



Preface

This preface describes the audience, organization, and conventions of the *Cisco MDS 9000 Family NX-OS Quality of Service Configuration Guide*. It also provides information on how to obtain related documentation.

Audience

This guide is for experienced network administrators who are responsible for configuring and maintaining the Cisco MDS 9000 Family of multilayer directors and fabric switches.

Document Organization

This document is organized as follows:

Chapter	Title	Description
Chapter 1	QoS Overview	Provides an overview Quality of Services.
Chapter 2	Configuring QoS	Provides information about configuring QoS.
Chapter 3	Configuring Priority Flow Control	Provides information about configuring priority flow control.

Document Conventions

Command descriptions use these conventions:

boldface font	Commands and keywords are in boldface.
<i>italic font</i>	Arguments for which you supply values are in italics.
[]	Elements in square brackets are optional.
[x y z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.

Screen examples use these conventions:

<code>screen font</code>	Terminal sessions and information the switch displays are in screen font.
boldface screen font	Information you must enter is in boldface screen font.
<i>italic screen font</i>	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.

Related Documentation

The documentation set for the Cisco MDS 9000 Family includes the following documents. To find a document online, use the Cisco MDS NX-OS Documentation Locator at:

http://www.cisco.com/en/US/docs/storage/san_switches/mds9000/roadmaps/doclocator.htm

Release Notes

- *Cisco MDS 9000 Family Release Notes for Cisco MDS NX-OS Releases*
- *Cisco MDS 9000 Family Release Notes for MDS SAN-OS Releases*
- *Cisco MDS 9000 Family Release Notes for Cisco MDS 9000 EPLD Images*
- *Cisco DCNM Release Notes*

Regulatory Compliance and Safety Information

- *Regulatory Compliance and Safety Information for the Cisco MDS 9000 Family*

Compatibility Information

- *Cisco Data Center Interoperability Support Matrix*
- *Cisco MDS 9000 NX-OS Hardware and Software Compatibility Information and Feature Lists*
- *Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide*

Hardware Installation

- *Cisco MDS 9700 Series Hardware Installation Guide*
- *Cisco MDS 9500 Series Hardware Installation Guide*
- *Cisco MDS 9250i Multiservice Switch Hardware Installation Guide*
- *Cisco MDS 9200 Series Hardware Installation Guide*
- *Cisco MDS 9100 Series Hardware Installation Guide*
- *Cisco MDS 9148 Multilayer Fabric Switch Quick Start Guide*
- *Cisco MDS 9148S Multilayer Fabric Switch Quick Start Guide*
- *Cisco MDS 9124 and Cisco MDS 9134 Multilayer Fabric Switch Quick Start Guide*

Software Installation and Upgrade

- *Cisco MDS 9000 NX-OS Software Upgrade and Downgrade Guide*

Cisco NX-OS

- *Cisco MDS 9000 Family NX-OS Licensing Guide*
- *Cisco MDS 9000 Family NX-OS Fundamentals Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Quality of Service Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Security Configuration Guide*
- *Cisco MDS 9000 Family NX-OS IP Services Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Intelligent Storage Services Configuration Guide*
- *Cisco MDS 9000 Family NX-OS High Availability and Redundancy Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Inter-VSAN Routing Configuration Guide*
- *Cisco MDS 9000 Family Cookbook for Cisco MDS SAN-OS*

Cisco DCNM-SAN

- *Cisco DCNM Fundamentals Guide, Release 5.x*
- *System Management Configuration Guide, Cisco DCNM for SAN, Release 5.x*
- *Interfaces Configuration Guide, Cisco DCNM for SAN, Release 5.x*
- *Fabric Configuration Guide, Cisco DCNM for SAN, Release 5.x*
- *Quality of Service Configuration Guide, Cisco DCNM for SAN, Release 5.x*
- *Security Configuration Guide, Cisco DCNM for SAN, Release 5.x*
- *IP Services Configuration Guide, Cisco DCNM for SAN, Release 5.x*
- *Intelligent Storage Services Configuration Guide, Cisco DCNM for SAN, Release 5.x*

