Cisco IOS(TM) Software Quality of Service Solutions

Executive Summary

Today's networks are carrying more data in the form of bandwidth-intensive, real-time voice, video and data, which stretch network capability and resources. Cisco IOS® software provides a toolbox full of quality of service (QoS) solutions to help you solve problems caused by increasing traffic demands on the network.

The goal of QoS is to provide better and more predictable network service by providing dedicated bandwidth, controlled jitter and latency, and improved loss characteristics. QoS achieves these goals by providing tools for managing network congestion, shaping network traffic, using expensive wide-area links more efficiently, and setting traffic policies across the network.

Internet service providers, small and medium-sized business networks, as well as large enterprise networks can take advantage of the solutions provided by the Cisco IOS QoS software features.

This white paper presents the Cisco IOS QoS implementation in the following sections:

- Introduction
- QoS Framework
- The Cisco QoS Toolkit
- QoS Signaling
- Cisco QoS Policy, Management, and Accounting Capabilities
- Network QoS Application Examples
- QoS Looking Forward

In addition, reference matrices and commonly requested QoS definitions are included in the appendices.

Introduction

Cisco Systems, the worldwide leader in networking for the Internet, provides products and services that give people access to the information they need by connecting information devices through intelligent, secure, and reliable networks. As the leader in global networking, Cisco's breadth of product covers virtually every market segment. Cisco is, therefore, in a unique position to coordinate and deliver end-to-end QoS across varying network technologies and architectures to achieve end-to end QoS solutions.

Networking users in general span three major market segments: major enterprises, network service providers, and the small and medium-sized business segment. Each segment has its own QoS requirements, but they also have many overlapping needs.

Network managers in today's enterprise networks must contend with numerous and diverse system platforms, network architectures, and protocols. Providing end-to-end QoS solutions across the various platforms often requires more than just linking them together; it also requires a different approach for each technology. Enterprises are increasingly depending on their networks

to carry complex mission-critical applications and databases such as SAP, PeopleSoft, and Oracle. These networks are also experiencing increased traffic from Web and multimedia applications. QoS prioritizes this traffic to ensure that mission-critical applications get the service they require, while simultaneously servicing these newer multimedia applications.

Internet service providers (ISPs) require assured scalability and performance. The ISP marketplace is also highly competitive and characterized by phenomenal growth. ISPs, who have traditionally offered best-effort IP connectivity, are now planning networks to transport voice, video, and other real-time, critical application data. ISPs need a model that will allow them to offer differentiated services to their customers, yet allow them to remain profitable. QoS provides the basis for a new business model by allowing ISPs to differentiate traffic from various customers or applications.

In the small and medium-sized business segment, managers are experiencing first hand the rapid growth of business on the Internet. Not so long ago, the "global networked business" concept was just that. Every day we witness more and more businesses participating in the reality. Besides the increased demands of Internet traffic, small and medium-sized business networks must also handle increasingly complex business applications. QoS lets the network handle the difficult task of utilizing an expensive wide-area network connection in the most efficient way for business applications.

With these increasing demands, it is important to find ways of utilizing and expanding upon existing network resources. Cisco IOS software allows the addition of QoS capabilities to the network primarily through software upgrades, helping to preserve valuable investments in network equipment, while meeting constantly growing needs.

This white paper provides an overview of the QoS architectural framework that explains how QoS applies in the network, and it provides details on technologies that Cisco IOS software provides in each piece of the architecture. The paper concludes with some examples of how these pieces work together to provide QoS services that help you get the most from scarce network resources. Tables summarizing feature capabilities and availability are contained in the appendix, and a glossary of QoS terms is also provided.

Network QoS Defined

QoS refers to the ability of a network to provide better service to selected network traffic over various technologies, including Frame Relay, Asynchronous Transfer Mode (ATM), Ethernet and 802.1 networks, SONET, as well as IP-routed networks that may use any or all of these underlying technologies. Primary goals of QoS include dedicated bandwidth, controlled jitter and latency, (required by some real-time and interactive traffic), and improved loss characteristics. QoS technologies provide the elemental building blocks that will be used for future business applications in campus, WAN, and service provider networks.

The Cisco IOS QoS software enables complex networks to control and predictably service a variety of networked applications and traffic types. Almost any network can take advantage of QoS for optimum efficiency, whether it is a small corporate network, Internet service provider, or enterprise network. The Cisco IOS QoS software provides these benefits:

- *Control over resources*—You have control over which resources (bandwidth, equipment, wide-area facilities, and so on) are being used. As an example, you can limit the bandwidth consumed over a backbone link by FTP transfers or give priority to an important database access.
- *More efficient use of network resources*—Using Cisco's network analysis management and accounting tools, you will know what your network is being used for and that you are servicing the most important traffic to your business.
- *Tailored services*—The control and visibility provided by QoS enables Internet service providers to offer carefully tailored grades of service differentiation to their customers.
- *Coexistence of mission-critical applications*—Cisco's QoS technologies 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.
- *Foundation for a fully integrated network in the future*—Implementing Cisco QoS technologies in your network now is a good first step toward the fully integrated multimedia network needed in the near future. For example, you can implement weighted fair queuing today and get its immediate benefit of increasing service predictability and IP Precedence signaling for traffic

differentiation. You reap additional benefits in the future, because weighted fair queuing is resource ReSerVation Protocol (RSVP) enabled, thereby allowing you to take advantage of dynamically signaled QoS from the inevitable coming wave of RSVP-enabled applications.

The following sections more fully describe the Cisco QoS architecture and the QoS tools that are provided in the Cisco IOS QoS software.

QoS Framework

This section describes the basic framework for providing QoS in the network. The "Basic QoS Architecture" section describes the three components necessary to deliver QoS across a network comprising heterogeneous technology (IP, ATM, LAN switches, and so on). The "End-to-End QoS Service Levels" section describes the three basic levels of QoS service that can be provided across a heterogeneous infrastructure.

