CHAPTER 9

Configuring Queuing and Scheduling on F1 Series Modules

This chapter describes how to configure the QoS queuing and scheduling features on the F1 Series module of the Cisco NX-OS device. This chapter includes the following sections:

Information About Queuing and Scheduling

On an F1 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 all QoS policy templates.

Note 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— Dynamic 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.

This section includes the following topics:

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 default-4q-8e-in-policy, default-4q-7e-in-policy, default-4q-6e-in-policy, and default-4q-4e-in-policy.

The predefined class map names (queue names) for ingress queuing are described in Table 9-1.

Table 9-1 Predefined Class Maps for Ingress Queuing

Ingress Policy Maps	Ingress Class Map Names
default-nq-8e-in-policy	2q4t-8e-in-q1 and 2q4t-8e-in-q-default
default-nq-7e-in-policy	4q4t-7e-in-q1, 4q4t-7e-in-q-default, 4q4t-7e-in-q3, and 4q4t-7e-in-q4
default-nq-6e-in-policy	4q4t-6e-in-q1, 4q4t-6e-in-q-default, 4q4t-6e-in-q3, and 4q4t-6e-in-q4.
default-nq-4e-in-policy	4q4t-4e-in-q1, 4q4t-4e-in-q-default, 4q4t-4e-in-q3, and 4q4t-4e-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.

The default queuing policy maps the priority and non-priority CoS values of the same class into different queues. However, you can configure a different CoS to queue mapping as needed. You can configure the CoS to queue map on the default VDC but the configuration is applicable system wide.

For example, in an 8e template, the network-admin may want to treat CoS 7 as a non-priority CoS and map it to the non-priority queue.

In the ingress class-map, you can modify the match cos value using the **match cos** command to configure the desired cos to queue mapping.

Note Modifying the default queuing policy maps is a disruptive operation and it may can 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, Release 5.x*.

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, Release 5.x*.

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

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 F1 VDCs. It is not supported in F1/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 an (DWRR) queue.
- For all the non 8e templates, second level scheduling is used.

Note

- Egress queues are of fixed size and are not user configurable.
- There are four queues in the egress port.

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, and default-4q-4e-out-policy. The flexible egress queues configuration is based on these queue types—1p3q1t-8e, 1p3q1t-7e, 3p1q1t-6e, and 2p2q1t-4e.

The predefined class map names (queue names) for egress queuing are described in Table 9-2.

Table 9-2 Predefined Class Maps for Egress Queuing

Egress Policy Names	Egress Class Map Names
default-nq-8e-out-policy	1p3q1t-8e-out-pq1, 1p3q1t-8e-out-q2, 1p3q1t-8e-out-q3, and 1p3q1t-8e-out-q-default
default-nq-7e-out-policy	1p3q1t-7e-out-pq1, 1p3q1t-7e-out-q2, 1p3q1t-7e-out-q3, and 1p3q1t-7e-out-q-default
default-nq-6-out-policy	3p1q1t-6e-out-pq1, 3p1q1t-6e-out-pq2, 3p1q1t-6e-out-pq3, and 3p1q1t-6e-out-q-default
default-nq-4e-out-policy	2p2q1t-4e-out-pq1, 2p2q1t-4e-out-pq2, 2p2q1t-4e-out-q3, and 2p2q1t-4e-out-q-default

You can modify an egress CoS to queue map irrespective of the ingress CoS to queue map. 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 on page 9-2. For a given network qos template, the egress queuing configuration (the number of DWRR queues, number of priority queues, 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, Release 5.x*.

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, Release 5.x*.

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. In order 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, Release 5.x*.

Note A shaper delays excess traffic, which do not confirm to the profile, by queuing it in a buffer, to shape the flow.

Licensing Requirements for Queuing and Scheduling

The following table shows the licensing requirements for this feature:

Product	License Requirement
Cisco NX-OS	The QoS feature does not a require license. Any feature not included in a license package is bundled with the Cisco NX-OS system images and is provided at no extra charge to you. For a complete explanation of the Cisco NX-OS licensing scheme, see the *Cisco NX-OS Licensing Guide*.

However, using VDCs requires an Advanced Services license.

Prerequisites for Queuing and Scheduling

Queuing and scheduling have the following prerequisites:

- You must be familiar with Chapter 3, "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 F1 Series cards 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 (this includes all F1 module ports), which is unlike the M1 series configuration where the default-in-policy is attached exclusively to each port.
- A queuing policy attached to a given port, overrides the system queuing policy on that port.
- F1 Series modules do not support the following in a QoS policy:

  - set discard-class or match discard-class
  - set qos-group or match qos-group
- F1 Series modules do not support WRED in ingress queuing policies.
- F1 Series modules support shaping in the priority queue. M1 modules do not support shaping in the priority queue.

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 their internal policies for the hierarchical queuing policy. Copying shorten 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 Chapter 3, "Using Modular QoS CLI."

