

Cisco MDS 9000 Series Quality of Service Configuration Guide, Release 9.x

First Published: 2022-09-02

Last Modified: 2025-08-08

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

© 2021–2025 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface v

Audience v

Document Conventions v

Documentation Feedback vii

Related Documentation vii

Communications, Services, and Additional Information vii

CHAPTER 1

New and Changed Information 1

New and Changed Information 1

CHAPTER 2

QoS Overview 3

QoS 3

QoS in Differentiated Service 3

Applying QoS to Traffic 4

QoS Configuration 4

QoS Licensing 5

Port Tracking 5

CHAPTER 3

Configuring QoS 7

Information About Control Traffic 7

Enabling or Disabling Control Traffic 8

Displaying Control Traffic Information 8

Information About Data Traffic 9

Comparing VSAN Versus Zone-Based QoS 10

Configuring Data Traffic 10

QoS Initiation for Data Traffic 11

Information About Class Map Creation 11
Creating a Class Map 12
Information About Service Policy Definition 13
Specifying Service Policies 13
About Service Policy Enforcement 14
Applying Service Policies 14
About the DWRR Traffic Scheduler Queue 15
Changing the Weight in a DWRR Queue 15
Displaying Data Traffic Information Examples 15
Configuration Examples for QoS 17
Example: Traffic Prioritization 18
Example: Address Congestion 19
Static Ingress Port Rate Limiting 20

CHAPTER 4 **Configuring Priority Flow Control** 21
Information About Priority Flow Control 21
Licensing Requirements for Priority Flow Control 21
Guidelines and Limitations for Priority Flow Control 22
Default Settings for Priority Flow Control 22
Configuring Priority Flow Control 22

CHAPTER 5 **Configuring Long Distance** 25
Information About Long Distance 25
Guidelines and Limitations for Long Distance 26
Default Settings for Long-Distance 26
Configuring Long Distance 26



Preface

This preface describes the audience, organization of, and conventions used in the Cisco MDS 9000 Series Configuration Guides. It also provides information on how to obtain related documentation, and contains the following sections:

- Audience, on page v
- Document Conventions, on page v
- Documentation Feedback, on page vii
- Related Documentation, on page vii
- Communications, Services, and Additional Information, on page vii

Audience

This publication is for network administrators who install, configure, and maintain Cisco MDS 9000 Series Switches.

Document Conventions

Command descriptions use these conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
Italic	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element (keyword or argument).
[x | y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x | y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.

Convention	Description
[x {y | z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
`variable`	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

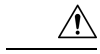
Examples use these conventions:

Convention	Description
`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
italic screen font	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:



Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.



Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.



Warning IMPORTANT SAFETY INSTRUCTIONS

This warning symbol means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device.

SAVE THESE INSTRUCTIONS

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to mds-docfeedback@cisco.com. We appreciate your feedback.

Related Documentation

The entire Cisco MDS 9000 Series switches documentation set is available at the following URL:

https://www.cisco.com/c/en/us/support/storage-networking/mds-9000-nx-os-san-os-software/series.html

Documentation Roadmap

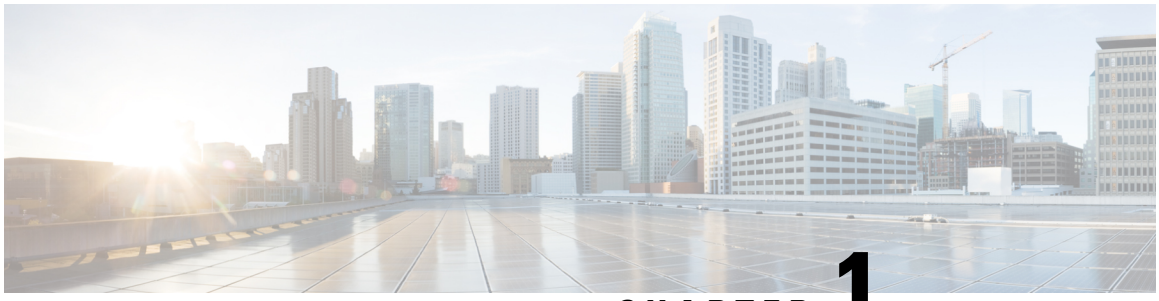
https://www.cisco.com/c/en/us/td/docs/storage/san_switches/mds9000/roadmaps/rel90.html

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

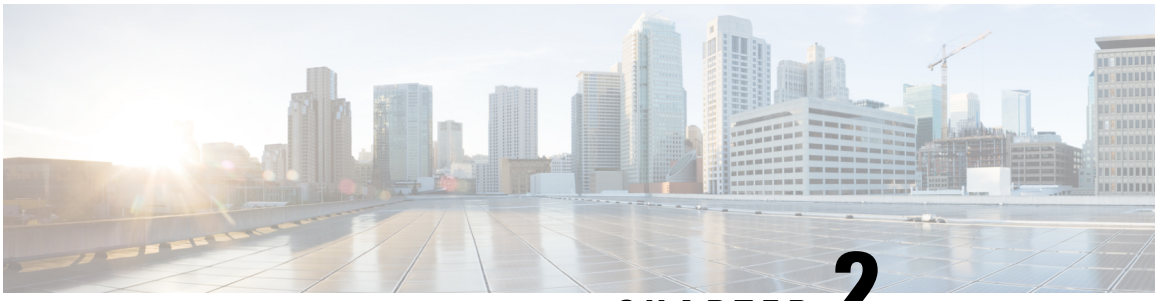
New and Changed Information

• New and Changed Information, on page 1

New and Changed Information

There are no new features in the Cisco MDS 9000 Series Quality of Service Configuration Guide for Cisco MDS NX-OS Release 9.x

New and Changed Information



CHAPTER 2

QoS Overview

The Cisco MDS 9000 NX-OS software offers traffic management features such as fabric-wide quality of service (QoS). These advanced capabilities are integrated into MDS 9000 Series switches to simplify deployment and to provide optimization of large-scale fabrics.