Basic QoS Architecture

The basic architecture introduces the three fundamental pieces for QoS implementation (see Figure 1):

- QoS within a single network element (for example, queuing, scheduling, and traffic shaping tools)
- QoS signaling techniques for coordinating QoS from end to end between network elements
- QoS policy, management, and accounting functions to control, and administer end-to-end traffic across a network

Figure 1 Basic QoS Architecture

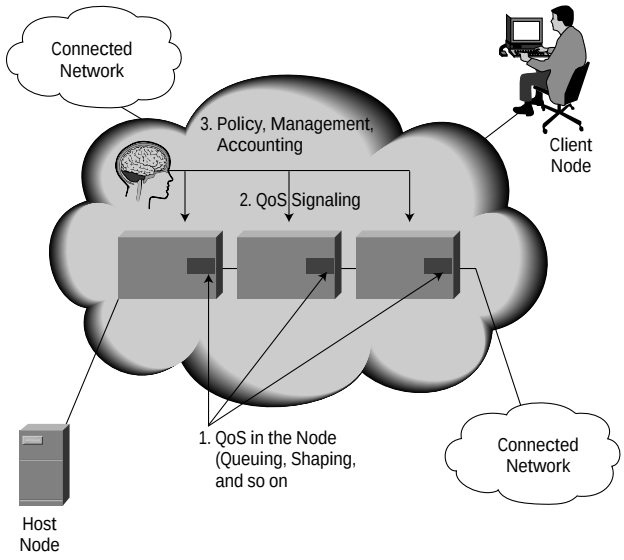

End-to-End QoS Service Levels

Service levels refer to the actual end-to-end QoS capabilities, meaning the ability of a network to deliver service needed by specific network traffic from end-to-end or edge-to-edge. The QoS services differ in their level of "QoS strictness," which describes how tightly the service can be bound by specific bandwidth, delay, jitter, and loss characteristics.

There are three basic levels of end-to-end QoS service that can be provided across a heterogeneous network, as shown in Figure 2.

- *Best Effort Service*—Also known as lack of QoS, best effort service is basic connectivity with no guarantees.
- *Differentiated Service (also Called Soft QoS)*—Some traffic is treated better than the rest (faster handling, more bandwidth on average, lower loss rate on average). This is a statistical preference, not a hard and fast guarantee.
- *Guaranteed Service (also Called Hard QoS)*—An absolute reservation of network resources for specific traffic.

These terms are defined in more detail in Appendix 2.

Figure 2 End-to-End QoS Service Levels

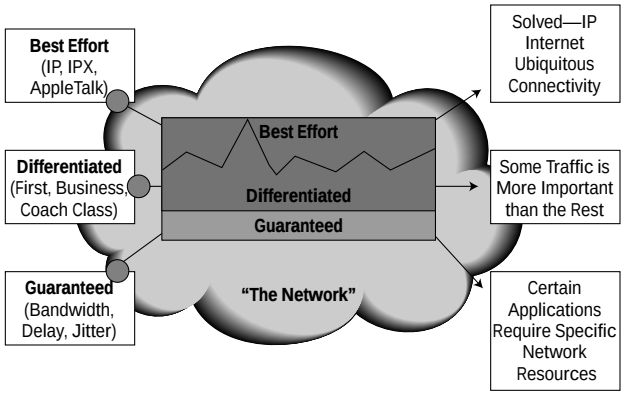

Deciding on which type of service is appropriate to deploy in the network depends on several factors:

- The application or problem the customer is trying to solve. Each of the three types of service is appropriate for certain applications. This does not imply that a customer must migrate to differentiated and then to guaranteed service (although we believe that many eventually will). A differentiated service—or even best effort service—may be appropriate depending on the customer application requirements.
- The rate at which customers can realistically upgrade their infrastructures. There is a natural upgrade path from the technology needed to provide differentiated services to that needed to provide guaranteed services, which is a superset of that needed for differentiated services.
- The cost-implementing and deploying guaranteed service is likely to be more expensive than doing so for a differentiated service.

The next three sections describe the tools that Cisco IOS provides in each section of the architecture, which, when combined, can create end-to-end QoS or simply solve specific problems at various points in the network.

The Cisco QoS Toolkit

Cisco IOS software provides a variety of QoS tools to provide the service levels described above. These tools are typically used within a single network element, as shown in the basic architecture depicted in Figure 1. Typically, these tools are turned on at an interface to provide the right QoS characteristics for a specific network application. The Cisco IOS QoS tools provide three major functions—congestion management (queuing and scheduling), congestion avoidance, and traffic shaping and policy making. In addition, Cisco IOS tools provides link efficiency mechanisms that integrate with the other three functions to provide additional improved QoS service.

Congestion Management Tools

One way that network elements handle an overflow of arriving traffic is to use a queuing algorithm to sort the traffic, then determine some method of prioritizing it onto an output link. Cisco IOS software includes the following queuing tools:

- First In, First Out (FIFO) Queuing
- Priority Queuing (PQ)
- Custom Queuing (CQ)
- Weighted Fair Queuing (WFQ)

Each queuing algorithm was designed to solve a specific network traffic problem and has a particular effect on network performance, as described in the following sections.

FIFO Provides Basic Store-and-Forward Capability

In its simplest form, FIFO queuing involves storing packets when the network is congested and forwarding them in order of arrival when the network is no longer congested. FIFO is the default queuing algorithm in some instances, thus requiring no configuration, but it has several shortcomings. Most importantly, FIFO queuing makes no decision about packet priority; the order of arrival determines bandwidth, promptness, and buffer allocation. Nor does it provide protection against ill-behaved applications (sources). Bursty sources can cause high delays in delivering time-sensitive application traffic, and potentially to network control and signaling messages. FIFO queuing was a necessary first step in controlling network traffic, but today's intelligent networks need more sophisticated algorithms. Cisco IOS software implements queuing algorithms that avoid the shortcomings of FIFO queuing.

PQ Prioritizes Traffic

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

In PQ, each packet is placed in one of four queues—High, Medium, Normal, or Low—based on an assigned priority. Packets that are not classified by this priority-list mechanism fall into the Normal queue; see Figure 3. During transmission, the algorithm gives higher-priority queues absolute preferential treatment over low-priority queues. This is a simple and intuitive approach but can cause queuing delays that the higher-priority traffic might have experienced to be randomly transferred to the lower-priority traffic, increasing jitter on the lower-priority traffic. Higher-priority traffic can be rate limited to avoid this problem.

