Send documentation comments to mdsfeedback-doc@cisco.com

Cisco Fabric Manager Quality of Services Configuration Guide

Cisco MDS NX-OS Release 4.2(1)
Cisco MDS 9000 FabricWare Release 4.x
August 2009

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-19778-01

Cisco Fabric Manager Quality of Services Configuration Guide

CONTENTS

New and Changed Information

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.

Some information from the *Cisco MDS 9000 Family CLI Configuration Guide* and the *Cisco MDS 9000 Family Fabric Manager Configuration Guide* now appears in the following guides that are common among products that run the Nexus operating system:

- *Cisco NX-OS Family Licensing Guide* – Explains the licensing model and describes the feature licenses.
- *Cisco NX-OS Fundamentals Configuration Guide* – Describes the switch setup utility and includes general CLI, file system, and configuration information.

For a complete list of document titles, see the list of Related Documentation in the “Preface.”

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 9: Traffic Management of the *Cisco MDS 9000 Family Fabric Manager Configuration Guide.*

There are no new or changed Fabric Manager features for quality of service in MDS NX-OS Release 4.2(1).

Preface

This preface describes the audience, organization, and conventions of the *Cisco MDS 9000 Family Cisco Fabric Manager Quality of Services 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.

Organization

This Guide is organized as follows: :

Chapter	Title	Description
Chapter 1	QoS Overview	Provides an overview Quality of Services.
Chapter 2	Configuring Fabric Congestion Control and QoS	Provides details on the QoS and FCC features provided in all switches
Chapter 3	Configuring Port Tracking	Provides information about a port tracking feature that provides a faster recovery from link failures.

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 Storage Services Interface Images*
- *Cisco MDS 9000 Family Release Notes for Cisco MDS 9000 EPLD Images*
- *Release Notes for Cisco MDS 9000 Family Fabric Manager*

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*

Send documentation comments to mdsfeedback-doc@cisco.com

- *Cisco MDS NX-OS Release Compatibility Matrix for Storage Service Interface Images*
- *Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide*
- *Cisco MDS NX-OS Release Compatibility Matrix for IBM SAN Volume Controller Software for Cisco MDS 9000*
- *Cisco MDS SAN-OS Release Compatibility Matrix for VERITAS Storage Foundation for Networks Software*

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 Release 4.1(x) and SAN-OS 3(x) Software Upgrade and Downgrade Guide*
- *Cisco MDS 9000 Family Storage Services Interface Image Install and Upgrade Guide*
- *Cisco MDS 9000 Family Storage Services Module Software Installation and Upgrade 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 Fabric Manager

- *Cisco Fabric Manager Fundamentals Configuration Guide*
- *Cisco Fabric Manager System Management Configuration Guide*
- *Cisco Fabric Manager Interfaces Configuration Guide*

- *Cisco Fabric Manager Fabric Configuration Guide*
- *Cisco Fabric Manager Quality of Service Configuration Guide*
- *Cisco Fabric Manager Security Configuration Guide*
- *Cisco Fabric Manager IP Services Configuration Guide*
- *Cisco Fabric Manager Intelligent Storage Services Configuration Guide*
- *Cisco Fabric Manager High Availability and Redundancy Configuration Guide*
- *Cisco Fabric Manager Inter-VSAN Routing Configuration Guide*
- *Cisco Fabric Manager Online Help*
- *Cisco Fabric Manager Web Services Online Help*

Command-Line Interface

- *Cisco MDS 9000 Family Command Reference*

Intelligent Storage Networking Services Configuration Guides

- *Cisco MDS 9000 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*
- *Cisco MDS 9000 Family Secure Erase Configuration Guide*
- *Cisco MDS 9000 Family Cookbook for Cisco MDS SAN-OS*

Troubleshooting and Reference

- *Cisco NX-OS System Messages Reference*
- *Cisco MDS 9000 Family NX-OS Troubleshooting Guide*
- *Cisco MDS 9000 Family NX-OS MIB Quick Reference*
- *Cisco MDS 9000 Family NX-OS SMI-S Programming Reference*
- *Cisco MDS 9000 Family Fabric Manager Server Database Schema*

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.

