



Quality of Service Overview

This chapter explains quality of service (QoS) and the service models that embody it. It also suggests benefits that you can gain from implementing Cisco IOS QoS in your network. Then it focuses on the Cisco IOS QoS features and the technologies that implement them.

This chapter contains the following sections:

- [What Is Quality of Service?](#)
- [About QoS Architecture](#)
- [Who Could Benefit from Using Cisco IOS QoS?](#)
- [Why Deploy Cisco IOS QoS?](#)
- [End-to-End QoS Models](#)
- [Cisco IOS QoS Features](#)

What Is Quality of Service?

QoS refers to the ability of a network to provide improved service to selected network traffic over various underlying technologies including Frame Relay, ATM, Ethernet and 802.1 networks, SONET, and IP-routed networks. In particular, QoS features provide improved and more predictable network service by implementing the following services:

- Supporting guaranteed bandwidth
- Improving loss characteristics
- Avoiding and managing network congestion
- Shaping network traffic
- Setting traffic priorities across the network



Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

© 2007 Cisco Systems, Inc. All rights reserved.

About QoS Architecture

You configure QoS features throughout a network to provide for end-to-end QoS delivery. The following three components are necessary to deliver QoS across a heterogeneous network:

- QoS within a single network element, which includes queuing, scheduling, and traffic shaping features.
- QoS signaling techniques for coordinating QoS for end-to-end delivery between network elements.
- QoS policing and management functions to control and administer end-to-end traffic across a network.

Not all QoS techniques are appropriate for all network routers. Because edge routers and backbone routers in a network do not necessarily perform the same operations, the QoS tasks that they perform might differ as well. To configure an IP network for real-time voice traffic, for example, you would need to consider the functions of both edge and backbone routers in the network and then select the appropriate QoS feature or features.

In general, edge routers perform the following QoS functions:

- Packet classification and marking
- Admission control
- Configuration management

In general, backbone routers perform the following QoS functions:

- Congestion management
- Congestion avoidance

Who Could Benefit from Using Cisco IOS QoS?

All networks can take advantage of aspects of QoS for optimum efficiency, whether the network is for a small corporation, an enterprise, or an Internet service provider (ISP). Different categories of networking users—such as major enterprises, network service providers, and small- and medium-sized businesses—have their own QoS requirements; in many areas, however, these requirements overlap. The Cisco IOS QoS features described in the [“Cisco IOS QoS Features” section on page 5](#) address these diverse and common needs.

Enterprise networks, for example, must provide end-to-end QoS solutions across the various platforms that comprise the network. Providing solutions for heterogeneous platforms often requires that you take a different QoS configuration approach for each technology. As enterprise networks carry more complex, mission-critical applications and experience increased traffic from web multimedia applications, QoS serves to prioritize this traffic to ensure that each application gets the service that it requires.

ISPs require assured scalability and performance. For example, ISPs that have long offered best-effort IP connectivity now also transfer voice, video, and other real-time critical application data. QoS answers the scalability and performance needs of these ISPs to distinguish different kinds of traffic, thereby enabling them to offer service differentiation to their customers.

In the small- and medium-sized business segment, managers are experiencing firsthand the rapid growth of business on the Internet. These business networks must also handle increasingly complex business applications. QoS lets the network handle the difficult task of utilizing an expensive WAN connection in the most efficient way for business applications.

Why Deploy Cisco IOS QoS?

The Cisco IOS QoS features enable networks to control and predictably service a variety of networked applications and traffic types. Implementing Cisco IOS QoS in your network has the following advantages:

- Control over resources. You have control over which resources (bandwidth, equipment, wide-area facilities, and so on) are being used. For example, you can limit bandwidth consumed over a backbone link by FTP transfers or give priority to an important database access.
- Tailored services. If you are an ISP, the control and visibility provided by QoS enables you to offer carefully tailored grades of service differentiation to your customers.
- Coexistence of mission-critical applications. Cisco IOS QoS features ensures following conditions:
 - That your WAN is used efficiently by mission-critical applications that are most important to your business.
 - That bandwidth and minimum delays required by time-sensitive multimedia and voice applications are available.
 - That other applications using the link get their fair service without interfering with mission-critical traffic.