- *High Availability and Redundancy Configuration Guide, Cisco DCNM for SAN, Release 5.x*
- *Inter-VSAN Routing Configuration Guide, Cisco DCNM for SAN, Release 5.x*
- *SMI-S and Web Services Programming Guide, Cisco DCNM for SAN, Release 5.x*

Command-Line Interface

- *Cisco MDS 9000 Family Command Reference*

Intelligent Storage Networking Services Configuration Guides

- *Cisco MDS 9000 Family I/O Acceleration Configuration Guide*
- *Cisco MDS 9000 Family SANTap Deployment Guide*
- *Cisco MDS 9000 Family Data Mobility Manager Configuration Guide*
- *Cisco MDS 9000 Family Storage Media Encryption Configuration Guide*

Troubleshooting and Reference

- *Cisco MDS 9000 Family and Nexus 7000 Series System Messages Reference*
- *Cisco MDS 9000 Family SAN-OS Troubleshooting Guide*
- *Cisco MDS 9000 Family NX-OS MIB Quick Reference*
- *Cisco DCNM for SAN Database Schema Reference*

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

<http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html>

- Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.



New and Changed Information

As of Cisco DCNM Release 5.2, Cisco Fabric Manager and Cisco Data Center Network Manager for LAN are merged into one unified product called Cisco Data Center Network Manager (DCNM) that can manage both LAN and SAN environments. As a part of this product merger, the name Cisco DCNM for SAN replaces the name Cisco Fabric Manager.

The following documentation changes support the merged Cisco DCNM product:

- Cisco DCNM product documentation for Cisco DCNM Release 5.2 is retitled with the name Cisco DCNM for LAN.
- Cisco Fabric Manager product documentation for Cisco DCNM Release 5.2 is retitled with the name Cisco DCNM for SAN.
- Cisco DCNM for SAN product documentation is now published to the Data Center Network Manager listing page on Cisco.com:
http://www.cisco.com/en/US/products/ps9369/tsd_products_support_configure.html
This URL is also the listing page for Cisco DCNM for LAN product documentation.
- Cisco Fabric Manager documentation for software releases earlier than Cisco DCNM Release 5.2, retains the name Cisco Fabric Manager and remains available at its current Cisco.com listing page:
http://www.cisco.com/en/US/products/ps10495/tsd_products_support_configure.html
You should continue to use the Cisco Fabric Manager documentation if you are using a release of Cisco Fabric Manager software that is earlier than Cisco DCNM Release 5.2.
- The name DCNM-SAN is used in place of Cisco DCNM for SAN in the user interface of Cisco Data Center Network Manager; likewise, the name DCNM-LAN is used in place of Cisco DCNM for LAN in the user interface. To match the user interface, the product documentation also uses the names DCNM-SAN and DCNM-LAN.
- The following new publications support both Cisco DCNM for LAN and DCNM for SAN, and address the new licensing model, the new installation process, and the new features of Cisco DCNM:
 - *Cisco DCNM Installation and Licensing Guide*
 - *Cisco DCNM Release Notes*
- For a complete list of Cisco DCNM documentation, see the “Related Documentation” section in the Preface.

As of Cisco MDS NX-OS Release 4.2(1), software configuration information is available in new feature-specific configuration guides for the following information:

- System management
- Interfaces

- Fabric
- Quality of service
- Security
- IP services
- High availability and redundancy

The information in these new guides previously existed in the *Cisco MDS 9000 Family CLI Configuration Guide* and in the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*. Those configuration guides remain available on Cisco.com and should be used for all software releases prior to MDS NX-OS Release 4.2(1). Each guide addresses the features introduced in or available in a particular release. Select and view the configuration guide that pertains to the software installed in your switch.

To find additional information about Cisco MDS NX-OS Release 4.2(x), see the *Cisco MDS 9000 Family Release Notes* available at the following Cisco Systems website:

http://www.cisco.com/en/US/products/ps5989/prod_release_notes_list.htm

About this Guide

The information in the new *Cisco Fabric Manager Quality of Service Configuration Guide* previously existed in *Part 5: QoS of the Cisco MDS 9000 Family Fabric Manager Configuration Guide*.

Table 1 lists the New and Changed features for this guide, starting with MDS NX-OS Release 6.2(9).

