Send documentation comments to dcnm-san-docfeedback@cisco.com

Quality of Service Configuration Guide, Cisco DCNM for SAN

Cisco DCNM for SAN, Release 5.x
July 2011

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

Text Part Number: OL-25337-01

Quality of Service Configuration Guide, Cisco DCNM for SAN

CONTENTS

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.

Preface

This preface describes the audience, organization, and conventions of the *Quality of Service Configuration Guide, Cisco DCNM for SAN*. 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 folows:

Chapter	Title	Description
Chapter 1	QoS Overview	Provides an overview Quality of Services.

Document Conventions

Command descriptions use these conventions:

boldface font	Commands and keywords are in boldface.
italic font	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:

`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.

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/doclocater.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 9500 Series Hardware Installation Guide*

- *Cisco MDS 9200 Series Hardware Installation Guide*
- *Cisco MDS 9100 Series Hardware Installation 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 System Management 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*

Send documentation comments to dcnm-san-docfeedback@cisco.com

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.

CHAPTER 1

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

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 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

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 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 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 2, “Configuring QoS”.

CHAPTER 2

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:

Information About QoS

QoS implementation in the Cisco MDS 9000 Family follows the differentiated services (DiffServ) model. The DiffServ standard is defined in RFCs 2474 and 2475. 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. By default, the QoS feature for certain critical control traffic is enabled. These critical control frames are assigned the highest (absolute) priority. 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.

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.

Configuring QoS

This section includes the following topics:

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

Toenable or disable the high priority assignment for control traffic using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **FC Services** and then select **QoS** in the Physical Attributes pane.

The QoS control traffic information is displayed in the Information pane. The **Control** tab is default.

Step 2 Select the switch on which you want to enable or disable control traffic.

Step 3 In the Command column, click the drop-down menu and select **enable** or **disable**.

Step 4 Click **Apply Changes** to save your changes.

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

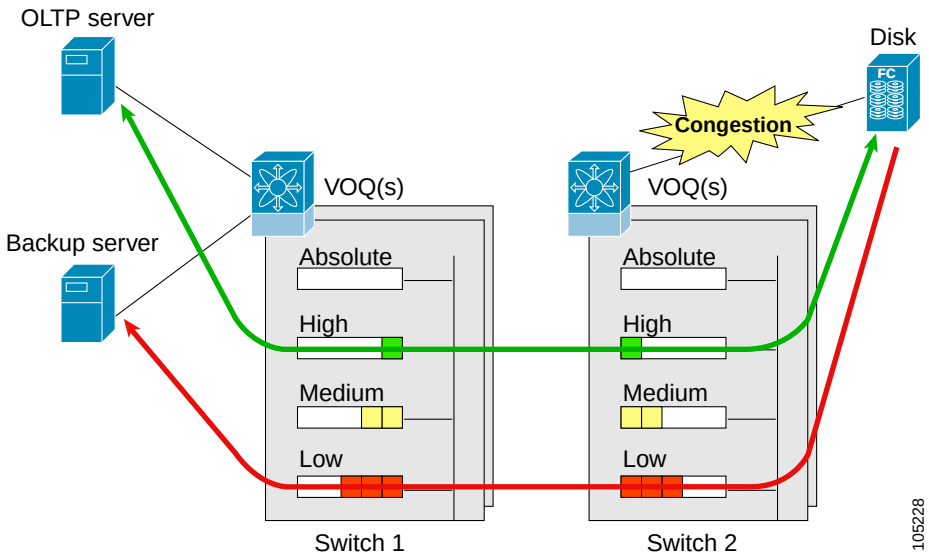

In Figure 2-1, the OLTP traffic arriving at Switch 1 is marked with a high priority level of throughput classification (class map) and marking (policy map). Similarly, the backup traffic is marked with a low priority level. The traffic is sent to the corresponding priority queue within a virtual output 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 using Fabric Manager, 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.

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

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

Detailed Steps To create a class map using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **FC Services** and then select **QoS** in the Physical Attributes pane.

The QoS information is displayed in the Information pane. The **Control** tab is the default.

Step 2 In the **Class Maps** tab, click **Create Row** to create a new class map.

You see the Create Class Maps dialog box.

Step 3 Select the switches for the class map.

Step 4 Enter the source ID or the destination **ID** in the field.

Step 5 Enter a name for the class map.

Step 6 Select a Match mode.You can either match **any** or **all** criterion with one match statement from the class map configuration mode.