- QoS, on page 3
- Port Tracking, on page 5

QoS

QoS monitors the ability of a network to provide better service to selected network traffic over various underlying technologies including Frame Relay, Asynchronous Transfer Mode (ATM), Ethernet and 802.1 networks, SONET, and IP-routed networks. QoS features provide better and more predictable network service with these functions:

- Supporting dedicated bandwidth
- Improving loss characteristics
- Avoiding and managing network congestion
- Shaping network traffic
- Setting traffic priorities across the network

QoS-enabled switches provided traffic differentiation and prioritization, enabling latency-sensitive applications such as Online Transaction Processing (OLTP) to share common storage resources alongside throughput-intensive applications such as data warehousing.

QoS can be used alongside other traffic engineering features such as FCC and ingress port-rate limiting and can be configured to apply different policies at different times of day using the command scheduler built into Cisco MDS 9000 NX-OS software.

QoS in Differentiated Service

A service model, also called a level of service, describes a set of end-to-end QoS capabilities. End-to-end QoS is the ability of the network to deliver service required by specific network traffic from one end of the network to another.

The QoS implementation in the Cisco MDS 9000 Series switch follows the differentiated services (DiffServ) model.

Differentiated service is a multiple service model that can satisfy differing QoS requirements. However, unlike the integrated service model, an application using differentiated service does not explicitly signal the router before sending data.

For differentiated service, the network tries to deliver a particular kind of service based on the QoS specified by each packet. This specification can occur in different ways, for example, using the IP Precedence bit settings in IP packets or source and destination addresses. The network uses the QoS specification to classify, shape, and police traffic, and to perform intelligent queueing.

Applying QoS to Traffic

QoS provides service differentiation in the fabric by applying different service levels to different traffic. The service differentiation can perform the following operations:

- Provide relative bandwidth guarantees to application traffic
- Control latency experienced by application traffic
- Prioritize one application traffic over another

QoS is accomplished by combining traffic classification and Virtual Output Queuing (VOQ). Data traffic is classified at ingress ports as low, medium, or high priority. Classified frames are queued in the appropriate location based on the traffic type and QoS priority.

Traffic is processed based on how you classify it and the policies that you create and apply to traffic classes.

You can classify data traffic based on the following criterion:

- VSAN ID
- Source or destination N port WWN
- Fibre Channel ID (FCID)
- Zone

Four distinct QoS priority levels are available: three for Fibre Channel data traffic and one for Fibre Channel control traffic. Control traffic is assigned the highest QoS priority automatically to accelerate convergence of fabric-wide protocols such as Fabric Shortest Path First (FSPF), zone merges, and principal switch selection.

QoS requires FCC to be enabled in the fabric to provide the configured bandwidth guarantees.

QoS Configuration

QoS configuration should be consistent across multiple switches to help ensure that all switches are enforcing a common policy for traffic in both send and receive directions.

QoS is configured in an identical manner regardless of whether the switch has first generation, second generation, or third generation modules present. QoS can be deployed in any one of three ways depending on the complexity of the QoS policy desired:

- Virtual SAN (VSAN)-based QoS—VSAN-based QoS enables QoS priority to be assigned on a per-VSAN basis.
- Zone-based QoS—QoS priority can be assigned on a per-zone basis when a more granular QoS is required.
- Individual QoS policies matching individual devices—QoS policy can be defined on a per-device basis, with individual policies applied to different devices and VSANs when maximum flexibility is required.

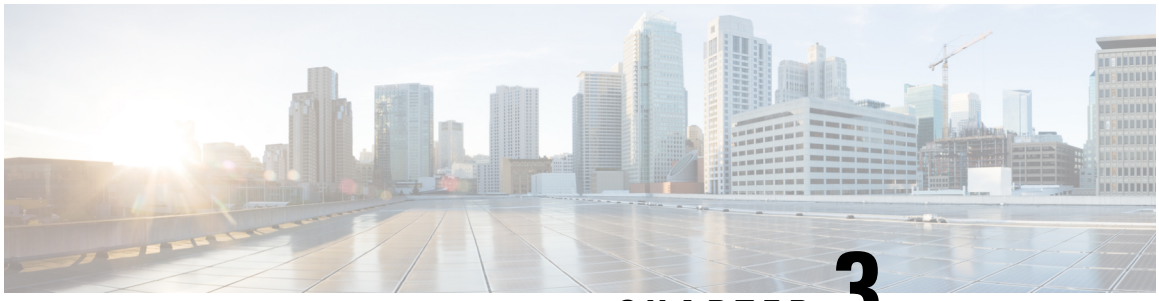
QoS Licensing

QoS is a licensed feature and requires an Enterprise Package license installed on all switches where QoS is enabled. However, you do not need a license to provide QoS for internally generated control traffic. You can also explicitly enable QoS by using the **qos enable** command.

Port Tracking

The port tracking feature in the Cisco MDS NX-OS software provides a resilient SAN extension.

If a switch detects a WAN or metropolitan-area network (MAN) link failure, it brings down the associated disk-array link when port tracking is configured. The array can redirect a failed I/O operation to another link without waiting for an I/O timeout. Otherwise, disk arrays must wait seconds for an I/O timeout to recover from a network link failure.