Figure 3 Priority Queuing

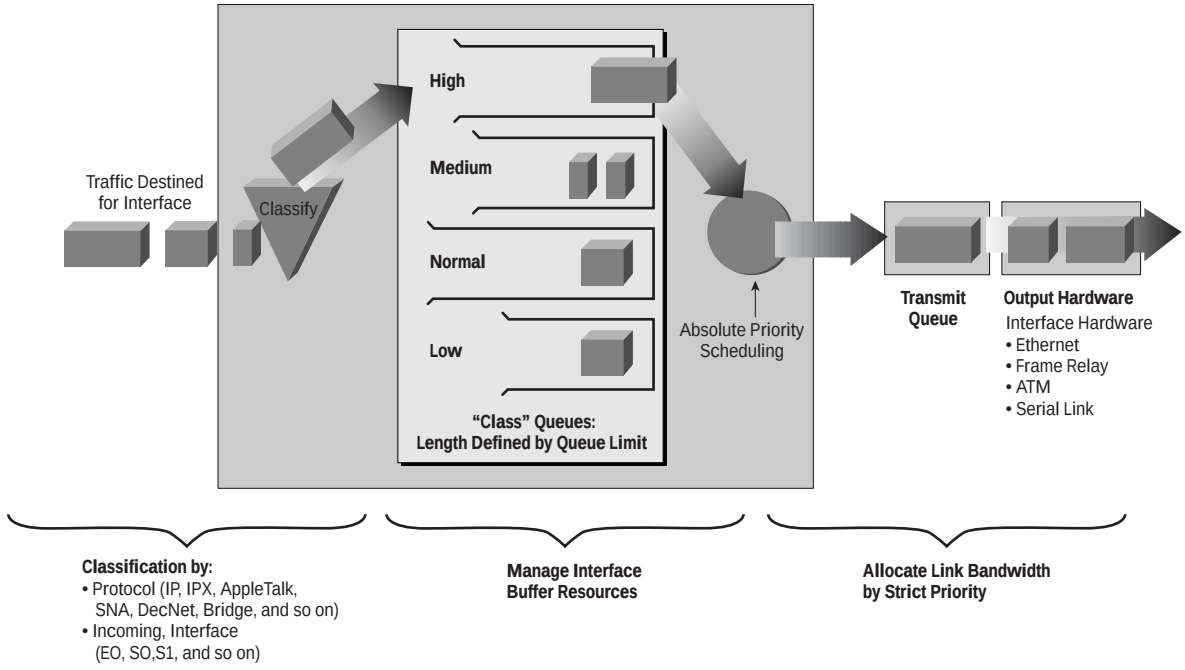

*For information on specific interface support, see Appendix 1.

PQ is useful for making sure that mission-critical traffic traversing various WAN links gets priority treatment. For example, Cisco uses PQ to ensure that important Oracle-based sales reporting data gets to its destination ahead of other less critical traffic. PQ currently uses static configuration and thus does not automatically adapt to changing network requirements.

CQ Guarantees Bandwidth

Custom queuing (CQ) was designed to allow various applications or organizations to share the network among applications with specific minimum bandwidth or latency requirements. In these environments, bandwidth must be shared proportionally between applications and users. You can use the Cisco CQ feature to provide guaranteed bandwidth at a potential congestion point, assuring the specified traffic a fixed portion of available bandwidth and leaving the remaining bandwidth to other traffic. Custom queuing handles traffic by assigning a specified amount of queue space to each class of packets and then servicing the queues in a round-robin fashion; see Figure 4.

Figure 4 Custom Queuing

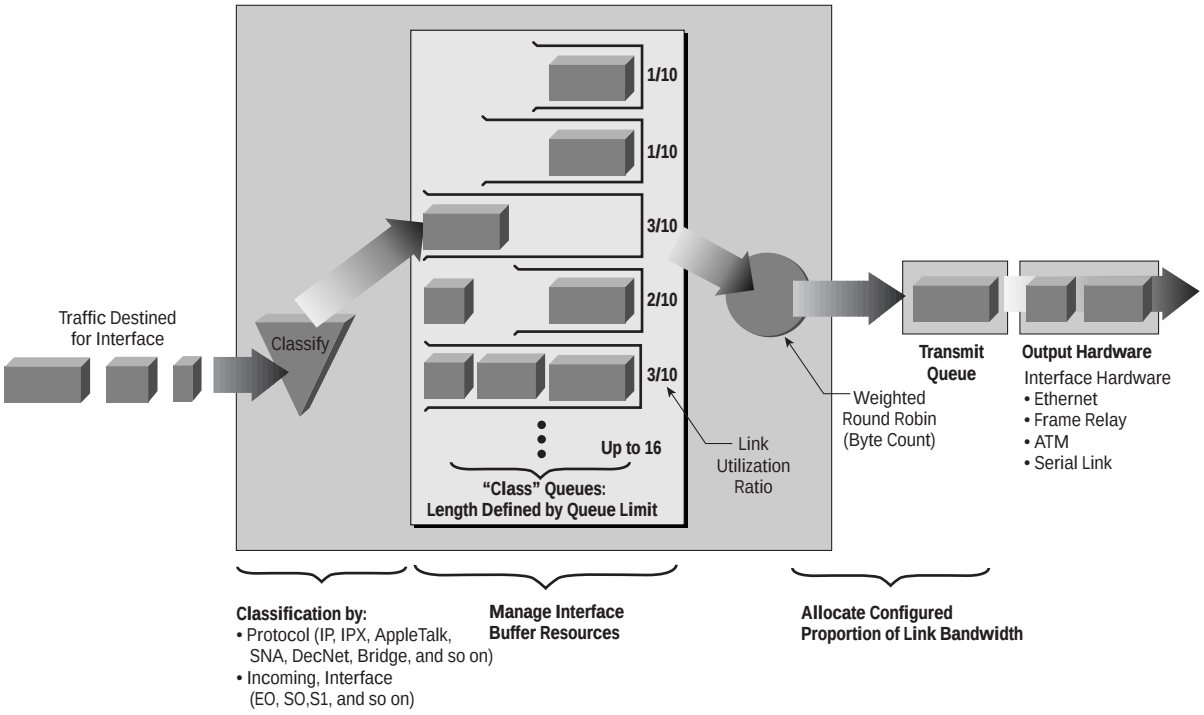

As an example, encapsulated SNA requires a guaranteed minimum level of service. You could reserve half of available bandwidth for SNA data, allowing the remaining half to be used by other protocols such as IP and IPX.

The queuing algorithm places the messages in one of 17 queues (queue 0 holds system messages such as keep-alives, signaling, and so on), and is emptied with weighted priority. The router services queues 1 through 16 in round-robin order, dequeuing a configured byte count from each queue in each cycle. This feature ensures that no application (or specified group of applications) achieves more than a predetermined proportion of overall capacity when the line is under stress. Like PQ, CQ is statically configured and does not automatically adapt to changing network conditions.

For situations in which it is desirable to provide consistent response time to heavy and light network users alike without adding excessive bandwidth, the solution is WFQ. WFQ is one of Cisco's premier queuing techniques. It is a flow-based queuing algorithm that does two things simultaneously: It schedules interactive traffic to the front of the queue to reduce response time, and it fairly shares the remaining bandwidth between high bandwidth flows.

WFQ ensures that queues do not starve for bandwidth, and that traffic gets predictable service. Low-volume traffic streams—which comprise the majority of traffic—receive preferential service, transmitting their entire offered loads in a timely fashion. High-volume traffic streams share the remaining capacity proportionally between them, as shown in Figure 5.

WFQ is designed to minimize configuration effort and adapts automatically to changing network traffic conditions. In fact, WFQ does such a good job for most applications that it has been made the default queuing mode on most serial interfaces configured to run at or below E1 speeds (2.048 Mbps).

