



Configuring Queuing and Scheduling on F-Series I/O Modules

This chapter describes how to configure the QoS queuing and scheduling features on the F-Series I/O module of the Cisco NX-OS device.

- [Finding Feature Information, page 1](#)
- [Information About Queuing and Scheduling, page 1](#)
- [Prerequisites for Queuing and Scheduling, page 11](#)
- [Guidelines and Limitations, page 11](#)
- [Configuring Queuing and Scheduling, page 14](#)
- [Verifying the Queuing and Scheduling Configuration, page 18](#)
- [Configuration Examples for Queuing and Scheduling on F-Series Modules, page 19](#)
- [Feature History for Queuing and Scheduling for F-Series Modules, page 21](#)

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

Information About Queuing and Scheduling

On an F-Series module, a queuing policy is closely coupled with the network qos policy. For each network qos policy that is activated, its corresponding default queuing policy is automatically selected for the system target. In the ingress direction, either two or four queues (buffer pools) are formed depending on the policy template. In the egress direction, there are four physical queues for qos policy templates on Cisco Nexus 7000 Series devices, except on the Cisco Nexus 7710 switch and Cisco Nexus 7718 switch, where, beginning with Cisco Release 6.2(2), there is support for eight physical queues.

The system queuing policy applied by default can be overridden on a per-port basis. In general, the user-configured queuing policies are per virtual device context (VDC).

Ingress queuing determines the following attributes:

- Queue-limit—Amount of buffers to be allocated for a class of service (CoS).
- Bandwidth—Priority grouping and its bandwidth allocation advertised using the Data Center Bridging Capability Exchange Protocol (DCBXP).
- Set CoS—Untrusted port default CoS (similar to the M1 modules).

Egress queuing determines the following attributes:

- Bandwidth—Differential Weighted Round Robin (DWRR) bandwidth for a given queue and the group.
- Priority level—The priority level of the queue.
- Shape—The shaper for the queue.

Ingress Queuing

You use the ingress queuing to partition the port ingress buffers that are 1.25 MB and an additional 256 KB (a total of 1.5 MB) to absorb the frames in transit after pause has been sent. This buffer is partitioned among the eight CoS values. The number of partitions is fixed for a given network qos template. The incoming CoS values are mapped to each partition. Each buffer partition is considered as an ingress queue.

There is a high threshold and a low threshold at which the pause or resume frames are generated when a threshold is met. This requirement is applicable to the no-drop CoS only. The frames that are in transit are absorbed by a skid buffer after a pause is generated. If the number of frames exceed the skid buffer threshold, the frames are tail dropped. There are three thresholds for drop eligible (DE), non-DE, and Bridge Protocol Data Unit (BPDU) frames for dropping. For the drop CoS, the high and low thresholds are the same.

The default policy ingress queues are created as follows:

- Different queues per drop class:
Drop queue = 70% buffers; no-drop queue = 30% buffers
- Different queues for priority and nonpriority CoS in a given drop class:
Nonpriority queue = 90% buffers; priority queue = 10% buffers

Each network qos policy has a corresponding default ingress queuing policy (template) and is automatically activated for the system. They are the default-4q-8e-in-policy, default-4q-7e-in-policy, default-4q-6e-in-policy, default-4q-4e-in-policy, default-8e-4q8q-in-policy, default-7e-4q8q-in-policy, default-6e-4q8q-in-policy, default-4e-4q8q-in-policy, and default-8e-4q4q-in-policy.

The predefined class map names (queue names) for ingress queuing are described in the table below.

Table 1: Predefined Class Maps for Ingress Queuing

Ingress Policy Maps	Ingress Class Map Names
default-4q-8e-in-policy	2q4t-8e-in-q1 and 2q4t-8e-in-q-default

Ingress Policy Maps	Ingress Class Map Names
default-4q-7e-in-policy	4q4t-7e-in-q1, 4q4t-7e-in-q-default, 4q4t-7e-in-q3, and 4q4t-7e-in-q4
default-4q-6e-in-policy	4q4t-6e-in-q1, 4q4t-6e-in-q-default, 4q4t-6e-in-q3, and 4q4t-6e-in-q4
default-4q-4e-in-policy	4q4t-4e-in-q1, 4q4t-4e-in-q-default, 4q4t-4e-in-q3, and 4q4t-4e-in-q4
default-8e-4q4q-in-policy	4q1t-8e-4q4q-in-q1, 4q1t-8e-4q4q-in-q-default, 4q1t-8e-4q4q-in-q3, and 4q1t-8e-4q4q-in-q4
default-8e-4q8q-in-policy (on Cisco Nexus 7710 / 7718 switches only)	8e-4q8q-in-q1, 8e-4q8q-in-q-default, 8e-4q8q-in-q3, and 8e-4q8q-in-q4
default-7e-4q8q-in-policy (Cisco Nexus 7710 / 7718 switches only) <ul style="list-style-type: none"> • default-7e-4q8q-drop-in-policy • default-7e-4q8q-ndrop-in-policy 	c-7e-4q8q-drop-in, c-7e-4q8q-ndrop-in 7e-4q8q-in-q1, 7e-4q8q-in-q-default and 7e-4q8q-in-q3 7e-4q8q-in-q4
default-6e-4q8q-in-policy (Cisco Nexus 7710 / 7718 switches only) <ul style="list-style-type: none"> • default-6e-4q8q-drop-in-policy • default-6e-4q8q-ndrop-in-policy 	c-6e-4q8q-drop-in and c-6e-4q8q-ndrop-in 6e-4q8q-in-q1 and 6e-4q8q-in-q-default 6e-4q8q-in-q3 and 6e-4q8q-in-q4
default-4e-4q8q-in-policy (Cisco Nexus 7710 / 7718 switches only) <ul style="list-style-type: none"> • default-4e-4q8q-drop-in-policy • default-4e-4q8q-ndrop-in-policy 	c-4e-4q8q-drop-in and c-4e-4q8q-ndrop-in 4e-4q8q-in-q1 and 4e-4q8q-in-q-default 4e-4q8q-in-q3 and 4e-4q8q-in-q4

