History Table

Feature Name	Release Information	Feature Description
Global Statistics Counters for Priority Flow Control and Priority Flow Control Watchdog	Release 7.5.5 Release 24.2.11	<p>You can now view statistics for Priority Flow Control (PFC) and PFC Watchdog for all interfaces in a consolidated, compact, tabular, and easy-to-read format.</p> <p>We've also made the display of these global statistics faster by ensuring data is collected from all line cards for their interfaces and cumulatively sent to the local statistics infrastructure from where the show commands collect the data.</p> <p>Previously, you could view statistics only per interface for PFC and PFC Watchdog, wherein the show commands get the data from each interface.</p> <p>This feature modifies the following command:</p> <ul style="list-style-type: none"> • show controllers

Previously, you could view per interface statistics using the following command forms:

- PFC: **show controllers priority-flow-control statistics**
- PFC Watchdog: **show controllers priority-flow-control watchdog statistics**

With Release 7.5.5 (and from Release 24.2.11 onwards), you can view the statistics for all interfaces using the following command forms:

- PFC: **show controllers all priority-flow-control statistics location all**
- PFC Watchdog: **show controllers all priority-flow-control watchdog-stats location all**

Running the **clear controller priority-flow-control statistics** and **clear controller priority-flow-control watchdog-stats** commands also clears the counters for the corresponding interfaces that the **show controllers all priority-flow-control statistics location all** and **show controllers all priority-flow-control watchdog-stats location all** commands display.

View Global Statistics for Priority Flow Control and Priority Flow Control Watchdog

Run the **show controllers all priority-flow-control statistics location all** command to view the following PFC details for all interfaces:

- Interface: Interface name
- TC: Traffic Class

- Rx PFC: PFC frames received for the traffic class
- Tx PFC: PFC frames transmitted for the traffic class
- Rx Dropped: Received PFC frames that are dropped. (0 indicates no frames were dropped, while NA indicates the information is 'Not Available'— the line card hasn't sent the data.

Router#show controllers all priority-flow-control statistics location all

Interface	TC	Rx PFC	TxFPC	RxDropped
FourHundredGigE0/0/0/0	0	0	0	NA
FourHundredGigE0/0/0/0	7	0	0	NA
FourHundredGigE0/0/0/0	all	0	0	0
FourHundredGigE0/0/0/1	0	0	0	NA
FourHundredGigE0/0/0/1	7	0	0	NA
FourHundredGigE0/0/0/1	all	0	0	0
FourHundredGigE0/0/0/2	0	0	0	NA
FourHundredGigE0/0/0/2	7	0	0	NA
FourHundredGigE0/0/0/2	all	0	0	0
FourHundredGigE0/0/0/3	0	0	0	NA

Run the **show controllers all priority-flow-control watchdog-stats location all** command to view the following PFC Watchdog details for all interfaces:

- Interface: interface name
- TC: traffic class
- Watchdog events: the number of times the Watchdog module receives notifications that excess PFC frames are received.
- Shutdown events: The number of times the PFC Watchdog moves the queue to the shutdown state.
- Auto restore Events: The number of times the PFC Watchdog sets the auto-restore timer.
- Total Dropped: Cumulative VOQ drops + output drops in all the watchdog events from the time you ran the clear controllers priority-flow-control watchdog statistics command to clear the statistics counter.
- Dropped: Cumulative VOQ drops + output queue drops in the most recent Watchdog shutdown event.

Router#show controllers all priority-flow-control watchdog-stats location all

Interface	TC	Watchdog	Shutdown	Auto Restore	SAR	SAR Instantaneous
Total Dropped		Events	Events	Events	Events	Events
FourHundredGigE0/0/0/0	0	0	0	0	0	0
0						
FourHundredGigE0/0/0/0	7	0	0	0	0	0
0						
FourHundredGigE0/0/0/1	0	0	0	0	0	0
0						
FourHundredGigE0/0/0/1	7	0	0	0	0	0
0						
FourHundredGigE0/0/0/2	0	0	0	0	0	0
0						
FourHundredGigE0/0/0/2	7	0	0	0	0	0
0						
FourHundredGigE0/0/0/3	0	0	0	0	0	0
0						
FourHundredGigE0/0/0/3	7	0	0	0	0	0
0						

```
FourHundredGigE0/0/0/4      0      0      0      0      0      0
0                               0
```



Note Disregard the SAR Events and SAR Instantaneous Events entries because those numbers have no bearing on your operations.

PFC Watchdog Statistics to Monitor Traffic Drops

Table 38: Feature History Table

Feature Name	Release Information	Feature Description
PFC Watchdog Statistics to Monitor Traffic Drops	Release 7.5.4	<p>We have enhanced the PFC Watchdog counters to ensure you get an accurate view of the packets dropped across queues. You now get separate statistic counters for the number of packets dropped and the total number of packets dropped.</p> <p>By providing deeper insights into Watchdog packet drops, we aim to help you plan your traffic flows better and avoid potential PFC storms.</p> <p>Earlier releases didn't support counters for packet drops.</p> <p>You can view the counters using:</p> <ul style="list-style-type: none"> • YANG data model (at Github under the 753 folder): Cisco-IOS-XR-ofa-npu-pfc-oper • CLI: <code>show controllers priority-flow-control watchdog-stats</code> command form.

Here's how the PFC Watchdog works when it detects that a queue has stalled:

1. It begins dropping packets that arrive at the stalled queue. (See [Priority Flow Control Watchdog Overview, on page 140](#) for details.)
2. After the congestion clears and no more PFC frames arrive at the queue, the Watchdog restores the queue and flushes all packets that are stuck in the Virtual Output Queue (VOQ) and the output queue.

This functionality with an enhanced statistics view provides two new counters for all the dropped packets.

PFC Watchdog Statistics for Two Counters

With Release 7.5.4, the PFC Watchdog statistics account for two counters:

- Dropped Packets = Cumulative VOQ drops + output queue drops in the most recent Watchdog shutdown event
- Total Dropped Packets = Cumulative VOQ drops + output drops in all the watchdog events from the time you ran the **clear controller priority-flow-control watchdog-stats** command to clear the statistics counter.

Thus, when you run the **show controllers priority-flow-control watchdog-stats** command, the output displays the statistics for dropped packets and total dropped packets.

PFC Watchdog Statistics Benefits

- With the additional insights the two counters provide, you can study traffic drop patterns across longer timelines.
- You can analyze traffic drops for every congestion recovery that a queue makes.
- With access to such deep analysis, you can plan your traffic flows better. You can adjust your traffic flow rates and prevent potential PFC storms based on previous drop patterns and trends.

View PFC Watchdog Statistics

```
Router#show controllers hundredGigE 0/1/0/43 priority-flow-control watchdog-stats
```

Priority flow control information for interface HundredGigE0/1/0/43:

Priority flow control watchdog statistics:

SAR: Auto restore and shutdown

Traffic Class	:	0	1	2	3	4	5	6
7								
Watchdog Events	:	0	0	0	3	3	0	0
0								
Shutdown Events	:	0	0	0	3	3	0	0
0								
Auto Restore Events	:	0	0	0	3	3	0	0
0								
SAR Events	:	0	0	0	3510	3510	0	0
0								
SAR Instantaneous Events	:	0	0	0	1172	1172	0	0
0								
Total Dropped Packets	:	0	0	0	941505767	941488166	0	0
0								
Dropped Packets	:	0	0	0	314855466	314887161	0	0
0								



Note Disregard the **SAR Events** and **SAR Instantaneous Events** entries because those numbers have no bearing on your operations.

Detect High Bandwidth Memory Congestion

Table 39: Feature History Table

Feature Name	Release Information	Feature Description
Detect High Bandwidth Memory Congestion	Release 7.5.3	<p>We provide detailed insights into congestion on the High Bandwidth Memory (HBM), such as the devices on which congestion has occurred, the time stamps, and when the device returned to its normal state. With such details, you can investigate the cause of congestion and identify the source ports causing congestion for future preventive actions.</p> <p>You must configure PFC in the buffer-extended mode for this option.</p> <p>The feature introduces the following to enable the option to detect HBM congestion:</p> <ul style="list-style-type: none"> • YANG data model (at Github under the 753 folder): Cisco-IOS-XR-um-8000-hw-module-profile-cfg • CLI: hw-module profile npu memory buffer-extended location bandwidth-congestion-detection enable <p>It also introduces the following to view the congestion and memory usage details:</p> <ul style="list-style-type: none"> • YANG data model (at Github under the 753 folder): Cisco-IOS-XR-8000-platforms-npu-memory-oper • CLI: show controllers npu packet-memory

Here's how congestion build-up in queues drives the usage of memory on your router:

1. Under normal queue conditions, packets egressing from an interface are enqueued into the Shared Memory System (SMS) or internal buffer memory.
2. When congestion occurs, the SMS remains the buffer for packets until the congestion exceeds the per-queue usage criteria. This usage criterion is a combination of the threshold buffer space per VOQ and the age of the packet (defined in units of milliseconds).