CHAPTER 3

Configuring QoS

This chapter provides details on the QoS features provided in all switches.

Quality of service (QoS) offers the following advantages:

- Provides relative bandwidth guarantee to application traffic.
- Controls latency experienced by application traffic.
- Prioritizes one application over another (for example, prioritizing transactional traffic over bulk traffic) through bandwidth and latency differentiation.

- Information About Control Traffic, on page 7
- Enabling or Disabling Control Traffic, on page 8
- Displaying Control Traffic Information, on page 8
- Information About Data Traffic, on page 9
- Comparing VSAN Versus Zone-Based QoS, on page 10
- Configuring Data Traffic, on page 10
- QoS Initiation for Data Traffic, on page 11
- Information About Class Map Creation, on page 11
- Creating a Class Map, on page 12
- Information About Service Policy Definition, on page 13
- Specifying Service Policies, on page 13
- About Service Policy Enforcement, on page 14
- Applying Service Policies, on page 14
- About the DWRR Traffic Scheduler Queue, on page 15
- Changing the Weight in a DWRR Queue, on page 15
- Displaying Data Traffic Information Examples, on page 15
- Configuration Examples for QoS, on page 17
- Static Ingress Port Rate Limiting, on page 20

Information About Control Traffic

The Cisco MDS 9000 Series supports QoS for internally and externally generated control traffic. Within a switch, control traffic is sourced to the supervisor module and is treated as a high priority frame. A high priority status provides absolute priority over all other traffic and is assigned in the following cases:

- Internally generated time-critical control traffic (mostly Class F frames).

- Externally generated time-critical control traffic entering a switch in the Cisco MDS 9000 Series from a another vendor's switch. High priority frames originating from other vendor switches are marked as high priority as they enter a switch in the Cisco MDS 9000 Series.

Enabling or Disabling Control Traffic

By default, the QoS feature for certain critical control traffic is enabled. These critical control frames are assigned the highest (absolute) priority.



Tip We do not recommend disabling this feature as all critical control traffic is automatically assigned the lowest priority once you issue this command.

To disable the high priority assignment for control traffic, follow these steps:

Procedure

Step 1 Enters configuration mode.

switch# **configure terminal**

Step 2 Enables the control traffic QoS feature.

switch(config)# **no qos control priority 0**

Step 3 Disables the control traffic QoS feature.

switch(config)# **qos control priority 0**

Displaying Control Traffic Information

Use the **show qos statistics** command to view the current state of the QoS configuration for critical control traffic. This command displays the current QoS settings along with the number of frames marked high priority. The count is only for debugging purposes and cannot be configured.

The following example displays Current QoS Settings

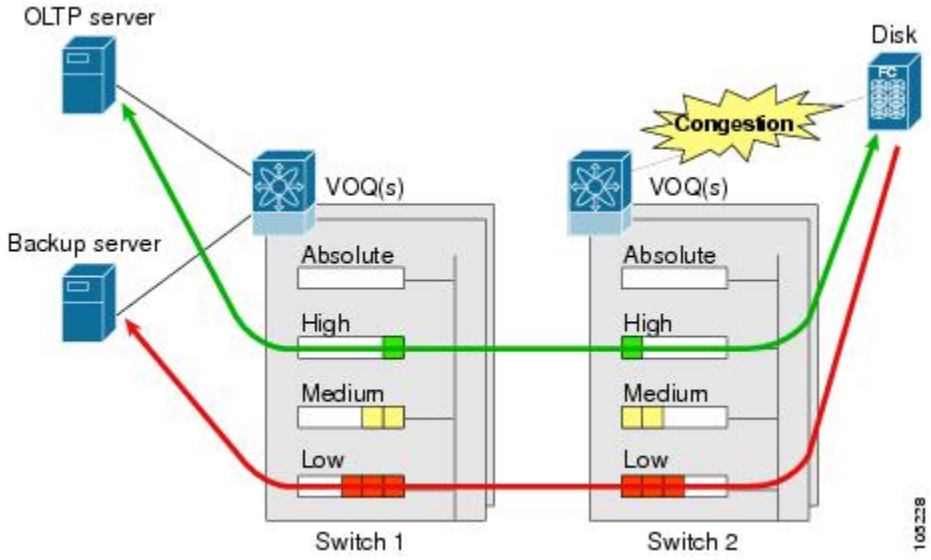
```
switch# show qos statistics
Total number of FC frames transmitted from the Supervisor= 15767
Number of highest-priority FC frames transmitted         = 8224
Current priority of FC control frames = 0    (0 = lowest; 7 = highest)
```

Information About Data Traffic

Online transaction processing (OLTP), which is a low volume, latency sensitive application, requires quick access to requested information. Backup processing application require high bandwidth but are not sensitive to latency. In a network that does not support service differentiation, all traffic is treated identically—they experience similar latency and are allocated similar bandwidths. The QoS feature in the Cisco MDS 9000 Series switches provides these guarantees.

Data traffic can be prioritized in distinct levels of service differentiation: low, medium, or high priority. You can apply QoS to ensure that Fibre Channel data traffic for your latency-sensitive applications receive higher priority over throughput-intensive applications such as data warehousing.

Figure 1: Prioritizing Data Traffic



In the above figure, the OLTP traffic arriving at Switch 1 is marked with a high priority level of throughput classification (class map) and marking (policy map). Similarly, the backup traffic is marked with a low priority level. The traffic is sent to the corresponding priority queue within a virtual output queue (VOQ).