Send documentation comments to mdsfeedback-doc@cisco.com

Send documentation comments to mdsfeedback-doc@cisco.com

CHAPTER 1

QoS Overview

The Cisco MDS 9000 NX-OS software offers traffic management features such as fabric-wide quality of service (QoS) and Fibre Channel Congestion Control (FCC). 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, FCC, 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 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 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 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.

For information on configuring QoS, refer to Chapter 2, “Configuring Fabric Congestion Control and QoS”.

FCC

FCC provides an innovative, end-to-end congestion-control mechanism that augments the standard Fibre Channel buffer-to-buffer credit mechanism to provide enhanced traffic management. A switch experiencing congestion explicitly signals this condition to the ingress switch (the entry point for traffic into the fabric that is causing congestion). Upon receipt of an explicit notification, the ingress switch throttles the N port or NL port traffic by reducing the buffer-to-buffer credits.

This section covers the following topics:

Causes of Congestion

FCC is used to prevent congestion occurring on an output port caused by sending devices transmitting traffic to receiving devices.

Congestion occurs due to the following two common causes:

- Receivers Unable To Maintain Sustained Performance
- Speed Mismatch Between Senders and Receivers in a Fabric

Receivers Unable To Maintain Sustained Performance

Congestion occurs if multiple senders are contending with a smaller number of receivers. If the aggregate rate of traffic transmitted by senders exceeds the size of the connection to the receivers, blocking occurs as shown in Figure 1-1.

Figure 1-1 Congestion Caused by Transmitters Outnumbering Receivers

Speed Mismatch Between Senders and Receivers in a Fabric

When there is a speed mismatch between senders and receivers, buffering occurs. Buffers are a finite resource on switches, typically in the range of 16 buffers (32 KB) to 255 buffers (512 KB) per port. When these buffers are full, blocking occurs and leads to congestion as shown in Figure 1-2.

Figure 1-2 Congestion Caused by Speed Mismatch between Senders and Receivers

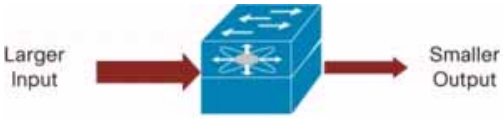

Using FCC To Reduce Congestion and Blocking

FCC reduces the congestion and blocking from spreading by using signaling to notify devices in the path to slow down the sending device that is causing the congestion.

When FCC detects congestion on an output port, it generates an FCC edge quench frame with a destination address of the device causing the congestion (the sender). Switches that receive FCC edge quench inspects the FCC frame and determine if it is targeted at a directly attached device. If it is, the switch instigates rate limiting on the input port attached to the sending device causing congestion and rate-limits the sender to the rate at which the receiver can receive.

FCC works on an active-loop feedback system. When there is congestion, the switch with the congested output port continues to send FCC edge quench frames. The switch connected to the sending device, causing the congestion, continues to rate-limit the flow until the congestion is minimized.

FCC and FCC edge quench frames are completely compatible with the existing Fibre Channel standards and fabrics and can be used in mixed-vendor fabrics.

For information on configuring FCC, refer to *Chapter 2, "Configuring Fabric Congestion Control 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 Fabric Congestion Control and QoS"*.

CHAPTER 2

Configuring Fabric Congestion Control and QoS

Fibre Channel Congestion Control (FCC) is a Cisco proprietary flow control mechanism that alleviates congestion on Fibre Channel networks.

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 provides details on the QoS and FCC features provided in all switches. It includes the following sections:

FCC

FCC reduces the congestion in the fabric without interfering with the standard Fibre Channel protocols. This section contains the following topics:

About FCC

The FCC protocol increases the granularity and the scale of congestion control applied to any class of traffic (see Figure 2-1).