WFQ is efficient in that it will use whatever bandwidth is available to forward traffic from lower priority flows if no traffic from higher priority flows is present. This is different from Time Division Multiplexing (TDM) which simply carves up the bandwidth and lets it go unused if no traffic is present for a particular traffic type. WFQ works with both of Cisco's primary QoS signaling techniques, IP Precedence and RSVP, described later in this white paper, to help provide differentiated QoS as well as guaranteed QoS services.

Figure 5 Weighted Fair Queuing

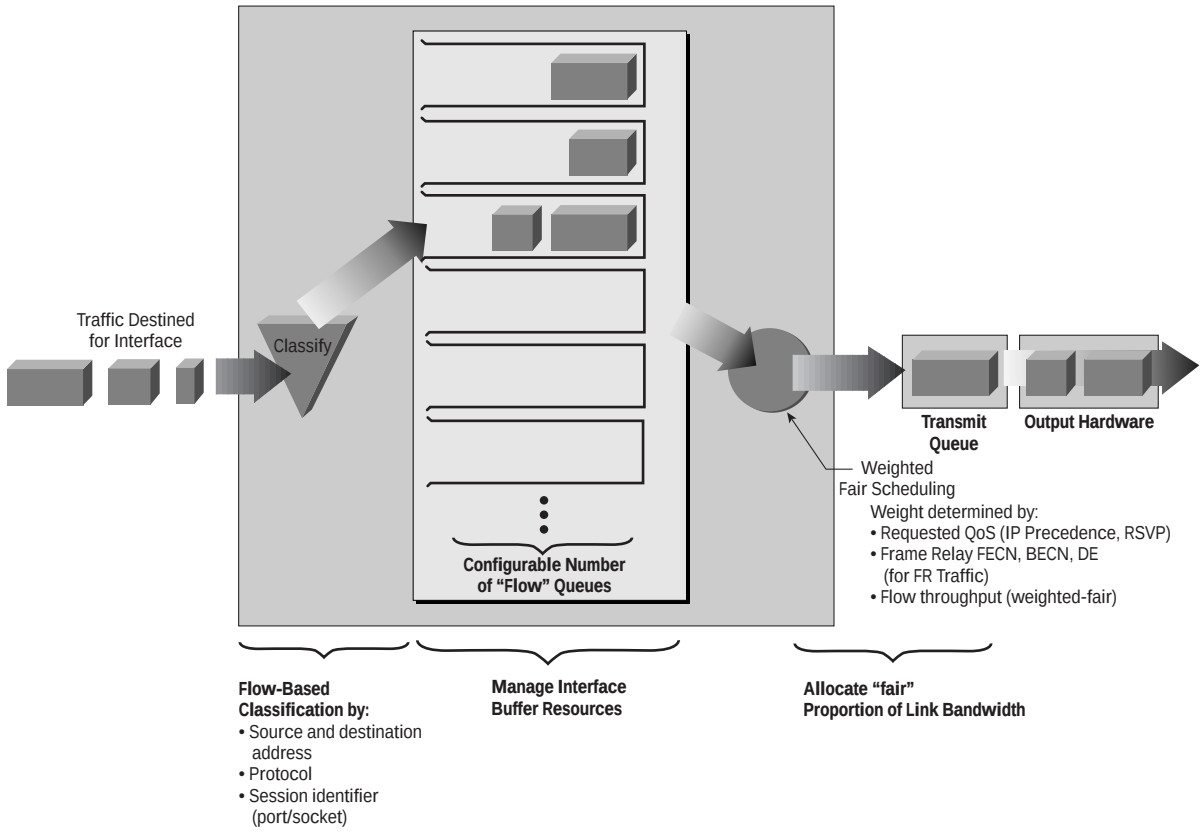

The WFQ algorithm also addresses the problem of round-trip delay variability. If multiple high-volume conversations are active, their transfer rates and interarrival periods are made much more predictable. WFQ greatly enhances algorithms such as the SNA Logical Link Control (LLC) and the Transmission Control Protocol (TCP) congestion control and slow-start features. The result is more predictable throughput and response time for each active flow, as shown in Figure 6.

Figure 6 Interactive Traffic Delay (128-Kbps Frame Relay WAN Link)

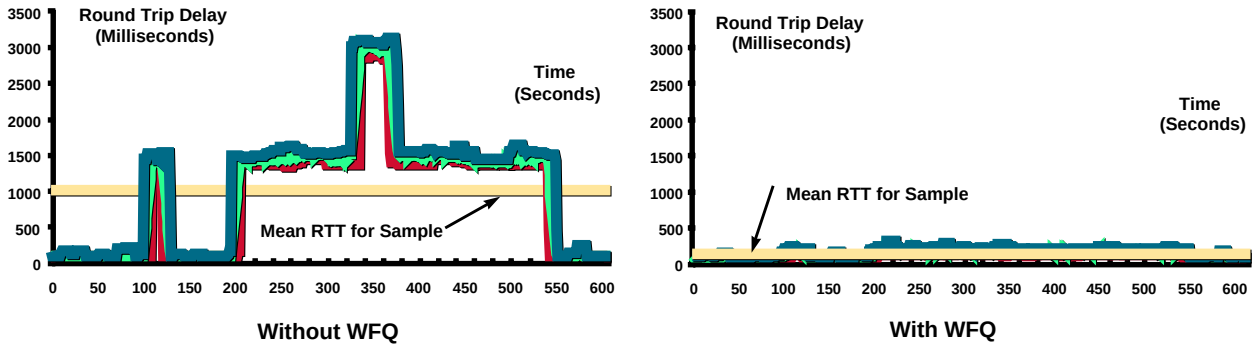

Cooperation between WFQ and QoS Signaling Technologies

WFQ is IP Precedence-aware, that is, it is able to detect higher priority packets marked with precedence by the IP Forwarder and can schedule them faster, providing superior response time for this traffic. The IP Precedence field has values between 0 (the default) and 7. As the precedence value increases, the algorithm allocates more bandwidth to that conversation to make sure that it gets served more quickly when congestion occurs. WFQ assigns a weight to each flow, which determines the transmit order for queued packets. In this scheme, lower weights are served first. IP Precedence serves as a divisor to this weighting factor. For instance, traffic with an IP Precedence field value of 7 gets a lower weight than traffic with an IP Precedence field value of 3, and thus has priority in the transmit order.

An example: If you have one flow at each precedence level on an interface, each flow will get precedence+1 parts of the link, as follows:

```
1+2+3+4+5+6+7+8 = 36
```

and the flows will get 8/36, 7/36, 6/36, and 5/36 of the link, and so on. However, if you have 18 precedence-1 flows and one of each of the others, the formula looks like this:

```
1+18*2+3+4+5+6+7+8 = 36-2+18*2 = 70
```

and the flows will get 8/70, 7/70, 6/70, 5/70, 4/70, 3/70, 2/70, and 1/70 of the link, and 18 of the flows will get approximately 2/70 of the link.