A deficit weighted round robin (DWRR) scheduler configured in the first switch ensures that high priority traffic is treated better than low priority traffic. For example, DWRR weights of 70:20:10 implies that the high priority queue is serviced at 7 times the rate of the low priority queue. This guarantees lower delays and higher bandwidths to high priority traffic if congestion sets in. A similar configuration in the second switch ensures the same traffic treatment in the other direction.

If the ISL is congested when the OLTP server sends a request, the request is queued in the high priority queue and is serviced almost immediately since the high priority queue is not congested. The scheduler assigns its priority over the backup traffic in the low priority queue.

Note When the high priority queue does not have traffic flowing through, the low priority queue uses all the bandwidth and is not restricted to the configured value.

A similar occurrence in Switch 2 sends a response to the transaction request. The round trip delay experienced by the OLTP server is independent of the volume of low priority traffic or the ISL congestion. The backup traffic uses the available ISL bandwidth when it is not used by the OLTP traffic.

Comparing VSAN Versus Zone-Based QoS

While you can configure both zone-based QoS and VSAN-based QoS configurations in the same switch, both configurations have significant differences. The following table highlights the differences between configuring QoS priorities based on VSANs versus zones.

Table 1: QoS Configuration Differences

VSAN-Based QoS	Zone-Based QoS
If you configure the active zone set on a given VSAN and also configure QoS parameters in any of the member zones, you cannot associate the policy map with the VSAN.	You cannot activate a zone set on a VSAN that already has a policy map associated.
If the same flow is present in two class maps associated to a policy map, the QoS value of the class map attached first takes effect.	If the same flow is present in two zones in a given zone set with different QoS values, the higher QoS value is considered.
—	During a zone merge, if the Cisco NX-OS software detects a mismatch for the QoS parameter, the link is isolated.
Takes effect only when QoS is enabled.	Takes effect only when QoS is enabled.

Configuring Data Traffic

To configure QoS, follow these steps:

Procedure

Step 1 Enable the QoS feature.
Step 2 Create and define class maps.
Step 3 Define service policies.
Step 4 Apply the configuration.

QoS Initiation for Data Traffic

By default, the QoS data traffic feature is disabled for data traffic. To configure QoS for data traffic, you must first enable the data traffic feature in the switch.



Tip QoS is supported in interoperability mode. For more information, refer to the *Cisco MDS 9000 Series Switch-to-Switch Interoperability Configuration Guide* .

To enable the QoS data traffic feature, follow these steps:

Procedure

Step 1 Enters configuration mode.

switch# **configure terminal**

Step 2 Enables QoS. You can now configure data traffic parameters.

switch(config)# **qos enable**

Information About Class Map Creation

Use the class map feature to create and define a traffic class with match criteria to identify traffic belonging to that class. The class map name is restricted to 63 alphanumeric characters and defaults to the match-all option. Flow-based traffic uses one of the following values:

- WWN—The source WWN or the destination WWN.
- Fibre Channel ID (FC ID) —The source ID (SID) or the destination ID (DID). The possible values for mask are FFFFFF (the entire FC ID is used—this is the default), FFFF00 (only domain and area FC ID is used), or FF0000 (only domain FC ID is used).



Note An SID or DID of 0x000000 is not allowed.

- Source interface—The ingress interface.



Tip The order of entries to be matched within a class map is not significant.

Creating a Class Map

Use the **class-map** command to create and define a traffic class with match criteria to identify traffic belonging to that class. Define each match criterion with one match statement from the class map configuration (switch(config-cmap)) mode.

Note The enhanced mode for the **source-device-alias** or **destination-device-alias** option is not supported.

Note The QoS attribute with IVR zone set and VSAN is not supported.

- Use the **source-wwn** option to specify the source WWN or the **destination-wwn** option to specify the destination WWN.
- Use the **source-address** option to specify the source ID (SID) or the **destination-address** option to specify the destination ID (DID).
- Use the **input-interface** option to specify the ingress interface.
- Use the **destination-device-alias** option to specify the distributed device alias.

To create a class map, follow these steps:

Procedure

Step 1 Specifies a logical AND operator for all matching statements in this class. If a frame matches all (default) configured criteria, it qualifies for this class. This is the default.

switch(config)# **qos class-map MyClass match-all**

Step 2 Specifies a logical OR operator for all matching statements in this class. If a frame matches any one configured criteria, it qualifies for this class.

switch(config)# **qos class-map MyClass match-any**

Step 3 Specifies a destination address match for frames with the specified destination FC ID.

switch(config-cmap)# **match destination-address 0x12ee00**

Step 4 Specifies a source address and mask match for frames with the specified source FC ID.

switch(config-cmap)# **match source-address 0x6d1090 mask 0xFFFFFF**

Step 5 Specifies a destination WWN to match frames.

switch(config-cmap)# **match destination-wwn 20:01:00:05:30:00:28:df**

Step 6 Specifies a source WWN to match frames.

switch(config-cmap)# **match source-wwn 23:15:00:05:30:00:2a:1f**

Step 7 Specifies a destination device alias to match frames.

switch(config-cmap)# **match destination-device-alias DocDeviceAlias**

Step 8 Specifies a source device alias to match frames.

switch(config-cmap)# **match source-device-alias DocDeviceAliase**

Step 9 Specifies a source interface to match frames.

switch(config-cmap)# **match input-interface fc 2/1**

Step 10 Removes a match based on the specified source interface.

switch(config-cmap)# **no match input-interface fc 3/5**

Information About Service Policy Definition

Service policies are specified using policy maps. Policy maps provide an ordered mapping of class maps to service levels. You can specify multiple class maps within a policy map, and map a class map to a high, medium, or low service level. The default priority is low. The policy map name is restricted to 63 alphanumeric characters.

