

# Introduction: Quality of Service Overview

---

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

## What Is Quality of Service?

QoS refers to the ability of a network to provide better service to selected network traffic over various underlying technologies including Frame Relay, Asynchronous Transfer Mode (ATM), Ethernet and 802.1 networks, SONET, and IP-routed networks. In particular, QoS features provide better and more predictable network service by:

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

## 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 queueing, scheduling, and traffic shaping features.
- QoS signalling techniques for coordinating QoS from end-to-end 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 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, then select the appropriate QoS feature or features.

In general, edge routers perform the following QoS functions:

- Packet classification
- 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 business networking users—have their own QoS requirements; in many areas, however, these requirements overlap. The Cisco IOS QoS features described in the section “Cisco’s QoS Features” later in this chapter address these diverse and common needs.

Enterprise networks, for example, must provide end-to-end QoS solutions across the various platforms comprising the network; providing solutions for heterogeneous platforms often requires taking 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 it requires.

ISPs require assured scalability and performance. For example, ISPs that long have 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 promotes:

- 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 File Transfer Protocol (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’s QoS features make certain 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; and 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 Quality of Service 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 from one another in how they enable applications to send data and in the ways in which the network attempts to deliver that data. For instance, a different service model applies to real-time applications, such as audio and video conferencing and IP telephony, from the one that applies to 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 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 ability 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 doing so for a best-effort service.

This section describes these service models:

- Best-Effort Service
- Integrated Service
- Differentiated Service

The features in Cisco IOS QoS software address the requirements for these service models.

### 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) queueing. 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 sending data. The request is made by explicit signalling; 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, based on 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 these features that provide controlled load service, which is a kind of integrated service:

- Resource Reservation Protocol (RSVP) can be used by applications to signal their QoS requirements to the router.
- Intelligent queuing mechanisms 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 32 Mbps 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 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, 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 these features that support the differentiated service model:

- Committed access rate (CAR), which performs packet classification through IP Precedence and QoS group settings. CAR 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 VIP-Distributed WRED and VIP-Distributed WFQ. These features can be used with CAR to deliver differentiated services.

## Cisco's QoS Features

The Cisco IOS QoS software provides these major features, some of which have been previously mentioned, and all of which are briefly introduced in this chapter and then described more fully in the overview chapters of this book:

- Classification
- Congestion Management
- Congestion Avoidance
- Policing and Shaping
- Signalling
- Link Efficiency Mechanisms

### Classification

Packet classification features provide the capability to partition network traffic into multiple priority levels or classes of service. For example, using the three precedence bits in the type of service (ToS) field of the IP packet header—two of the values are reserved for other purposes—you can categorize packets into a limited set of up to six traffic classes. After you classify packets, you can utilize other QoS features to assign the appropriate traffic handling policies including congestion management, bandwidth allocation, and delay bounds for each traffic class.

Packets can also be classified by external sources, that is, by a customer or by a downstream network provider. You can either allow the network to accept the classification or override it and reclassify the packet according to a policy that you specify.

Packets can be classified based on policies specified by the network operator. Policies can be set that include classification based on physical port, source or destination IP or MAC address, application port, IP protocol type, and other criteria that you can specify by access lists or extended access lists.

You can use Cisco IOS QoS policy-based routing (PBR) and the classification features of Cisco IOS QoS CAR to classify packets. You can use Border Gateway Protocol (BGP) Policy Propagation to propagate destination-based packet classification policy throughout a large network via BGP routing updates. This section gives a brief description of these features.

For more complete conceptual information on packet classification, see the chapter “Classification Overview” in this book.

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

- “Configuring Policy-Based Routing”
- “Configuring QoS Policy Propagation via Border Gateway Protocol”
- “Configuring Committed Access Rate”

For complete command syntax information, see the *Quality of Service Solutions Command Reference*.

### IP Precedence

IP Precedence allows you to specify a packet's class of service using the three precedence bits in the IPv4 header's ToS field. Other features configured throughout the network can then use these bits to determine how to treat the packet in regard to the type of service to grant it. For example, although IP Precedence is not a queueing method, other queueing methods such as WFQ can use the IP Precedence setting of the packet to prioritize traffic.

### Policy-Based Routing

Cisco IOS QoS PBR allows you to do the following:

- Classify traffic based on extended access list criteria.
- Set IP Precedence bits.
- Route specific traffic to engineered paths, which may be required to allow a specific QoS service through the network.

