



CHAPTER 1

Overview

This chapter describes the configurable Cisco NX-OS quality of service (QoS) features on the Cisco NX-OS device.

QoS allows you to classify the network traffic, police and prioritize the traffic flow, and help avoid traffic congestion in a network.

This chapter includes the following sections:

- [Information About QoS Features, page 1-1](#)
- [High Availability Requirements for QoS Features, page 1-5](#)
- [QoS Feature Configuration with MQC, page 1-5](#)
- [QoS Statistics, page 1-5](#)
- [QoS Policies on Fabric Extenders, page 1-6](#)

Information About QoS Features

You use the QoS features to provide the most desirable flow of traffic through a network. QoS allows you to classify the network traffic, police and prioritize the traffic flow, and provide congestion avoidance. The control of traffic is based on the fields in the packets that flow through the system. You use the Modular QoS CLI (MQC) to create the traffic classes and policies of the QoS features.

QoS features are applied using QoS policies and queuing policies are as follows:

- QoS policies include the policing feature and the marking features.
- Queuing policies use the queuing and scheduling features as well as a limited set of the marking feature.



Note

The system-defined QoS features and values that are discussed in [Chapter 3, “Using Modular QoS CLI”](#) apply globally to the entire switch and cannot be modified. For complete information on virtual device contexts (VDCs), see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 5.x*.

Send document comments to nexus7k-docfeedback@cisco.com.

**Caution**

Before you attempt a downgrade from Cisco NX-OS Release 5.2(x) to any release prior to Release 5.2(1), you should clear the QoS MIB and MPLS QoS defaults by using the **clear qos mpls-snmp** command. The downgrade might fail if the defaults are not cleared.

Before you downgrade from Cisco NX-OS Release 5.2(x) or 5.1(x) to Cisco NX-OS Release 5.0(x) or an earlier release, remove all system QoS and QoS policies configured on F1-series modules. Use the **clear qos policies** command to remove the defaults for F1-series modules. An internal process failure can result if the QoS policies are not removed prior to the downgrade.

This section includes the following topics:

- [Using QoS, page 1-2](#)
- [Classification, page 1-2](#)
- [Marking, page 1-3](#)
- [Mutation, page 1-3](#)
- [Policing, page 1-3](#)
- [Queuing and Scheduling, page 1-3](#)
- [Sequencing of QoS Actions, page 1-4](#)

Using QoS

Traffic is processed based on how you classify it and the policies that you create and apply to traffic classes.

To configure QoS features, you use the following steps:

1. Create traffic classes by classifying the incoming and outgoing packets that match criteria such as IP address or QoS fields.
2. Create policies by specifying actions to take on the traffic classes, such as limiting, marking, or dropping packets.
3. Apply policies to a port, port channel, VLAN, or a subinterface.

You use MQC to create the traffic classes and policies of the QoS features. For more information, see [Chapter 3, “Using Modular QoS CLI.”](#)

**Note**

The queuing and scheduling operations of the overall QoS feature are applicable to both IPv4 and IPv6.

Classification

You use classification to partition traffic into classes. You classify the traffic based on the port characteristics (class of service [CoS] field) or the packet header fields that include IP precedence, Differentiated Services Code Point (DSCP), Layer 2 to Layer 4 parameters, and the packet length.

The values used to classify traffic are called match criteria. When you define a traffic class, you can specify multiple match criteria, you can choose to not match on a particular criterion, or you can determine the traffic class by matching any or all criteria.

Traffic that fails to match any class is assigned to a default class of traffic called class-default.

Send document comments to nexus7k-docfeedback@cisco.com.

For more information about configuring classification, see [Chapter 2, “Configuring Classification.”](#)

Marking

Marking is the setting of QoS information that is related to a packet. You can set the value of a standard QoS field IP precedence, DSCP and CoS, and internal labels that can be used in subsequent actions. Marking is used to identify the traffic type for policing, queuing, and scheduling traffic (only CoS is used in scheduling).

For more information about configuring marking, see [Chapter 4, “Configuring Marking.”](#)

Mutation

Mutation is the changing of packet header QoS fields. You can map IP precedence, DSCP, or CoS values to all incoming or outgoing packets. You can use mutation in policies that contain policing commands, but you cannot use mutation in queuing and scheduling commands. You use configurable, user-defined table maps for mutation.

For more information about configuring mutation, see [Chapter 5, “Configuring Mutation Mapping.”](#)

Policing

Policing is the monitoring of data rates for a particular class of traffic. The device can also monitor associated burst sizes.

Three colors, or conditions, are determined by the policer depending on the data rate parameters supplied: conform (green), exceed (yellow), or violate (red). You can configure only one action for each condition. When the data rate exceeds the user-supplied values, packets are either marked down or dropped. You can define single-rate, dual-rate, and color-aware policers.

Single-rate policers monitor the specified committed information rate (CIR) of traffic. Dual-rate policers monitor both CIR and peak information rate (PIR) of traffic. Color-aware policers assume that traffic has been previously marked with a color.

For more information about configuring policing, see [Chapter 6, “Configuring Policing.”](#)

Queuing and Scheduling

The queuing and scheduling process allows you to control the bandwidth allocated to traffic classes, so you achieve the desired trade-off between throughput and latency.

