CHAPTER 25

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. It includes the following sections:

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 (see Figure 25-1).

Figure 25-1 FCC Mechanisms

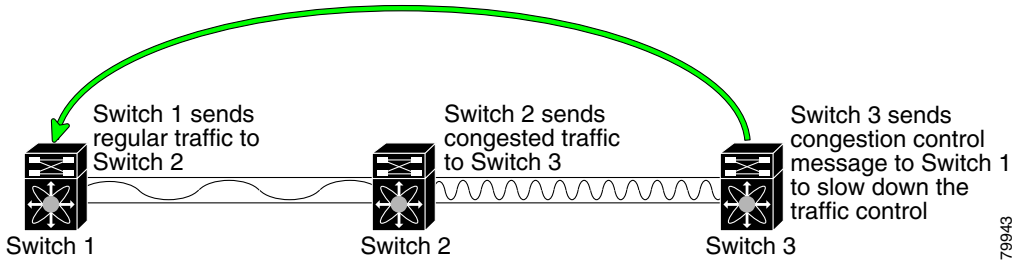

Edge quench congestion control provides feedback to the source about the rate at which frames should be injected into the network (frame intervals).

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 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.

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`	Enters configuration mode.
Step 2	`switch(config)# fcc`	Enables FCC in this switch.
	`switch(config)# no fcc`	Disables FCC in this switch (default).

Assigning FCC Priority

To assign FCC priority, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
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** command to view FCC settings (see Example 25-1).

Example 25-1 Displays Configured FCC Information

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

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, page 25-4
- Data Traffic, page 25-5

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:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# no qos control priority 0`	Enables the control traffic QoS feature.
	`switch(config)# qos control priority 0`	Disables the control traffic QoS feature.

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 25-2).

Example 25-2 Displays Current QoS Settings

```
switch# 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

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 as of Cisco MDS SAN-OS Release 1.3.

Prior versions of the Cisco SAN-OS software only differentiated traffic priority based on control traffic. Cisco MDS SAN-OS Release 1.3 enables you to take full advantage of the QoS capabilities. 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 like data warehousing (see Figure 25-2).

Figure 25-2 Prioritizing Data Traffic

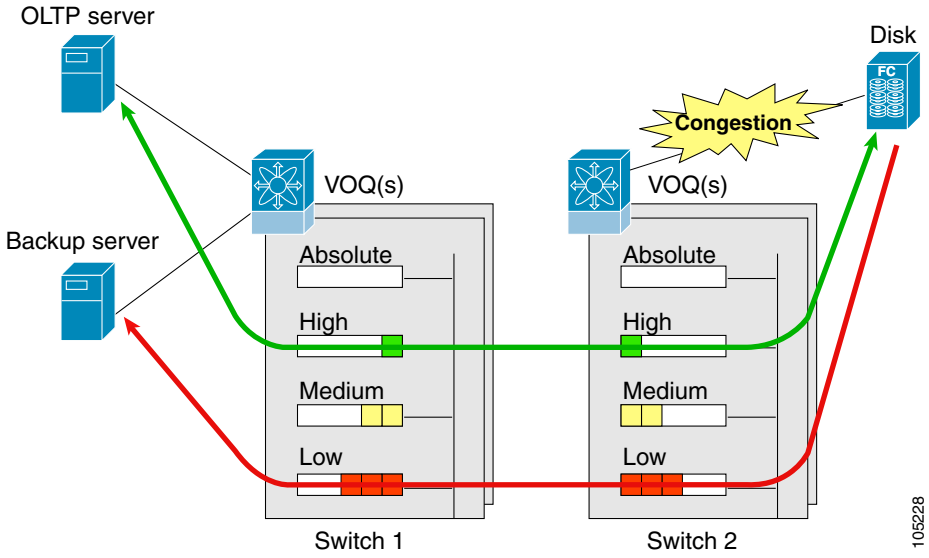

In Figure 25-2, the OLTP traffic arriving at Switch 1 is marked with a high priority level of through classification (class map) and marking (policy map). Similarly, the backup traffic is marked with a low priority level. The traffic is sent to the corresponding priority queue within a virtual output 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.

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 25-3).

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.

To enable the QoS data traffic feature, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# qos enable`	Enables QoS. You can now configure data traffic parameters.
	`switch(config)# no qos enable`	Removes the currently applied QoS configuration and disables QoS. You can no longer configure data traffic parameters.

Tip QoS is supported in interoperability mode—its effectiveness depends on the location of Cisco 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 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 **source-address** or **destination-address** of 0x000000 is not allowed.

- Source interface—Use the **input-interface** option to specify the ingress interface.

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

To create a class map, follow these steps:

	Command	Purpose
Step 1	`switch(config)# qos class-map MyClass` `switch(config-cmap)#`	Creates a class map called MyClass and places you in the class-map submode to match all criteria specified for this class.
	`switch(config)# qos class-map MyClass match-all` `switch(config-cmap)#`	Specifies a logical AND operator for all matching statements in this class. If a frame matches all (default) configured criteria, it qualifies for this class. This is the default.
	`switch(config)# qos class-map MyClass match-any` `switch(config-cmap)#`	Specifies a logical OR operator for all matching statements in this class. If a frame matches any one configured criteria, it qualifies for this class.
Step 2	`switch(config-cmap)# match destination-address 0x12ee00`	Specifies a destination address match for frames with the specified destination FC ID.
	`switch(config-cmap)# match source-address 0x6d1090 mask 0xFFFFFF`	Specifies a source address and mask match for frames with the specified source FC ID.
Step 3	`switch(config-cmap)# match destination-wwn 20:01:00:05:30:00:28:df`	Specifies a destination WWN to match frames.
	`switch(config-cmap)# match source-wwn 23:15:00:05:30:00:2a:1f`	Specifies a source WWN to match frames.
Step 4	`switch(config-cmap)# match input-interface fc 2/1`	Specifies a source interface to match frames.
Step 5	`switch(config-cmap)# no match input-interface fc 3/5`	Removes a match based on the specified source interface.

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.

Note Refer to http://www.cisco.com/warp/public/105/dscpvalues.html#dscpandassuredforwardingclasses 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.

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

To specify a service policy, follow these steps:

	Command	Purpose
Step 1	`switch(config)# qos policy-map MyPolicy` `switch(config-pmap)#`	Creates a policy map called MyPolicy and places you in the policy-map submode.
	`switch(config)# no qos policy-map OldPolicy` `switch(config)#`	Deletes the policy map called OldPolicy and places you in the policy-map submode.
Step 2	`switch(config-pmap)# class MyClass` `switch(config-pmap-c)#`	Specifies the name of a predefined class and places you at the policy-map submode for that class.
	`switch(config-pmap)# no class OldClass`	Removes the class map called OldClass from the policy map.
Step 3	`switch(config-pmap-c)# priority high`	Specifies the priority to be given for each frame matching this class.
	`switch(config-pmap-c)# no priority high`	Deletes a previously assigned priority and reverts to the default value of low.
Step 4	`switch(config-pmap-c)# dscp 2`	Specifies the DSCP value to mark each frame matching this class.
	`switch(config-pmap-c)# no dscp 60`	Deletes a previously assigned DSCP value and reverts to the factory default of 0.

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

To apply a service policy, follow these steps:

	Command	Purpose
Step 1	`switch(config)# qos service policy MyPolicy vsan 3`	Applies a configured policy to VSAN 3.
	`switch(config)# no qos service policy OldPolicy vsan 7`	Deletes a configured policy that was applied to VSAN 7.

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

Scheduling Traffic

The Cisco SAN-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 queues:
  - Use the **dwrr-q high** option to schedule high priority traffic.
  - Use the **dwrr-q medium** option to schedule medium priority traffic.
  - Use the **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 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).

To associate a weight with a DWRR queue, follow these steps:

	Command	Purpose
Step 1	`switch(config)# qos dwrr-q high weight 10`	Associates a relative weight (10) to a specified queue (default queue).
	`switch(config)# no dwrr-q low weight 51`	Restores the default weight of 20.

Displaying Data Traffic Information

The **show qos** commands display the current QoS settings for data traffic (see Examples 25-3 to 25-11).

Example 25-3 Displays the Contents of all Class Maps