As an alternative, you can map a class map to a differentiated services code point (DSCP).The DSCP is an indicator of the service level for a specified frame. The DSCP value ranges from 0 to 63, and the default is 0. A DSCP value of 46 is disallowed.

The order of the class maps within a policy map is important to determine the order in which the frame is compared to class maps. The first matching class map has the corresponding priority marked in the frame.

Note For more information on implementing QoS DSCP values, see the following document: *Implementing Quality of Service Policies with DSCP*.

Note Class maps are processed in the order in which they are configured in each policy map.

Specifying Service Policies

To specify a service policy, follow these steps:

Procedure

Step 1 Creates a policy map called MyPolicy and places you in the policy-map submode.

switch(config)# **qos policy-map MyPolicy**

switch(config-pmap)#

Step 2 Deletes the policy map called OldPolicy and places you in the policy-map submode.

switch(config)# **no qos policy-map OldPolicy**

Step 3 Specifies the name of a predefined class and places you at the policy-map submode for that class.

switch(config-pmap)# **class MyClass**

switch(config-pmap-c)#

Step 4 Removes the class map called OldClass from the policy map.

switch(config-pmap)# **no class OldClass**

Step 5 Specifies the priority to be given for each frame matching this class.

switch(config-pmap-c)# **priority high**

Step 6 Deletes a previously assigned priority and reverts to the default value of low.

switch(config-pmap-c)# **no priority high**

Step 7 Specifies the DSCP value to mark each frame matching this class.

switch(config-pmap-c)# **dscp 2**

Step 8 Deletes a previously assigned DSCP value and reverts to the factory default of 0.

switch(config-pmap-c)# **no dscp 60**

About Service Policy Enforcement

When you have configured a QoS data traffic policy, you must enforce the data traffic configuration by applying that policy to the required VSAN(s). If you do not apply the policy to a VSAN, the data traffic configuration is not enforced. You can only apply one policy map to a VSAN.

Note You can apply the same policy to a range of VSANs.

Applying Service Policies

To apply a service policy, follow these steps:

Procedure

Step 1 Applies a configured policy to VSAN 3.

switch(config)# **qos service policy MyPolicy vsan 3**

Step 2 Deletes a configured policy that was applied to VSAN 7.

switch(config)# **no qos service policy OldPolicy vsan 7**

About the DWRR Traffic Scheduler Queue

The Cisco NX-OS software supports four scheduling queues:

- Strict priority queues are queues that are serviced in preference to other queues—it is always serviced if there is a frame queued in it regardless of the state of the other queues.
- QoS assigns all other traffic to the DWRR scheduling high, medium, and low priority traffic queues.

The DWRR scheduler services the queues in the ratio of the configured weights. Higher weights translate to proportionally higher bandwidth and lower latency. The default weights are 50 for the high queue, 30 for the medium queue, and 20 for the low queue. Decreasing order of queue weights is mandated to ensure the higher priority queues have a higher service level, though the ratio of the configured weights can vary (for example, one can configure 70:30:5 or 60:50:10 but not 50:70:10).



Note Generation 1 and Generation 2 modules are not supported from Cisco MDS NX-OS Release 6.x and later. Generation 3 and Generation 4 modules are not supported from Cisco MDS NX-OS Release 8.x and later.

For more information on the modules and ports supported, refer the Cisco MDS 9000 Series Interface Configuration Guide, Release 8.x guide.

Changing the Weight in a DWRR Queue

To associate a weight with a DWRR queue, follow these steps:

Procedure

Step 1 Associates a relative weight (10) to a specified queue (default queue).

switch(config)# **qos dwrr-q high weight 10**

Step 2 Restores the default weight of 20.

switch(config)# **no qos dwrr-q low weight 51**

Displaying Data Traffic Information Examples

The **show qos** commands display the current QoS settings for data traffic (see the following examples).

Example: Class Maps

The following example displays the Contents of all Class Maps

```
switch# show qos class-map
qos class-map MyClass match-any
    match destination-wwn 20:01:00:05:30:00:28:df
    match source-wwn 23:15:00:05:30:00:2a:1f
    match input-interface fc2/1
qos class-map Class2 match-all
    match input-interface fc2/14
qos class-map Class3 match-all
    match source-wwn 20:01:00:05:30:00:2a:1f
```

Example: Specified Class Map

The following example displays the Contents of a Specified Class Map:

```
switch# show qos class-map name MyClass
qos class-map MyClass match-any
    match destination-wwn 20:01:00:05:30:00:28:df
    match source-wwn 23:15:00:05:30:00:2a:1f
    match input-interface fc2/1
```

Example: All Configured Policy Maps

The following example displays All Configured Policy Maps:

```
switch# show qos policy-map
qos policy-map MyPolicy
    class MyClass
    priority medium
qos policy-map Policy1
    class Class2
    priority low
```

Example: Specified Policy Map

The following example displays a Specified Policy Map:

```
switch# show qos policy-map name MyPolicy
qos policy-map MyPolicy
    class MyClass
        priority medium
```

Example: Scheduled DWRR Configurations

The following example displays Scheduled DWRR Configurations:

```
switch# show qos dwrr
qos dwrr-q high weight 50
qos dwrr-q medium weight 30
qos dwrr-q low weight 20
```

Example: All Applied Policy Maps

The following example displays All Applied Policy Maps:

```
switch# show qos service policy
qos service policy MyPolicy vsan 1
qos service policy Policy1 vsan 4
```

Example: Policy Map Associated with a Specified VSAN

The following example displays the Policy Map Associated with a Specified VSAN:

```
switch# show qos service policy vsan 1
qos policy-map pmap1
   class cmap1
       priority medium
   class cmap2
       priority high
```