You can apply weighted random early detection (WRED) to a class of traffic, which allows packets to be dropped based on the CoS field. The WRED algorithm allows you to perform proactive queue management to avoid traffic congestion.

You can schedule traffic by imposing a maximum data rate on a class of traffic so that excess packets are retained in a queue to smooth (constrain) the output rate.

For information about configuring queuing and scheduling, see [Chapter 7, “Configuring Queuing and Scheduling.”](#)

Send document comments to nexus7k-docfeedback@cisco.com.

Sequencing of QoS Actions

The following are the three types of policies:

- **network qos**—Defines the characteristics of QoS properties network wide.
- **qos**—Defines MQC objects that you can use for marking and policing.
- **queuing**—Defines MQC objects that you can use for queuing and scheduling as well as a limited set of the marking objects.



Note

The default type of policy is **qos**.

The Cisco NX-OS device processes the QoS policies that you define based on whether they are applied to ingress or egress packets. The system performs actions for QoS policies only if you define them under the type **qos** service policies.



Note

You can apply only ingress traffic actions for QoS policies on Layer 2 interfaces. You can apply both ingress and egress traffic actions on Layer 3 interfaces.

This section includes the following topics:

- [Sequencing of Ingress Traffic Actions, page 1-4](#)
- [Sequencing of Egress Traffic Actions, page 1-4](#)

Sequencing of Ingress Traffic Actions

The sequence of QoS actions on ingress traffic is as follows:

1. Queuing and scheduling
2. Mutation
3. Classification
4. Marking
5. Policing

Sequencing of Egress Traffic Actions

The sequencing of QoS actions on egress traffic is as follows:

1. Classification
2. Marking
3. Policing
4. Mutation
5. Queuing and scheduling



Note

Mutation occurs much closer to the beginning of the traffic actions on the ingress packets, and any further classification and policing is based on the changed QoS values. Mutation occurs at the end of the traffic actions on the egress packets, right before queuing and scheduling.

Send document comments to nexus7k-docfeedback@cisco.com.

High Availability Requirements for QoS Features

The Cisco NX-OS QoS software recovers its previous state after a software restart, and it is capable of a switchover from the active supervisor to the standby supervisor without a loss of state.



Note

For complete information on high availability, see the *Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide, Release 5.x*.

QoS Feature Configuration with MQC

You use MQC to configure QoS features. The MQC configuration commands are shown in [Table 1-1](#).

Table 1-1 MQC Configuration Commands

MQC Command	Description
class-map	Defines a class map that represents a class of traffic.
table-map	Defines a table map that represents a mapping from one set of field values to another set of field values. You can reference a table map from a policy map.
policy-map	Defines a policy map that represents a set of policies to be applied to a set of class maps. Policy maps can reference table maps.

You can modify or delete MQC objects, except system-defined objects, when the objects are not associated with any interfaces. For information on system-defined MQC objects, see [Chapter 3, “Using Modular QoS CLI.”](#)

After a QoS policy is defined, you can attach the policy map to an interface by using the interface configuration command shown in [Table 1-2](#).

Table 1-2 Interface Command to Attach a Policy Map to an Interface

Interface Command	Description
service-policy	Applies the specified policy map to input or output packets on the interface.

For information on how to use MQC, see [Chapter 3, “Using Modular QoS CLI.”](#)

QoS Statistics

Statistics are maintained for each policy, class action, and match criteria per interface. You can enable or disable the collection of statistics, you can display statistics using the **show policy-map** interface command, and you can clear statistics based on an interface or policy map with the **clear qos statistics** command. Statistics are enabled by default and can be disabled globally.

For information about monitoring QoS statistics, see [Chapter 11, “Monitoring QoS Statistics.”](#)

[Send document comments to nexus7k-docfeedback@cisco.com.](mailto:nexus7k-docfeedback@cisco.com)

Default QoS Behavior

The QoS queuing features are enabled by default. Specific QoS-type features, policing and marking, are enabled only when a policy is attached to an interface. Specific policies are enabled when that policy is attached to an interface.

By default, the device always enables a system default queuing policy, or system-defined queuing policy map, on each port and port channel. When you configure a queuing policy and apply the new queuing policy to specified interfaces, the new queuing policy replaces the default queuing policy and those rules now apply. For more information on the system-defined, default queuing policies and the default values that apply to each interface, see [Chapter 3, “Using Modular QoS CLI.”](#)

The device enables other QoS features, policing and marking, only when you apply a policy map to an interface.

QoS Policies on Fabric Extenders

The Cisco Nexus 2000 Series Fabric Extender (FEX) is a remote line card that you can connect to the Cisco Nexus 7000 Series switch. The FEX has 48 1-Gbps front-panel or server-facing ports, which are satellite ports. The FEX has four uplink ports that you can use to connect it to the Cisco Nexus 7000 Series switch. The four ports on the Cisco Nexus 7000 Series switch that connect to the uplink ports are fabric ports. Only QoS policies can be configured on the server-facing FEX ports. Currently, queuing on the FEX interfaces is not supported.

For more information on FEX, see the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 5.x*, *Cisco Nexus 7000 Series NX-OS Fundamentals Configuration Guide, Release 5.x*, and *Cisco Nexus 7000 Series NX-OS Fundamentals Command Reference, Release 5.x*.