**Note**

- The naming conventions of the queue are similar to the M1 modules. Also, the process for referring to queuing class maps and changing CoS to queue maps is also similar to M1 modules.
- When a port becomes part of a port channel, the port inherits the policy of the port channel. When the port is moved out of the port channel, the default system queuing policy gets activated on the port.

By default, the queuing policy maps the priority CoS values (CoS 5-7) and nonpriority CoS values (CoS 0-4) into different ingress queues (IVL). CoS to ingress queue mapping is configured from the default VDC and the configuration is applied system wide. A network administrator user role is required to change CoS to IVL.

Starting with the Cisco NX-OS 6.1 release, DSCP to IVL is supported on F2 modules, in the ingress direction, using the `match dscp` command with the `2q4t-8e-in-q1` class map and the `2q4t-8e-in-q-default` class map.

**Note**

Starting with the Cisco NX-OS 6.1(2) release, DSCP to IVL is supported on IPV6 using F2e modules.

Guidelines for the **match dscp** command are as follows:

- The **match dscp** command is applicable only to queues that have at least one CoS value associated with it. If all DSCP values are not mapped to a nondefault ingress queue, the default queue should have the CoS values associated with it.
- DSCP queuing is automatically disabled when the user removes all **match dscp** commands (using **no match** statements).
- If the **match dscp** command is used in the `2q4t-8e-in-q1` class map to set some DSCP values, all remaining DSCP values are automatically mapped to the default queue.

Bridged Traffic	<p>When DSCP to ingress queue is enabled, the ingress queue selection is based on the DSCP value. However, the egress queue selection is based on the CoS value of the packet. If a CoS value does not exist, then all packets are accepted as CoS 0.</p> <p>To override this behavior and use DSCP for egress queue selection, an ingress QoS policy has to be applied that matches the DSCP value and uses the set dscp value command to set the DSCP to the same value as the matched DSCP on the class map.</p> <p>When CoS to ingress queue is enabled, the ingress queue selection is based on the CoS value. The egress queue selection is based on the CoS value of the packet, including non-IP packets. If a CoS value does not exist, then all packets are accepted as CoS 0.</p>
-----------------	---

<p>Routed Traffic without Proxy Mode (native F-Series modules)</p>	<p>For port type Layer 3-to-Layer L3/SVI— When the DSCP-to-ingress queue is enabled, the ingress queue selection is based on the DSCP. However, the egress queue selection is based on the DSCP value when the Ingress Port is a Layer 3 type.</p> <p>When CoS-to-ingress queue is enabled for subinterfaces, the ingress queue selection is based on the CoS. However, the egress queue selection is based on the DSCP value when the ingress port is a Layer 3 type(subinterface) because the packets are routed and the DSCP-to-egress queue takes place.</p> <p>For port type Layer 2 SVI-to-SVI Layer 2— The egress queue selection is based on the DSCP value and does not rely on whether DSCP-to-ingress queue is enabled in ingress or not.</p> <p>When the CoS-to-ingress queue is enabled, the ingress queue selection is based on the CoS. The egress queue selection is based on the CoS value, or 0 for no CoS value, because the forwarding decision uses the CoS value to drive the egress queue. To use the DSCP/CoS value for egress queue selection, apply the type QoS policy to mutate the DSCP values one-to-one or set the DSCP values to one-to-one.</p>
<p>Routed Traffic with Proxy Mode (mixed F-Series and M-Series modules)</p>	<p>When the DSCP-to-ingress queue is enabled, the ingress queue selection is based on the DSCP. However, the egress queue selection is based on the DSCP value because packets are proxy routed on the M-Series module and the DSCP-to-egress queue takes place.</p> <p>When CoS-to-ingress queue is enabled, the ingress queue selection is based on the CoS. However, the egress queue selection is based on the DSCP value because packets are proxy routed on the M-Series module and the DSCP-to-egress queue takes place.</p>

The following table contains an example of when the match dscp command is used in the 2q4t-8e-in-q1 class map to set specific DSCP values.