Feature	New or Changed Topics	Changed in Release	Where Documented
FCoE long-distance SAN extension	Provides information about the FCoE long-distance SAN extension feature that you can enable on certain ports where peer switches are at a longer distance.	6.2(9)	Chapter 3, “Configuring Priority Flow Control,”



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 Family switches to simplify deployment and to provide optimization of large-scale fabrics.

This chapter describes the QoS, and port-tracking features on the Cisco MDS 9000 switches and includes the following sections:

- [QoS, page 1-1](#)
- [Port Tracking, page 1-3](#)

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.

This section covers the following topics:

- [QoS in Differentiated Service, page 1-2](#)
- [Applying QoS to Traffic, page 1-2](#)
- [QoS Configuration, page 1-2](#)
- [QoS Licensing, page 1-3](#)

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 Family 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 queuing.

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.

For information on configuring QoS, refer to *Chapter 1, “Configuring Fabric Congestion and QoS”*.

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.

For information on configuring port tracking, refer to *Chapter 2, “Configuring Port Tracking”*.



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.

This chapter includes the following topics:

- [Configuring QoS, page 2-5](#)
- [Configuration Examples for QoS, page 2-14](#)
- [Ingress Port Rate Limiting, page 2-16](#)

Configuring QoS

This section includes the following topics:

- [Information About Control Traffic, page 2-6](#)
- [Enabling or Disabling Control Traffic, page 2-6](#)
- [Displaying Control Traffic Information, page 2-6](#)
- [Information About Data Traffic, page 2-7](#)
- [Comparing VSAN Versus Zone-Based QoS, page 2-8](#)
- [Configuring Data Traffic, page 2-8](#)
- [QoS Initiation for Data Traffic, page 2-8](#)
- [Information About Class Map Creation, page 2-9](#)
- [Creating a Class Map, page 2-9](#)
- [Information About Service Policy Definition, page 2-10](#)
- [Specifying Service Policies, page 2-11](#)
- [About Service Policy Enforcement, page 2-11](#)
- [Applying Service Policies, page 2-12](#)
- [About the DWRR Traffic Scheduler Queue, page 2-12](#)

- [Changing the Weight in a DWRR Queue, page 2-12](#)
- [Displaying Data Traffic Information, page 2-12](#)

Information About Control Traffic

The Cisco MDS 9000 Family 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 Family 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 Family.

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.

Detailed Steps

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

	Command	Purpose
Step 1	switch# confi g t	Enters configuration mode.
Step 2	switch(config)# no qos control priority 0	Enables the control traffic QoS feature.
	switch(config)# qos control priority 0	Disables the control traffic QoS feature.

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 (see [Example 2-1](#)).

Example 2-1 Displays Current QoS Settings

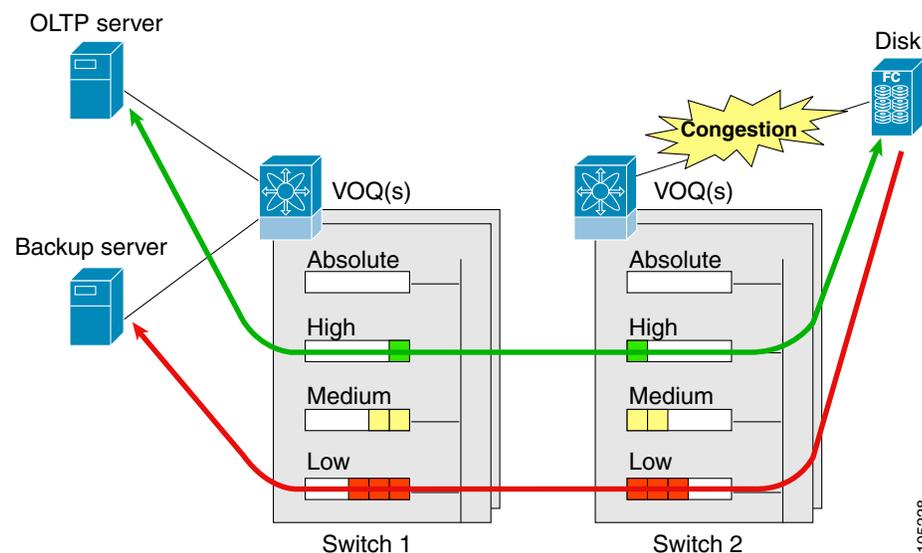
```
switch# show qos statistics
Total number of FC frames transmitted from the Supervisor= 15767
Number of highest-priority FC frames transmitted           = 8224
Current priority of FC control frames = 0      (0 = lowest; 7 = highest)
```

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 Family 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 (see [Figure 2-1](#)).