Moreover, in implementing QoS features in your network, you put in place the foundation for a future fully integrated network.

End-to-End QoS Models

A service model, also called a level of service, describes a set of end-to-end QoS capabilities. End-to-end QoS is the ability of the network to deliver service required by specific network traffic from one end of the network to another. Cisco IOS QoS software supports three types of service models: best effort, integrated, and differentiated services.



Note

QoS service models differ in how they enable applications to send data and in the ways in which the network attempts to deliver that data. For instance, one service model can be used for real-time applications, such as audio and video conferencing and IP telephony, while another service model can be used for file transfer and e-mail applications.

Consider the following factors when deciding which type of service to deploy in the network:

- The application or problem that you are trying to solve. Each of the three types of service—best effort, integrated, and differentiated—is appropriate for certain applications.
- The kind of capability that you want to allocate to your resources.
- Cost-benefit analysis. For example, the cost of implementing and deploying differentiated service is certain to be more expensive than the cost for a best-effort service.

The following sections describe the service models that are supported by features in Cisco IOS software:

- [Best-Effort Service](#)
- [Integrated Service](#)
- [Differentiated Service](#)

Best-Effort Service

Best effort is a single service model in which an application sends data whenever it must, in any quantity, and without requesting permission or first informing the network. For best-effort service, the network delivers data if it can, without any assurance of reliability, delay bounds, or throughput.

The Cisco IOS QoS feature that implements best-effort service is first-in, first-out (FIFO) queuing. Best-effort service is suitable for a wide range of networked applications such as general file transfers or e-mail.

Integrated Service

Integrated service is a multiple service model that can accommodate multiple QoS requirements. In this model the application requests a specific kind of service from the network before it sends data. The request is made by explicit signaling; the application informs the network of its traffic profile and requests a particular kind of service that can encompass its bandwidth and delay requirements. The application is expected to send data only after it gets a confirmation from the network. It is also expected to send data that lies within its described traffic profile.

The network performs admission control on the basis of information from the application and available network resources. It also commits to meeting the QoS requirements of the application as long as the traffic remains within the profile specifications. The network fulfills its commitment by maintaining per-flow state and then performing packet classification, policing, and intelligent queuing based on that state.

Cisco IOS QoS includes the following features that provide controlled load service, which is a kind of integrated service:

- The Resource Reservation Protocol (RSVP), which can be used by applications to signal their QoS requirements to the router.
- Intelligent queuing mechanisms, which can be used with RSVP to provide the following kinds of services:
 - Guaranteed rate service, which allows applications to reserve bandwidth to meet their requirements. For example, a Voice over IP (VoIP) application can reserve the required amount of bandwidth end-to-end using this kind of service. Cisco IOS QoS uses weighted fair queuing (WFQ) with RSVP to provide this kind of service.
 - Controlled load service, which allows applications to have low delay and high throughput even during times of congestion. For example, adaptive real-time applications, such as playback of a recorded conference, can use this kind of service. Cisco IOS QoS uses RSVP with Weighted Random Early Detection (WRED) to provide this kind of service.

Differentiated Service

Differentiated service is a multiple service model that can satisfy differing QoS requirements. However, unlike in the integrated service model, an application using differentiated service does not explicitly signal the router before sending data.

For differentiated service, the network tries to deliver a particular kind of service based on the QoS specified by each packet. This specification can occur in different ways, for example, using the IP Precedence bit settings in IP packets or source and destination addresses. The network uses the QoS specification to classify, mark, shape, and police traffic and to perform intelligent queuing.

The differentiated service model is used for several mission-critical applications and for providing end-to-end QoS. Typically, this service model is appropriate for aggregate flows because it performs a relatively coarse level of traffic classification.

Cisco IOS QoS includes the following features that support the differentiated service model:

- Committed access rate (CAR), which performs metering and policing of traffic, providing bandwidth management.
- Intelligent queuing schemes such as WRED and WFQ and their equivalent features on the Versatile Interface Processor (VIP), which are distributed WRED (DWRED) and distributed WFQ. These features can be used with CAR to deliver differentiated services.

For more information on how to implement differentiated services using the components of Cisco IOS software, see the [“Overview of DiffServ for Quality of Service”](#) chapter.