Commands	Description
<pre>class-map type queuing match-any 2q4t-8e-in-q1 match cos 5-7 match dscp 40-45</pre>	<p>The values set by the match dscp command are displayed by the show run command.</p>

Commands	Description
<pre>class-map type queuing match-any 2q4t-8e-in-q-default match dscp 0-39,46-63.</pre>	<p>The remaining DSCP values (0-39, 46-63) are automatically mapped to the default queue.</p> <p>The values associated with the default queue are not displayed by the show run command. These values are implicitly programmed in the hardware.</p>
<pre>class-map type queuing match-any 2q4t-8e-in-q-default match dscp 40-45</pre>	<p>When specific DSCP values are mapped to the default queue (2q4t-8e-in-q-default), the remaining DSCP values are automatically mapped to the default queue.</p> <p>There is no restriction when specifying all of the remaining DSCP values in the default queue.</p> <p>The values set by the match dscp command are displayed by the show run command.</p>
<pre>class-map type queuing match-any 2q4t-8e-in-q-default match dscp 0-39,46-63</pre>	<p>The DSCP values (0-39, 46-63) are automatically mapped to the default queue (2q4t-8e-in-q-default).</p> <p>The values associated with the default queue are not displayed by the show run command. These values are implicitly programmed in the hardware.</p>

**Note**

Modifying the default queuing policy maps is a disruptive operation that might cause frame drops.

You can assign a bandwidth percentage to each ingress queue. The CoS values (priority group) of each queue and its bandwidth are relayed to the peer using the DCBXP.

With the Enhanced Transmission Selection (ETS; specifies scheduling of queues based on priority) implementation, when you define both the drop and no-drop classes in a non-8e network qos policy template, the queuing follows a hierarchical pattern. In a hierarchical queuing pattern, queues within a class are configured with respect to the buffer at the first level, and buffers across the queuing groups are configured at the second level.

You use the **queue-limit** command to tune the ingress queue sizes (buffers). You can define the percentage of the total buffer to be allocated to the queue. For more information about the queue-limit command, see the *Cisco Nexus 7000 Series NX-OS Quality of Service Command Reference*.

You use the **bandwidth** command to control the bandwidth allocated to the traffic classes (CoS) in the ingress queue. The bandwidth allocated to a traffic class in the ingress queue does not impact the switch. Instead, it sends the bandwidth information to the peer as an indication of the bandwidth for the traffic classes (CoS) that the peer sends. For more information about the bandwidth command, see the *Cisco Nexus 7000 Series NX-OS Quality of Service Command Reference*.

You use the **set cos** command only on the default queue to make a port that is untrusted on the default queue.

Starting with Cisco NXOS 6.2(2) Release, default dscp values are provided for all the following five templates on F-Series Modules:

- default-nq-4e-policy template 4e

- default-nq-6e-policy template 6e
- default-nq-7e-policy template 7e
- default-nq-8e-policy template 8e
- default-nq-8e-4q4q-policy template 8e-4q4q

The following table lists the default dscp values for 4q mode templates:

Ingress Queue	DSCP Map Value
Template: default-nq-4e-policy template 4e	
4q4t-4e-in-q-default	0-39
4q4t-4e-in-q1	40-63
4q4t-4e-in-q3	—
4q4t-4e-in-q4	—
default-nq-6e-policy template 6e	
4q4t-6e-in-q-default	0-39
4q4t-6e-in-q1	40-63
4q4t-6e-in-q3	—
4q4t-6e-in-q4	—
default-nq-7e-policy template 7e	
4q4t-7e-in-q-default	0-15
4q4t-7e-in-q1	40-63
4q4t-7e-in-q3	16-39
4q4t-7e-in-q4	
default-nq-8e-policy template 8e	
2q4t-8e-4q4q-in-q-default	0-39
2q4t-8e-4q4q-in-q1	40-63
default-nq-8e-4q4q-policy template 8e-4q4q	
4q1t-8e-4q4q-in-q-default	0-15

Ingress Queue	DSCP Map Value
4q1t-8e-4q4q-in-q1	40-63
4q1t-8e-4q4q-in-q3	24-39
4q1t-8e-4q4q-in-q4	16-23

Similarly, the default dscp values are mapped for ingress queues for Cisco 7710/7718 switches.

Egress Queuing

You use egress queuing to determine how to schedule the traffic from the egress queues out of a port. The class map names represent queues and match cos represents the CoS values mapped to them. You can modify the egress class map and match cos to achieve the desired CoS-to-queue mapping.



Note

CoS remapping is supported only in strict F-Series VDCs. It is not supported in F-Series/M1 mixed VDCs.

Each egress port has about 0.7 MB of buffers that are distributed equally among the 8 CoS values. A CoS has approximately 0.1 MB of buffers.

The default policy egress queues are created as follows:

- The drop and no-drop CoS must be mapped to different queues.
- The priority CoS is mapped to a strict priority (SP) queue. All the nonpriority CoS values are mapped to a DWRR queue.
- For all the non-8e templates, second level scheduling is used.



Note

- Egress queues have a fixed size and are not user configurable.
- The egress port has four queues, except for the Cisco Nexus 7710 switch and the Cisco Nexus 7718 switch, whereby, beginning with Cisco Release 6.2(2), has support for eight queues (4q8q mode).

Each network qos policy has a corresponding default egress queuing policy (template) and is automatically activated for the system. They are the default-4q-8e-out-policy, default-4q-7e-out-policy, default-4q-6e-out-policy, default-4q-4e-out-policy, default-8e-4q8q-out-policy, default-7e-4q8q-out-policy, default-6e-4q8q-out-policy, default-4e-4q8q-out-policy and the default-8e-4q4q-out-policy. The flexible egress queues configuration is based on these queue types— 1p7qlt-8e, 1p7qlt-7e, 1p3qlt-8e, 1p3qlt-7e, 2p2qlt-4e, 2p6qlt-4e, 3p1qlt-6e, and 3p5qlt-6e.