3. At this stage, the packets are evicted to the High Bandwidth Memory (HBM) or external buffer memory. This provides additional buffer memory to absorb short bursts of congestion.
4. However, when packets from multiple VOQs are evicted to the HBM, the HBM bandwidth begins to experience congestion. The HBM congestion causes data loss for lossless traffic classes.

It's for insights into these extreme events that we've provided the option to detect congestion in the HBM. This information is meant for post-event analysis and reporting. You can enable this functionality by configuring the command **hw-module profile npu memory buffer-extended bandwidth-congestion-detection enable**. By doing so, you can view information such as:

- Devices on which congestion has occurred and the time stamps
- The current buffer memory usage and the highest memory usage watermark reached since the last reading

Benefits of Detecting High Bandwidth Memory Congestion

When you enable the ability to detect HBM congestion, the **show controllers npu packet-memory** command displays:

- Details of congestion along with the affected devices.
- A snapshot of the maximum memory usage with time stamps.

Using this information, you can conduct post-event analysis such as:

- investigate the cause of the congestion.
- identify the ports and flows that are the source of this congestion.

High Bandwidth Memory Congestion Detection: Guidelines and Limitations

- You must configure PFC in the buffer-extended mode for this functionality.
- There's no requirement for a line card reload after you configure **hw-module profile npu memory buffer-extended bandwidth-congestion-detection enable**.
- This functionality is supported on:
 - Cisco Silicon One Q200-based routers and line cards
 - 8201-32FH routers
 - 88-LC0-48TH-MO line cards
 - 88-LC0-36FH line cards
 - 88-LC0-36FH-M line cards

Configure High Bandwidth Memory Congestion Detection

To enable the detection of HBM congestion, configure the **hw-module profile npu memory buffer-extended bandwidth-congestion-detection enable** command.

```

Router#config
Router(config)#hw-module profile npu buffer-extended location 0/6/CPU0
bandwidth-congestion-detection enable
Router(config)#commit
Router(config)#exit

```



Note There's no requirement to reload the line card after the configuration.

Verification

The following table lists the various verification commands you can run depending on your required information. The outputs for these commands are available after the table.

Table 40: Verification Commands

Type of Information	Commands	Details
Buffer usage	<ul style="list-style-type: none"> show controllers npu packet-memory usage instance all location all show controllers npu packet-memory usage verbose instance all location all 	<ul style="list-style-type: none"> These commands provide data for the current use and the highest watermark reached since the last reading for both SMS and HBM. The refresh interval for the information is 30 seconds.
HBM bandwidth congestion	<ul style="list-style-type: none"> show controllers npu packet-memory congestion instance all location all show controllers npu packet-memory congestion detail instance all location all show controllers npu packet-memory congestion verbose instance all location all 	<ul style="list-style-type: none"> These commands provide data for when the HBM congestion occurred or is about to happen. The output maintains a history of the last 120 events regardless of the elapsed time. The refresh interval for new events to be added is 30 seconds.

Run the **show controllers npu packet-memory usage instance all location all** command to view the following details:

- the timestamp at which data is sampled
- the network processor name (**Device**)
- packet memory usage for that timestamp for SMS (**Buff-int Usage** in units of buffers) and HBM (**Buff-ext Usage** in units of 8 KB blocks)
- highest maximum watermark reached for SMS (**Buff-int Max WM**) and HBM (**Buff-ext Max WM**) since the last reading


```
Router#show controllers npu packet-memory usage instance all location all
HW memory Information For Location: 0/6/CPU0
```

Timestamp(msec)	Device	Buff-int Usage	Buff-int Max WM	Buff-ext Usage	Buff-ext Max WM
Wed 2022-09-21 01:54:11.154 UTC	0	7	8	0	0
Wed 2022-09-21 01:54:12.154 UTC	0	7	8	0	0
Wed 2022-09-21 01:54:24.023 UTC	1	22	22	0	0
Wed 2022-09-21 01:54:34.088 UTC	2	11	12	0	0
Wed 2022-09-21 01:54:35.088 UTC	2	11	12	0	0
Wed 2022-09-21 01:54:36.088 UTC	2	11	12	0	0
Wed 2022-09-21 01:54:37.088 UTC	2	11	12	0	0
Wed 2022-09-21 01:54:38.089 UTC	2	11	12	0	0
Wed 2022-09-21 01:54:39.089 UTC	2	11	12	0	0
Wed 2022-09-21 01:54:40.089 UTC	2	11	12	0	0

In addition, the **show controllers npu packet-memory usage verbose instance all location all** command displays timestamp in milliseconds:

```
RP/0/RP0/CPU0:Router#show controllers npu packet-memory usage verbose instance all location all
```

```
HW memory Information For Location: 0/RP0/CPU0
```

```
* Option 'verbose' formatted data is for internal consumption.
```

Timestamp(msec)	Device	Buff-int Usage	Buff-int Max WM	Buff-ext Usage	Buff-ext Max WM
1663958881006	0	2455	2676	637	640
1663958882007	0	2461	2703	635	640
1663958883007	0	2364	2690	635	640
1663958884007	0	71603	75325	3183	18336
1663958885008	0	2458	2852	1275	1279
1663958886008	0	2484	2827	1275	1279

Run the **show controllers npu packet-memory congestion instance all location all** command to view if HBM congestion has occurred and the timestamp of the congestion state.

- Each row indicates:
 - the timestamp at which data is sampled
 - the network processor name (**Device**)
 - the event type (**Normal** or **Congest**)
- The data displayed is for the last 120 events, and new events get added every 30 seconds. To view the updated data, re-run the command.
- After 120 events, the latest entry replaces the oldest entry. You can't clear the events from the list.

```
Router#show controllers npu packet-memory congestion instance all location all
HW memory Information For Location: 0/6/CPU0
```

Timestamp(msec)	Device	Buff-ext Event Type
Wed 2022-09-21 02:14:41.709 UTC	1	Congest

```

Wed 2022-09-21 02:14:41.959 UTC          Congest      1
Wed 2022-09-21 02:14:42.960 UTC          Congest      1
Wed 2022-09-21 02:14:43.960 UTC          Congest      1
Wed 2022-09-21 02:14:45.210 UTC          Congest      1
Wed 2022-09-21 02:14:45.710 UTC          Congest      1
Wed 2022-09-21 02:14:47.711 UTC          Normal       1

```

Run the **show controllers npu packet-memory congestion detail instance all location all** command to view the following details:

- the timestamp at which data is sampled
- the event type (**Normal** or **Congest**)
- the network processor name (**Device**) and the slice number (**Slice**) for that device. Every network processor has a fixed number of slices, and each slice, in turn, has a set number of ports.
- single VOQ buffer and aggregated SMS VOQ buffers
- packet memory usage for that timestamp for SMS (**Buff-int Usage** in units of buffers) and HBM (**Buff-ext Usage** in units of 8 KB blocks)
- highest maximum watermark reached for SMS (**Buff-int Max WM**) and HBM (**Buff-ext Max WM**) since the last reading

```
RRP/0/RP0/CPU0:ios#show controllers npu packet-memory congestion detail instance all location all
```

```
Fri Sep 23 18:49:50.640 UTC
```

```
HW memory information for location: 0/RP0/CPU0
```

* Option 'detail' formatted data is for internal consumption.

Timestamp (msec)	Evicted-buff	Buff-int	Buff-int	Buff-ext	Device	Slice	VOQ	VOQ-buff
	int-WM	UC-WM	Usage	Event Type	Usage	Max WM		int-WM
Fri 2022-09-23 18:42:30.349 UTC				Congest		0	5	16011
63969		65451	70410		34405	34405	534	
Fri 2022-09-23 18:42:31.101 UTC				Normal		0	5	0
0		900	2440		0	0	534	
Fri 2022-09-23 18:42:37.354 UTC				Congest		0	5	16011
63984		65493	70573		34408	34408	534	
Fri 2022-09-23 18:42:38.354 UTC				Normal		0	5	0
0		915	2455		0	0	534	
Fri 2022-09-23 18:42:44.606 UTC				Congest		0	5	16011
64002		65520	70081		34532	34532	534	

```
... ..
```

Run the **show controllers npu packet-memory congestion verbose instance all location all** command to view the following details:

- timestamp in milliseconds when the data is sampled
- the network processor name (**Device**) and the slice number (**Slice**) for that device. Every network processor has a fixed number of slices, and each slice, in turn, has a set number of ports
- event type, where 0 = single VOQ-based congestion and 1 = single VOQ-based congestion backoff (**VOQ-buff int-WM**), 2 = congestion in aggregated SMS buffers for VOQ and 3 = congestion backoff in aggregated SMS buffers for VOQ (**Evicted-buff int-WM**)
- buffer internal for unicast, which is for information only.

- packet memory usage for that timestamp for SMS (**Buff-int Usage** in units of buffers) and HBM (**Buff-ext Usage** in units of 8 KB blocks)
- highest maximum watermark reached for SMS (**Buff-int Max WM**) and HBM (**Buff-ext Max WM**) since the last reading

Router#show controllers npu packet-memory congestion verbose instance all location all
 HW memory Information For Location: 0/RP0/CPU0

* Option 'verbose' formatted data is for internal consumption.