Cisco IOS QoS Features

The Cisco IOS QoS software provides the major features described in the following sections. Some of these features have been previously mentioned, and all of them are briefly introduced in this chapter.

- [Classification](#)
- [Congestion Management](#)
- [Congestion Avoidance](#)
- [Policing and Shaping](#)
- [Signaling](#)
- [Link Efficiency Mechanisms](#)
- [QoS Solutions](#)
- [Modular QoS Command-Line Interface](#)
- [Security Device Manager](#)
- [AutoQoS](#)

The features listed are described more fully in the overview chapters of this book, which is organized into parts, one for each of the major features listed. Each book part contains an overview chapter and one or more configuration chapters.

Classification

Classifying network traffic allows you to organize traffic (that is, packets) into traffic classes or categories on the basis of whether the traffic matches specific criteria. Classifying network traffic (used in conjunction with marking network traffic) is the foundation for enabling many QoS features on your network.

For more conceptual information about classification, see the [“Classification Overview”](#) chapter.

For more information about classifying network traffic, see the [“Classifying Network Traffic”](#) chapter.

For more information about classifying network traffic using Network-Based Application Recognition (NBAR), see the [“Classifying Network Traffic Using NBAR”](#) chapter.

For more information about marking network traffic, see the [“Marking Network Traffic”](#) chapter.

Congestion Management

Congestion management features operate to control congestion once it occurs. One way that network elements handle an overflow of arriving traffic is to use a queuing algorithm to sort the traffic and then determine some method of prioritizing it onto an output link. Each queuing algorithm is designed to solve a specific network traffic problem and has a particular effect on network performance.

The Cisco IOS software congestion management, or queuing, features include the following:

- FIFO queuing
- Priority queuing (PQ)
- Frame Relay permanent virtual circuit (PVC) interface priority queuing (FR PIPQ)
- Custom queuing (CQ)
- Weighted fair queuing (WFQ) and distributed WFQ (DWFQ)
- Class-based WFQ (CBWFQ) and Distributed CBWFQ (DCBWFQ)
- IP RTP Priority
- Frame Relay IP RTP Priority
- Low Latency Queuing (LLQ)
- Distributed LLQ (DLLQ)
- LLQ for Frame Relay

For more complete conceptual information on congestion management, see the [“Congestion Management Overview”](#) chapter.

For information on how to configure the various protocols that implement congestion management, see the following chapters:

- [“Configuring Weighted Fair Queueing”](#)
- [“Configuring Custom Queueing”](#)
- [“Configuring Priority Queueing”](#)

For complete command syntax information, see the [Cisco IOS Quality of Service Solutions Command Reference](#).

What Is Congestion in Networks?

To give you a more definite sense of congestion in networks, this section briefly describes some of its characteristics, drawing on the explanation presented by V. Paxson and S. Floyd in a paper titled *Wide Area Traffic: The Failure of Poisson Modeling*.

What does congestion look like? Consideration of the behavior of congested systems is not simple and cannot be dealt with in a simplistic manner, because traffic rates do not simply rise to a level, stay there a while and then subside. Periods of traffic congestion can be quite long, with losses that are heavily concentrated. In contrast to Poisson traffic models, linear increases in buffer size do not result in large decreases in packet drop rates; a slight increase in the number of active connections can result in a large increase in the packet loss rate. This understanding of the behavior of congested networks suggests that because the level of busy period traffic is not predictable, it would be difficult to efficiently size networks to reduce congestion adequately. Observers of network congestion report that in reality, traffic “spikes,” which causes actual losses that ride on longer-term ripples, which in turn ride on still longer-term swells.

FIFO Queuing

FIFO provides basic store-and-forward capability. FIFO is the default queuing algorithm in some instances, thus requiring no configuration. See the “[FIFO Queuing](#)” section on page 7 for a complete explanation of default configuration.

PQ

Designed to give strict priority to important traffic, PQ ensures that important traffic gets the fastest handling at each point where PQ is used. PQ can flexibly prioritize according to network protocol (such as IP, IPX, or AppleTalk), incoming interface, packet size, source/destination address, and so on.