For the Cisco Nexus 7710 switch and the Cisco Nexus 7718 switch, a hierarchical scheduling pattern is followed on the 7e-4q8q, 6e-4q8q, and 4e-4q8q templates.

The predefined class map names (queue names) for egress queuing are described in the table below.

Table 2: Predefined Class Maps for Egress Queuing

Egress Policy Names	Egress Class Map Names
default-4q-8e-out-policy	1p3q1t-8e-out-pq1, 1p3q1t-8e-out-q2, 1p3q1t-8e-out-q3, and 1p3q1t-8e-out-q-default
default-4q-7e-out-policy	1p3q1t-7e-out-pq1, 1p3q1t-7e-out-q2, 1p3q1t-7e-out-q3, and 1p3q1t-7e-out-q-default
default-4q-6e-out-policy	3p1q1t-6e-out-pq1, 3p1q1t-6e-out-pq2, 3p1q1t-6e-out-pq3, and 3p1q1t-6e-out-q-default
default-4q-4e-out-policy	2p2q1t-4e-out-pq1, 2p2q1t-4e-out-pq2, 2p2q1t-4e-out-q3, and 2p2q1t-4e-out-q-default
default-8e-4q4q-out-policy	1p3q1t-8e-4q4q-out-pq1, 1p3q1t-8e-4q4q-out-q2, 1p3q1t-8e-4q4q-out-q3, and 1p3q1t-8e-4q4q-out-q-default
default-8e-4q8q-out-policy (Cisco Nexus 7710 / 7718 switches only)	8e-4q8q-out-q1 (priority queue), 8e-4q8q-out-q2, 8e-4q8q-out-q3, 8e-4q8q-out-q4, 8e-4q8q-out-q5, 8e-4q8q-out-q6, 8e-4q8q-out-q7, and 8e-4q8q-out-q-default
default-7e-4q8q-out-policy (Cisco Nexus 7710 / 7718 switches only) <ul style="list-style-type: none"> • default-7e-4q8q-drop-out-policy • default-7e-4q8q-ndrop-out-policy 	c-7e-4q8q-drop-out and c-7e-4q8q-ndrop-out 7e-4q8q-out-q1 (priority queue), 7e-4q8q-out-q2, 7e-4q8q-out-q3, 7e-4q8q-out-q4, 7e-4q8q-out-q6, 7e-4q8q-out-q7, and 7e-4q8q-out-q-default 7e-4q8q-out-q5
default-6e-4q8q-out-policy (Cisco Nexus 7710 / 7718 switches only) <ul style="list-style-type: none"> • default-6e-4q8q-drop-out-policy • default-6e-4q8q-ndrop-out-policy 	c-6e-4q8q-drop-out and c-6e-4q8q-ndrop-out 6e-4q8q-out-q1 (priority queue), 6e-4q8q-out-q2, 6e-4q8q-out-q3, 6e-4q8q-out-q6, 6e-4q8q-out-q7, and 6e-4q8q-out-q-default 6e-4q8q-out-q4 (priority queue) and 6e-4q8q-out-q5 (priority queue)
default-4e-4q8q-out-policy (Cisco Nexus 7710 / 7718 switches only) <ul style="list-style-type: none"> • default-4e-4q8q-drop-out-policy • default-4e-4q8q-ndrop-out-policy 	c-4e-4q8q-drop-out and c-4e-4q8q-ndrop-out 4e-4q8q-out-q1 (priority queue), 4e-4q8q-out-q2, 4e-4q8q-out-q3, and 4e-4q8q-out-q-default 4e-4q8q-out-q4 (priority queue), 4e-4q8q-out-q5, and 4e-4q8q-out-q6, 4e-4q8q-out-q7

You can modify an egress CoS to queue map irrespective of the ingress CoS to queue map by using the **match cos** command to configure the desired CoS to queue mapping.

An egress queue follows a hierarchical scheduling pattern when both drop classes are present. For more information, see the “Ingress Queuing” section. For a given network qos template, the egress queuing configuration (the number of DWRR queues, number of priority queues, and the scheduling hierarchy) are fixed. You can modify the bandwidth percentage, priority level, and shaper for a given port.

You use the **bandwidth** command to control the bandwidth allocated to an egress queue (traffic class). For more information about the bandwidth command, see the *Cisco Nexus 7000 Series NX-OS Quality of Service Command Reference*.



Note Bandwidth and priority are mutually exclusive on a class map (queue).

You use the **priority** command to specify that a class of traffic has low latency requirements with respect to other classes. You can configure the priority level to a traffic queue as high or low. Use the priority command to define multiple levels of a strict priority service model. For more information about the **priority** command, see the *Cisco Nexus 7000 Series NX-OS Quality of Service Command Reference*.

The shaper can be configured with a percentage value and it can be enabled on any queue. You use the **shape** command to specify that a class of traffic has a maximum rate imposed on it and the outgoing traffic has a smooth output rate. To achieve a smooth output rate, the excess packets are retained in the queue and then scheduled for transmission later. For more information about the **shape** command, see the *Cisco Nexus 7000 Series NX-OS Quality of Service Command Reference*.



Note A shaper delays excess traffic that does not conform to the profile by queuing it in a buffer to shape the flow.

Shared Buffer Queuing on the F3 Series Module



Note This feature is available only on the F3 Series modules. If you attempt to configure shared buffer queuing on other modules, the switch returns an error.