Example: Class Map Associated with a Specified Interface

The following example displays the Class Map Associated with a Specified Interface:

```
switch# show qos service policy interface fc3/10
qos policy-map pmap1
   class cmap3
       priority high
   class cmap4
       priority low
```

Example: QoS Statistics

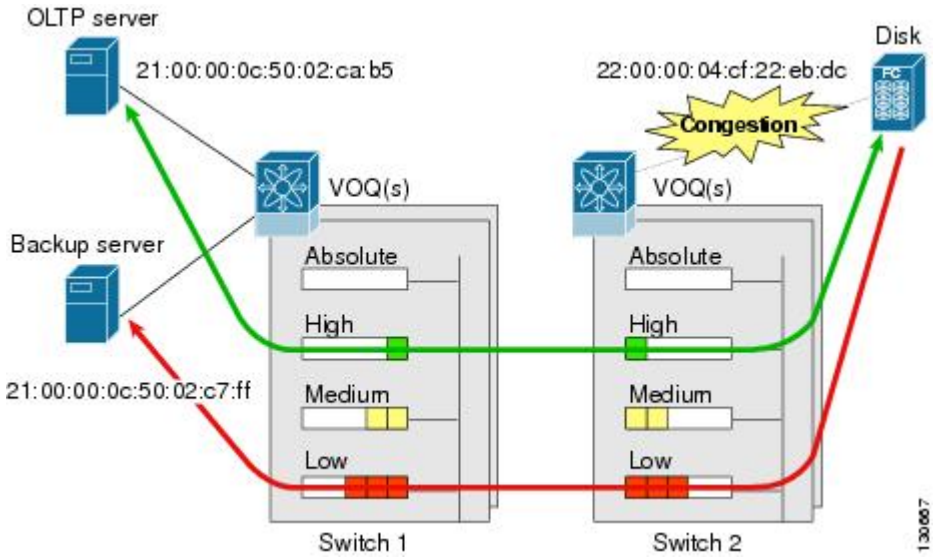
The following example displays QoS Statistics:

```
switch# show qos statistics
Total number of FC frames transmitted from the Supervisor= 301431
Number of highest-priority FC frames transmitted         = 137679
Current priority of FC control frames = 7    (0 = lowest; 7 = highest)
```

Configuration Examples for QoS

This section describes a configuration example for the application illustrated in the following figure.

Figure 2: Example Application for Traffic Prioritization



Both the OLTP server and the backup server are accessing the disk. The backup server is writing large amounts of data to the disk. This data does not require specific service guarantees. The volumes of data generated by the OLTP server to the disk are comparatively much lower but this traffic requires faster response because transaction processing is a low latency application.

The point of congestion is the link between Switch 2 and the disk, for traffic from the switch to the disk. The return path is largely uncongested as there is little backup traffic on this path.

Service differentiation is needed at Switch 2 to prioritize the OLTP-server-to-disk traffic higher than the backup-server-to-disk traffic.

Example: Traffic Prioritization

To configure traffic prioritization for the example application, follow these steps:

Procedure

Step 1 Create the class maps.

```
Switch 2# config t
Switch 2(config)# qos class-map jc1 match-all
Switch 2(config-cmap)# match source-wwn 21:00:00:0c:50:02:ca:b5
Switch 2(config-cmap)# match destination-wwn 22:00:00:04:cf:22:eb:dc
Switch 2(config-cmap)# exit
Switch 2(config)# qos class-map jc2 match-all
Switch 2(config-cmap)# match source-wwn 21:00:00:0c:50:02:c7:ff
Switch 2(config-cmap)# match destination-wwn 22:00:00:04:cf:22:eb:dc
Switch 2(config-cmap)# exit
Switch 2(config)#
```

Step 2 Create the policy map.

```
Switch 2(config)# qos policy-map jp1
Switch 2(config-pmap)# class jc1
Switch 2(config-pmap-c)# priority high
Switch 2(config-pmap-c)# exit
Switch 2(config-pmap)# class jc2
Switch 2(config-pmap-c)# priority low
Switch 2(config-pmap-c)# exit
Switch 2(config-pmap)# exit
Switch 2(config)#
```

Step 3 Assign the service policy.

```
Switch 2(config)# qos service policy jp1 vsan 1
```

Step 4 Assign the weights for the DWRR queues.

```
Switch 2(config)# qos dwrr-q high weight 50
Switch 2(config)# qos dwrr-q medium weight 30
Switch 2(config)# qos dwrr-q low weight 20
```

Step 5 Repeat Step 1 through Step 4 on Switch 1 to address forward path congestion at both switches.

Example: Address Congestion

Congestion could occur anywhere in the example configuration. To address congestion of the return path at both switches, you need to create two more class maps and include them in the policy map as follows:

Procedure

Step 1 Create two more class maps.

```
Switch 2(config)# qos class-map jc3 match-all
Switch 2(config-cmap)# match source-wwn 22:00:00:04:cf:22:eb:dc
Switch 2(config-cmap)# match destination-wwn 21:00:00:0c:50:02:ca:b5
Switch 2(config-cmap)# exit
Switch 2(config)# qos class-map jc4 match-all
Switch 2(config-cmap)# match source-wwn 22:00:00:04:cf:22:eb:dc
Switch 2(config-cmap)# match destination-wwn 21:00:00:0c:50:02:c7:ff
Switch 2(config-cmap)# exit
Switch 2(config)#
```

Step 2 Assign the class maps to the policy map.

```
Switch 2(config)# qos policy-map jp1
Switch 2(config-pmap)# class jc3
Switch 2(config-pmap-c)# priority high
Switch 2(config-pmap-c)# exit
Switch 2(config-pmap)# class jc4
```

```
Switch 2(config-pmap-c)# priority low
Switch 2(config-pmap-c)# exit
Switch 2(config-pmap)# exit
Switch 2(config)#
```

Step 3 Repeat Step 1 through Step 2 on Switch 1 to address return path congestion at both switches.

