



CHAPTER 1

Overview

This chapter describes the quality of service (QoS) features that you can use on Cisco Nexus 1000V to prevent traffic congestion in your network.

Information About Quality of Service

You can use QoS to provide the most desirable flow of traffic through a network. QoS allows you to classify your network traffic, police and prioritize the traffic flow, and provide congestion avoidance. Traffic is processed based on how you classify it and the QoS policies that you put in place.

You can implement a QoS policy using the following steps:

1. Define a traffic class by using the **class-map** command. For more information, see [Chapter 2, “Configuring QoS Classification.”](#)
2. Create a traffic policy by using the **policy-map** command. A traffic policy defines how specific traffic is to be acted upon to improve the quality of service. For more information, see [Chapter 3, “Configuring QoS Marking Policies.”](#)
3. Attach the traffic policy to an interface or port profile by using the **service-policy** command. For more information, see the [“Creating Ingress and Egress Policies”](#) section on page 3-10.
4. Police the traffic. For more information, see [Chapter 4, “Configuring QoS Policing.”](#)

Traffic Classification and Marking

You can use traffic classification and marking to sort and modify traffic for the best quality of service. [Table 1-1](#) describes these processes.

Table 1-1 Traffic Classification and Traffic Marking

QoS Method	Description	Command	Mechanism
Traffic Classification	Groups network traffic based on defined criteria.	match	class maps
Traffic Marking	Modifies traffic attributes by class.	set	policy maps

Traffic Classification

Traffic classification allows you to organize traffic (packets) into traffic classes or categories on the basis of whether the traffic matches the criteria you specify. 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. For more information, see [Chapter 2, “Configuring QoS Classification.”](#)

Traffic Marking

Marking is the process of assigning a priority and involves setting the fields, such as class of service or DSCP, in a packet. The traffic is then marked accordingly as it comes into the device on an ingress interface. The markings are used to treat the traffic as it leaves the device on the egress interface. For more information about configuring marking, see [Chapter 3, “Configuring QoS Marking Policies.”](#)

Policing

Policing is the monitoring of data rates for a particular class of traffic. The Cisco Nexus 1000V can also monitor associated burst sizes.

Three conditions, are determined by the policer depending on the data rate parameters supplied: conform, exceed, or violate. 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 or dual-rate 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.

For more information about configuring policing, see [Chapter 4, “Configuring QoS Policing.”](#)

QoS Commands

QoS configuration commands are shown in [Table 1-2](#).

Table 1-2 QoS Configuration Commands

Global Configuration Commands	Class Map QoS Configuration Command	Policy Map QoS Configuration Command	Interface or Port Profile Configuration Command	Policy Map Class QoS Configuration 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.

Table 1-2 QoS Configuration Commands

Global Configuration Commands	Class Map QoS Configuration Command	Policy Map QoS Configuration Command	Interface or Port Profile Configuration Command	Policy Map Class QoS Configuration Command	Description
<code>policy-map</code>	—	—	—	—	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.
—	—	—	<code>service-policy</code>	—	Applies a specified policy map to input or output packets on interfaces configured as follows: <ul style="list-style-type: none"> • inherited from a port-profile¹ • port-channel • Ethernet • VEthernet
—	—	—	—	<code>police</code>	Defines the action to take regarding packet data rates.
—	<code>match</code>	—	—	—	Defines the criteria for a class map.
—	—	<code>set</code>	—	—	Defines the packet header values for a policy map.

1. For information about port profiles, see the *Cisco Nexus 1000V Port Profile Configuration Guide, Release 4.0(4)SV1(3)*.

QoS Statistics

Statistics are maintained for each policy, class action, and match criteria per interface. For information about monitoring QoS statistics, see [Chapter 5, “Monitoring QoS Statistics.”](#)

Default QoS Behavior

QoS has no default behavior. Policing and prioritization of traffic are only implemented when you apply a policy map to an interface. The only exception is that, by default, the CoS value for control and packet VLAN traffic is set to 6. This value can be overridden with an explicit QoS policy that is configured on the interface that carries the control and packet VLAN traffic.

However, when designing your QoS and ACL policies, note that access control lists (ACLs) that are referenced within a QoS policy are processed as follows as part of the QoS policy:

- QoS ingress processing follows ACL processing.
- QoS egress processing precedes ACL egress processing.

Supported RFCs

Table 1-3 lists RFCs that are supported by QoS.

Table 1-3 **Supported RFCs**

Number	Title
RFC 2475	Architecture for Differentiated Services
RFC 2697	A Single Rate Three Color Marker
RFC 2698	A Dual Rate Three Color Marker
RFC 3289	Management Information Base for the Differentiated Services Architecture
RFC 3550	RTP: A Transport Protocol for Real-Time Applications

High Availability Requirements for QoS Features

QoS recovers its previous state after a software restart, and it is able to switch over from the active supervisor to the standby supervisor without a loss of state.