Beginning with Cisco NX-OS Release 6.2(10) you can split QoS buffers into dedicated and shared buffers. With only dedicated buffers based on the CoS value, one queue may have very high traffic even though memory associated with some of the other queues may be lying idle. The shared buffer pools address this problem. The shared buffer is between ports in a port group.

The default is disabled for shared buffer queuing.

When you enable this feature, you enable it for each specific module. After you have enabled shared buffer queuing, the queue is, by default, divided equally into dedicated and shared buffer pools, 50:50 for the specified module. The dedicated buffer pools continue to function as they always did.

Then, if you want a different ratio, you modify the ratio per port group on the specified module, using the Command Line Interface (CLI). First you specify the port group for the given module and then you can modify the default queue limit ratio for that port group. If you want to change the default queue limit ration for other port groups on that module, you must enter the command for each port group separately.

Finally, you can apply a custom queuing policy to the specified port group.

**Note**

If global level shared buffer queuing is disabled, then shared buffer for all port-groups in that VDC are disabled. If the global level is enabled, all port groups in the VDC are enabled.

Shared buffer queuing is applicable only to port groups of physical interfaces. Shared buffer queuing on port groups is independent of membership in port channels. Thus, members of a port channel may have different shared buffer queuing configurations.

The command is applicable only to the ports in the VDC in which you are working. When you move any of the port groups from this VDC, the shared buffer queuing feature returns to the default disabled state. When you move a port group into the VDC, the port group assumes the global shared buffer configuration of that VDC (for example, if shared buffering is enabled in the VDC, it will also be enabled for the newly moved port group). Finally, if all the port groups are removed from a given VDC with this feature enabled, the shared buffering for that VDC with no ports is returned to the default disabled state.

After you enable shared buffering, the shared buffer pools are configured, using the active template, on the default ingress queuing policy. If you change the template, the setting for shared buffering remains as you last set this feature, either enabled or disabled. Then, the shared buffers are reconfigured based on the ingress queuing policy of the new template. If no user-defined policies are attached to a port group, the same default ingress queuing policy is applied to the shared buffer pool as to the dedicated buffer pool.

**Note**

You cannot change the template if you have any user-defined policies attached to the port group. You must first remove all user-defined queuing policies from the port group, and then you can change the template.

If you enable and configure shared buffer queuing, the shared buffer pools are used before the dedicated buffer queues. Any necessary dropped packets come from the dedicated buffer pools and pause is always honored from the dedicated buffer pools.

**Note**

The cos2q maps and the dscp2 maps are global in scope. The same maps are applied to the shared buffer pools.

Prerequisites for Queuing and Scheduling

Queuing and scheduling have the following prerequisites:

- You must be familiar with “Using Modular QoS CLI.”
- You are logged on to the switch.
- You are in the correct VDC. A VDC is a logical representation of a set of system resources. You can use the **switchto vdc** command with a VDC number.

Guidelines and Limitations

Queuing and scheduling of F-Series modules have the following configuration guidelines and limitations:

- A queuing policy that is being activated should be consistent with the system network qos policy.
- The default queuing policy is attached to the system target (includes all F Series module ports), which is unlike the M1 series configuration where the default-in-policy is attached exclusively to each port.
- A queuing policy that is attached to a given port, overrides the system queuing policy on that port.
- The DSCP to egress queue selection for DSCP values 2-7 are set to be the same as the values for CoS 2-7. To change this setting, access the type QoS policy and use the set cos command to change the selected egress queue (applicable for all types of interfaces, such as access, trunk, routed, and so on).
- Egress policies on VLAN configurations do not support set match on CoS.
- Egress policies on VLAN configurations do not support set QoS group or discard class.
- The ingress type QoS policy supports **set dscp/cos** and **set qos-group** commands. You can configure either **set dscp/cos** or **set qos-group** command but not both. It is possible to migrate between these configurations at any time.
- F-Series modules do not support the following commands in a QoS policy:
 - **set discard-class** or **match discard-class**
 - **set qos-group** or **match qos-group**
- F Series modules do not support WRED in ingress queuing policies.
- F2 modules do not support CoS-to-queue mapping changes when M1 modules are also installed in the switch.
- F Series modules and M2 modules support shaping in the priority queue. M1 modules do not support shaping in the priority queue.
- F-Series modules and M2 modules support shaping in the priority queue. M1 modules do not support shaping in the priority queue.
- See the following information about the default-nq-8e-4q4q-policy template that support four ingress buffers:
 - The default-nq-8e-4q4q-policy template is supported only with F2 modules.
 - When F1 modules are online, the default-nq-8e-4q4q-policy template cannot be attached to the system qos.
 - When the default-nq-8e-4q4q-policy template is attached to system qos, F1 modules are allowed to come online. However, all interfaces of the F1 modules go to the unallocated pool of the corresponding VDC.
 - To make software downgrades nondisruptive, the following is required before the software downgrade:
 - All user defined and cloned 8e-4q4q template queuing policies should be detached manually from all interfaces in each VDC.
 - The default-nq-8e-4q4q-policy or the user defined/cloned 8e-4q4q template network-qos policy should be detached from the system qos.
 - All user defined and cloned 8e-4q4q template network-qos policies should be removed manually from the default VDC.

- All user defined 8e-4q4q template queuing policies should be removed manually from all VDCs.
- Use the **clear qos policies 8e-4q4q** command in the default VDC to clear the default 8e-4q4q template policies. This command clears PPF (Policy Propagation Facility) nodes of 8e-4q4q template policies.
- After executing **clear qos policies 8e-4q4q** command, you must perform an in-service software downgrade (ISSD). If an ISSD is not performed, unexpected results might occur.