Timestamp(msec)	Event	Device	Slice	VOQ	VOQ-buff	Evicted-buff	Buff-int
Buff-int	Buff-int	Buff-ext	Buff-ext				
		Type					
Usage	Max WM	Usage	Max WM		int-WM	int-WM	UC-WM
1663958550349		0	0	5	534	16011	63969
70410	70410	34405	34405				
1663958551101		1	0	5	534	0	900
2440	2440	0	0				
1663958557354		0	0	5	534	16011	63984
70573	70573	34408	34408				
1663958558354		1	0	5	534	0	915
2455	2455	0	0				
1663958564606		0	0	5	534	16011	64002
70081	70081	34532	34532				
1663958565356		1	0	5	534	0	915
2417	2417	0	0				

Available Shared Memory System and High Bandwidth Memory

Table 41: Feature History Table

Feature Name	Release Information	Feature Description
Available Shared Memory System and High Bandwidth Memory Buffers	Release 7.5.5 Release 24.1.1	<p>You can now view buffer availability for Shared Memory System (SMS) and High Bandwidth Memory (HBM) with higher accuracy without any lag between the minimum and maximum watermark readings, especially when the packet buffers are used and released rapidly. This is possible because we've enabled the instantaneous display of available or free SMS and HBM.</p> <p>Previously, you could view details only for the highest watermark readings for SMS and HBM.</p> <p>You must configure PFC in the buffer-extended mode for this option, and this functionality is available only for Cisco Silicon One Q200-based routers and line cards.</p> <p>This functionality modifies the following:</p> <ul style="list-style-type: none"> • CLI: show controllers npu packet-memory • YANG Data Model: Cisco-IOS-XR-8000-platforms-npu-memory-oper

From Release 7.5.5Release 24.1.1 onwards, you can additionally view details of the buffer available for SMS (**Buff-int-free Min WM**) and HBM (**Buff-ext-free Min WM**) when you run the **show controllers npu packet-memory usage instance all location all** command.

These counters are accounted for as follows:

- Available SMS buffer (**Buff-int-free Min WM**) at a given instant= Maximum SMS buffer - highest maximum watermark reached for SMS (**Buff-int Max WM**)
- Available HBM buffer (**Buff-ext-free Min WM**) at a given instant = Maximum HBM buffer - highest maximum watermark reached for HBM (**Buff-ext Max WM**)

This additional detail, available only for Cisco Silicon One Q200-based routers and line cards, provides even better insights to analyze the congestion affecting buffer occupancy, especially when packet buffers are used and released quickly.

```
Router#show controller npu packet-memory usage instance all location all
HW memory Information For Location: 0/6/CPU0
```

Timestamp(msec)	Device	Buff-int	Buff-int	Buff-ext	Buff-ext	
Buff-int-free	Buff-ext-free	Usage	Max WM	Usage	Max WM	Min
WM	Min WM					
Wed 2023-08-30 23:47:40.918 UTC	0	1518	6668	17154	17656	

```

293394          982846
Wed 2023-08-30 23:47:41.918 UTC      0      1227      5631      16010      16427
293685          983990
Wed 2023-08-30 23:47:42.919 UTC      0      1398      8295      15041      15734
293514          984959
Wed 2023-08-30 23:47:43.919 UTC      0      1765      8892      14744      15678
293147          985256
Wed 2023-08-30 23:47:41.011 UTC      1      10380     12419     37532     38165
284532          962468
Wed 2023-08-30 23:47:42.011 UTC      1      10463     11977     37315     38326
284449          962685
Wed 2023-08-30 23:47:43.013 UTC      1        9145     12604     37714     38242
285767          962286
Wed 2023-08-30 23:47:44.013 UTC      1      10996     13272     37429     38051
283916          962571

```

Global Pause Frames for High Bandwidth Memory Congestion

Table 42: Feature History Table

Feature Name	Release Information	Feature Description
Global Pause Frames for High Bandwidth Memory Congestion	Release 7.5.4	<p>We ensure no packet drops on PFC-enabled queues due to High Bandwidth Memory (HBM) congestion. Such prevention of drops is possible because we have enabled the triggering of global pause frames (X-Off) whenever there's HBM congestion.</p> <p>This functionality is disabled by default. You have the following options to enable it:</p> <ul style="list-style-type: none"> • CLI: hw-module profile npu memory buffer-extended bandwidth-congestion-protect enable • YANG Data Model: Cisco-IOS-XR-um-8000-hw-module-profile-cfg (see GitHub, YANG Data Models Navigator) <p>This feature introduces the show hw-module bandwidth-congestion-protect command to view the status of the global X-Off configuration.</p>

When High Bandwidth Memory congestion occurs (see [Detect High Bandwidth Memory Congestion](#), on page 148 for details), global pause frames (X-Off) are triggered for all PFC-enabled queues, regardless of whether those queues are the aggressor queues, hence the name 'global.' The only queues that don't transmit the X-Off trigger are those that don't receive any traffic. Such action ensures no packet drops on lossless queues, allowing you to meet your traffic bandwidth commitments for specific customers and requirements. It also ensures that the X-Off isn't triggered prematurely, which would have affected uncongested queues as well, causing a drop in performance.

To enable this functionality, configure the **hw-module profile npu memory buffer-extended bandwidth-congestion-protect enable** command.

Global Pause Frames for High Bandwidth Memory Congestion: Guidelines and Limitations

- This functionality isn't supported for the buffer-extended mode where the devices are more than 0.5 km apart.
- Configuring the **hw-module profile npu memory buffer-extended bandwidth-congestion-protect enable** command for line cards where you've configured headroom values exceeding 6144000 bytes could result in a commit error or the feature not being enabled.
- You must reload the line card for the **hw-module profile npu memory buffer-extended bandwidth-congestion-protect enable** command to take effect.
- This functionality is supported on:
 - 88-LC0-36FH line cards
 - 88-LC0-36FH-M line cards

Configure Global Pause Frames for High Bandwidth Memory Congestion

To enable triggering of global pause frames (X-off) whenever there's HBM congestion in the buffer-extended mode, configure the **hw-module profile npu memory buffer-extended bandwidth-congestion-protect enable** command.

```
Router#config
Router(config)#hw-module profile npu memory buffer-extended location 0/1/CPU0
bandwidth-congestion-protect enable
Router(config)#commit
```

Verification

Run the **show hw-module bandwidth-congestion-protect** command to view details about global X-off.


```
RP/0/RP1/CPU0:router#show hw-module bandwidth-congestion-protect location 0/1/CPU0
```

Location	Configured	Applied	Action
0/1/CPU0	Yes	No	Reload

The table lists the various possibilities for the command output based on your activity.

If you...	Configured field displays...	Applied field displays...	Action field displays...
Configure the hw-module profile npu memory buffer-extended command	Yes	No	Reload
Use the no form of the hw-module profile npu memory buffer-extended command after configuring it, but before reloading the line card	No	No	N/A

If you...	Configured field displays...	Applied field displays...	Action field displays...
Configure the hw-module profile npu memory buffer-extended command for a supported variant and reload the line card	Yes	Yes, Active Note Yes indicates that the configuration is programmed to the hardware, Active indicates that the global X-off functionality is active on the hardware.	N/A
Use the no form of the hw-module profile npu memory buffer-extended command when it is active, and commit the no form but don't reload the line card	No Note At this stage, the output displays the user action and not the hardware status.	No Note At this stage, the output displays the user action and not the hardware status.	Reload
Reload the line card after committing the no form of the hw-module profile npu memory buffer-extended command	No Note At this stage, the output displays the hardware status.	No Note At this stage, the output displays the hardware status.	N/A



CHAPTER 8

Congestion Management

- [Congestion Management Overview, on page 159](#)
- [Important Notes About Congestion Management, on page 159](#)
- [Ingress Traffic Management Model, on page 160](#)
- [Low Latency Queueing with Strict Priority Queueing, on page 161](#)
- [Committed Bursts, on page 175](#)
- [Excess Bursts, on page 176](#)
- [Two-Rate Policer Details, on page 176](#)
- [References for Modular QoS Management, on page 177](#)

Congestion Management Overview

Congestion management features allow you to control congestion by determining the order in which a traffic flow (or packets) is sent out an interface based on priorities assigned to packets. Congestion management entails the creation of queues, assignment of packets to those queues based on the classification of the packet, and scheduling of the packets in a queue for transmission.