Figure 2-1 FCC Mechanisms

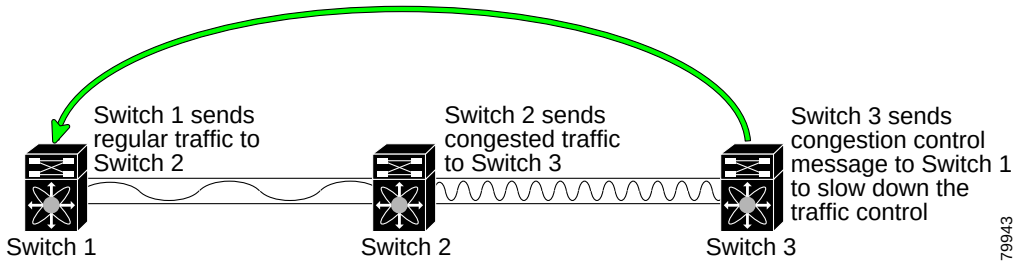

Edge quench congestion control provides feedback to the source about the rate at which frames should be injected into the network (frame intervals).

Note FCC is not supported on the Cisco Fabric Switch for HP c-Class BladeSystem and Cisco Fabric Switch for IBM BladeCenter.

FCC Process

When a node in the network detects congestion for an output port, it generates an edge quench message. These frames are identified by the Fibre Channel destination ID (DID) and the source ID. A switch from other vendors simply forwards these frames.

Any receiving switch in the Cisco MDS 9000 Family handles frames in one of these ways:

- It forwards the frame.
- It limits the rate of the frame flow in the congested port.

The behavior of the flow control mechanism differs based on the Fibre Channel DID:

- If the Fibre Channel DID is directly connected to one of the switch ports, the input rate limit is applied to that port.
- If the destination of the edge quench frame is a Cisco domain or the next hop is a Cisco MDS 9000 Family switch, the frame is forwarded.
- If neither of these mechanisms is true, then the frame is processed in the port going towards the FC DID.

All switches (including the edge switch) along the congested path process path quench frames. However, only the edge switch processes edge quench frames.

Enabling FCC

By default, the FCC protocol is disabled. FCC can only be enabled for the entire switch.

Tip If you enable FCC, be sure to enable it in all switches in the fabric.

To enable or disable the FCC feature using Fabric Manager, follow these steps:

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

The FCC information is displayed in the Information pane. The **General** tab is the default.

Step 2 Select the switch on which you want to enable FCC.

Step 3 Check the **Enable** check box.

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

Assigning FCC Priority

To assign FCC priority using Fabric Manager, follow these steps:

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

The FCC information is displayed in the Information pane. The **General** tab is the default.

Step 2 Select the switch for which you want to assign the FCC priority.

Step 3 Enter the priority in the **Priority** column.

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

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.

All switches support the following types of traffic:

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.

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.

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 2-2).