FR PIPQ

FR PIPQ provides an interface-level PQ scheme in which prioritization is based on destination PVC rather than on packet contents. For example, FR PIPQ allows you to configure PVC transporting voices traffic to have absolute priority over a PVC transporting signaling traffic and a PVC transporting signaling traffic to have absolute priority over a PVC transporting data.

FR PIPQ provides four levels of priority: high, medium, normal, and low. The Frame Relay packet is examined at the interface for the data-link connection identifier (DLCI) value. The packet is then sent to the correct priority queue on the basis of the priority level configured for that DLCI.

CQ

CQ reserves a percentage of the available bandwidth of an interface for each selected traffic type. If a particular type of traffic is not using the bandwidth reserved for it, then other traffic types may use the remaining reserved bandwidth.

WFQ and DWFQ

WFQ applies priority (or weights) to identified traffic to classify traffic into conversations and determine how much bandwidth each conversation is allowed relative to other conversations. WFQ classifies traffic into different flows on the basis of such characteristics as source and destination address, protocol, and port and socket of the session.

To provide large-scale support for applications and traffic classes that require bandwidth allocations and delay bounds over the network infrastructure, Cisco IOS QoS includes a version of WFQ that runs only in distributed mode on VIPs. This version is called distributed WFQ (DWFQ). It provides increased flexibility in terms of traffic classification, weight assessment, and discard policy, and delivers Internet-scale performance on the Cisco 7500 series platforms.

For serial interfaces at E1 (2.048 Mbps) and below, WFQ is used by default. When no other queuing strategies are configured, all other interfaces use FIFO by default.

CBWFQ and DCBWFQ

The CBWFQ and DCBWFQ features extend the standard WFQ functionality to provide support for user-defined traffic classes. They allow you to specify the exact amount of bandwidth to be allocated for a specific class of traffic. Taking into account available bandwidth on the interface, you can configure up to 64 classes and control distribution among them.

DCBWFQ is intended for use on the VIP-based Cisco 7000 series routers with the Route Switch Processors (RSPs) and on the Cisco 7500 series routers.

IP RTP Priority

The IP RTP Priority feature provides a strict PQ scheme that allows delay-sensitive data such as voice to be dequeued and sent before packets in other queues are dequeued. This feature can be used on serial interfaces and Frame Relay PVCs in conjunction with either WFQ or CBWFQ on the same outgoing interface. In either case, traffic matching the range of UDP ports specified for the priority queue is guaranteed strict priority over other CBWFQ classes or WFQ flows; packets in the priority queue are always serviced first.

Frame Relay IP RTP Priority

The Frame Relay IP RTP Priority feature provides a strict PQ scheme on a Frame Relay PVC for delay-sensitive data such as voice. Voice traffic can be identified by its Real-Time Transport Protocol (RTP) port numbers and can be classified into a priority queue configured by the **frame-relay ip rtp priority** command. With this feature, voice traffic receives preferential treatment over nonvoice traffic.

LLQ

LLQ provides strict PQ on ATM VCs and serial interfaces. This feature allows you to configure the priority status for a class within CBWFQ, and it is not limited to UDP port numbers, as is IP RTP Priority. LLQ and IP RTP Priority can be configured at the same time, but IP RTP Priority takes precedence.

Additionally, the functionality of LLQ has been extended to allow you to specify the committed burst (Bc) size in LLQ and to change (or vary) the number of packets contained in the hold queue per-VC (on ATM adapters that support per-VC queuing). For more information, see the [“Congestion Management Overview”](#) chapter.

DLLQ

The DLLQ feature provides the ability to specify low-latency behavior for a traffic class on a VIP-based Cisco 7500 series router. DLLQ allows delay-sensitive data such as voice to be dequeued and sent before packets in other queues are dequeued.

The DLLQ feature also introduces the ability to limit the depth of a device transmission ring.

LLQ for Frame Relay

LLQ for Frame Relay provides strict PQ for voice traffic and WFQs for other classes of traffic. Before the release of this feature, LLQ was available at the interface and ATM VC levels. It is now available at the Frame Relay VC level when Frame Relay Traffic Shaping is configured.