Static Ingress Port Rate Limiting

A static port rate limiting feature helps control the bandwidth for individual Fibre Channel ports using the **switchport ingress-rate** *limit* command. Port rate limiting is also referred to as ingress rate limiting because it controls ingress traffic into a Fibre Channel port. The feature controls traffic flow by slowing the rate of B2B credits transmitted from the FC port to the adjacent device. Port rate limiting works on all Fibre Channel ports. Prior to Cisco MDS NX-OS Release 8.5(1), the rate limit ranges from 1 to 100%. From Cisco MDS NX-OS Release 8.5(1), the limit ranges from 0.0126 to 100%. The default rate limit is 100%.

Starting from Cisco MDS NX-OS Release 8.5(1), the FPM feature needs to be configured before configuring the dynamic or static ingress port rate limiting feature on all Cisco MDS switches except Cisco MDS 9250i and MDS 9148S switches. Prior to Cisco MDS NX-OS Release 8.5(1) or on Cisco MDS 9250i and MDS 9148S switches, static ingress port rate limiting can be configured on all Cisco MDS switches and modules only if the QoS feature is enabled.

To configure the port rate limiting value, follow these steps:

Procedure

Step 1 Enters the configuration mode.

switch # **configure terminal**

switch(config)#

Step 2 Selects the interface to specify the ingress port rate limit.

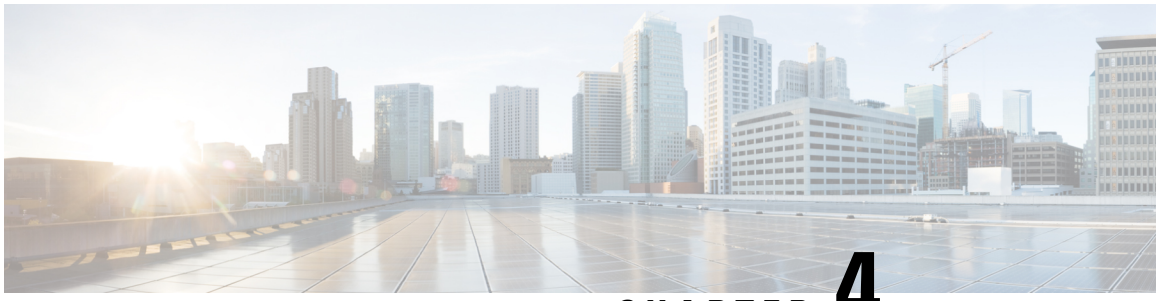
switch(config)# **interface fc 1/1**

Step 3 Configures a 50% port rate limit for the selected interface.

switch(config-if)# **switchport ingress-rate 50**

Step 4 Reverts a previously configured rate to the factory default of 100%.

switch(config-if)# **no switchport ingress-rate 50**



CHAPTER 4

Configuring Priority Flow Control

- Information About Priority Flow Control, on page 21
- Licensing Requirements for Priority Flow Control, on page 21
- Guidelines and Limitations for Priority Flow Control, on page 22
- Default Settings for Priority Flow Control, on page 22
- Configuring Priority Flow Control, on page 22

Information About Priority Flow Control

Priority flow control (PFC; IEEE 802.1bb), 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 802.3x Flow Control (pause frames) or link-level flow control (LFC). However, PFC functions on a per class-of-service (CoS) basis.

When a buffer threshold is exceeded due to congestion, LFC sends a pause frame to its peer to pause all data transmission on the link for a specified period of time. When the congestion is mitigated (traffic comes under the configured threshold), a resume frame is generated to restart data transmission on the link.

In contrast, during congestion, PFC sends a pause frame that indicates which CoS value needs to be paused. A PFC pause frame contains a 2-octet timer value for each CoS that indicates the length of time that the traffic needs to be paused. The unit of time for the timer is specified in pause quanta. A quanta is the time that is required for transmitting 512 bits at the speed of the port. The range is from 0 to 65535. A pause frame with a pause quanta of 0 indicates a resume frame to restart the paused traffic.



Note Only certain classes of service of traffic can be flow controlled while other classes are allowed to operate normally.

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 that is not forwarded when received by the peer. When the congestion is mitigated, PFC can request the peer to restart transmitting frames.

Licensing Requirements for Priority Flow Control

The following table shows the licensing requirements for this feature:

Product	License Requirement
Cisco MDS NX-OS	The QoS feature does not require a license. Any feature not included in a license package is bundled with the Cisco MDS NX-OS system images and is provided at no extra charge to you. For a complete explanation of the Cisco MDS NX-OS licensing scheme, see the *Cisco MDS NX-OS Licensing Guide*.

Guidelines and Limitations for Priority Flow Control

PFC has the following configuration guidelines and limitations:

- If PFC is enabled on a port or a port channel, it does not cause a port flap.
- A flap occurs when both the PFC and LFC are enabled and PFC is disabled before LFC is configured.
- PFC configuration enables PFC in both the send (Tx) and receive (Rx) direction.
- PFC on mode is used to support the hosts that support PFC but do not support the Data Center Bridging Capability Exchange Protocol (DCBXP).
- Only an exact match of the no-drop CoS is considered as a successful negotiation of PFC by the DCBXP.