- The **clear qos policies 8e-4q4q** command is only supported in the default VDC. Using this command in the default VDC also clears the 8e-4q4q policy-maps in non-default VDCs.
- Reloading an F2 module brings up all the cleared default 8e-4q4q template related policy-maps by using the **clear qos policies 8e-4q4q** command.
- The default 8e-4q4q-policy template is published when a software upgrade is completed.

- See the following information about the Cisco 7710/7718 switches and the four default 4p8q policy templates that support eight egress queues on these switches:
 - The default 4q8q-policy templates are supported and enabled by default on the Cisco Nexus 7710 switch and Cisco Nexus 7718 switch only.
 - The default 4q8q-policy templates are supported on F2e modules only.
 - DSCP queuing is enabled by default on the Cisco Nexus 7710/7718 switches. You must use the **no hardware qos dscp-to-queue** command to disable DSCP queuing on the switch. You can use the **hardware qos dscp-to-queue command module type** command to reenables DSCP queuing.

- See the following information about the **match dscp** command:
 - Supports only the ingress queues for F2 modules for the 8E template. (It does not support egress queues, M1 queues, or fabric-qos queues.)
 - Supports only ingress queues that have at least one CoS value associated with it without any restriction on which CoS value is used.
 - Cannot be used in user-defined class maps.
 - Cannot be used in a user configuration session.
 - Must be disabled for ISSD. (If it is not disabled, the ISSD is disruptive).
 - DSCP to IVL mapping is disabled by default.
 - The **queue-limit** command cannot be specified based on CoS or DSCP values. The configured queue-limit sizes are applicable for both the DSCP and CoS values.
 - No additional statistics are generated to differentiate how many packets are matched on DSCP or CoS.
 - When DSCP to IVL is enabled, an interface uses the DSCP value as trusted for IP packets and the CoS value is trusted for non-IP packets.
 - DSCP to IVL mapping is enabled by default on the Cisco Nexus 7710/7718 switches. You must use the **no hardware qos dscp-to-queue** command to disable DSCP to IVL mapping.
 - DSCP to IVL mapping for FabricPath interfaces is not supported.

- DSCP to IVL mapping for IPv6 packets is not supported.
 - DSCP to IVL mapping change is a disruptive operation and might cause BFD/routing protocols to flap.
- Shared buffer queuing between ports in a port group is available only on the F3 Series modules.
 - Shared buffering is supported only in 8e and 8e-4q4q templates.
 - Break-out ports do not support shared buffering.
 - The M1, M2, F1, F2 and F2e modules do not support shared buffering.
 - Ports in a port channel with a user-defined policy attached should have this same user-defined policy attached to the port groups.
 - When a user-defined policy map is attached to a port group, the set cos and bandwidth commands are not applicable to the port group.
 - Shared buffer queuing does not apply on the FEX Hif ports.
 - When changing egress Class of Service (CoS) to queue mapping, ensure that you specify 2 or 3 seconds as the minimum time limit between changes. Otherwise, continuous traffic drop might occur.
 - The M3 modules do not support per-queue counters for egress drops (multicast, unknown unicast, or broadcasts). The egress drops will be per port and per Q-Default counter.
 - Only 8e templates are supported on M3 modules.
 - The M3 module supports only network-qos template with default-nq-8e-4q8q-policy. It does not support default-nq-4e-4q8q-policy, default-nq-6e-4q8q-policy, default-nq-7e-4q8q-policy, and default-nq-8021qav-4q8q-policy.
 - The M3 module supports only the network-qos template. This template contains all the CoS values that match the MTU size.
 - All data traffic will be enqueued to the default queue of dot1q-tunnel port because this port is untrusted by default.
 - Starting with Cisco NX-OS Release 8.0(1), the dscp-to-queue mapping for M3-Series modules is enabled by using the **hardware qos dscp-to-queue ingress module type f-series** command.

Configuring Queuing and Scheduling

You configure queuing and scheduling by creating policy maps of type queuing that you apply to either traffic direction of an interface. You can configure a queuing policy by following one of these methods:

- Copying predefined policy—You can copy a queuing policy template and modify it as needed.



Note

When you copy an ingress or egress queuing policy, you are also copying the internal policies for the hierarchical queuing policy. Copying shortens the default policy name by stripping the *default* and *policy* substrings from it.

- User-defined policy—You can create a queuing policy that conforms to one of the system-defined queuing policy templates.

For information about configuring policy maps and class maps, see “Using Modular QoS CLI.”

Configuring an Ingress Queuing Policy

You must modify the ingress queuing policy only if you want to change the default policy that the port inherited from the system default.