Strict PQ improves QoS by allowing delay-sensitive traffic such as voice to be pulled from the queue and sent before other classes of traffic.

LLQ for Frame Relay allows you to define classes of traffic according to protocol, interface, or access lists. You can then assign characteristics to those classes, including priority, bandwidth, queue limit, and WRED.

Congestion Avoidance

Congestion avoidance techniques monitor network traffic loads in an effort to anticipate and avoid congestion at common network and internetwork bottlenecks before it becomes a problem. These techniques are designed to provide preferential treatment for premium (priority) class traffic under congestion situations while concurrently maximizing network throughput and capacity utilization and minimizing packet loss and delay. WRED and DWRED are the Cisco IOS QoS congestion avoidance features.

Router behavior allows output buffers to fill during periods of congestion, using the tail drop feature to resolve the problem when WRED is not configured. During tail drop, a potentially large number of packets from numerous connections are discarded because of lack of buffer capacity. This behavior can result in waves of congestion followed by periods during which the transmission link is not fully used. WRED obviates this situation proactively by providing congestion avoidance. That is, instead of waiting for buffers to fill before dropping packets, the router monitors the buffer depth and performs early discards on selected packets sent over selected connections.

WRED is the Cisco implementation of the RED class of congestion avoidance algorithms. When RED is used and the source detects the dropped packet, the source slows its transmission. RED is primarily designed to work with TCP in IP internetwork environments.

WRED can also be configured to use the DSCP value when it calculates the drop probability of a packet, enabling WRED to be compliant with the DiffServ standard being developed by the Internet Engineering Task Force (IETF).

For more complete conceptual information, see the [“Congestion Avoidance Overview”](#) chapter.

For information on how to configure WRED, DWRED, flow-based WRED, and DiffServ compliant WRED, see the [“Configuring Weighted Random Early Detection”](#) chapter.

For complete command syntax information, see the [Cisco IOS Quality of Service Solutions Command Reference](#).

WRED

WRED, the Cisco implementation of RED, combines the capabilities of the RED algorithm with IP Precedence to provide preferential traffic handling for higher priority packets. It can selectively discard lower priority traffic when the interface begins to get congested and provide differentiated performance characteristics for different classes of service. WRED is also RSVP-aware. WRED is available on the Cisco 7200 series Route Switch Processor (RSP).

DWRED

DWRED is the Cisco high-speed version of WRED. The DWRED algorithm was designed with ISP providers in mind; it allows an ISP to define minimum and maximum queue depth thresholds and drop capabilities for each class of service. DWRED, which is available on the Cisco 7500 series routers or the Cisco 7000 series router with RSPs, is analogous in function to WRED, which is available on the Cisco 7200 series RSP.

Flow-Based WRED

The Flow-Based WRED feature forces WRED to afford greater fairness to all flows on an interface in regard to how packets are dropped. To provide fairness to all flows, the Flow-Based WRED feature has the following functionality:

- It ensures that flows that respond to WRED packet drops by backing off packet transmission are protected from flows that do not respond to WRED packet drops.
- It prohibits a single flow from monopolizing the buffer resources at an interface.

DiffServ Compliant WRED

The DiffServ Compliant WRED feature extends the functionality of WRED to enable support for Differentiated Services (DiffServ) and Assured Forwarding (AF) Per Hop Behavior (PHB). This feature enables customers to implement AF PHB by coloring packets according to DSCP values and then assigning preferential drop probabilities to those packets.

The DiffServ and the AF PHB standards are supported by this feature.

Policing and Shaping

For traffic policing, Cisco IOS QoS includes traffic policing capabilities implemented through the rate-limiting aspects of CAR and the Traffic Policing feature.

For traffic shaping, Cisco IOS QoS includes Generic Traffic Shaping (GTS), Class-Based Shaping, and Frame Relay Traffic Shaping (FRTS). These features allow you to regulate packet flow (that is, the flow of traffic) on your network.

For more complete conceptual information about traffic policing and traffic shaping, see the [“Policing and Shaping Overview”](#) chapter.

Signaling

Cisco IOS QoS signaling provides a way for an end station or network node to signal its neighbors to request special handling of certain traffic. QoS signaling is useful for coordinating the traffic-handling techniques provided by other QoS features. It plays a key role in configuring successful overall end-to-end QoS service across your network.