Classification of traffic through PBR allows you to identify traffic for different classes of service at the perimeter of the network and then implement QoS defined for each class of service in the core of the network using priority, custom, or weighted fair queueing techniques. This process obviates the need to classify traffic explicitly at each WAN interface in the core-backbone network.

Some possible applications for policy routing are to provide equal access, protocol-sensitive routing, source-sensitive routing, routing based on interactive versus batch traffic, or routing based on dedicated links.

### Committed Access Rate (Packet Classification)

CAR is the main feature supporting packet classification. CAR uses the ToS bits in the IP header to classify packets. You can use the CAR classification commands to classify and reclassify a packet.

Here are some example packet classification policies:

- All packets received on a particular T1 line are classified as high priority (port-based classification).
- All HTTP traffic is classified as medium priority (application classification).
- Video traffic from a specified IP address is classified as medium priority.
- Packets bound for particular destinations are classified as high priority traffic (for example, international traffic or traffic bound for a premium customer).
- Some packets are classified for subrate IP services. The network operator delivers a physical T1/E1 or T3/E3 line to the customer, but offers a less expensive subrate service, for example, 1 Mbps on an E1 line or 10 Mbps on a T3 line. The customer pays for the subrate bandwidth and may be upgraded to additional access bandwidth over time based on demand. CAR limits the traffic rate available to the customer and delivered to the network to the agreed upon rate limit (with the ability to temporarily burst over the limit). The network operator may upgrade the service without any physical network arrangement.
- Traffic is classified for exchange point traffic control. An ISP offers transit services to downstream ISPs via exchange point connectivity provided by a Layer 2 switch. The upstream provider utilizes MAC-address rate limits provided by CAR to enforce bandwidth usage limitations on the downstream ISPs.

---

**Note** CAR also implements rate-limiting services, described briefly later in this chapter.

---

## BGP Policy Propagation

BGP provides a powerful, scalable means of utilizing attributes, such as community values, to propagate destination-based packet classification policy throughout a large network via BGP routing updates. Packet classification policy can be scalably propagated via BGP without writing and deploying complex access lists at each of a large number of routers. BGP ensures that return traffic to customers is handled as premium traffic by the network.

## 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 queueing algorithm to sort the traffic, then determine some method of prioritizing it onto an output link. Each queueing algorithm was designed to solve a specific network traffic problem and has a particular effect on network performance. The Cisco IOS software congestion management, or queueing, features include FIFO, priority queueing (PQ), custom queueing (CQ), and WFQ (and VIP-Distributed WFQ).

For more complete conceptual information on packet classification, see the chapter “Congestion Management Overview” in this book.

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 *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, as traffic rates do not simply rise to a level, stay there a while, 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, and they in turn ride on still longer-term swells.

## FIFO Queueing

FIFO provides basic store and forward capability. FIFO is the default queueing algorithm in some instances, thus requiring no configuration. See “WFQ and VIP-Distributed WFQ” later in this chapter for 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 forth.

### CQ

CQ reserves a percentage of an interface's available bandwidth 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 VIP-Distributed WFQ

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 based on 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 requiring 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 VIP-Distributed WFQ. 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.

## 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 its counterpart for the VIP, VIP-Distributed WRED, are Cisco's IOS QoS congestion avoidance features.

Router behavior allows output buffers to fill during periods of congestion, using tail drop 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 transmitted over selected connections.

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

For more complete conceptual information, see the chapter "Congestion Avoidance Overview" in this book.

For information on how to configure WRED and VIP-Distributed WRED, see the chapter “Configuring Weighted Random Early Detection” in this book.

For complete command syntax information, see the *Quality of Service Solutions Command Reference*.

## WRED

Cisco’s implementation of RED, called WRED, 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 RSP processors.

## VIP-Distributed WRED

VIP-Distributed WRED is Cisco’s high-speed version of WRED. The VIP-Distributed WRED 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.

VIP-Distributed WRED, which is available on the Cisco 7500 series VIPs, is analogous in function to WRED, which is available on the 7200 series RSP processors.

## Policing and Shaping

Cisco IOS QoS includes traffic policing capabilities implemented through the rate-limiting aspects of CAR and traffic shaping capabilities provided by the Generic Traffic Shaping (GTS) and Frame Relay Traffic Shaping (FRTS) protocols.

For more complete conceptual information, see the chapter “Policing and Shaping Overview” in this book.