This section includes the following topics:

Configuring an Ingress Queuing Policy

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

You can configure an ingress queuing policy by following steps given below.

The example in this section assumes that you are copying an 8e template queuing policy.

SUMMARY STEPS

1. **qos copy policy type queuing default-4q-8e-in-policy** {**prefix** *prefix* | **suffix** *suffix*}
2. (Optional) **show policy-map type queuing** [*policy-map-name*]
3. **configure terminal**
4. **policy-map type queuing** [*policy-map-name*]
5. **class type queuing** [**2q4t-8e-in-q-default** | **2q4t-8e-in-q1**]
6. **queue-limit percent** [**1-100**]
7. **bandwidth percent** [**1-100**]
8. **exit**

9. **service-policy type queuing input** [*policy-map-name*]
10. (Optional) **show policy-map type queuing** [*policy-map-name*]
11. (Optional) **show policy-map interface ethernet** [*slot/port*]

DETAILED STEPS

	Command	Purpose
Step 1	**qos copy policy type queuing** [**default-4q-8e-in-policy** {**prefix** *prefix* | **suffix** *suffix*} **Example:** switch# qos copy policy type queuing default-4q-8e-in-policy prefix my_ switch#	Copies a system-defined queuing policy and renames it with a prefix or suffix. The policy is renamed as my_4q-8e-in-policy in this example.
Step 2	**show policy-map type queuing** [*policy-map-name*] **Example:** switch# show policy-map type queuing my_4q-8e-in	(Optional) Displays the queuing policy that you copied and renamed.
Step 3	**configure terminal** **Example:** switch# configure terminal switch(config)#	Enters configuration mode.
Step 4	**policy-map type queuing** [*policy-map-name*] **Example:** switch(config)# policy-map type queuing my_4q-8e-drop-in switch(config-pmap-que)#	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	**class type queuing** [**2q4t-8e-in-q-default** | **2q4t-8e-in-q1**] **Example:** switch(config)# class type queuing 2q4t-8e-in-q-default	Configures the class map of type queuing and then enters policy-map class queuing mode.
Step 6	**queue-limit percent** [**1-100**] **Example 1:** switch(config)# class type queuing 2q4t-8e-in-q-default switch(config-pmap-c-que)# queue-limit 40 **Example 2:** switch(config)# class type queuing 2q4t-8e-in-q1 switch(config-pmap-c-que)# queue-limit 60	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.

	Command	Purpose
Step 7	**bandwidth percent** [**1-100**] **Example:** switch(config-pmap-c-que)# bandwidth percent 10 switch(config-pmap-c-que)#	Allocates the bandwidth to the CoS values mapped to the queues for exchanging with the peer. The range is from 1 to 100. The bandwidth is set to 10 percent in this example.
Step 8	**exit** **Example:** switch(config-cmap-que)# exit switch(config)#	Exits policy-map queue mode and enters configuration mode.
Step 9	**service-policy type queuing input** [*policy-map-name*] **Example:** switch(config)# interface ethernet 5/5 switch(config-if)# service-policy type queuing input my_4q-8e-in switch(config-if)#	Applies a policy to an interface.
Step 10	**show policy-map type queuing** [*policy-map-name*] **Example:** switch# show policy-map type queuing default-4q-4e-in-policy	(Optional) Displays information about all configured policy maps or a selected policy map of type queuing.
Step 11	**show policy-map interface ethernet** [*slot/port*] **Example:** switch# show policy-map interface ethernet 1/5	(Optional) Displays information about the service policy on an Ethernet interface.

Configuring an Egress Queuing Policy

You can configure an egress queuing policy.

The example in this section assumes that you are copying a 4e template queuing policy.

SUMMARY STEPS

1. **qos copy policy type queuing default-4q-8e-out-policy** | {**prefix** *prefix* | **suffix** *suffix*}
2. **show policy-map type queuing** [*policy-map-name*]
3. **configure terminal**
4. **policy-map type queuing** [*policy-map-name*]
5. **class type queuing** [**1p3q1t-8e-out-pq1** | **1p3q1t-8e-out-q-default** | **1p3q1t-8e-out-q2** | **1p3q1t-8e-out-q3**]
6. **bandwidth** [**percent** {**1-100**} | **remaining**]
7. **priority level**{**1** | **2**}
8. **shape** [**average** | **percent** {**1-100**}]
9. **exit**

10. **service-policy type queuing output** [*policy-map-name*]
11. (Optional) **show policy-map type queuing** [*policy-map-name*]

DETAILED STEPS

	Command	Purpose

Step 1	**qos copy policy type queuing default-4q-8e-out-policy** {**prefix** *prefix* | **suffix** *suffix*}
	
	Example:
	`switch# qos copy policy type queuing default-4q-8e-out-policy prefix my_`
	`switch#`
		Copies a system-defined queuing policy and renames it with a prefix or suffix.
		
