



Configuring Traffic Management

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.

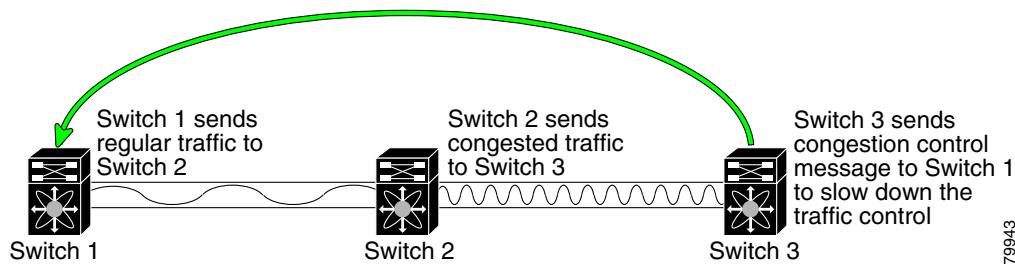
This chapter provides details on the QoS and FCC features provided in all switches.

This chapter contains the following topics:

- [FCC, page 27-1](#)
- [QoS, page 27-2](#)
- [Control Traffic, page 27-3](#)
- [Data Traffic, page 27-3](#)
- [Ingress Port Rate Limiting, page 27-6](#)
- [Default Settings, page 27-6](#)

FCC

FCC reduces the congestion in the fabric without interfering with the standard Fibre Channel protocols. The FCC protocol increases the granularity and the scale of congestion control applied to any class of traffic.

Figure 27-1 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).

FCC Process

When a node in the network detects a 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 quest 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 quest frames. However, only the edge switch processes edge quest frames.

Enabling FCC

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

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

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:

- Control Traffic
- DataTraffic

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 another vendor 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.

Disabling Control Traffic

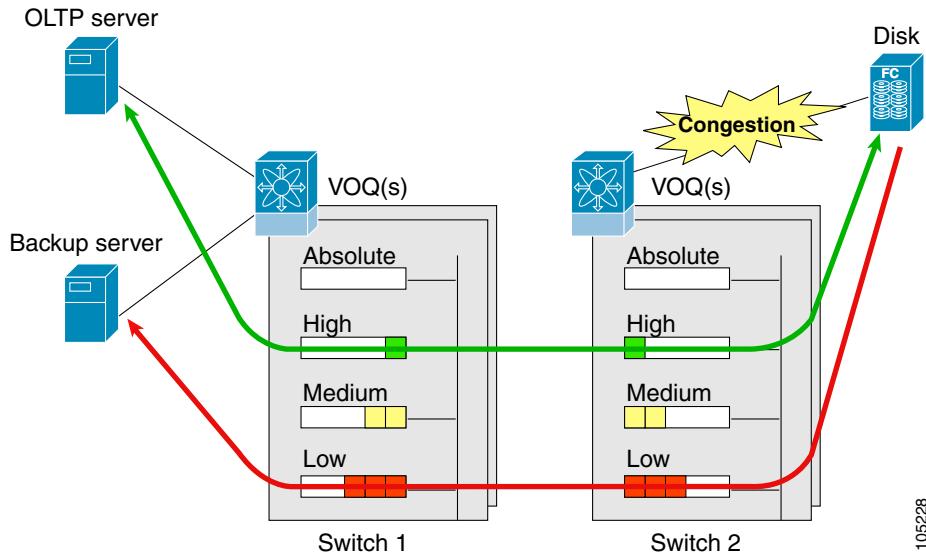
By default, the QoS feature for certain critical control traffic is enabled. These critical control frames are assigned the highest (absolute) priority.

We do not recommended disabling this feature as all critical control traffic will automatically be assigned the lowest priority once you issue this command. You can view the current state of the QoS configuration for critical control traffic using the **show qos statistics** command.

Data Traffic

Transaction processing, a low volume, latency sensitive application, requires quick access to requested information. Backup processing requires high bandwidth but is not sensitive to latency. In a network that does not support service differentiation, all traffic is treated identically; they experience similar latency and get similar bandwidths. The QoS feature in all switches in the Cisco MDS 9000 Family provides these guarantees from SAN-OS Release 1.3(x).

Prior versions of the SAN-OS software only differentiated traffic priority based on control traffic. SAN-OS Release 1.3(x) enables you to take full advantage of the QoS capabilities. Data traffic can now be prioritized in four distinct levels of service differentiation: low, medium, high, or absolute priority. You can apply QoS to ensure that Fibre Channel data traffic for your latency-sensitive applications receive higher priority over throughput-intensive applications like data warehousing.

Figure 27-2 Prioritizing Data Traffic

In Figure 27-2, the OLTP traffic arriving at Switch 1 is marked with a **High** priority level of through 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 as the high priority queue is not congested. The scheduler assigns it priority over the backup traffic in the low priority queue.

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.

To achieve this traffic differentiation, be sure to enable FCC.

Configuring Data Traffic

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.

Enabling QoS 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.

QoS is supported in interoperability mode. The effectiveness of the setting depends on the location of MDS switches in the fabric relative to the location of the source or destination of the prioritized devices.

Creating Class Maps

Use the **class-map** option 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. 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—Use the **source-wwn** option to specify the source WWN or the **destination-wwn** option to specify the destination WWN.
- Fibre Channel ID (FC ID)—Use the **source-address** option to specify the source ID (SID) or the **destination-address** option to specify the destination ID (DID). The possible values for mask are:
 - FFFFFF—The entire FCID is used. This is the default.
 - FFFF00—Only domain and area FCID is used.
 - FF0000—Only domain FCID is used.A source-address or destination-address of 0x000000 is not allowed.
- Source interface—Use the **input-interface** option to specify the ingress interface.

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

Defining Service Policies

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.

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. See <http://www.cisco.com/warp/public/105/dscvalues.html#dscpanassuredforwardingclasses> for further information on implementing QoS DSCP values.

Use the **policy-map** option to specify the class of service. The policy map name is restricted to 63 alphanumeric characters.

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

Applying a Service Policy

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 will not be enforced. You can only apply one policy map to a VSAN.

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

Scheduling Traffic

The SAN-OS software supports four scheduling queues:

- Strict-priority queues are queues that are serviced in preference to other queues. A strict-priority queue 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 queues:
 - Use **dwrr-q high** option to schedule high priority traffic.
 - Use **dwrr-q medium** option to schedule medium priority traffic.
 - Use **dwrr-q low** option to schedule low priority traffic.

Use the **qos dwrr-q** command to associate a weight with a DWRR queue.

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 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).

Ingress Port Rate Limiting

A port rate limiting feature is available in SAN-OS 1.3(x). This feature helps control the bandwidth for individual FC ports. Port rate limiting is also referred to as ingress rate limiting because it controls ingress traffic into a FC 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.

Port rate limiting can only be configured in switches in the Cisco MDS 9100 Series.

This command can only be configured if the following conditions hold true:

- The QoS feature is enabled using the **qos enable** command.
- The command is issued in a Cisco MDS 9100 series switch.

The rate limit ranges from 1 to 100% and the default is 100%.

Default Settings

[Table 27-1](#) lists the default settings for FCC, QoS, and rate limiting features:

Table 27-1 Default FCC, QoS, and Rate Limiting Settings

Parameters	Default
FCC protocol	Disabled.
QoS control traffic	Enabled.
QoS data traffic	Disabled.
Rate limit	100%