Cisco IOS QoS signaling takes advantage of IP. Either in-band (IP Precedence, 802.1p) or out-of-band (RSVP) signaling is used to indicate that a particular QoS service is desired for a particular traffic classification. Together, IP Precedence and RSVP provide a robust combination for end-to-end QoS signaling: IP Precedence signals for differentiated QoS, and RSVP signals for guaranteed QoS.

Cisco IOS software offers the following features and functionality associated with signaling:

- ATM User Network Interface (UNI) signaling and Frame Relay Local Management Interface (LMI)
Achieves the end-to-end benefits of IP Precedence and RSVP signaling, and provides signaling into their respective backbone technologies.
- Common Open Policy Service (COPS) with RSVP
Achieves centralized monitoring and control of RSVP signaling.
- Subnetwork Bandwidth Manager (SBM)
Enables admission control over IEEE 802-styled networks.
- RSVP-ATM QoS Interworking feature
Provides support for Controlled Load Service using RSVP over an ATM core network.
- RSVP support for Low Latency Queuing (LLQ) and Frame Relay.

For more complete conceptual information, see the [“Signalling Overview”](#) chapter.

For information on how to configure the various protocols that implement signaling, see the following chapters:

- [“Configuring RSVP”](#)
- [“Configuring RSVP Support for LLQ”](#)
- [“Configuring RSVP Support for Frame Relay”](#)
- [“Configuring COPS for RSVP”](#)
- [“Configuring Subnetwork Bandwidth Manager”](#)
- [“Configuring RSVP-ATM QoS Interworking”](#)

For complete command syntax information, see the [Cisco IOS Quality of Service Solutions Command Reference](#).

Link Efficiency Mechanisms

Cisco IOS software offers a number of link-layer efficiency mechanisms or features designed to reduce latency and jitter for network traffic. These link efficiency mechanisms include the following:

- Multilink PPP (MLP)
- Frame Relay Fragmentation
- Header Compression

These mechanisms work with queuing and fragmentation to improve the efficiency and predictability of the application service levels.

For more complete conceptual information, see the [“Link Efficiency Mechanisms Overview”](#) chapter.

Multilink PPP

At the highest level, MLP provides packet interleaving, packet fragmentation, and packet resequencing across multiple logical data links. The packet interleaving, packet fragmentation, and packet resequencing are used to accommodate the fast transmission times required for sending real-time packets (for example, voice packets) across the network links. MLP is especially useful over slow network links (that is, a network link with a link speed less than or equal to 768 kbps).

For more conceptual information about MLP, see [“Reducing Latency and Jitter for Real-Time Traffic Using Multilink PPP”](#) module.

Frame Relay Fragmentation

Cisco has developed the following a number of methods of performing Frame Relay fragmentation, including the following:

- End-to-end FRF.12 (and higher) fragmentation
- Frame Relay fragmentation using FRF.11 Annex C (and higher)
- Cisco proprietary encapsulation

For more information about Frame Relay fragmentation, see the [“Frame Relay Queuing and Fragmentation at the Interface”](#) module.

Header Compression

Header compression is a mechanism that compresses the IP header in a packet before the packet is transmitted. Header compression reduces network overhead and speeds up the transmission of Real-Time Transport Protocol (RTP) and Transmission Control Protocol (TCP) packets. Header compression also reduces the amount of bandwidth consumed when the RTP or TCP packets are transmitted.

For more information about header compression, see the [“Link Efficiency Mechanisms Overview”](#) chapter.

QoS Solutions

The Cisco IOS QoS software includes a number of features collectively referred to as “QoS solutions.” These software features include the following:

- IP to ATM CoS
- QoS features for voice
- Differentiated services implementations
- QoS: Classification, Policing, and Marking on a Layer 2 Tunneling Protocol (L2TP) Access Concentrator (LAC)
- QoS Bandwidth Estimation

IP to ATM CoS

IP to ATM CoS is a feature suite that maps QoS characteristics between IP and ATM, making it possible to support differentiated services in network service provider environments.