Step 7 Click **Create** to proceed with creating the class map.

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.

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.

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).

Table 2-2 describes the QoS behavior for Generation 1, Generation 2, and Generation 3 switching modules.

Table 2-2 QoS Behavior for Generation 1 and Generation 2 Switching Modules

Source Module Type	Destination Module Type	QoS Behavior Description
Generation 1	Generation 1	QoS behavior reflects the DWRR configuration for traffic coming in through a given port and queued to the same egress port. All the other traffic share equal bandwidth.
Generation 1	Generation 2 or Generation 3	QoS behavior reflects the DWRR configuration for traffic coming in through a given port and queued to the same egress port. All the other streams share equal bandwidth.

Table 2-2 QoS Behavior for Generation 1 and Generation 2 Switching Modules

Source Module Type	Destination Module Type	QoS Behavior Description
Generation 2 or Generation 3	Generation 1	Bandwidth partitioning is equal for all the traffic.
Generation 2 or Generation 3	Generation 2 or Generation 3	QoS behavior reflects the DWRR weights configuration for all possible streams.

Changing the Weight in a DWRR Queue

Detailed Steps To change the weight in a DWRR queue using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **FC Services** and then select **QoS** in the Physical Attributes pane.

The QoS control traffic information is displayed in the Information pane. The default is the **Control** tab.

Step 2 Click the **DWRR** tab.

You see the queue status and weight.

Step 3 Select a switch and change the weight.

Step 4 Click the **Apply Changes** icon to save your changes.

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

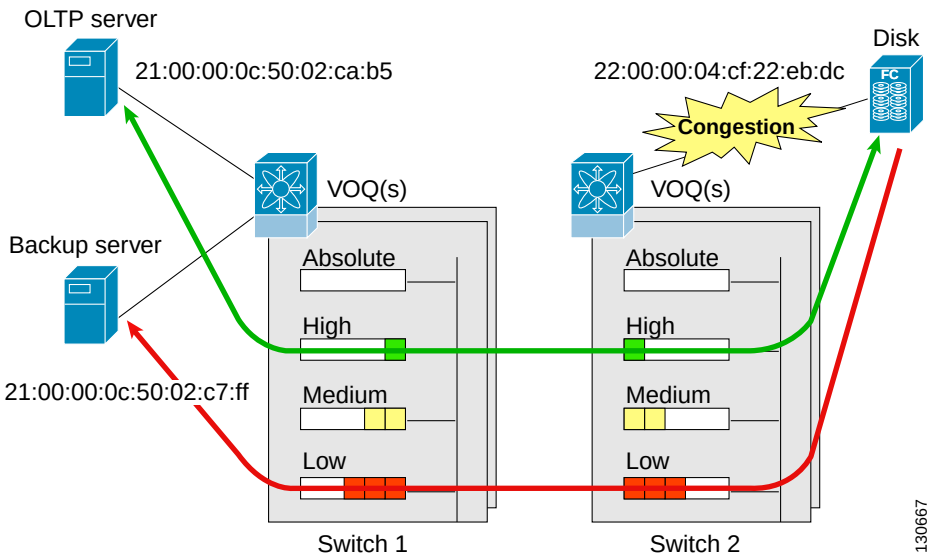

Both the OLTP server and the backup server are accessing the disk. The backup server is writing large amounts of data to the disk. This data does not require specific service guarantees. The volumes of data generated by the OLTP server to the disk are comparatively much lower but this traffic requires 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.

Step 2 Create the policy map.

Step 3 Assign the service policy.

Step 4 Assign the weights for the DWRR queues.

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.

Step 2 Assign the class maps to the policy map.

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

Limiting 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 only be configured on Cisco MDS 9100 Series switches, Cisco MDS 9216i switches, and MPS-14/2 modules.

This feature can only be configured if the QoS feature is enabled and if this configuration is performed on a Cisco MDS 9100 series switch, Cisco MDS 9216i switch, or MPS-14/2 module.

Detailed Steps

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

Step 1 Expand **Switches**, expand **FC Services** and then select **QoS** in the Physical Attributes pane.

The QoS control traffic information is displayed in the Information pane. The default is the **Control** tab.

Step 2 Click the **Rate Limit** tab.

Step 3 Select the switch whose port rate limit you want to change.

Step 4 Enter the desired port rate limit in the Percent column.

Step 5 Click the **Apply Changes** icon to save your changes.

INDEX

V