



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, page 1-1](#)
- [FCC, page 1-3](#)
- [Port Tracking, page 1-4](#)

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 1-3](#)

Send documentation comments to mdsfeedback-doc@cisco.com

- [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.

Send documentation comments to mdsfeedback-doc@cisco.com

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 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, page 1-3](#)
- [Using FCC To Reduce Congestion and Blocking, page 1-4](#)

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

Send documentation comments to mdsfeedback-doc@cisco.com

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 inspect 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 1, “Configuring Fabric Congestion Control and QoS”](#).

Port Tracking

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

Send documentation comments to mdsfeedback-doc@cisco.com

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 1, “Configuring Fabric Congestion Control and QoS”*.

Send documentation comments to mdsfeedback-doc@cisco.com