Figure 2-2 Prioritizing Data Traffic

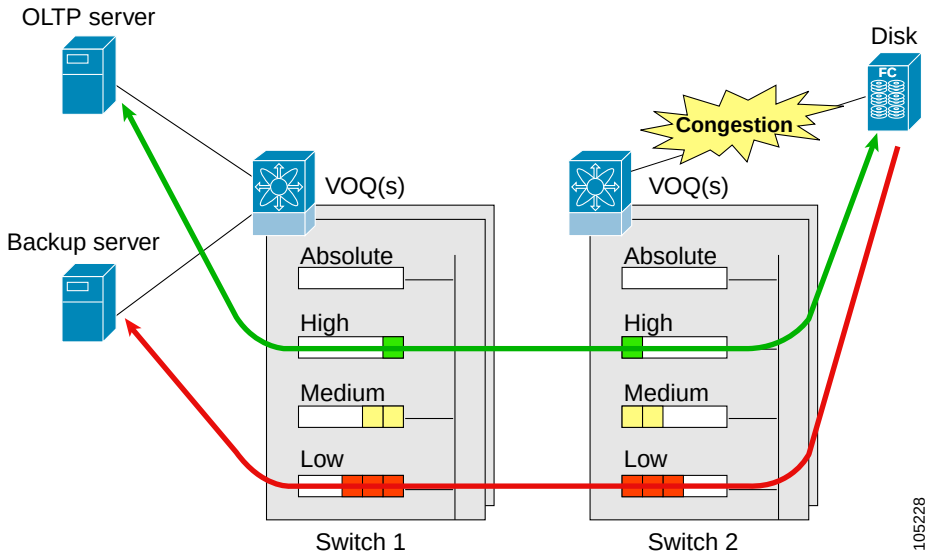

In Figure 2-2, the OLTP traffic arriving at Switch 1 is marked with a high priority level of throughput classification (class map) and marking (policy map). Similarly, the backup traffic is marked with a low priority level. The traffic is sent to the corresponding priority queue within a virtual output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

Tip To achieve this traffic differentiation, be sure to enable FCC (see the "Enabling FCC" section on page 2-2).

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

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*.

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

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 shown in Figure 2-3. The **Control** tab is the default.

Figure 2-3 Quality of Service Control Tab

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

You see the Create Class Maps dialog box shown in Figure 2-4.

Figure 2-4 Create Class Maps Dialog Box

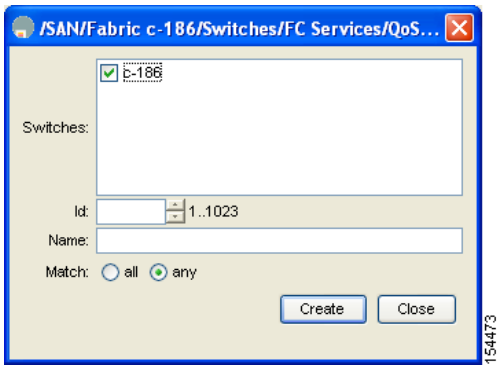

Step 3 Select the switches for the class map.

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.

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.
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

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 shown in Figure 2-5. The default is the **Control** tab.

Figure 2-5 Quality of Service Control Tab

Step 2 Click the **DWRR** tab.

You see the queue status and weight (see Figure 2-6).

Figure 2-6 QoS Queue Status and Weight

Step 3 Select a switch and change the weight.

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

Example Configuration

This section describes a configuration example for the application illustrated in Figure 2-7.