Figure 2-1 Prioritizing Data Traffic



In [Figure 2-1](#), 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. [Table 2-1](#) highlights the differences between configuring QoS priorities based on VSANs versus zones.

Table 2-1 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

Detailed Steps

To configure QoS, follow these steps:

-
- 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 Family Switch-to-Switch Interoperability Configuration Guide*.

Detailed Steps

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

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# qos enable	Enables QoS. You can now configure data traffic parameters.
	switch(config)# no qos enable	Removes the currently applied QoS configuration and disables QoS. You can no longer configure data traffic parameters.

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 FFFFFFFF (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.

Detailed Steps

To create a class map, follow these steps:

	Command	Purpose
Step 1	<code>switch(config)# qos class-map MyClass match-all switch(config-cmap)#</code>	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.
	<code>switch(config)# qos class-map MyClass match-any switch(config-cmap)#</code>	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.
Step 2	<code>switch(config-cmap)# match destination-address 0x12ee00</code>	Specifies a destination address match for frames with the specified destination FC ID.
	<code>switch(config-cmap)# match source-address 0x6d1090 mask 0xFFFFF</code>	Specifies a source address and mask match for frames with the specified source FC ID.
	<code>switch(config-cmap)# match destination-wwn 20:01:00:05:30:00:28:df</code>	Specifies a destination WWN to match frames.
	<code>switch(config-cmap)# match source-wwn 23:15:00:05:30:00:2a:1f</code>	Specifies a source WWN to match frames.
	<code>switch(config-cmap)# match destination-device-alias DocDeviceAlias</code>	Specifies a destination device alias to match frames.
	<code>switch(config-cmap)# match source-device-alias DocDeviceAliase</code>	Specifies a source device alias to match frames.
	<code>switch(config-cmap)# match input-interface fc 2/1</code>	Specifies a source interface to match frames.
	<code>switch(config-cmap)# no match input-interface fc 3/5</code>	Removes a match based on the specified source interface.

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

Refer to http://www.cisco.com/en/US/tech/tk543/tk757/technologies_tech_note09186a00800949f2.shtml for further information on implementing QoS DSCP values.



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

Specifying Service Policies

Detailed Steps

To specify a service policy, follow these steps:

	Command	Purpose
Step 1	<code>switch(config)# qos policy-map MyPolicy</code> <code>switch(config-pmap)#</code>	Creates a policy map called MyPolicy and places you in the policy-map submode.
	<code>switch(config)# no qos policy-map OldPolicy</code> <code>switch(config)#</code>	Deletes the policy map called OldPolicy and places you in the policy-map submode.
Step 2	<code>switch(config-pmap)# class MyClass</code> <code>switch(config-pmap-c)#</code>	Specifies the name of a predefined class and places you at the policy-map submode for that class.
	<code>switch(config-pmap)# no class OldClass</code>	Removes the class map called OldClass from the policy map.
Step 3	<code>switch(config-pmap-c)# priority high</code>	Specifies the priority to be given for each frame matching this class.
	<code>switch(config-pmap-c)# no priority high</code>	Deletes a previously assigned priority and reverts to the default value of low.
Step 4	<code>switch(config-pmap-c)# dscp 2</code>	Specifies the DSCP value to mark each frame matching this class.
	<code>switch(config-pmap-c)# no dscp 60</code>	Deletes a previously assigned DSCP value and reverts to the factory default of 0.

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

Detailed Steps

To apply a service policy, follow these steps:

	Command	Purpose
Step 1	<code>switch(config)# qos service policy MyPolicy vsan 3</code>	Applies a configured policy to VSAN 3.
	<code>switch(config)# no qos service policy OldPolicy vsan 7</code>	Deletes a configured policy that was applied to 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.