Procedure

	Command or Action	Purpose
Step 1	switch# qos copy policy type queuing default-4q-8e-in-policy {prefix <i>prefix</i> suffix <i>suffix</i> }	Copies a system-defined queuing policy and renames it with a prefix or suffix.
Step 2	switch# show policy-map type queuing [<i>policy-map-name</i>]	(Optional) Displays the queuing policy that you copied and renamed.
Step 3	switch# configure terminal	Enters global configuration mode.
Step 4	switch(config)# policy-map type queuing [<i>policy-map-name</i>]	Configures the policy map of type queuing and enters policy-map mode for the policy map name that you specify. The policy map names can contain alphabetic, hyphen, or underscore characters, are case sensitive, and can be up to 40 characters.
Step 5	switch(config)# class type queuing [2q4t-8e-in-q-default 2q4t-8e-in-q1]	Configures the class map of type queuing and then enters policy-map class queuing mode.
Step 6	switch(config)# queue-limit percent [1-100]	Sets the queue limit for the queue. The range is from 1 to 100. Note The total queue limit for all the queues in the policy cannot exceed 100. In this example, the queue limit is set to 40 percent in the 2q4t-8e-in-q-default and 60 percent in 2q4t-8e-in-q1.
Step 7	switch(config-pmap-c-que)# bandwidth percent [1-100]	Allocates the bandwidth to the CoS values mapped to the queues for exchanging with the peer. The range is from 1 to 100.
Step 8	switch(config-pmap-c-que)# exit	Exits policy-map queue mode and enters configuration mode.
Step 9	switch(config)# service-policy type queuing input [<i>policy-map-name</i>]	Applies a policy to an interface.

	Command or Action	Purpose
Step 10	switch(config)# show policy-map type queuing [<i>policy-map-name</i>]	(Optional) Displays information about all configured policy maps or a selected policy map of type queuing.
Step 11	switch(config)# show policy-map interface ethernet [<i>slot/port</i>]	(Optional) Displays information about the service policy on an Ethernet interface.

Configuring an Egress Queuing Policy

Procedure

	Command or Action	Purpose
Step 1	switch# qos copy policy type queuing default-4q-8e-in-policy { <i>prefix prefix</i> <i>suffix suffix</i> }	Copies a system-defined queuing policy and renames it with a prefix or suffix.
Step 2	switch# show policy-map type queuing [<i>policy-map-name</i>]	(Optional) Displays the queuing policy that you copied and renamed.
Step 3	switch# configure terminal	Enters global configuration mode.
Step 4	switch(config)# policy-map type queuing [<i>policy-map-name</i>]	Configures the policy map of type queuing and enters policy-map mode for the policy map name that you specify. The policy map names can contain alphabetic, hyphen, or underscore characters, are case sensitive, and can be up to 40 characters.
Step 5	switch(config)# class type queuing [1p3q1t-8e-out-pq1 1p3q1t-8e-out-q-default 1p3q1t-8e-out-q2 1p3q1t-8e-out-q3]	Configures the class map of type queuing and then enters policy-map class queuing mode.
Step 6	switch(config-pmap-c-que)# bandwidth percent [<i>1-100</i>]	Allocates the bandwidth in all ingress packets to the value specified. The range is from 1 to 100. Alternatively, absolute values in Gbps, Mbps, Kbps can also be specified.
Step 7	switch(config-cmap-que)# priority level { 1 2 }	Marks the priority level of the traffic queue. 1 stands for the highest priority and 2 stands for the lowest priority.

	Command or Action	Purpose
Step 8	switch(config-cmap-que)# shape [average percent {1-100}]	Shapes the traffic rate from a queue. The range is from 80000 bits per second to 10 Gigabytes per second.
Step 9	switch(config-pmap-que)# exit	Exits policy-map queue mode and enters configuration mode.
Step 10	switch(config)# service-policy type queuing input [<i>policy-map-name</i>]	Applies a policy to an interface.
Step 11	switch(config)# show policy-map type queuing [<i>policy-map-name</i>]	(Optional) Displays information about all configured policy maps or a selected policy map of type queuing.

Enabling DSCP to Queue Mapping

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# hardware qos dscp-to-queue ingress module type {all f-series m-series }	Enables the dscp-to-queue mapping on the specified module(s). Note Starting with Cisco NX-OS Release 8.0(1), use the f-series keyword to enable dscp-to-queue mapping on M3-Series I/O modules also.
Step 3	switch(config)# show hardware qos dscp-to-queue ingress	(Optional) Displays information about the status of dscp-to-queue mapping in ingress direction.
Step 4	switch(config)# copy running-config startup-config	(Optional) Saves the running configuration to the startup configuration.

Configuring Shared Buffer Queuing

You enable or disable shared buffer queuing per module. You then specify the port group to change from the default queue limit ration of 50:50 for dedicated and shared pools.

Currently, the configuring shared buffer queuing functionality is not supported on M3 modules.

The default value is no shared buffer queuing.



Note This command is applicable only on an F3 Series module.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# hardware qos shared-buffer module module-number	Enables the shared buffer queuing for the specified module. This command enables shared buffer queuing in the default ratio of 50:50 for dedicated:shared queues. The default value for shared buffer queuing is disabled. Use the no form of this command to disable shared buffer queuing on the specified module.
Step 3	switch(config)# hardware module module-number port-group port-group-number	(Optional) Enters the configuration for the specified port group on the module. If you want to change the default queue limit ratio, you do it by port group on the specified module.
Step 4	switch(config-port-group)# qos shared-buffer queue-limit percent	(Optional) Sets the queue limit for the shared buffer queue for the specified port group. The range is from 10 to 80 percent. Note The total queue limit for all the queues in the policy cannot exceed 100.
Step 5	switch(config)# copy running-config startup-config	(Optional) Saves the running configuration to the startup configuration.

Verifying the Queuing and Scheduling Configuration

To display the queuing policy configuration, perform one of the following tasks:



Note The **show** commands display only the default policies that correspond to the active template.