Figure 2-7 Example Application for Traffic Prioritization

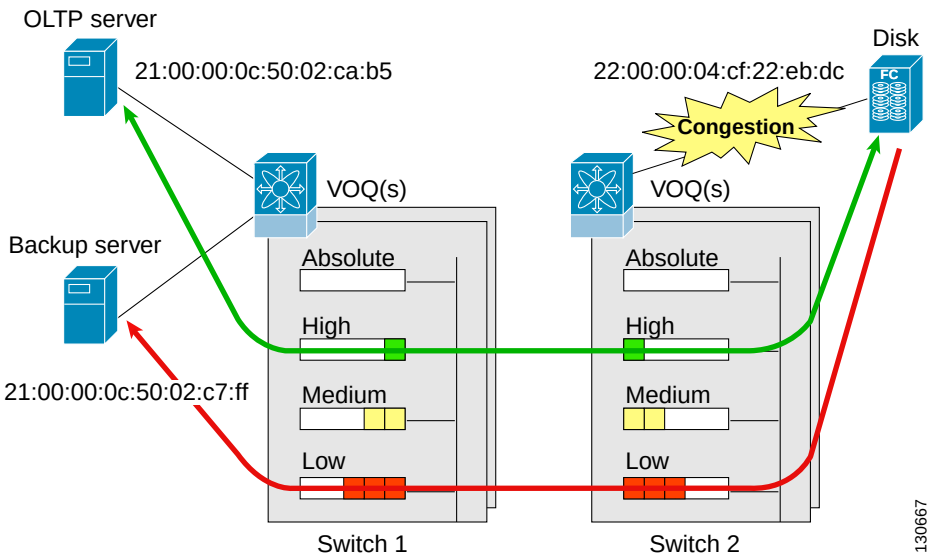

Both the OLTP server and the backup server are accessing the disk. The backup server is writing large amounts of data to the disk. This data does not require specific service guarantees. The volumes of data generated by the OLTP server to the disk are comparatively much lower but this traffic requires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

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.

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.

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 shown in Figure 2-8. The default is the **Control** tab.

Figure 2-8 Quality of Service Control Tab

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

You see the information shown in Figure 2-9.

Figure 2-9 Rate Limits for Switch Interfaces

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.

Default Settings

Table 2-3 lists the default settings for FCC, QoS, and rate limiting features.

Table 2-3 Default FCC, QoS, and Rate Limiting Settings

Parameters	Default
FCC protocol	Disabled.
QoS control traffic	Enabled.
QoS data traffic	Disabled.
Zone-based QoS priority	Low.
Rate limit	100%

CHAPTER 3

Configuring Port Tracking

The port tracking feature is unique to the Cisco MDS 9000 Family of switches. This feature uses information about the operational state of the link to initiate a failure in the link that connects the edge device. This process of converting the indirect failure to a direct failure triggers a faster recovery process towards redundant links. When enabled, the port tracking feature brings down the configured links based on the failed link and forces the traffic to be redirected to another redundant link.

This chapter includes the following sections:

About Port Tracking

Generally, hosts can instantly recover from a link failure on a link that is immediately (direct link) connected to a switch. However, recovering from an indirect link failure between switches in a WAN or MAN fabric with a keep-alive mechanism is dependent on several factors such as the time out values (TOVs) and on registered state change notification (RSCN) information.

In Figure 3-1, when the direct link 1 to the host fails, recovery can be immediate. However, when the ISL 2 fails between the two switches, recovery depends on TOVs, RSCNs, and other factors.

Figure 3-1 Traffic Recovery Using Port Tracking

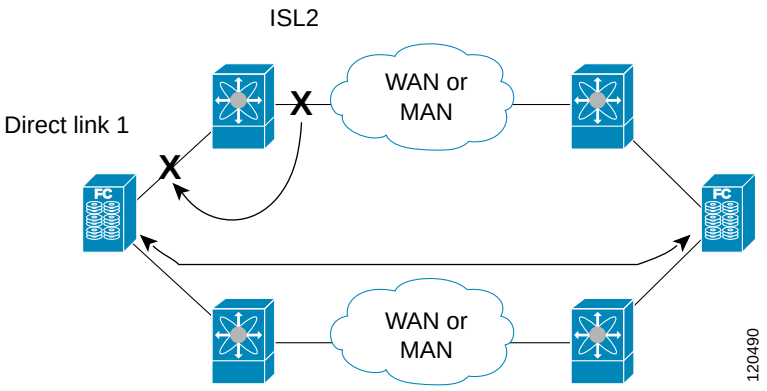

The port tracking feature monitors and detects failures that cause topology changes and brings down the links connecting the attached devices. When you enable this feature and explicitly configure the linked and tracked ports, the Cisco NX-OS software monitors the tracked ports and alters the operational state of the linked ports on detecting a link state change.

The following terms are used in this chapter:

- Tracked ports—A port whose operational state is continuously monitored. The operational state of the tracked port is used to alter the operational state of one or more ports. Fibre Channel, VSAN, PortChannel, FCIP, or a Gigabit Ethernet port can be tracked. Generally, ports in E and TE port modes can also be Fx ports.
- Linked ports—A port whose operational state is altered based on the operational state of the tracked ports. Only a Fibre Channel port can be linked.