The types of traffic regulation mechanisms supported are:

- [Low Latency Queueing with Strict Priority Queueing, on page 161](#)
- [Traffic Shaping, on page 167](#)
- [Traffic Policing, on page 169](#)

Important Notes About Congestion Management

The following points are applicable to the line card 88-LC1-36EH:

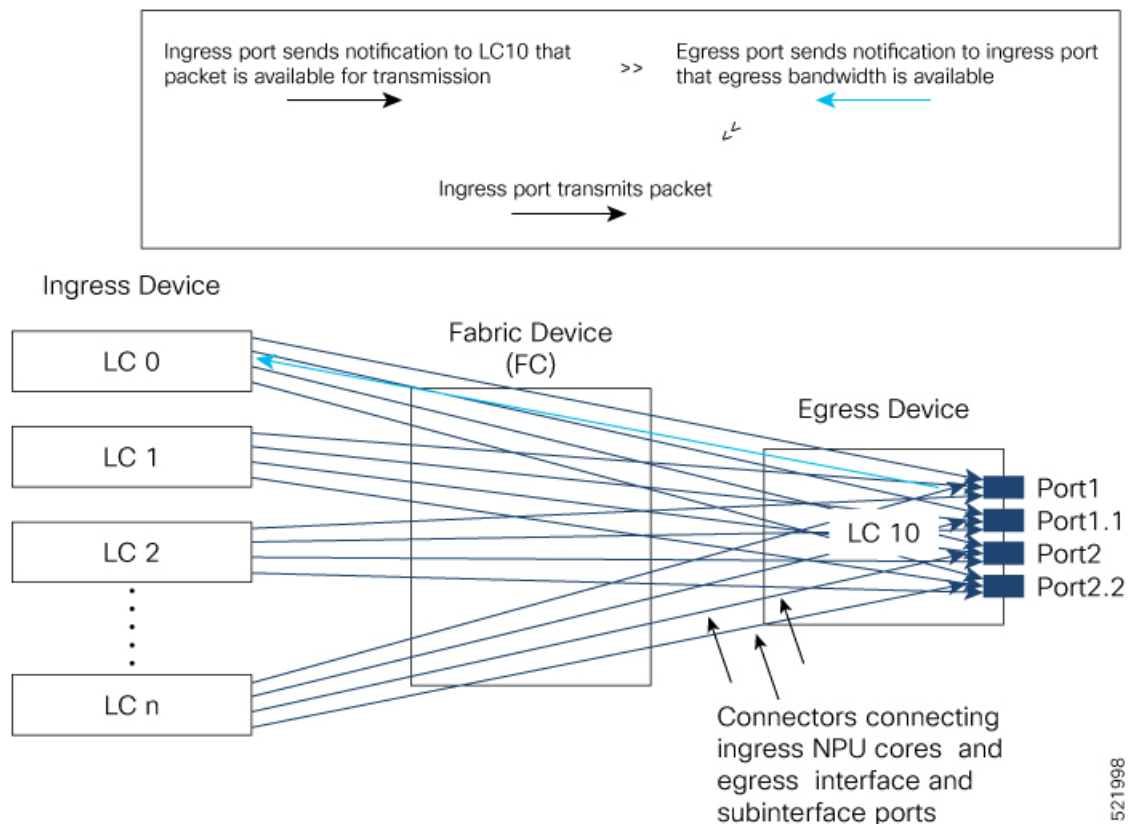
- The default scheduling hierarchy operates in strict-priority mode, with traffic prioritized on classes TC7 > TC6 > TC5 > TC4 > TC3 > TC2 > TC1 > TC0. You can change this default behaviour of strict-priority mode by applying an egress queuing policy map.
- Priority flow control (PFC), explicit congestion notification (ECN), and 4 VOQ mode aren't supposed.
- Egress queuing policy map on the subinterface (physical subinterface or bundle subinterface) isn't supported.

Ingress Traffic Management Model

The ingress traffic management model relies on packet queueing on the egress interface using Virtual Output Queueing (VOQ) on the ingress. In this model, buffering takes place at ingress. Here's how the VOQ process works.

Your routers support up to eight output queues per main interface or physical port. For every egress output queue, the VOQ model earmarks buffer space on every ingress pipeline. This buffer space is in the form of dedicated VOQs. These queues are called virtual because the queues physically exist on the ingress interface only when the line card actually has packets enqueued to it. To support the modular model of packet distribution, each network processing unit (NPU) core at the ingress needs connectors to every egress main interface and subinterface. The ingress traffic management model thus requires a mesh of connectors to connect the ingress NPU cores to the egress interfaces, as shown in **The Ingress Traffic Management Model**.

Figure 10: The Ingress Traffic Management Model



In the figure, every ingress interface (LC 0 through LC n) port has eight VOQs for the single egress line card LC 10.

Here's how packet transmission takes place:

1. When a packet arrives at ingress port (say on LC 0), the forwarding lookup on ingress line card points to the egress interface. Based on the egress interface (say it is on LC10), the packet is enqueued to the VOQ of LC 10. The egress interface is always mapped to a physical port.

2. Once egress bandwidth is available, the LC 10 ports ready to receive the packets (based on the packet marking and distribution model) send grants to the ingress ports via the connectors. (The figure shows a separate line for the grant for the sake of visual representation. In reality, the same connector is used for requests, grants, and transmission between an NPU core at the ingress and the egress port on LC 10.)
3. The ingress ports respond to this permission by transmitting the packets via FC to the LC 10 ports. (The time it takes for the ingress ports to request for egress port access, the egress port to grant access, and the packet to travel across FC is the round-trip time.)

The VOQ model thus operates on the principle of storing excess packets in buffers at ingress until bandwidth becomes available. Based on the congestion that builds up and the configured threshold values, packets begin to drop at the ingress itself, instead of having to travel all the way to the egress interface and then getting dropped.

Hardware Limitation:

In a scale scenario where 1000+ VoQs (created using egress QoS policies) store packets due to active traffic flows and may consume all the available on-chip buffer (OCB), unexpected traffic drops will be seen even though the traffic rate at the VoQ level is less than that of the VoQ shaper.

Low Latency Queueing with Strict Priority Queueing

Table 43: Feature History Table

Feature Name	Release Information	Feature Description
Low Latency Queueing with Strict Priority Queueing	Release 24.4.1	<p>Introduced in this release on: Fixed Systems (8700) (select variants only*)</p> <p>This feature ensures that time-sensitive traffic receives dedicated priority, ensuring minimal latency and jitter while congestion for other types of traffic is managed fairly, guaranteeing high performance for critical applications and efficient overall traffic management.</p> <p>*This functionality is now supported on Cisco 8712-MOD-M routers.</p>

Feature Name	Release Information	Feature Description
Priority Propagation on Egress Queues	Release 24.3.1	<p>Introduced in this release on: Fixed Systems (8200, 8700); Centralized Systems (8600); Modular Systems (8800 [LC ASIC: P100])</p> <p>You can now ensure that high-priority traffic is consistently prioritized across multiple sub-interfaces on the same network port, enhancing the performance of critical applications.</p> <p>This is achieved by enabling support for egress queuing policy maps that allow high-priority P1 traffic (on class TC7) to take precedence over non-P1 traffic across sub-interfaces. This feature ensures that critical traffic, such as voice or real-time video, is always given the highest priority, regardless of the sub-interface it is associated with, thereby maintaining optimal network performance and reducing latency for essential services.</p>

Feature Name	Release Information	Feature Description
Low Latency Queueing with Strict Priority Queueing	Release 24.3.1	<p>Introduced in this release on: Modular Systems (8800 [LC ASIC: P100]) (select variants only*), Fixed Systems (8200) (select variants only*), Fixed Systems (8700 (P100, K100)) (select variants only*)</p> <p>This feature ensures that time-sensitive traffic receives dedicated priority, ensuring minimal latency and jitter while congestion for other types of traffic is managed fairly, guaranteeing high performance for critical applications and efficient overall traffic management.</p> <p>*This feature is supported on:</p> <ul style="list-style-type: none">• 88-LC1-12TH24FH-E• 88-LC1-52Y8H-EM• 8212-48FH-M• 8711-32FH-M

Feature Name	Release Information	Feature Description
Low Latency Queueing with Strict Priority Queueing	Release 24.2.1	<p>Introduced in this release on: Modular Systems (8800 [LC ASIC: P100])(select variants only*)</p> <p>With this feature, time-sensitive traffic receives dedicated priority, ensuring minimal latency and jitter while congestion for other types of traffic is managed fairly, guaranteeing high performance for critical applications and efficient overall traffic management.</p> <p>You can also configure low latency queueing with strict priority queueing at the subinterface level, ensuring that the strict priority policy applies only to traffic on that specific subinterface, not the entire physical interface. This configuration is useful in scenarios where multiple services share the same physical link but have different QoS requirements.</p> <p>*This feature is supported on 88-LC1-36EH.</p>

The Low-Latency Queueing (LLQ) feature brings strict priority queueing (PQ) to the CBWFQ scheduling mechanism. Priority Queueing (PQ) in strict priority mode ensures that one type of traffic is sent, possibly at the expense of all others. For PQ, a low-priority queue can be detrimentally affected, and, in the worst case, never allowed to send its packets if a limited amount of bandwidth is available or the transmission rate of critical traffic is high. Strict PQ allows delay-sensitive data, such as voice, to be de-queued and sent before packets in other queues are de-queued.

What is the priority propagation on egress queues functionality all about?

The egress scheduling hierarchy operates in strict-priority mode by default, which means traffic on classes TC7 > TC6 > TC5 > TC4 > TC3 > TC2 > TC1 > TC0. This default behaviour of strict priority mode can be changed to the required configuration by applying the egress queuing policy map.