See “IP Precedence Signals Differentiated QoS” on page 20 for more on this.

WFQ is also RSVP aware; RSVP (see “RSVP Guarantees QoS” on page 21) uses WFQ to allocate buffer space and schedule packets, and guarantees bandwidth for reserved flows.

Additionally, in a Frame Relay network, the presence of congestion is flagged by the forward explicit congestion notification (FECN) and backward explicit congestion notification (BECN) bits. WFQ weights are affected by Frame Relay discard eligible (DE), FECN, and BECN bits when the traffic is switched by the Frame Relay switching module. Once congestion is flagged, the weights used by the algorithm are altered so that the conversation encountering the congestion transmits less frequently.

D-WFQ-A High-Speed Version for the 7500 Platform

Cisco IOS software also provides distributed weighted fair queuing (D-WFQ), a special high-speed version of WFQ designed initially for IP-only networks. D-WFQ is currently available only on VIP processors and only in Cisco IOS release 11.1cc, a special version for 7500 VIP processors. The 11.1cc functionality was initially distributed to a select set of ISP customers, but will be released with an upcoming version of Cisco IOS software for enterprise customers as well (see the "Cisco IOS QoS Capabilities Matrix" in the appendices for more details).

Congestion Avoidance Tools

Congestion avoidance techniques monitor network traffic loads in an effort to anticipate and avoid congestion at common network bottlenecks, as opposed to congestion management techniques that operate to control congestion once it occurs. The primary Cisco IOS congestion avoidance tool is Weighted Random Early Detection (WRED), which is described next.

WRED Avoids Congestion

The Random Early Detection (RED) class of algorithms are designed to avoid congestion in internetworks before it becomes a problem. RED works by monitoring traffic load at points in the network and stochastically discarding packets if the congestion begins to increase. The result of the drop is that the source detects the dropped traffic and slows its transmission. RED is primarily designed to work with TCP in IP internetwork environments.

WRED is Cisco's implementation. In early documentation it was simply called RED, but the name has been changed to WRED to better reflect its capabilities, as described below.

Figure 7 Weighted Random Early Detection

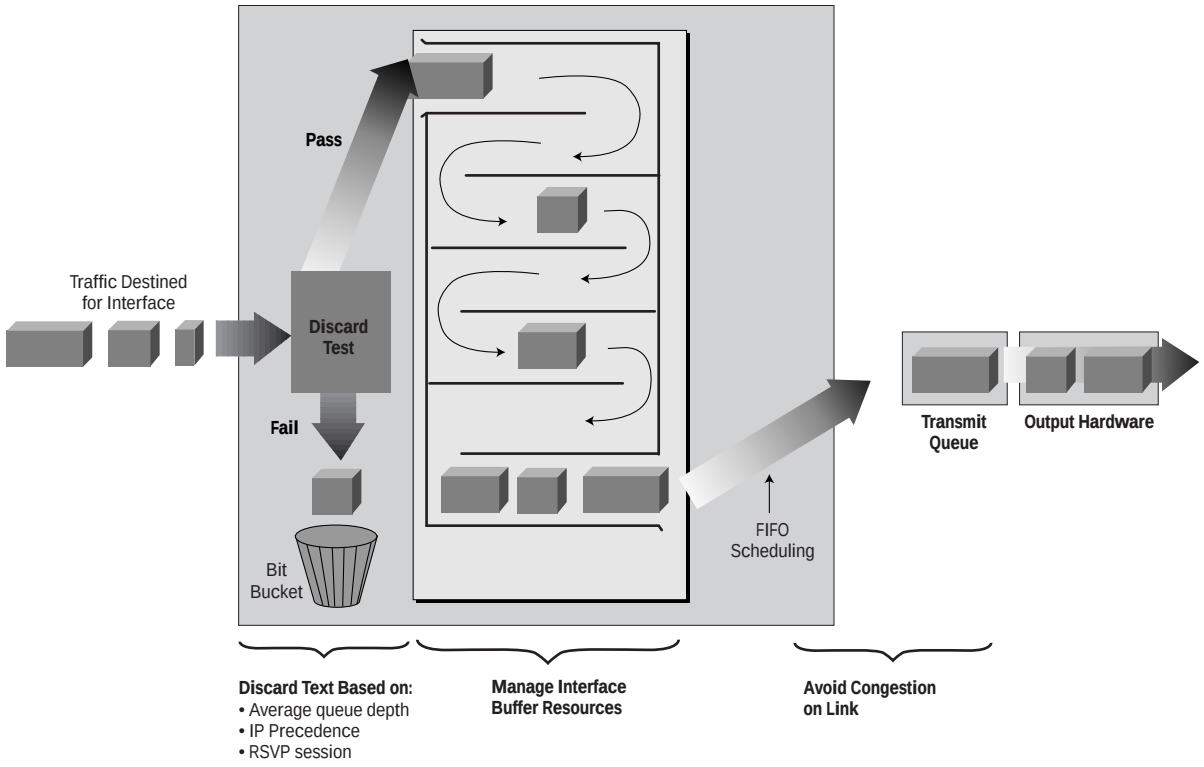

WRED Cooperation with QoS Signaling Technologies

WRED combines the capabilities of the RED algorithm with IP Precedence. This combination provides for preferential traffic handling for higher-priority packets. It can selectively discard lower-priority traffic when the interface starts to get congested and provide differentiated performance characteristics for different classes of service. See Figure 7.

WRED is also RSVP-aware, and can provide an integrated services controlled-load QoS service.

D-WRED Delivers High-Speed Differentiated Traffic on the 7500 Platform

Cisco IOS software also provides Distributed Weighted Random Early Detection (D-WRED), a high-speed version of WRED that runs on VIP-distributed processors. The D-WRED algorithm provides functionality beyond what WRED provides, such as minimum and maximum queue depth thresholds and drop capabilities for each class of service. D-WRED is currently available only in the Cisco IOS 11.1cc release, a special image for IP-only applications. The 11.1cc functionality was initially distributed to a select set of ISP customers, but will released with an upcoming version of Cisco IOS software for enterprise customers as well (see the "Cisco IOS QoS Capabilities Matrix" in the appendices for more details).

Traffic Shaping and Policing Tools

Cisco's QoS software solutions include two traffic shaping tools—Generic Traffic Shaping (GTS) and Frame Relay Traffic Shaping (FRTS)—to manage traffic and congestion on the network.

