



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, page 47-2](#)

[QoS, page 47-3](#)

[Control Traffic, page 47-4](#)


[Data Traffic, page 47-5](#)

[Ingress Port Rate Limiting, page 47-13](#)

[Default Settings, page 47-14](#)

FCC

Figure 47-1 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 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, follow these steps:

	Command	Purpose
Step 1	switch# config t	
Step 2	switch(config)# fcc	
	switch(config)# no fcc	

Assigning FCC Priority

To assign FCC priority, follow these steps:

	Command	Purpose
Step 1	switch# config t	
Step 2	switch(config)# fcc priority 2	Defines the FCC priority threshold to have a priority of 2—0 is the lowest priority and 7 is the highest priority.

Displaying FCC

Use the **show fcc**

Example 47-1 Displays Configured FCC Information

```
switch# show fcc
fcc is disabled
fcc is applied to frames with priority up to 4
```

QoS

-
-

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.

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.

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

switch# config t	
switch(config)# no qos control priority 0	
switch(config)# qos control priority 0	

Displaying Control Traffic Information

Use the **show qos statistics** command to view the current state of the QoS configuration for critical control traffic. This command displays the current QoS settings along with the number of frames marked high priority. The count is only for debugging purposes and cannot be configured (see [Example 47-2](#)).

Example 47-2 Displays Current QoS Settings

```

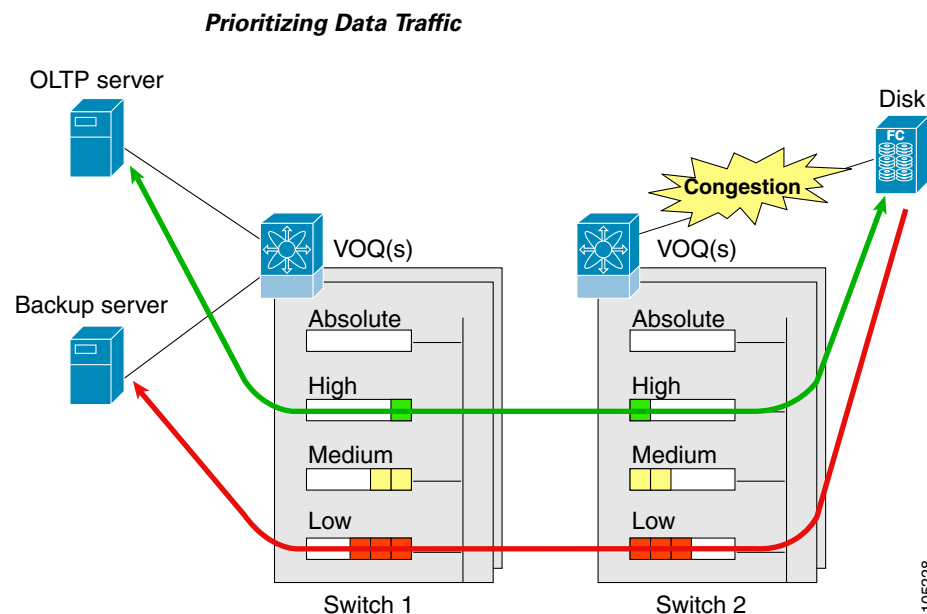
show qos statistics
Total number of FC frames transmitted from the Supervisor= 15767
Number of highest-priority FC frames transmitted           = 8224
Current priority of FC control frames = 0    (0 = lowest; 7 = highest)

```

Data Traffic

Online transaction processing, 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 47-2](#)).



In [Figure 47-2](#), 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.



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.



[“Enabling FCC” section on page 47-3](#)).

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 47-1](#) highlights the differences between configuring QoS priorities based on VSANs versus zones.

Table 47-1 QoS Configuration Differences

VSAN-Based QoS	Zone-Based QoS

See the [“Zone-Based Traffic Priority” section on page 19-15](#) for details on configuring a zone-based QoS policy.

To configure QoS, follow these steps.

Enable the QoS feature.

Create and define class maps.

Step 3

Step 4

QoS Initiation for Data Traffic



Tip

<code>qos enable</code>	
<code>no qos enable</code>	

Class Map Creation

belonging to that class. The class map name is restricted to 63 alphanumeric characters and defaults to the match-all option. Flow-based traffic uses one of the following values:

- WWN—The source WWN or the destination WWN.
- Fibre Channel ID (FC ID) —The source ID (SID) or the destination ID (DID). The possible values for mask are FFFFFFFF (the entire FC ID is used—this is the default), FFFF00 (only domain and area FC ID is used), or FF0000 (only domain FC ID is used).



Note A 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.