Changing the Weight in a DWRR Queue

Detailed Steps

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

	Command	Purpose
Step 1	<code>switch(config)# qos dwrr-q high weight 10</code>	Associates a relative weight (10) to a specified queue (default queue).
	<code>switch(config)# no qos dwrr-q low weight 51</code>	Restores the default weight of 20.

Displaying Data Traffic Information

Examples

The `show qos` commands display the current QoS settings for data traffic (see Examples 2-2 to 2-10).

Example 2-2 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 2-3 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 2-4 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 2-5 Displays a Specified Policy Map

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

Example 2-6 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 2-7 Displays All Applied Policy Maps

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

Example 2-8 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 2-9 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 2-10 Displays QoS Statistics

```

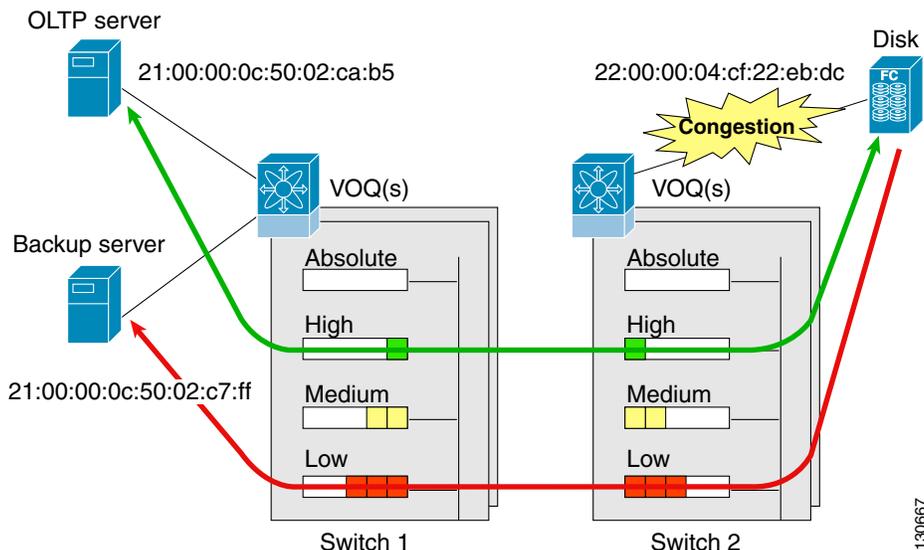
switch# show qos statistics
Total number of FC frames transmitted from the Supervisor= 301431
Number of highest-priority FC frames transmitted           = 137679
Current priority of FC control frames = 7 (0 = lowest; 7 = highest)

```

Configuration Examples for QoS

This section describes a configuration example for the application illustrated in [Figure 2-2](#).

Figure 2-2 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.

Detailed Steps

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

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.

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:

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.

Ingress Port Rate Limiting

A port rate limiting feature helps control the bandwidth for individual Fibre Channel ports. 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 limiting the number of frames that are transmitted out of the exit point on the MAC. Port rate limiting works on all Fibre Channel ports. The rate limit ranges from 1 to 100% and the default is 100%.



Note

Port rate limiting can be configured on all Cisco MDS switches and linecards only if the QoS feature is enabled.

Detailed Steps

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

	Command	Purpose
Step 1	switch # config t switch(config)#	Enters the configuration mode.
Step 2	switch(config)# interface fc 1/1	Selects the interface to specify the ingress port rate limit.
Step 3	switch(config-if)# switchport ingress-rate 50	Configures a 50% port rate limit for the selected interface.
	switch(config-if)# no switchport ingress-rate 50	Reverts a previously configured rate to the factory default of 100%.



Configuring Priority Flow Control

This chapter describes how to configure priority flow control (PFC) on the Cisco MDS 9000 switches. This chapter includes the following sections:

- [Information About Priority Flow Control, page 3-1](#)
- [Licensing Requirements for Priority Flow Control, page 3-2](#)
- [Guidelines and Limitations for Priority Flow Control, page 3-2](#)
- [Default Settings for Priority Flow Control, page 3-2](#)
- [Configuring Priority Flow Control, page 3-2](#)
- [Information About Long Distance, page 3-4](#)
- [Guidelines and Limitations for Long Distance, page 3-4](#)
- [Configuring Long Distance, page 3-4](#)
- [Configuration Examples for Priority Flow Control, page 3-5](#)

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

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

Table [Table 3-1](#) lists the default setting for PFC.