GTS Controls Outbound Traffic Flow

Generic Traffic Shaping (GTS) provides a mechanism to control the traffic flow on a particular interface. It reduces outbound traffic flow to avoid congestion by constraining specified traffic to a particular bit rate (also known as the token bucket approach), while queuing bursts of the specified traffic. Thus, traffic adhering to a particular profile can be shaped to meet downstream requirements, eliminating bottlenecks in topologies with data-rate mismatches. Figure 8 illustrates GTS.

Figure 8 Generic Traffic Shaping

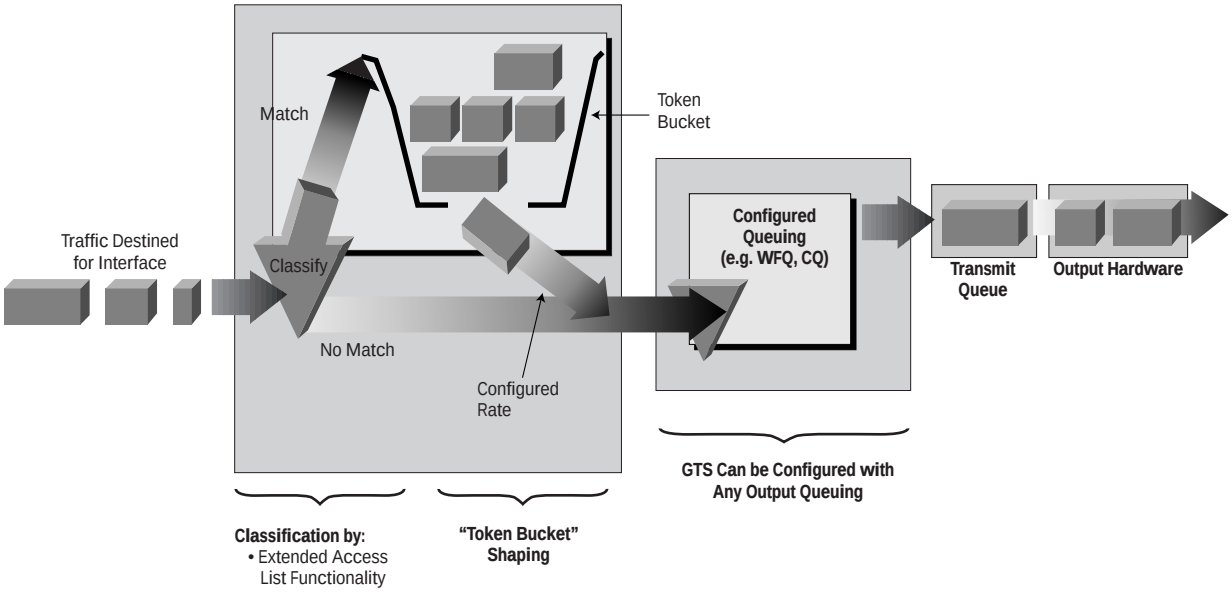

GTS applies on a per-interface basis, can use access lists to select the traffic to shape, and works with a variety of Layer 2 technologies, including Frame Relay, ATM, Switched Multimegabit Data Service (SMDS), and Ethernet.

On a Frame Relay subinterface, GTS can be set up to adapt dynamically to available bandwidth by integrating BECN signals, or set up simply to shape to a prespecified rate. GTS can also be configured on an ATM/AIP interface card to respond to RSVP signaled over statically configured ATM permanent virtual circuits (PVCs).

FRTS Manages Frame Relay Traffic

Frame Relay Traffic Shaping (FRTS) provides parameters that are useful for managing network traffic congestion. These include committed information rate (CIR), forward and backward explicit congestion notification (FECN/BECN), and the discard eligibility (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 with additional capabilities that improve the scalability and performance of a Frame Relay network increasing the density of virtual circuits and improving response time.

For example, you can configure rate enforcement—a peak rate configured to limit outbound traffic—to either the CIR or some other defined value, such as the excess information rate (EIR), on a per-virtual-circuit (VC) basis.

You can also define priority and custom queuing at the VC or subinterface level. This allows for finer granularity in the prioritization and queuing of traffic and provides more control over the traffic flow on an individual VC. If you combine CQ with the per-VC queuing and rate enforcement capabilities, you enable Frame Relay VCs to carry multiple traffic types such as IP, Systems Network Architecture (SNA), and Internetwork Packet Exchange (IPX), with bandwidth guaranteed for each traffic type.

FRTS can eliminate bottlenecks in Frame Relay networks with high-speed connections at the central site and low-speed connections at the branch sites. You can configure rate enforcement to limit the rate at which data is sent on the VC at the central site. You can also use rate enforcement with the existing DLCI Prioritization feature to further improve performance in this situation.

FRTS applies only to Frame Relay Permanent Virtual Connections (PVCs) and Switched Virtual Connections (SVCs).

Using information contained in BECN-tagged packets received from the network, FRTS can also dynamically throttle traffic. With BECN-based throttling, packets are held in the router's buffers to reduce the data flow from the router into the Frame Relay network. The throttling is done on a per-VC basis and the transmission rate is adjusted based on the number of BECN-tagged packets received.

FRTS also provides a mechanism for sharing media by multiple VCs. Rate enforcement allows the transmission speed used by the router to be controlled by criteria other than line speed, such as the CIR or EIR. The rate enforcement feature can also be used to preallocate bandwidth to each VC, creating a virtual time division multiplexing network.

And finally, with the Cisco's FRTS feature, you can integrate StrataCom ATM Foresight closed loop congestion control to actively adapt to downstream congestion conditions.

Link Efficiency Mechanisms

Currently, Cisco IOS software offers two link efficiency mechanisms—Real Time Protocol Header Compression (RTP-HC) and Link Fragmentation and Interleaving (LFI)—which work with queuing and traffic shaping to improve the efficiency and predictability of the application service levels.

LFI Fragments and Interleaves IP Traffic

Interactive traffic (Telnet, voice on IP, and the like) is susceptible to increased latency and jitter when the network processes large packets, (LAN-to-LAN FTP transfers traversing a WAN link, for example), especially as they are queued on slower links. The Cisco IOS Link Fragmentation and Interleaving (LFI) feature reduces delay and jitter on slower-speed links by breaking up large datagrams and interleaving low delay traffic packets with the resulting smaller packets; see Figure 9.