Network managers can use existing features such as CAR or PBR to classify and mark different IP traffic by modifying the IP Precedence field in the IPv4 packet header. Subsequently, WRED or DWRED can be configured on a per-VC basis so that the IP traffic is subject to different drop probabilities (and therefore priorities) as IP traffic coming into a router competes for bandwidth on a particular VC.

IP to ATM CoS provides support for ATM VC bundle management, allowing you to configure multiple VCs that have different QoS characteristics between any pair of ATM-connected routers.

IP to ATM CoS also provides for per-VC WFQ and CBWFQ, which allows you to apply CBWFQ functionality—normally applicable at the interface or subinterface levels only—to an individual VC configured for IP to ATM CoS. You can use this feature to apply either CBWFQ or flow-based WFQ on a per-VC basis.

For more complete conceptual information, see the [“IP to ATM Class of Service Overview”](#) chapter.

For information on how to configure IP to ATM CoS, see the [“Configuring IP to ATM Class of Service”](#) chapter.

QoS Features for Voice

Many of the QoS features already mentioned in this chapter are useful for voice applications. For a high-level overview of Cisco IOS QoS features for voice, see the [“Introduction to QoS Features for Voice”](#) chapter.

Differentiated Services Implementations

Many of the QoS features can be used to implement Differentiated Services on your network. For a high-level overview of how to use the Cisco IOS components to implement Differentiated Services, see the [“Overview of DiffServ for Quality of Service”](#) chapter.

QoS: Classification, Policing, and Marking on a LAC

The QoS: Classification, Policing, and Marking on a Layer 2 Tunneling Protocol (L2TP) Access Concentrator (LAC) feature allows service providers to classify packets based upon the IP type of service (ToS) bits in an embedded IP packet. The classification will be used to police the incoming traffic according to the differentiated services code point (DSCP) value. The purpose of classifying the packet by examining its encapsulation is to simplify the implementation and configuration needed for a large number of Point-to-Point Protocol (PPP) sessions.

For more information about this feature, see the [“QoS: Classification, Policing, and Marking on a LAC”](#) module.

QoS Bandwidth Estimation

The QoS Bandwidth Estimation feature uses Corvil Bandwidth technology to allow you as a network manager to determine the bandwidth requirements to achieve user-specified QoS targets for networked applications.

For more information about the QoS Bandwidth Estimation feature, see the [QoS Bandwidth Estimation](#) module.

Modular QoS Command-Line Interface

The Modular Quality of Service Command-Line Interface (MQC) is a CLI structure that allows you to define a traffic class, create a traffic policy (policy map), and attach the traffic policy to an interface. The traffic policy contains the QoS feature that will be applied to the traffic class.

The MQC structure consists of the following three high-level steps:

-
- Step 1** Define a traffic class by using the **class-map** command. A traffic class is used to classify traffic.
 - Step 2** Create a traffic policy by using the **policy-map** command. (The terms *traffic policy* and *policy map* are often synonymous). A traffic policy (policy map) contains a traffic class and one or more QoS features that will be applied to the traffic class. The QoS features in the traffic policy determine how to treat the classified traffic.
 - Step 3** Attach the traffic policy (policy map) to the interface by using the **service-policy** command.
-

For concepts and tasks associated with the MQC, see the [“Applying QoS Features Using the MQC”](#) module.

Security Device Manager

The Cisco Router and Security Device Manager (SDM) provides an intuitive, graphical user interface for configuring and monitoring advanced IP-based QoS functionality within Cisco routers.

For a high-level overview of SDM, see the [“Security Device Manager Overview”](#) chapter.

AutoQoS

The AutoQoS feature allows you to automate the delivery of QoS on your network and provides a means for simplifying the implementation and provisioning of QoS.

For more information about AutoQoS, see the [“AutoQoS — VoIP”](#) module or the [“AutoQoS for the Enterprise”](#) module.

CCDE, CCENT, CCSI, Cisco Eos, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nurse Connect, Cisco Pulse, Cisco SensorBase, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco:Financed (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumin, Cisco Nexus, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Follow Me Browsing, GainMaker, iLynx, IOS, iPhone, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, SenderBase, SMARTnet, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0910R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

© 2007 Cisco Systems, Inc. All rights reserved.