Table 3-1 Default PFC Setting

Parameter	Default
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.

SUMMARY STEPS

1. **configure terminal**
2. **interface ethernet** [*slot/port-number*] or **interface ethernet port-channel** [*port-number*]
3. **[no] priority-flow-control mode** {**auto** | **off** | **on**}
4. **show interface priority-flow-control**

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface ethernet [<i>slot/port-number</i>] Example: switch(config)# interface ethernet 2/5 switch(config-if)#	Enters interface mode on the interface specified.
Step 3	priority-flow-control mode { auto off on }	Sets the PFC to the auto , off , or on mode. By default, PFC mode is set to auto on all ports.
Step 4	show interface priority-flow-control Example: switch# show interface priority-flow-control	Displays the status of PFC on all interfaces.

Verifying the Priority Flow Control Configuration

To display the PFC configuration, perform the following task:

Command	Purpose
show interface priority-flow-control	Displays the status of PFC on all interfaces.

For detailed information about the fields in the output from these commands, see the *Cisco MDS 9000 Family Command Reference*.

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 Family, 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%.

Configuring Long Distance

SUMMARY STEPS

1. **configure terminal**
2. **interface ethernet** *[slot/port-number]* or **interface ethernet-port-channel** *[port-number]*
3. **[no] priority-flow-control long-distance**
4. **show system internal eth-qos port-node ethernet** *[slot/port-number]*

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface ethernet [slot/port-number] Example: switch(config)# interface ethernet-port-channel 2/5 switch(config-if)#	Enters interface mode on the interface specified.
Step 3	priority-flow-control long-distance Example: switch(config-if)# priority-flow-control fcoe long-distance switch(config-if)# no priority-flow-control long-distance switch(config-if)# no priority-flow-control fcoe long-distance	Enables or disables the long-distance feature.
Step 4	show sys internal eth-qos port-node ethernet 1/1 Example: switch# show system internal eth-qos port-node ethernet 1/1	Displays the status of long-distance feature.

Default Settings for Long-Distance

Table [Table 3-1](#) lists the default setting for long-distance.

Table 3-2 Default PFC Setting

Parameter	Default
long-distance	Disabled.

Configuration Examples for Priority Flow Control

The following example shows how to configure PFC:

```
configure terminal
interface ethernet 5/5
priority-flow-control mode on
```




C

class maps

- configuring for data traffic [2-5](#)
- creating [2-5](#)

control traffic

- disabling QoS [2-2](#)
- enabling for QoS [2-2](#)

D

data traffic

- applying service policies [2-8](#)
- class maps [2-5](#)
- comparing VSANs and QoS [2-4](#)
- defining service policies [2-6](#)
- displaying information [2-9](#)
- DWRR queues [2-8](#)
- enabling QoS [2-4](#)
- enforcing service policies [2-7](#)
- example configuration [2-10](#)

deficit weighted round robin schedulers. See DWRR schedulers

DWRR queues

- changing weights [2-9](#)

DWRR schedulers

- description [2-3](#)

F

FCC

- benefits [2-1](#)

Fibre Channel Congestion Control. See FCC

G

Generation 1 switching modules

- QoS behavior [2-8](#)

Generation 2 switching modules

- QoS behavior [2-8](#)

P

port rate limiting

- configuring [2-13](#)
- description [2-12](#)
- hardware restrictions [2-12](#)

Q

QoS

- class maps [2-5](#)
- comparison with VSANs [2-4](#)
- control traffic support [2-2](#)
- creating class maps [2-5](#)
- data traffic support [2-3 to ??](#)
- description [2-1](#)
- displaying information [2-2, 2-9](#)
- DWRR queues [2-8](#)
- enabling control traffic [2-2](#)
- enabling for data traffic [2-4](#)
- example data traffic configuration [2-10](#)
- port rate limiting [2-12](#)
- service policies [2-6, 2-7](#)

S

service policies

 applying [2-8](#)

 defining [2-6](#)

 enforcement [2-7](#)

V

VSANs

 comparison with QoS [2-4](#)