Figure 9 Link Fragmentation and Interleaving

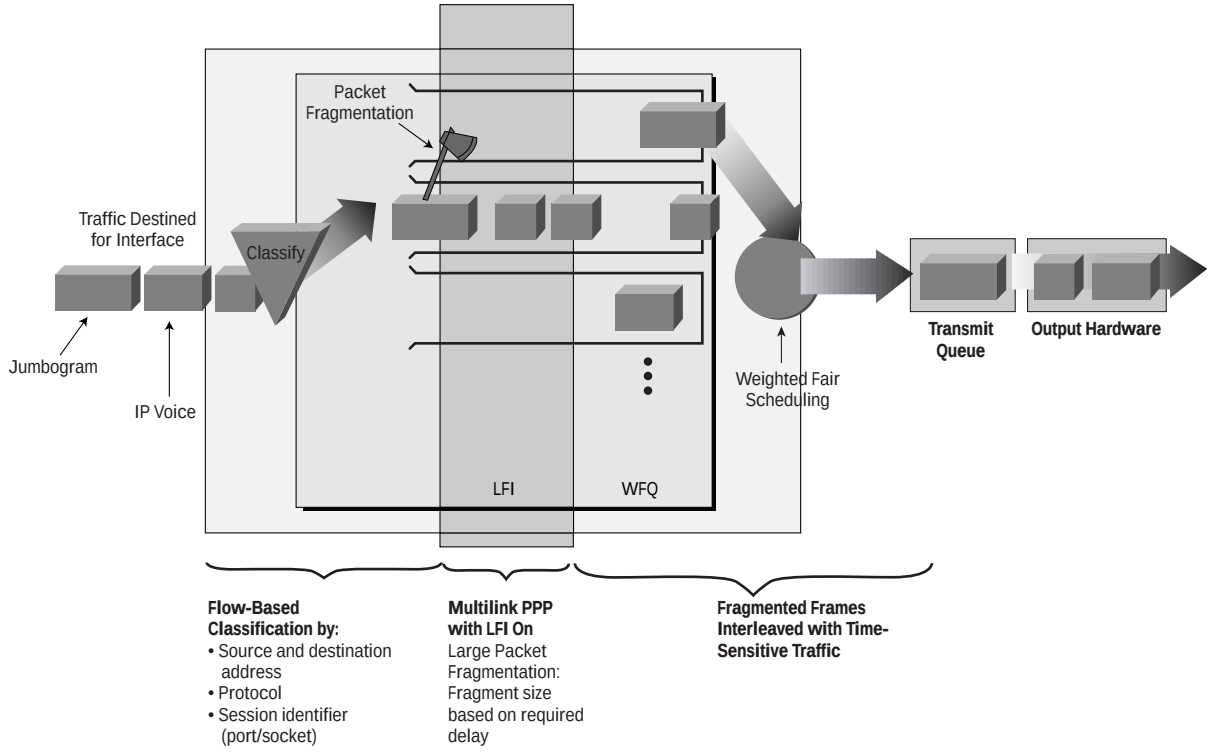

LFI was designed especially for lower-speed links where serialization delay is significant. LFI requires that multilink PPP be configured on the interface with interleaving turned on. A related IETF Draft, called Multiclass Extensions to Multilink PPP (MCML) implements almost the same function as LFI.

RTP Header Compression Increases Efficiency of Real-time Traffic

Real-time Transport 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 over multicast or unicast network services. RTP Header Compression increases efficiency for many of the newer voice-over-IP or multimedia applications that take advantage of RTP, especially on slow links. Figure 10 illustrates RTP header compression.

For compressed-payload audio applications, the RTP packet has a 40-byte header and typically a 20- to 150-byte payload. Given the size of the IP/UDP/RTP header combination, it is inefficient to transmit an uncompressed header. RTP header compression helps RTP run more efficiently—especially over lower-speed links—by compressing the RTP/UDP/IP header from 40 bytes to two to five bytes. This is especially beneficial for smaller packets (such as IP voice traffic) on slower links (385 kbps and below), where RTP header compression can reduce overhead and transmission delay significantly.

RTP header compression reduces line overhead for multimedia RTP traffic with a corresponding reduction in delay, especially for traffic that uses short packets relative to header length.

Figure 10 RTP Header Compression

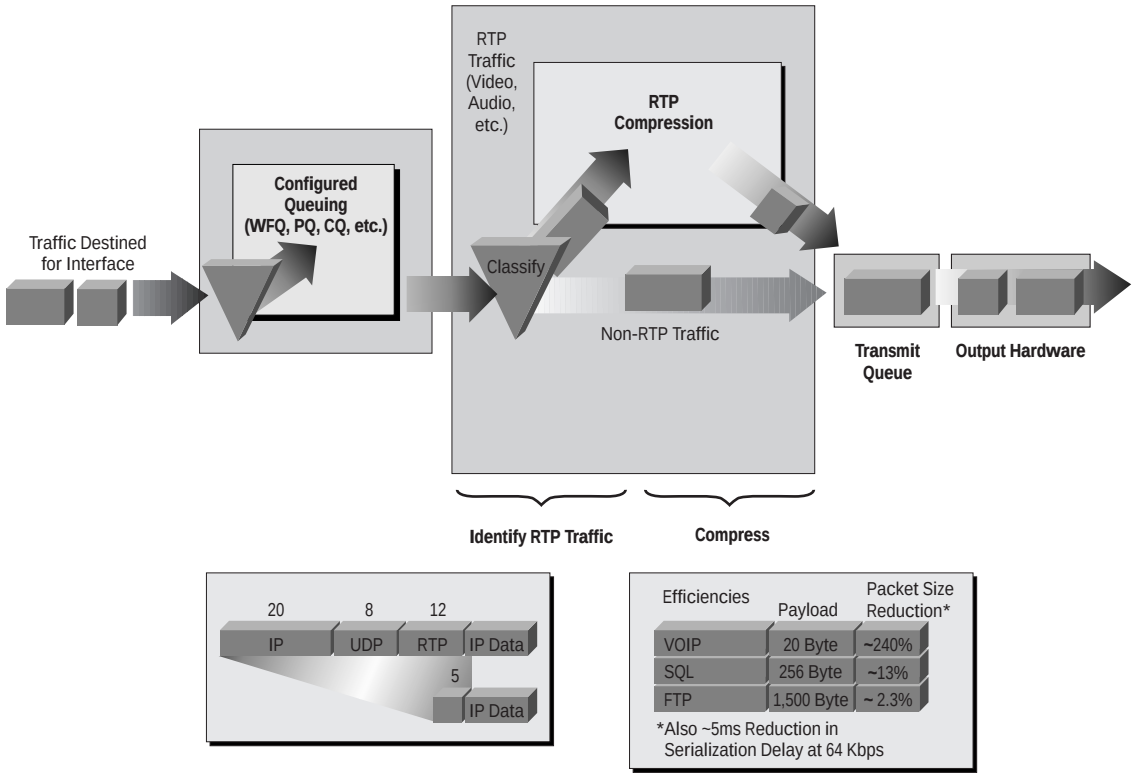

RTP header compression is supported on serial lines using Frame Relay, HDLC, or PPP encapsulation. It is also supported over ISDN interfaces. A related IETF Draft, called Compressed RTP (CRTP), defines essentially the same functionality.