Port Tracking

Before configuring port tracking, consider the following guidelines:

- Verify that the tracked ports and the linked ports are on the same Cisco MDS switch.
- Be aware that the linked port is automatically brought down when the tracked port goes down.
- Do not track a linked port back to itself (for example, Port fc1/2 to Port fc2/5 and back to Port fc1/2) to avoid recursive dependency.

This section includes the following topics:

About Port Tracking

Port tracking has the following features:

- The application brings the linked port down when the tracked port goes down. When the tracked port recovers from the failure and comes back up again, the tracked port is also brought up automatically (unless otherwise configured).
- You can forcefully continue to keep the linked port down, even though the tracked port comes back up. In this case, you must explicitly bring the port up when required.

Enabling Port Tracking

The port tracking feature is disabled by default in all switches in the Cisco 9000 Family. When you enable this feature, port tracking is globally enabled for the entire switch.

To configure port tracking, enable the port tracking feature and configure the linked port(s) for the tracked port.

To enable port tracking with Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **Interfaces**, and then select **Port Tracking** in the Physical Attributes pane.

The port tracking information is displayed in the Information pane shown in Figure 3-2. The default is the **Controls** tab.

Figure 3-2 Port Tracking

Step 2 Click in the Command column to **enable** or **disable** port tracking.

Depending on your selection the corresponding entry in the Status column changes.

Step 3 Click the **Apply Changes** icon to save your changes.

The entry in the Result column changes to **success**.

About Configuring Linked Ports

You can link ports using one of two methods:

- Operationally binding the linked port(s) to the tracked port (default).
- Continuing to keep the linked port down forcefully—even if the tracked port has recovered from the link failure.

Operationally Binding a Tracked Port

When you configure the first tracked port, operational binding is automatically in effect. When you use this method, you have the option to monitor multiple ports or monitor ports in one VSAN.

To operationally bind a tracked port, follow these steps:

Step 1 Expand **Switches**, expand **Interfaces**, and then select **Port Tracking** in the Physical Attributes pane.

The port tracking information is displayed in the Information pane. The default is the Controls tab.

Figure 3-3 Port Tracking Controls Tab

Step 2 Click the **Dependencies** tab.

Step 3 Click **Create Row**.

You see the Create Port Tracking Dependencies dialog box shown in Figure 3-4.

Figure 3-4 Create Port Tracking Dependencies Dialog Box

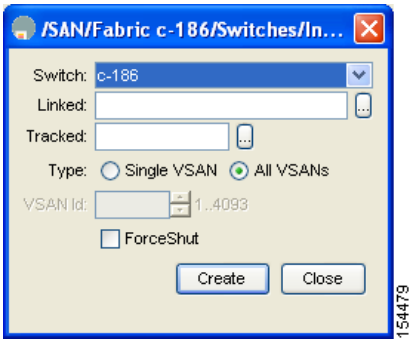

Step 4 Select the switch whose ports you want to track by and selecting a switch from the drop-down list.

Step 5 Select the linked port(s) that should be bound to the tracked port(s) by clicking the browse button and selecting from the list.

Step 6 Click the **Single VSAN** radio button if you want to track these ports only in one VSAN or click the **All VSANs** radio button if you want to track these ports in all the available VSANs.

See "About Monitoring Ports in a VSAN" section on page 3-6 for details.

Step 7 If you chose Single VSAN in the previous step, enter the ID of the VSAN where these ports will be monitored.

Step 8 Check the **Forceshut** check box if you want to forcefully shutdown the tracked port.

See "About Forceful Shutdown" section on page 3-6 for details.

Step 9 Click **Create** to proceed with creating this dependency.

If tracking is established, you see **Success** in the lower left corner of the dialog box (see Figure 3-5).