Configure Low Latency Queueing with Strict Priority Queueing

Configuring low latency queueing (LLQ) with strict priority queueing (PQ) allows delay-sensitive data such as voice to be de-queued and sent before the packets in other queues are de-queued.

Guidelines

- Only priority level 1 to 7 is supported, with 1 being the highest priority and 7 being the lowest. However, the default CoSQ 0 has the lowest priority among all.

- Egress policing is not supported. Hence, in the case of strict priority queueing, there are chances that the other queues do not get serviced.
- You can configure **shape average** and **queue-limit** commands along with **priority**.
- You can configure an optional shaper to cap the maximum traffic rate. This is to ensure the lower priority traffic classes are not starved of bandwidth.

Configuration Example

You have to accomplish the following to complete the LLQ with strict priority queueing:

1. Creating or modifying a policy-map that can be attached to one or more interfaces
2. Specifying the traffic class whose policy has to be created or changed
3. Specifying priority to the traffic class
4. (Optional) Shaping the traffic to a specific bit rate
5. Attaching the policy-map to an output interface

```
Router# configure
Router(config)# policy-map test-priority-1
Router(config-pmap)# class qos1
Router(config-pmap-c)# priority level 7
Router(config-pmap-c)# exit
Router(config-pmap)# exit
Router(config)# interface HundredGigE 0/6/0/18
Router(config-if)# service-policy output test-priority-1
Router(config-if)# no shutdown
Router(config-if)# commit
```

Running Configuration

```
policy-map strict-priority
class tc7
  priority level 1
  queue-limit 75 mbytes
!
class tc6
  priority level 2
  queue-limit 75 mbytes
!
class tc5
  priority level 3
  queue-limit 75 mbytes
!
class tc4
  priority level 4
  queue-limit 75 mbytes
!
class tc3
  priority level 5
  queue-limit 75 mbytes
!
class tc2
  priority level 6
  queue-limit 75 mbytes
```

```

!
class tc1
  priority level 7
  queue-limit 75 mbytes
!
class class-default
  queue-limit 75 mbytes
!
end-policy-map
!

class-map match-any tc1
match traffic-class 1
end-class-map
!
class-map match-any tc2
match traffic-class 2
end-class-map
!
class-map match-any tc3
match traffic-class 3
end-class-map
!
class-map match-any tc4
match traffic-class 4
end-class-map
!
class-map match-any tc5
match traffic-class 5
end-class-map
!
class-map match-any tc6
match traffic-class 6
end-class-map
!
class-map match-any tc7
match traffic-class 7
end-class-map
!

interface HundredGigE0/6/0/18
  service-policy input 100g-s1-1
  service-policy output test-priority-1
!

```

Verification

Router# **show qos int hundredGigE 0/6/0/18 output**

```

NOTE:- Configured values are displayed within parentheses
Interface HundredGigE0/6/0/18 ifh 0x3000220 -- output policy
NPU Id:                               3
Total number of classes:               3
Interface Bandwidth:                   100000000 kbps
VOQ Base:                             11176
VOQ Stats Handle:                     0x88550ea0
Accounting Type:                       Layer1 (Include Layer 1 encapsulation and above)
-----
Level1 Class (HP7)                     =   qos-1
Egress Queue ID                        =  11177 (HP7 queue)

```



```

TailDrop Threshold                = 125304832 bytes / 10 ms (default)
WRED not configured for this class

Levell Class (HP6)                = qos-2
Egressq Queue ID                  = 11178 (HP6 queue)
TailDrop Threshold                = 125304832 bytes / 10 ms (default)
WRED not configured for this class

Levell Class                      = class-default
Egressq Queue ID                  = 11176 (Default LP queue)
Queue Max. BW.                   = 101803495 kbps (default)
Queue Min. BW.                   = 0 kbps (default)
Inverse Weight / Weight          = 1 (BWR not configured)
TailDrop Threshold                = 1253376 bytes / 10 ms (default)
WRED not configured for this class

```

Related Topics

- [Congestion Management Overview, on page 159](#)
- [Traffic Shaping, on page 167](#)

Traffic Shaping

Traffic shaping allows you to control the traffic flow exiting an interface to match its transmission to the speed of the remote target interface and ensure that the traffic conforms to policies contracted for it. Traffic adhering to a particular profile can be shaped to meet downstream requirements, thereby eliminating bottlenecks in topologies with data-rate mismatches.



Note Traffic shaping is supported only in egress direction.

Configure Traffic Shaping

The traffic shaping performed on outgoing interfaces is done at the Layer 1 level and includes the Layer 1 header in the rate calculation.

Guidelines

- Only egress traffic shaping is supported.
- It is mandatory to configure all the eight qos-group classes (including class-default) for the egress policies.
- You can configure **shape average** command along with **priority** command.
- It is recommended that you configure all the eight <traffic-class classes> (including **class-default**) for the egress policies. A limited number of < traffic-class class> combinations are supported, but unicast and multicast traffic-class may not be handled consistently. Hence, such configurations are recommended, when the port egress has only unicast traffic.

Configuration Example

You have to accomplish the following to complete the traffic shaping configuration:

1. Creating or modifying a policy-map that can be attached to one or more interfaces
2. Specifying the traffic class whose policy has to be created or changed
3. Shaping the traffic to a specific bit rate
4. Attaching the policy-map to an output interface

```
Router# configure
Router(config)# policy-map egress_policy1
Router(config-pmap)# class c5
Router(config-pmap-c)# shape average percent 40
Router(config-pmap-c)# exit
Router(config-pmap)# exit
Router(config)# interface HundredGigE 0/1/0/0
Router(config-if)# service-policy output egress_policy1
Router(config-if)# commit
```

Running Configuration

```
policy-map egress_policy1
  class c5
    shape average percent 40
  !
  class class-default
  !
end-policy-map
!

interface HundredGigE0/6/0/18
  service-policy input 100g-s1-1
  service-policy output egress_policy1
!
```

Verification

```
Router# show qos interface hundredGigE 0/6/0/18 output
```

```
NOTE:- Configured values are displayed within parentheses
Interface HundredGigE0/6/0/18 ifh 0x3000220 -- output policy
NPU Id: 3
Total number of classes: 2
Interface Bandwidth: 100000000 kbps
VOQ Base: 11176
VOQ Stats Handle: 0x88550ea0
Accounting Type: Layer1 (Include Layer 1 encapsulation and above)
-----
Level1 Class = c5
Egressq Queue ID = 11177 (LP queue)
Queue Max. BW. = 40329846 kbps (40 %)
Queue Min. BW. = 0 kbps (default)
Inverse Weight / Weight = 1 (BWR not configured)
Guaranteed service rate = 40000000 kbps
TailDrop Threshold = 50069504 bytes / 10 ms (default)
WRED not configured for this class

Level1 Class = class-default
```

```

Egressq Queue ID           = 11176 (Default LP queue)
Queue Max. BW.             = 101803495 kbps (default)
Queue Min. BW.             = 0 kbps (default)
Inverse Weight / Weight    = 1 (BWR not configured)
Guaranteed service rate     = 50000000 kbps
TailDrop Threshold         = 62652416 bytes / 10 ms (default)
WRED not configured for this class

```

Related Topics

- [Congestion Management Overview, on page 159](#)

Egress class-level traffic shaping for subinterfaces

Table 44: Feature History Table

Feature Name	Release Information	Feature Description
Egress class-level traffic shaping for subinterfaces	Release 24.3.1	<p>Introduced in this release on: Modular Systems (8800 [LC ASIC: P100]) (select variants only*)</p> <p>You can now granularly control the traffic flow on each class in the egress queuing policy map ensuring optimal bandwidth allocation and improved network performance.</p> <p>This is achieved by introducing support for enabling class-level rate limiting traffic shapers for each class in the egress subinterface-level (main and LAG) queuing policy maps.</p> <p>*This feature is enabled by default and supported on:</p> <ul style="list-style-type: none"> • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM

- **What is this functionality all about?**—The subinterface queuing egress policy map can have class-level rate limiters (traffic shapers) on each class. This helps to granularly control the traffic flow on each class in the queuing policy map.

Traffic Policing

Traffic policing 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 class of service (CoS). Traffic policing manages the maximum rate of traffic through a token bucket algorithm. The token bucket algorithm uses user-configured values to determine the maximum rate of traffic allowed on an interface at a given moment in time. The token bucket

algorithm is affected by all traffic entering or leaving the interface (depending on where the traffic policy with traffic policing is configured) and is useful in managing network bandwidth in cases where several large packets are sent in the same traffic stream. By default, the configured bandwidth value takes into account the Layer 2 encapsulation that is applied to traffic leaving the interface.

The router supports the following traffic policing mode:

- Single-Rate Two-Color (SR2C) in color-blind mode. See [Single-Rate Policer, on page 170](#).

Restrictions

- 1R3C policers are not supported.
- Up to eight policers are supported per ingress policy.
- Policers are allocated in multiples of 2, so any request for allocating odd number of policers is internally rounded up by a factor of one.
- Only one conditional marking action is supported. You can set discard class to 0/1, that is used to set either virtual output queueing (VOQ) limits or Random Early Detection (RED) profiles.
- If you configure discard-class explicitly at the class level and not under policer, then the explicit mark action is applied to all the transmitted policer traffic.
- Egress policing is not supported.