Command	Purpose
show queuing interface ethernet	Displays information about whether the queuing policy is applied correctly to the module.
show class-map type queuing	Displays information about all configured class maps or a selected class map of type queuing.

Command	Purpose
show policy-map type queuing	Displays information about all configured policy maps or a selected policy map of type queuing.
show policy-map system	Displays information about the network qos and queuing policy-maps that are currently in effect on the system.
show hardware qos dscp-to-queue	Shows the status of DSCP queuing.

When you modify a network QoS template, remove queuing policies, if any, that are attached exclusively on an F-Series module and M3 interface because these policies will be inconsistent with the new network QoS template.

For more information about the fields in the output of these commands, see the *Cisco Nexus 7000 Series NX-OS Quality of Service Command Reference* document.

Configuration Examples for Queuing and Scheduling on F-Series Modules

Example: Ingress Queuing Policy Configuration

The following example shows how to configure an ingress queuing policy:

```

policy-map type queuing p-4que-7e-drop-in
  class type queuing 4q4t-7e-in-q1
    queue-limit percent 45
    bandwidth percent 25
  class type queuing 4q4t-7e-in-q2
    queue-limit percent 10
    bandwidth percent 25
  class type queuing 4q4t-7e-in-q3
    queue-limit percent 45
    bandwidth percent 25
policy-map type queuing p-4que-7e-ndrop-in
  class type queuing 4q4t-7e-in-q4
    queue-limit percent 100
    bandwidth percent 25
policy-map type queuing p-4que-7e-in
  class type queuing c-4q-7e-drop-in
    service-policy type queuing p-4que-7e-drop-in
    queue-limit percent 70
  class type queuing c-4q-7e-drop-in
    service-policy type queuing p-4que-7e-ndrop-in
    queue-limit percent 30

```

Example: Egress Queuing Policy Configuration

The following example shows how to configure an egress queuing policy:

```

policy-map type queuing p-4que-6e-drop-out

```

Example: Hierarchical Queuing Policy Configuration

```

class type queuing lq3plt-6e-out-pq1
  priority level 1
  shape average percent 50
class type queuing lq3plt-6e-out-q4
  bandwidth remaining percent 100
policy-map type queuing p-4que-6e-ndrop-out
class type queuing lq3plt-6e-out-pq2
  priority level 1
  shape average percent 50
class type queuing lq3plt-6e-out-pq3
  priority level 2
policy-map type queuing p-4que-6e-out
class type queuing c-4q-6e-drop-out
  service-policy type queuing p-4que-6e-drop-out
  bandwidth percent 70
class type queuing c-4q-6e-ndrop-out
  service-policy type queuing p-4que-6e-ndrop-out
  bandwidth percent 30

```

Example: Hierarchical Queuing Policy Configuration

The following example shows how to configure a hierarchical queuing policy:

```

policy-map type queuing inner-policy-1
  class type queuing lp3qlt-out-q1
    bandwidth percent 40
  class type queuing lp3qlt-out-q2
    bandwidth percent 60
policy-map type queuing inner-policy-2
  class type queuing lp3qlt-out-q3
    bandwidth percent 40
  class type queuing lp3qlt-out-q4
    bandwidth percent 60
class-map type queuing drop-class
  match class-map lp3qlt-out-q1
  match class-map lp3qlt-out-q2
class-map type queuing nodrop-class
  match class-map lp3qlt-out-q3
  match class-map lp3qlt-out-q4
policy-map type queuing example-hierarchical-policy
  class type queuing drop-class
    bandwidth percent 40
  service-policy type queuing inner-policy-1
    match class nodrop-class
      percent 60
  service-policy type queuing inner-policy-2

```

Example: Verifying the Status of DSCP-to-queue Mapping

The following sample output from the **show hardware qos dscp-to-queue ingress** command displays the status of DSCP-to-queue mapping enabled in ingress direction on F-series modules:

```

Switch# show hardware qos dscp-to-queue ingress

status: Enabled
module_type : f-series

```

Feature History for Queuing and Scheduling for F-Series Modules

The table below summarizes the new and changed features for this document and shows the releases in which each feature is supported. Your software release might not support all the features in this document. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release.

Table 3: Feature History for Queuing and Scheduling for F-Series Modules

Feature Name	Release	Feature Information
DSCP to Queue Mapping	8.0(1)	Use the hardware qos dscp-to-queue ingress module type f-series command to enable dscp-to-queue mapping on M3-Series I/O modules also.
Shared buffer queuing on the F3 Series modules	6.2(10)	Support for shared memory buffer queues on the F3 Series modules only.
DSCP to Queue Mapping	6.2(2)	Support for five default templates to enable DSCP to Queue Mapping on F-Series Modules. Support to enable DSCP to Queue mapping using hardware qos dscp-to-queue ingress module-type command.
Support for 4q8q policy templates	6.2(2)	Support for four 4q8q policy templates that provide eight egress queues on the Cisco Nexus 7710 switch and Cisco Nexus 7718 switch only.
Support for four ingress buffers	6.1(3)	Support for default-8e-4q4q-policy template that supports four ingress buffers.
DSCP mapping for F2 modules	6.1(1)	Support for DSCP mapping for F2 modules.
Scheduling and Queuing for F1 Series Modules	5.1(1)	This chapter was added. (Chapter title subsequently changed to accommodate other F-Series Modules.)