For information on how to configure these protocols, which implement policing and shaping, see the following chapters:

- “Configuring Generic Traffic Shaping”
- “Configuring Frame Relay and Frame Relay Traffic Shaping”

For complete command syntax information, see the *Quality of Service Solutions Command Reference*.

## CAR Rate Limiting

The rate-limiting feature of CAR provides the network operator with the means to define Layer 3 aggregate or granular access, or egress bandwidth rate limits, and to specify traffic handling policies when the traffic either conforms to or exceeds the specified rate limits. Aggregate access or egress matches all packets on an interface or subinterface. Granular access or egress matches a particular type of traffic based on precedence. You can designate CAR rate-limiting policies based on physical port, packet classification, IP address, MAC address, application flow, and other criteria specifiable by access lists or extended access lists. CAR rate limits may be implemented either on input or output interfaces or subinterfaces including Frame Relay and ATM subinterfaces.

An example of use of CAR’s rate-limiting capability is application-based rates limiting HTTP World Wide Web traffic to 50 percent of link bandwidth, which ensures capacity for non-Web traffic including mission-critical applications.

## Shaping

Cisco's IOS QoS software includes these two traffic shaping features, which manage traffic and congestion on the network:

- GTS, which provides a mechanism to control the flow of outbound traffic on a particular interface. It reduces outbound traffic flow to avoid congestion by constraining specified traffic to a particular bit rate. Traffic adhering to a particular profile can be shaped to meet downstream requirements, eliminating bottlenecks in topologies with data rate mismatches.
- FRTS, which provides parameters, such as the following, that are useful for managing network traffic congestion:
  - Committed information rate (CIR)
  - Forward and backward explicit congestion notification (FECN/BECN)
  - The discard eligible (DE) bit

For some time Cisco has provided support for FECN for DECnet and OSI, BECN for SNA traffic using direct LLC2 encapsulation via RFC 1490, and DE bit support. The FRTS feature builds upon this Frame Relay support by providing additional capabilities that improve the scalability and performance of a Frame Relay network by increasing the density of virtual circuits and improving response time.

FRTS applies only to Frame Relay permanent virtual circuits (PVCs) and switched virtual circuits (SVCs).

## Signalling

Cisco IOS QoS signalling provides a way for an end station or network node to signal its neighbors to request special handling of certain traffic. QoS signalling 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 signalling takes advantage of IP. Either in-band (IP Precedence, 802.1p) or out-of-band (RSVP) signalling 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 signalling: IP Precedence signals for differentiated QoS and RSVP for guaranteed QoS.

To achieve the end-to-end benefits of IP Precedence and RSVP signalling, Cisco IOS QoS software offers ATM User Network Interface (UNI) signalling and the Frame Relay Local Management Interface (LMI) to provide signalling into their respective backbone technologies.

For more complete conceptual information, see the chapter "Signalling Overview" in this book.

For information on how to configure RSVP, see the *Quality of Service Solutions Command Reference*.

## Link Efficiency Mechanisms

Cisco IOS QoS software offers these two link efficiency mechanisms that work in conjunction with queueing and traffic shaping to improve efficiency and predictability of the application services levels: Compressed Real-Time Protocol (CRTP) and Link Fragmentation and Interleaving (LFI).

For more complete conceptual information, see the chapter "Link Efficiency Mechanisms Overview" in this book.

For information on how to configure CRTP, see the chapter “Configuring Compressed Real-Time Protocol” in this book.

For information on how to configure LFI, see the chapter “Configuring Link Fragmentation and Interleaving for Multilink PPP” in this book.

For complete command syntax information, see the *Quality of Service Solutions Command Reference*.

## Compressed Real-Time Protocol

Real-Time Protocol (RTP) is a host-to-host protocol used for carrying newer multimedia application traffic, including packetized audio and video, over an IP network. RTP provides end-to-end network transport functions intended for applications transmitting real-time requirements, such as audio, video, or simulation data multicast or unicast network services.

To avoid the unnecessary consumption of available bandwidth, the RTP header compression feature—referred to as CRTP—is used on a link-by-link basis.

## Link Fragmentation and Interleaving

Interactive traffic, such as Telnet and Voice over IP, is susceptible to increased latency and jitter when the network processes large packets, such as LAN-to-LAN FTP Telnet transfers traversing a WAN link. This susceptibility increases as the traffic is queued on slower links. Cisco IOS QoS LFI reduces delay and jitter on slower speed links by breaking up large datagrams and interleaving low-delay traffic packets with the resulting smaller packets.