Committed Bursts and Excess Bursts

Unlike a traffic shaper, a traffic policer does not buffer excess packets and transmit them later. Instead, the policer executes a “send or do not send” policy without buffering. Policing uses normal or committed burst (bc) values and excess burst values (be) to ensure that the router reaches the configured committed information rate (CIR). Policing decides if a packet conforms or exceeds the CIR based on the burst values you configure. Burst parameters are based on a generic buffering rule for routers, which recommends that you configure buffering to be equal to the round-trip time bit-rate to accommodate the outstanding TCP windows of all connections in times of congestion. During periods of congestion, proper configuration of the excess burst parameter enables the policer to drop packets less aggressively.

For more details, see [Committed Bursts, on page 175](#) and [Excess Bursts, on page 176](#).

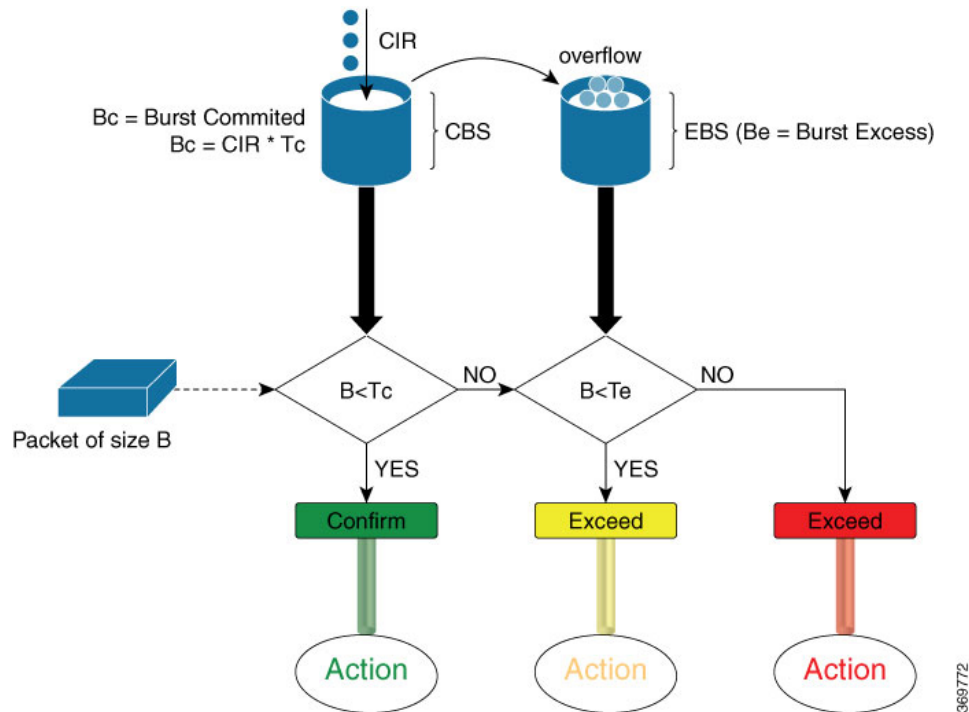
Single-Rate Policer

This section explains the concept of the single-rate two-color policer.

Single-Rate Two-Color Policer

A single-rate two-color (1R2C) policer provides one token bucket with two actions for each packet: a conform action and an exceed action.

Figure 11: Workflow of Single-Rate Two-Color Policy



Based on the committed information rate (CIR) value, the token bucket is updated at every refresh time interval. The Tc token bucket can contain up to the Bc value, which can be a certain number of bytes or a period of time. If a packet of size B is greater than the Tc token bucket, then the packet exceeds the CIR value and a default action is performed. If a packet of size B is less than the Tc token bucket, then the packet conforms and a different default action is performed.

Configure Traffic Policing (Single-Rate Two-Color)

Traffic policing is often configured on interfaces at the edge of a network to limit the rate of traffic entering or leaving the network. The default conform action for single-rate two color policer is to transmit the packet and the default exceed action is to drop the packet. Users cannot modify these default actions.

Configuration Example

You have to accomplish the following to complete the traffic policing configuration:

1. Creating or modifying a policy-map that can be attached to one or more interfaces
2. Specifying the traffic class whose policy has to be created or changed
3. (Optional) Specifying the marking action
4. Specifying the policy rate for the traffic
5. Attaching the policy-map to an input interface

```
Router# configure
Router(config)# policy-map test-police-1R2C
Router(config-pmap)# class dscp1
```

```

Router(config-pmap-c)# police rate 10 gbps
Router(config-pmap-c-police)# exit
Router(config-pmap-c)# exit
Router(config-pmap)# exit
Router(config)# interface hundredGigE 0/0/0/18
Router(config-if)# service-policy input test-police-1R2C
Router(config-if)# commit

```

Running Configuration

```

class-map match-any dscp1
  match dscp ipv4 1
end-class-map
!
!
policy-map test-police-1R2C
  class dscp1
    police rate 10 gbps
  !
  !
  class class-default
  !
  !
end-policy-map
!
!
interface HundredGigE0/0/0/8
service-policy input test-police-1R2C
!

```

Verification

```

Router# show qos interface hundredGigE 0/0/0/8 input
NOTE:- Configured values are displayed within parentheses
Interface HundredGigE0/0/0/8 ifh 0xf0001e8 -- input policy
NPU Id: 0
Total number of classes: 2
Interface Bandwidth: 100000000 kbps
Policy Name: test-police-1R2C
Accounting Type: Layer1 (Include Layer 1 encapsulation and above)
-----
Level1 Class = dscp1
Policer committed rate = 10000000 kbps (10 gbits/sec)
Policer conform burst = 1024000 bytes (default)
Policer conform action = Just TX
Policer exceed action = DROP PKT

Level1 Class = class-default
Policer not configured for this class

```

Related Topics

- [Traffic Policing, on page 169](#)


Two-Rate Policer

The two-rate policer manages the maximum rate of traffic by using two token buckets: the committed token bucket and the peak token bucket. The dual-token bucket algorithm uses user-configured values to determine the maximum rate of traffic allowed on a queue at a given moment. In this way, the two-rate policer can meter traffic at two independent rates: the committed information rate (CIR) and the peak information rate (PIR).

The dual-token bucket algorithm provides users with three actions for each packet—a conform action, an exceed action, and an optional violate action. Traffic entering a queue with the two-rate policer configured is placed into one of these categories. The actions are pre-determined for each category. The default conform and exceed actions are to transmit the packet, and the default violate action is to drop the packet.

This figure shows how the two-rate policer marks a packet and assigns a corresponding action to the packet.

Figure 12: Marking Packets and Assigning Actions—Two-Rate Policer



Also, see [Two-Rate Policer Details](#), on page 176.

The router supports Two-Rate Three-Color (2R3C) policer.

Configure Traffic Policing (Two-Rate Three-Color)

The default conform and exceed actions for two-rate three-color (2R3C) policer are to transmit the packet and the default violate action is to drop the packet. Users cannot modify these default actions.

Configuration Example

You have to accomplish the following to complete the two-rate three-color traffic policing configuration:

1. Creating or modifying a policy-map that can be attached to one or more interfaces
2. Specifying the traffic class whose policy has to be created or changed
3. Specifying the packet marking

4. Configuring two rate traffic policing
5. Attaching the policy-map to an input interface

```

Router# configure
Router(config)# policy-map test-police-2R3C
Router(config-pmap)# class dscp1
Router(config-pmap-c)# police rate 10 gbps peak-rate 20 gbps
Router(config-pmap-c-police)# exit
Router(config-pmap-c)# exit
Router(config-pmap)# exit
Router(config)# interface hundredGigE 0/0/0/8
Router(config-if)# service-policy input test-police-2R3C
Router(config-if)# commit

```

Running Configuration

```

class-map match-any dscp1
  match dscp ipv4 1
end-class-map
!
!
policy-map test-police-2R3C
  class dscp1
    police rate 10 gbps peak-rate 20 gbps
  !
  !
  class class-default
  !
  !
end-policy-map
!
!
interface HundredGigE0/0/0/8
service-policy input test-police-2R3C
!