Use the **class-map**

source-wwn

destination-wwn

source-address

destination-address

input-interface

destination-device-alias

	<code>qos class-map MyClass</code>	
	<code>qos class-map MyClass</code>	
	<code>match-all</code>	
	<code>qos class-map MyClass</code>	
	<code>match-any</code>	
	<code>match</code>	
	<code>destination-address 0x12ee00</code>	
	<code>match source-address</code>	
	<code>0x6d1090 mask 0xFFFFFF</code>	
	<code>match destination-wwn</code>	
	<code>20:01:00:05:30:00:28:df</code>	
	<code>match source-wwn</code>	
	<code>23:15:00:05:30:00:2a:1f</code>	
	<code>match</code>	
	<code>destination-device-alias DocDeviceAlias</code>	
Step 5	<code>match input-interface fc</code>	
	<code>2/1</code>	
Step 6		

Service Policy Definition



Note

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



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

	Command	Purpose
Step 1	<code>qos policy-map MyPolicy</code>	
	<code>no qos policy-map OldPolicy</code>	
	<code>class MyClass</code>	
	<code>no class OldClass</code>	
	<code>priority high</code>	
	<code>no priority high</code>	
	<code>dscp 2</code>	
	<code>no dscp 60</code>	



	<code>qos service policy MyPolicy</code>	
vsan 3		
	<code>no qos service policy OldPolicy</code>	
vsan 7		

DWRR Traffic Scheduler

-
-

qos dwrr-q
dwrr-q medium
dwrr-q high

dwrr-q low

switch(config)#	
switch(config)#	

show qos

Example 47-3 Displays the Contents of all Class Maps

```

qos class-map MyClass match-any
  match destination-wwn 20:01:00:05:30:00:28:df
  match source-wwn 23:15:00:05:30:00:2a:1f
  match input-interface fc2/1
qos class-map Class2 match-all
  match input-interface fc2/14
qos class-map Class3 match-all
  match source-wwn 20:01:00:05:30:00:2a:1f

```

```

switch#
qos class-map MyClass match-any
  match destination-wwn 20:01:00:05:30:00:28:df
  match source-wwn 23:15:00:05:30:00:2a:1f
  match input-interface fc2/1

```

Example 47-5 Displays All Configured Policy Maps

```

qos policy-map MyPolicy
  class MyClass
  priority medium
qos policy-map Policy1
  class Class2
  priority low

```

Example 47-6 Displays a Specified Policy Map

Example 47-7 Displays Scheduled DWRR Configurations

Example 47-8 Displays All Applied Policy Maps

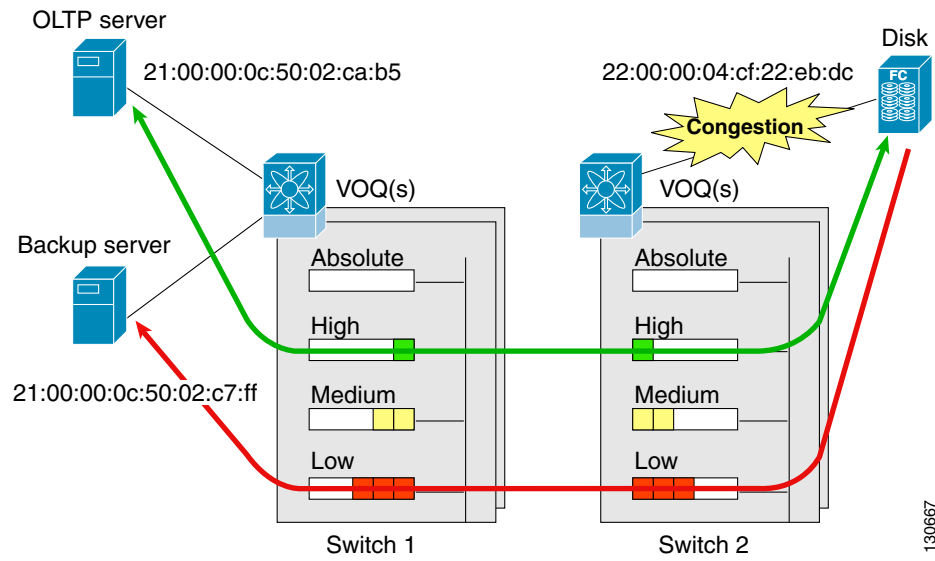
Example 47-9 Displays the Policy Map Associated with a Specified VSAN

Example 47-10 Displays the Class Map Associated with a Specified Interface

Example 47-11 Displays QoS Statistics

```
Number of highest-priority FC frames transmitted          = 137679
Current priority of FC control frames = 7      (0 = lowest; 7 = highest)
```

Example Configuration



Step 1

```
qos class-map jc1 match-all
  match source-wwn 21:00:00:0c:50:02:ca:b5
  match destination-wwn 22:00:00:04:cf:22:eb:dc
  exit
qos class-map jc2 match-all
  match source-wwn 21:00:00:0c:50:02:c7:ff
  match destination-wwn 22:00:00:04:cf:22:eb:dc
  exit

qos policy-map jp1
  class jc1
    priority high
  exit
  class jc2
```

Step 3

Step 4

Step 5

Step 1

Step 2

Step 3

Ingress Port Rate Limiting

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, follow these steps.

	Enters the configuration mode.
	Selects the interface to specify the ingress port rate limit.
	Configures a 50% port rate limit for the selected interface.
	Reverts a previously configured rate to the factory default of 100%.

Default Settings

Default FCC, QoS, and Rate Limiting Settings

Parameters	Default