```
switch# show qos class-map
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
```

Example 25-4 Displays the Contents of a Specified Class Map

```
switch# show qos class-map name MyClass
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 25-5 Displays All Configured Policy Maps

```
switch# show qos policy-map
qos policy-map MyPolicy
    class MyClass
    priority medium

qos policy-map Policy1
    class Class2
    priority low
```

Example 25-6 Displays a Specified Policy Map

```
switch# show qos policy-map name MyPolicy
qos policy-map MyPolicy
    class MyClass
        priority medium
```

Example 25-7 Displays Scheduled DWRR Configurations

```
switch# show qos dwrr
qos dwrr-q high weight 50
qos dwrr-q medium weight 30
qos dwrr-q low weight 20
```

Example 25-8 Displays All Applied Policy Maps

```
switch# show qos service policy
qos service policy MyPolicy vsan 1
qos service policy Policy1 vsan 4
```

Example 25-9 Displays the Policy Map Associated with a Specified VSAN

```
switch# show qos service policy vsan 1
```

```
qos policy-map pmap1
   class cmap1
       priority medium
   class cmap2
       priority high
```

Example 25-10 Displays the Class Map Associated with a Specified Interface

```
switch# show qos service policy interface fc3/10
qos policy-map pmap1
   class cmap3
       priority high
   class cmap4
       priority low
```

Example 25-11 Displays QoS Statistics

```
switch# show qos statistics
Total number of FC frames transmitted from the Supervisor= 301431
Number of highest-priority FC frames transmitted         = 137679
Current priority of FC control frames = 7    (0 = lowest; 7 = highest)
```

Ingress Port Rate Limiting

A port rate limiting feature is available in Cisco SAN-OS 1.3. 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.

Note 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 form 1 to 100% and the default is 100%.

To configure the port rate limiting value, follow these steps.

	Command	Purpose
Step 1	`switch # config t` `switch(config)#`	Enters the configuration mode.
Step 2	`switch(config)# interface fc 1/1`	Selects the interface to specify the ingress port rate limit.
Step 3	`switch(config-if)# switchport ingress-rate 50`	Configures a 50% port rate limit for the selected interface.
	`switch(config-if)# no switchport ingress-rate 50`	Reverts a previously configured rate to the factory default of 100%.

Default Settings

Table 25-1 lists the default settings for FCC, QoS, and rate limiting features:

Table 25-1 Default FCC, QoS, and Rate Limiting Settings

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