```

Verification

```
Router# show qos interface hundredGigE 0/0/0/8 input
```

```

NOTE:- Configured values are displayed within parentheses
Interface HundredGigE0/0/0/8 ifh 0xf0001e8 -- input policy
NPU Id: 0
Total number of classes: 2
Interface Bandwidth: 100000000 kbps
Policy Name: test-police-2R3C
Accounting Type: Layer1 (Include Layer 1 encapsulation and above)
-----
Level1 Class = dscp1
Policer committed rate = 10000000 kbps (10 gbits/sec)
Policer peak rate = 20000000 kbps (20 gbits/sec)
Policer conform burst = 1024000 bytes (default)
Policer exceed burst = 2048000 bytes (default)
Policer conform action = Just TX
Policer exceed action = DROP PKT
Policer violate action = DROP PKT

```



```

Level1 Class                               = class-default
Policer not configured for this class

Router# policy-map interface hundredGigE 0/0/0/8 input

HundredGigE0/0/0/8 input: test-police-2R3C

Class dscp1
  Classification statistics                (packets/bytes)      (rate - kbps)
  Matched                                :      289228439/289228439000      27734775
  Transmitted                            :      56422213/56422213000      5410359
  Total Dropped                          :      232806226/232806226000      22324416
  Policing statistics                    (packets/bytes)      (rate - kbps)
  Policed(conform)                       :      56422213/56422213000      5410359
  Policed(exceed)                        :      56422215/56422215000      5410358
  Policed(violate)                       :      176384011/176384011000      16914058
  Policed and dropped                    :      232806226/232806226000
Class class-default
  Classification statistics                (packets/bytes)      (rate - kbps)
  Matched                                :      61136620/61136620000           0
  Transmitted                            :      61136620/61136620000           0
  Total Dropped                          :              0/0              0
Policy Map Stats time: 1570155764000 [Local Time: 10/04/19 02:22:44.000]

```

Related Topics

- [Two-Rate Policer, on page 173](#)

Committed Bursts

The committed burst (bc) parameter of the police command implements the first, conforming (green) token bucket that the router uses to meter traffic. The bc parameter sets the size of this token bucket. Initially, the token bucket is full and the token count is equal to the committed burst size (CBS). Thereafter, the meter updates the token counts the number of times per second indicated by the committed information rate (CIR).

The following describes how the meter uses the conforming token bucket to send packets:

- If sufficient tokens are in the conforming token bucket when a packet arrives, the meter marks the packet green and decrements the conforming token count by the number of bytes of the packet.
- If there are insufficient tokens available in the conforming token bucket, the meter allows the traffic flow to borrow the tokens needed to send the packet. The meter checks the exceeding token bucket for the number of bytes of the packet. If the exceeding token bucket has a sufficient number of tokens available, the meter marks the packet.

Green and decrements the conforming token count down to the minimum value of 0.

Yellow, borrows the remaining tokens needed from the exceeding token bucket, and decrements the exceeding token count by the number of tokens borrowed down to the minimum value of 0.

- If an insufficient number of tokens is available, the meter marks the packet red and does not decrement either of the conforming or exceeding token counts.



Note When the meter marks a packet with a specific color, there must be a sufficient number of tokens of that color to accommodate the entire packet. Therefore, the volume of green packets is never smaller than the committed information rate (CIR) and committed burst size (CBS). Tokens of a given color are always used on packets of that color.

Excess Bursts

The excess burst (be) parameter of the police command implements the second, exceeding (yellow) token bucket that the router uses to meter traffic. The exceeding token bucket is initially full and the token count is equal to the excess burst size (EBS). Thereafter, the meter updates the token counts the number of times per second indicated by the committed information rate (CIR).

The following describes how the meter uses the exceeding token bucket to send packets:

- When the first token bucket (the conforming bucket) meets the committed burst size (CBS), the meter allows the traffic flow to borrow the tokens needed from the exceeding token bucket. The meter marks the packet yellow and then decrements the exceeding token bucket by the number of bytes of the packet.
- If the exceeding token bucket does not have the required tokens to borrow, the meter marks the packet red and does not decrement the conforming or the exceeding token bucket. Instead, the meter performs the exceed-action configured in the police command (for example, the policer drops the packets).

Important Points

- User configurable burst values—committed burst size (CBS) and excess burst size (EBS)—are not supported. Instead, they are derived using user configured rates—committed information rates (CIR) and peak information rates (PIR).
- The router supports two burst values: low (10 kilobytes) and high (1 megabyte). For a lower range of configured values (1.6 MBps to less than 1 GBps), the burst value is 10 KB. For a higher range of configured values (1 GBps to 400 GBps), the burst value is 1 MB.

Two-Rate Policer Details

The committed token bucket can hold bytes up to the size of the committed burst (bc) before overflowing. This token bucket holds the tokens that determine whether a packet conforms to or exceeds the CIR as the following describes:

- A traffic stream is conforming when the average number of bytes over time does not cause the committed token bucket to overflow. When this occurs, the token bucket algorithm marks the traffic stream green.
- A traffic stream is exceeding when it causes the committed token bucket to overflow into the peak token bucket. When this occurs, the token bucket algorithm marks the traffic stream yellow. The peak token bucket is filled as long as the traffic exceeds the police rate.

The peak token bucket can hold bytes up to the size of the peak burst (be) before overflowing. This token bucket holds the tokens that determine whether a packet violates the PIR. A traffic stream is violating when it causes the peak token bucket to overflow. When this occurs, the token bucket algorithm marks the traffic stream red.

For example, if a data stream with a rate of 250 kbps arrives at the two-rate policer, and the CIR is 100 kbps and the PIR is 200 kbps, the policer marks the packet in the following way:

- 100 kbps conforms to the rate
- 100 kbps exceeds the rate
- 50 kbps violates the rate

The router updates the tokens for both the committed and peak token buckets in the following way:

- The router updates the committed token bucket at the CIR value each time a packet arrives at the interface. The committed token bucket can contain up to the committed burst (bc) value.
- The router updates the peak token bucket at the PIR value each time a packet arrives at the interface. The peak token bucket can contain up to the peak burst (be) value.
- When an arriving packet conforms to the CIR, the router takes the conform action on the packet and decrements both the committed and peak token buckets by the number of bytes of the packet.
- When an arriving packet exceeds the CIR, the router takes the exceed action on the packet, decrements the committed token bucket by the number of bytes of the packet, and decrements the peak token bucket by the number of overflow bytes of the packet.
- When an arriving packet exceeds the PIR, the router takes the violate action on the packet, but does not decrement the peak token bucket.

See [Two-Rate Policer](#), on page 173.

References for Modular QoS Management

Read this section for more information on committed bursts, excess bursts, and two-rate policer.



CHAPTER 9

Configure Modular QoS on Link Bundles

- [QoS on Link Bundles, on page 179](#)

QoS on Link Bundles

A bundle is a group of one or more ports that are aggregated together and treated as a single link. The router supports Ethernet interfaces and VLAN interfaces (bundle sub-interfaces) bundles. All QoS features currently supported on physical interfaces, are also supported on all link bundle interfaces. Applying QoS on bundle members is not supported.

Restrictions for Link Bundles

- Only Ethernet link bundling is supported.
- Only relative (percentage-based) shaper values can be configured. You cannot configure absolute values
- A bundle interface can only contain physical interface.
- All links within a single bundle must be configured either to run 802.3ad (LACP) or EtherChannel (non-LACP). Mixed links within a single bundle are not supported.
- MAC accounting is not supported on Ethernet link bundles.
- Maximum number of links supported in each link bundle is 64.
- The maximum number of link bundles supported is 128.

Load Balancing

Load balancing function is a forwarding mechanism to distribute traffic over multiple links based on Layer 3 routing information in the router. Per-destination load balancing is only supported on the router, where the router is allowed to distribute packets over one of the links in the bundle. When the per-destination load balancing is enabled, all packets for a certain source-destination pair go through the same link, though there are multiple links available. In other words, per-destination load balancing can ensure that packets for a certain source-destination pair could arrive in order.

Layer 3 Load Balancing on Link Bundles

Layer 3 load balancing for link bundles is done on Ethernet Flow Points (EFPs) and is based on the IPv4 source and destination addresses in the packet. When Layer 3 service-specific load balancing is configured, all egress bundles are load balanced based on the IPv4 source and destination addresses. When packets do not have IPv4 addresses, default load-balancing (based on the MAC SA/DA fields in the packet header) is used.

Configure QoS on Link Bundles

QoS is configured on link bundles in the same way that it's configured on individual interfaces.

Guidelines

- When you apply a QoS policy on a bundle (ingress or egress direction), the queuing policy is applied at each member interface. The reference bandwidth that is used to calculate shaper or bandwidth values is applied as per the physical member interface bandwidth.
- If a QoS policy is not applied to a bundle interface, both the ingress and egress traffic use the default queue of the per link member port.
- The shape rate specified in the bundle policy-map is not an aggregate for all bundle members. The shape rate applied to the bundle depends on the load balancing of the links. For example, if a policy map with a shape rate of 10 Mbps is applied to a bundle with two member links, and if the traffic is always load-balanced to the same member link, then an overall rate of 10 Mbps applies to the bundle. However, if the traffic is load-balanced evenly between the two links, the overall shape rate for the bundle becomes 20 Mbps.
- If a member is deleted from a bundle, the total bundle statistics changes because the statistics that belongs to the detached link is lost.
- Add bundle members before you apply the QoS policy to the bundle. This action ensures that the QoS policy applies across the bundle and on the new member. If a QoS policy exists and you add a new member later, the policy may not apply on the bundle if the member interface has an incompatible configuration. This scenario may result in errors.
- **Queue-limit error:**
 - If you experience a queue-limit error, remove the existing QoS policy from the link bundle interfaces. However, this removal may take up to ten minutes to come into effect, depending on the time lapsed since the error displayed. To see the time lapsed, check the **retries performed** information in the subsequent syslog.
 - When you encounter an invalid queue-limit error, the syslog lists only one interface, though the error may occur across multiple interfaces. To know the number of affected interfaces, see **operations** in the syslog.