Figure 3-5 Successful Port Tracking Established

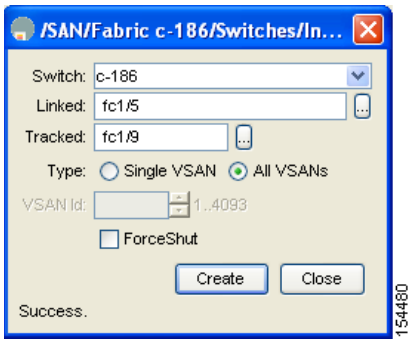

Step 10 Click **Close** to close the dialog box.

About Tracking Multiple Ports

You can control the operational state of the linked port based on the operational states of multiple tracked ports. When more than one tracked port is associated with a linked port, the operational state of the linked port will be set to down only if all the associated tracked ports are down. Even if one tracked port is up, the linked port will stay up.

In Figure 3-6, only if both ISLs 2 and 3 fail, will the direct link 1 be brought down. Direct link 1 will not be brought down if either 2 or 3 are still functioning as desired.

Figure 3-6 Traffic Recovery Using Port Tracking

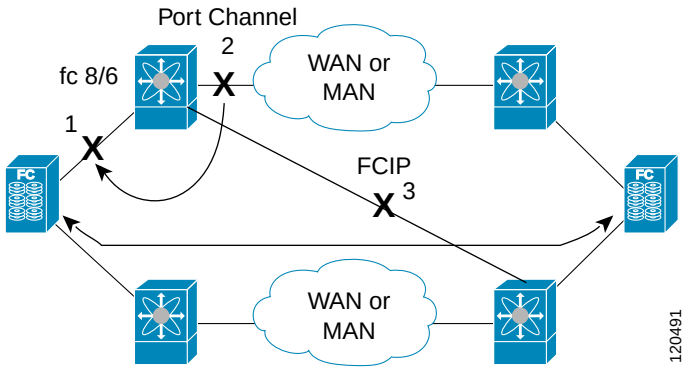

Tracking Multiple Ports

To track multiple ports, see "Operationally Binding a Tracked Port" section on page 3-3.

About Monitoring Ports in a VSAN

You can optionally configure one VSAN from the set of all operational VSANs on the tracked port with the linked port by specifying the required VSAN. This level of flexibility provides higher granularity in tracked ports. In some cases, when a tracked port is a TE port, the set of operational VSANs on the port can change dynamically without bringing down the operational state of the port. In such cases, the port VSAN of the linked port can be monitored on the set of operational VSANs on the tracked port.

If you configure this feature, the linked port is up only when the VSAN is up on the tracked port.

Tip The specified VSAN does not have to be the same as the port VSAN of the linked port.

Monitoring Ports in a VSAN

To monitor a tracked port in a specific VSAN, see "Operationally Binding a Tracked Port" section on page 3-3.

About Forceful Shutdown

If a tracked port flaps frequently, then tracking ports using the operational binding feature may cause frequent topology change. In this case, you may choose to keep the port in the down state until you are able to resolve the reason for these frequent flaps. Keeping the flapping port in the down state forces the traffic to flow through the redundant path until the primary tracked port problems are resolved. When the problems are resolved and the tracked port is back up, you can explicitly enable the interface.

Tip If you configure this feature, the linked port continues to remain in the shutdown state even after the tracked port comes back up. You must explicitly remove the forced shut state (by administratively bringing up this interface) of the linked port once the tracked port is up and stable.

Forcefully Shutting Down a Tracked Port

To forcefully shut down a tracked port, see "Operationally Binding a Tracked Port" section on page 3-3.

Default Port Tracking Settings

Table 3-1 lists the default settings for port tracking parameters.

Table 3-1 Default Port Tracking Parameters

Parameters	Default
Port tracking	Disabled.
Operational binding	Enabled along with port tracking.

INDEX

C

D

E

F

G

I

L

P

Q

R

S

T

V