		The policy is renamed as my_4q-8e-out-policy in this example.

Step 2	**show policy-map type queuing** [policy-map-name] **Example:** `switch# show policy-map type queuing my_4q-8e-out-policy`	(Optional) Displays the queuing policy that you copied and renamed.

Step 3	**configure terminal**
	
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters configuration mode.

Step 4	**policy-map type queuing** [*policy-map-name*]
	
	Example:
	`switch(config)# policy-map type queuing my_4q-8e-out-policy`
	`switch(config-pmap-que)#`
		Configures the policy map of type queuing and enters policy-map mode for the policy-map name that you specify. Policy-map names can contain alphabetic, hyphen, or underscore characters, are case sensitive, and can be up to 40 characters.

Step 5	**class type queuing** [**1p3q1t-8e-out-pq1** | **1p3q1t-8e-out-q-default** | **1p3q1t-8e-out-q2** | **1p3q1t-8e-out-q3**]
	
	Example:
	`switch(config)# class type queuing 1p3q1t-8e-out-pq1`
	`switch(config-pmap-c-que)#`
		Configures the class map of type queuing and then enters policy-map class queuing mode.

Step 6	**bandwidth** {**percent** {**1-100**}| **remaining**}
	
	Example:
	`switch(config-pmap-c-que)# bandwidth percent 25`
	`switch(config-pmap-c-que)#`
		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.
		
		The bandwidth is set to 25 percent in this example.

Step 7	**priority level** {**1** | **2**}
	
	Example:
	`switch(config-cmap-que)# priority level 1`
	`switch(config-cmap-que)#`
		Marks the priority level of the traffic queue. 1 stands for the highest priority and 2 stands for the lowest priority.
		
		The priority level is set to 1 in this example.

	Command	Purpose
Step 8	**shape percent** [**average** | **percent** {**1-100**}] **Example:** switch(config-cmap-que)# shape 50000 bps switch(config-cmap-que)#	Shapes the traffic rate from a queue. The range is from 80000 bits per second to 10 Gigabytes per second. The shaper is set to 50000 bits per second in this example.
Step 9	**exit** **Example:** switch(config-cmap-que)# exit switch(config)#	Exits policy-map queue mode, and enters configuration mode.
Step 10	**service-policy type queuing output** [*policy-map-name*] **Example:** switch(config)# interface ethernet 5/5 switch(config-if)# service-policy type queuing input my_4q-8e-out switch(config-if)#	Applies a policy to an interface.
Step 11	**show policy-map type queuing** [*policy-map-name*] **Example:** switch# show policy-map type queuing my_4q-8e-out-policy	(Optional) Displays information about all configured policy maps or a selected policy map of type queuing.

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.
show policy-map type queuing	Displays information about all configured policy maps or a selected policy map of type queuing.

When changing the network qos template, you must remove any queuing policy that is attached exclusively on an F1 module interface because the queuing policy would be inconsistent with the new network qos template.

For more information about the fields in the output from these commands, see the *Cisco Nexus 7000 Series NX-OS Quality of Service Command Reference, Release 5.x*.

Configuration Examples for Queuing and Scheduling on F1 Series Modules

In this section you can find examples of configuring queuing and scheduling for the F1 Series modules.

This section includes the following topics:

Ingress Queuing Policy Configuration Example

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
```

Egress Queuing Policy Configuration Example

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

```
policy-map type queuing p-4que-6e-drop-out
  class type queuing 1q3p1t-6e-out-pq1
    priority level 1
    shape average percent 50
  class type queuing 1q3p1t-6e-out-q4
    bandwidth remaining percent 100
policy-map type queuing p-4que-6e-ndrop-out
  class type queuing 1q3p1t-6e-out-pq2
    priority level 1
    shape average percent 50
  class type queuing 1q3p1t-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
```

Hierarchical Queuing Policy Configuration Example

```
policy-map type queuing inner-policy-1
  class type queuing 1p3q1t-out-q1
    bandwidth percent 40
  class type queuing 1p3q1t-out-q2
    bandwidth percent 60
policy-map type queuing inner-policy-2
   class type queuing 1p3q1t-out-q3
     bandwidth percent 40
   class type queuing 1p3q1t-out-q4
     bandwidth percent 60
   class-map type queuing drop-class
    match class-map 1p3q1t-out-q1
    match class-map 1p3q1t-out-q2
   class-map type queuing nodrop-class
    match class-map 1p3q1t-out-q3
    match class-map 1p3q1t-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
```

Feature History for Queuing and Scheduling for F1 Series Module

Table 9-3 lists the release history for this feature.

Table 9-3 Feature History for Queuing

Feature Name	Releases	Feature Information
Scheduling and Queuing for F1 Series Modules	5.1(1)	This chapter was added.