For example:

```
RP/0/RP0/CPU0:Aug 17 09:56:37.417 PDT: BM-DISTRIB[1156]:
%L2-BM-4-ERR_OP_RETRY_THRESHOLD :
Exceeded threshold for the number of retries for 4 operations on members.
First (and error) FortyGigE0/0/0/34 ('DPA_QOSEA' detected the 'warning'
condition 'Invalid Queue Limit.
Value should be in range 614400-390070272 bytes or equivalent value in usec or
msec based on the service rate of the queue'),
```

```
retries performed for 10 hrs 16 mins
```

```
RP/0/RP0/CPU0:Aug 17 10:06:40.388 PDT: BM-DISTRIB[1156]:
%L2-BM-4-ERR_OP_RETRY_THRESHOLD :
Exceeded threshold for the number of retries for 4 operations on members.
First (and error) FortyGigE0/0/0/34 ('DPA_QOSEA' detected the 'warning'
condition 'Invalid Queue Limit.
Value should be in range 614400-390070272 bytes or equivalent value in
usec or msec based on the service rate of the queue'),
retries performed for 10 hrs 26 mins
```

- The QoS policy applied on bundle is inherited to all its member links and the reference bandwidth used to calculate shaper/bandwidth is applied as per the physical member interface bandwidth, and not the bundle as a whole.
- Classification resources (TCAM/MAP) and counters are allocated per bundle interface, not per member.
- The policing policy can be applied on bundle main interface or sub interface, but not on both.
- Policer committed information rates (CIR) and peak information rates (PIR) can be configured only in percent units.
- Policers and policer counters are allocated per bundle interface but not per member.
- The queuing policy can be applied on bundle main or bundle sub interface, but not on both.
- Queuing resources (VoQs) and counters are allocated per bundle member interface.
- The classification and marking policies can be applied on the bundle main or the bundle sub interface, but not on both.

Configuration Example

You have to accomplish the following to complete the QoS configuration on link bundles:

1. Creating a class-map
2. Creating a policy-map and specifying the respective class-map
3. Specifying the action type for the traffic
Refer [Attach a Traffic Policy to an Interface, on page 27](#) for details on step 1, 2 and 3.
4. Creating a link bundle
5. Applying traffic policy to the link bundle

```
/* Configure Ether-Bundle and apply traffic policy */
Router(config)# interface Bundle-Ether 12000
Router(config-if)# mtu 9100
Router(config-if)# service-policy input ingress
Router(config-if)# service-policy output egress
Router(config-if)# ipv4 address 100.12.0.0 255.255.255.254
Router(config-if)# bundle maximum-active links 64
Router(config-if)# commit
```

Running Configuration

This example shows how a queuing policy is applied on an Ethernet link bundle. The policy is applied to all interfaces that are members of the Ethernet link bundle.

```
/* Policy-map */

policy-map ingress
  class inet4-classifier-af1
    set qos-group 1
  !
  class inet4-classifier-af2
    set qos-group 2
  !
  class inet4-classifier-af3
    set qos-group 3
  !
  class inet4-classifier-af4
    set qos-group 4
  !
  class inet4-classifier-be1
    set qos-group 5
  !
  class inet4-classifier-nc1
    set qos-group 6
  !
  class class-default
  !
end-policy-map
!

/* Ether Bundle */
interface Bundle-Ether12000
  mtu 9100
  service-policy input ingress
  service-policy output egress
  ipv4 address 100.12.0.0 255.255.255.254
  load-interval 30
!
```

Verification

- Verify that the bundle status is UP.

```
router# show bundle bundle-ether 1200
Wed Dec 16 19:55:49.974 PST

Bundle-Ether12000
  Status: Up
  Local links <active/standby/configured>: 35 / 0 / 35
  Local bandwidth <effective/available>: 35000000000 (35000000000) kbps
  MAC address (source): ea3b.745f.c4b0 (Chassis pool)
  Inter-chassis link: No
  Minimum active links / bandwidth: 1 / 1 kbps
  Maximum active links: 64
  Wait while timer: 2000 ms
  Load balancing: Default
  LACP: Operational
  Flap suppression timer: Off
  Cisco extensions: Disabled
```



```

Non-revertive: Disabled
mLACP: Not configured
IPv4 BFD: Not configured

```

Port	Device	State	Port ID	B/W, kbps
Hu0/4/0/0	Local	Active	0x8000, 0x0009	100000000
Link is Active				
Hu0/4/0/1	Local	Active	0x8000, 0x000a	100000000
Link is Active				
- - -				
- - -				
Hu0/4/0/35	Local	Active	0x8000, 0x002b	100000000
Link is Active				

- Verify the bundle statistics:

```
router# show policy-map interface bundle-ether 12000
```

```
Bundle-Ether12000 input: ingress
```

```

Class inet4-classifier-af1
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      : 4647401962/21236124455654      26403040
  Transmitted                   : 4647401962/21236124455654      26403040
  Total Dropped                 : 0/0                      0
Class inet4-classifier-af2
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      : 4502980177/20576584333939      25571493
  Transmitted                   : 4502980177/20576584333939      25571493
  Total Dropped                 : 0/0                      0
Class inet4-classifier-af3
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      : 4647404125/21236213667880      26389086
  Transmitted                   : 4647404125/21236213667880      26389086
  Total Dropped                 : 0/0                      0
Class inet4-classifier-af4
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      : 9291188840/42456120548683      52771168
  Transmitted                   : 9291188840/42456120548683      52771168
  Total Dropped                 : 0/0                      0
Class inet4-classifier-be1
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      : 4647413429/21235847852686      26393414
  Transmitted                   : 4647413429/21235847852686      26393414
  Total Dropped                 : 0/0                      0
Class inet4-classifier-nc1
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      : 9294887621/42473100149807      52778258
  Transmitted                   : 9294887621/42473100149807      52778258
  Total Dropped                 : 0/0                      0

Class class-default
  Classification statistics      (packets/bytes)      (rate - kbps)
  Matched                      : 0/0                      0
  Transmitted                   : 0/0                      0
  Total Dropped                 : 0/0                      0

```

```
Bundle-Ether12000 output: egress
```

```

Class c1
  Classification statistics      (packets/bytes)      (rate - kbps)

```

```

    Matched          :          16665494532/75878118942463      8760591
    Transmitted       :          16655834643/75834136022017      8760591
    Total Dropped     :          9659889/43982920446           0
  Queueing statistics
    Queue ID          : None (Bundle)
    Taildropped(packets/bytes) : 9659889/43982920446
Class c2
  Classification statistics          (packets/bytes)      (rate - kbps)
    Matched          :          16665421959/75877849543188      8718687
    Transmitted       :          16665421959/75877849543188      8718687
    Total Dropped     :              0/0                        0
  Queueing statistics
    Queue ID          : None (Bundle)
    Taildropped(packets/bytes) : 0/0
Class c3
  Classification statistics          (packets/bytes)      (rate - kbps)
    Matched          :          16665247833/75877509455458      8703470
    Transmitted       :          16665187414/75877234624197      8703470
    Total Dropped     :          60419/274831261                0
  Queueing statistics
    Queue ID          : None (Bundle)
    Taildropped(packets/bytes) : 60419/274831261
Class c4
  Classification statistics          (packets/bytes)      (rate - kbps)
    Matched          :          33330896131/151755393012945      17470745
    Transmitted       :          33330745421/151754709368565      17470745
    Total Dropped     :          150710/683644380                0
  Queueing statistics
    Queue ID          : None (Bundle)
    Taildropped(packets/bytes) : 150710/683644380
Class c5
  Classification statistics          (packets/bytes)      (rate - kbps)
    Matched          :          16878910340/76849791869834      8833394
    Transmitted       :          16878849464/76849514633309      8833394
    Total Dropped     :          60876/277236525                0
  Queueing statistics
    Queue ID          : None (Bundle)
    Taildropped(packets/bytes) : 60876/277236525
Class c6
  Classification statistics          (packets/bytes)      (rate - kbps)
    Matched          :          33330898844/151756094112925      17456785
    Transmitted       :          33330752668/151755427708382      17456785
    Total Dropped     :          146176/666404543                0
  Queueing statistics
    Queue ID          : None (Bundle)
    Taildropped(packets/bytes) : 146176/666404543
Class c7
  Classification statistics          (packets/bytes)      (rate - kbps)
    Matched          :          244106/79922040                74
    Transmitted       :          244106/79922040                74
    Total Dropped     :              0/0                        0
  Queueing statistics
    Queue ID          : None (Bundle)
    Taildropped(packets/bytes) : 0/0
Class class-default
  Classification statistics          (packets/bytes)      (rate - kbps)
    Matched          :          267075066180/1215993441123215      139917482
    Transmitted       :          267075066180/1215993441123215      139917482
    Total Dropped     :              0/0                        0
  Queueing statistics
    Queue ID          : None (Bundle)
    Taildropped(packets/bytes) : 0/0

```

Related Topics

- [QoS on Link Bundles, on page 179](#)