Default Settings for Priority Flow Control

The following table lists the default setting for PFC:

Parameter	Parameter
PFC	Auto

Configuring Priority Flow Control

You can configure PFC one of these three modes with the long-distance feature:

- auto—Enables the no-drop CoS values to be advertised by the DCBXP and negotiated with the peer. A successful negotiation enables PFC on the no-drop CoS. Any failures because of a mismatch in the capability of peers causes the PFC not to be enabled.
- on—Enables PFC on the local port regardless of the capability of the peers.
- off—Disables PFC on the local port.



Note You can also enable Link-level Flow Control (LFC) on the same port in which PFC is enabled. However, PFC, if enabled, always gets the priority.

Procedure

Step 1 Enter global configuration mode:

switch# **configure terminal**

Step 2 Enter interface configuration mode:

switch(config)# **interface ethernet** *slot/ethernet*

Step 3 Set the PFC to auto, off, or on mode. By default, PFC mode is set to auto on all ports.

switch(config-if)# **priority-flow-control mode { auto | off | on }**

Step 4 Exit interface and global configuration modes:

switch(config-if)# **exit**

switch(config)# **exit**

Step 5 (Optional) Verify the status of PFC on all interfaces:

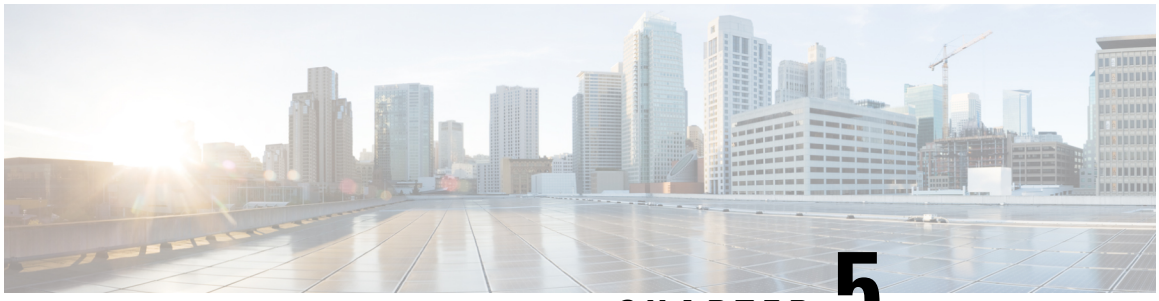
switch# **show interface priority-flow-control**

Example: Configuring Priority Flow Control

The following running configuration shows how to configure PFC. Replace the placeholders with relevant values for your setup.

```
configure terminal
 interface <ethernet 1/2>
  priority-flow-control mode on
  exit
 exit
```

Configuring Priority Flow Control



CHAPTER 5

Configuring Long Distance

- Information About Long Distance, on page 25
- Guidelines and Limitations for Long Distance, on page 26
- Default Settings for Long-Distance, on page 26
- Configuring Long Distance, on page 26

Information About Long Distance

In case of congestion, the long distance ingress ports sends a pause frame to its peer to pause all data transmission on the link for a specified period of time in order to absorb the in-flight packets.

When a pause frame is generated from a switch, it takes time to reach the peer switch. Until the packet reaches to peers all the in-flight packets need to be buffered and should not be dropped.

If the distance between the switches is longer, time taken by a pause frame to reach the peer switch is more and the number of the in-flight packets becomes higher. All these in-flights packets need to be buffered and should not be dropped. To absorb the large numbers of in-flight packets, the buffer size needs to be larger.

The ingress buffer is governed by these two thresholds:

- Lower threshold—Pause: The lower threshold is to trigger the generation of pause packets. When the buffer reaches the pause threshold, pause frames are generated.
- High threshold—Stop: This threshold is to stop generating the pause frames. When the buffer reaches the stop threshold, the pause frames stop getting generated.

To absorb a large number of in-flight packets, the difference between the pause threshold and the stop threshold should to be more. The pause threshold should be lower and the stop threshold should be higher. If the difference between the pause and stop thresholds are higher, the number of in-flight packets are absorbed in ingress buffer. If the pause threshold is lower, a pause frame is generated even for small buffer utilization.

On Cisco MDS 9000 switches, there is only one queue, the 7e template and all buffers are allocated to this queue. The pause threshold and the stop threshold are marked at same level. With this limitation, it cannot handle the long distance functionality.

With introduction of the long distance feature for Cisco MDS 9000 Series, you can enable the long distance feature on certain ports where peer switches are at a longer distance.

Guidelines and Limitations for Long Distance

The long distance feature has the following guidelines:

- If you use the long-distance command, a pause frame is generated even for a small buffer utilization.
- We recommend that you enable the long distance feature only on those ports where this feature is required.
- The internal buffer is currently set at 33%.

Default Settings for Long-Distance

The following table lists the default setting for long-distance.

Parameter	Default
long-distance	Disabled

Configuring Long Distance

Procedure

Step 1 Enter global configuration mode:

switch# **configure terminal**

Step 2 Enter interface configuration mode:

switch(config)# **interface ethernet** *slot/ethernet*

Step 3 Configure long distance:

switch(config-if)# **priority-flow-control long-distance**

Step 4 Exit interface and global configuration modes:

switch(config-if)# **exit**

switch(config)# **exit**

Step 5 (Optional) Verify the status of the long-distance feature:

switch# **show system internal eth-qos port-node ethernet** *slot/ethernet*

Example: Configuring Long Distance

The following running configuration shows how to configure long distance. Replace the placeholders with relevant values for your setup.

```
configure terminal
 interface <ethernet 1/2>
  priority-flow-control long-distance
  exit
 exit
```

Configuring Long Distance