QoS Signaling

Think of QoS signaling as a form of network communication. It provides a way for an end station or network element to signal certain requests to a neighbor. For example, an IP network can use part of the IP packet header to request special handling of priority or time-sensitive traffic. QoS signaling is useful for coordinating the traffic handling techniques described earlier in this paper and has a key role in configuring successful end-to-end QoS service across your network.

True end-to-end QoS requires that every element in the network path—switch, router, firewall, host, client, and so on—deliver its part of QoS, and it all must be coordinated with QoS signaling. However, the challenge is finding a robust QoS signaling solution that can operate end-to-end over heterogeneous network infrastructures. Although many viable QoS signaling solutions provide QoS at some places in the infrastructure, they often have limited scope across the network, as shown in Figure 11.

Figure 11 QoS Signaling Solutions

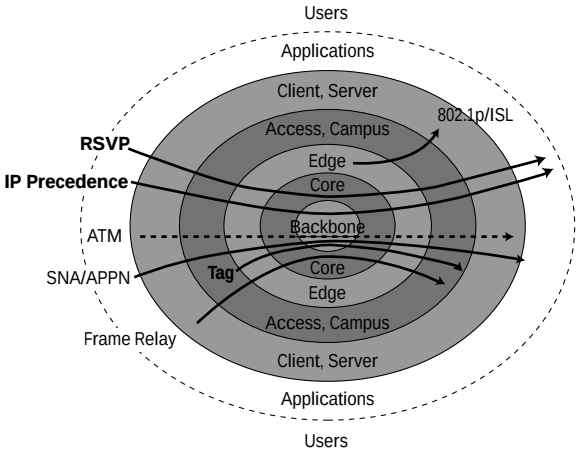

Cisco IOS software takes advantage of the end-to-end nature of IP to meet this challenge by overlaying Layer 2 technology-specific QoS signaling solutions with the Layer 3 IP QoS signaling methods of RSVP and IP Precedence.

This paper focuses on IP Precedence and RSVP, because both of these methods take advantage of the end-to-end nature of the IP protocol. As the majority of applications converge on the use of IP as the primary networking protocol, IP Precedence and RSVP provide a powerful combination for QoS signaling—IP Precedence signals for Differentiated QoS, and RSVP for Guaranteed QoS.

In addition to these mechanisms, Cisco leads the industry in QoS signaling integration, as shown in Figure 11. To achieve the end-to-end benefits of IP Precedence and RSVP signaling, Cisco IOS software offers ATM user to network interface (UNI) signaling and Frame Relay local management interface (LMI) to provide signaling into their ATM and Frame Relay backbone technologies. Cisco also provides similar priority signaling in its implementation of the IETF's multiprotocol label switching (MPLS), called Tag Switching.

With the Remote Switch Module (RSM) now running Cisco IOS software on the Catalyst® switch platforms, Cisco will soon support IEEE 802.1p for differentiated QoS, and Subnet Bandwidth Manager (SBM) for RSVP signaling of guaranteed QoS on switched internetworks. Using a feature called SNA type of service (ToS), Cisco IOS software also integrates native SNA class of service to provide the QoS required by mission-critical mainframe applications across IP-routed networks. All of these are standards-based mechanisms to integrate QoS functionality across heterogeneous networks; however, as previously mentioned, IP Precedence and RSVP are the two primary QoS signaling methods for the future.

IP Precedence Signals Differentiated QoS

IP Precedence utilizes the three precedence bits in the IPv4 header's ToS field to specify class of service for each packet, as shown in Figure 12. You can partition traffic in up to six classes of service using IP Precedence (two others are reserved for internal network use). The queuing technologies throughout the network can then use this signal to provide the appropriate expedited handling.

Figure 12 IP Precedence ToS Field

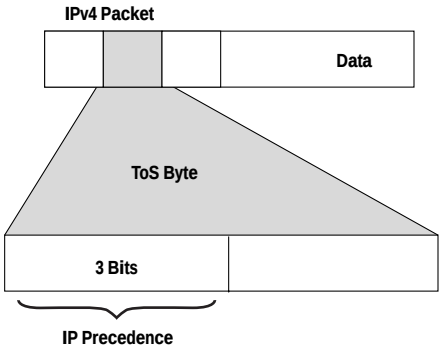

Features such as policy-based routing and CAR can be used to set precedence based on extended access-list classification. This allows considerable flexibility for precedence assignment, including assignment by application or user, or by destination and source subnet, and so on. Typically this functionality is deployed as close to the edge of the network (or administrative domain) as possible, so that each subsequent network element can provide service based on the determined policy.

IP Precedence can also be set in the host or network client, and this signaling can be used optionally; however, this can be overridden by policy within the network.

IP Precedence enables service classes to be established using existing network queuing mechanisms (for example, WFQ or WRED) with no changes to existing applications or complicated network requirements. Note that this same approach is easily extended to IPv6 using its priority field.

RSVP Guarantees QoS

RSVP is an IETF Internet Standard (RFC 2205) protocol for allowing an application to dynamically reserve network bandwidth. RSVP enables applications to request a specific QoS for a data flow, as shown in Figure 13. Cisco's implementation also allows RSVP to be initiated within the network using configured proxy RSVP. Using this capability, network managers can take advantage of the benefits of RSVP in the network, even for non-RSVP enabled applications and hosts.

Figure 13 Resource ReSerVation Protocol

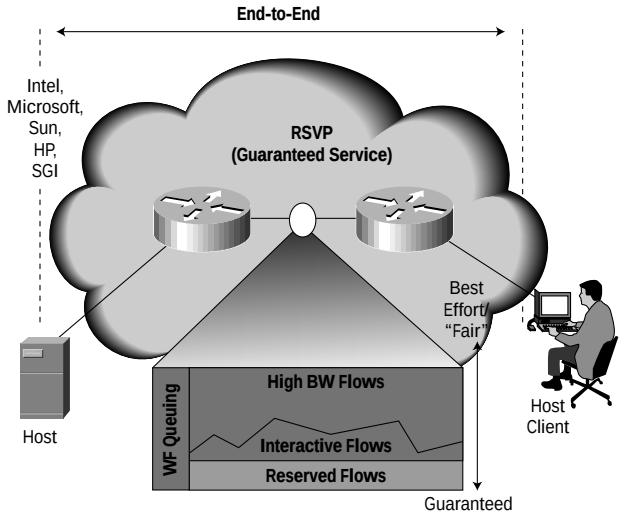

The press has written extensively about RSVP as a solution for guaranteeing bandwidth for new multimedia applications; however, RSVP's applicability is much broader than multimedia, as it is currently the only standard signaling protocol designed to guarantee network bandwidth from end to end for IP networks.

Hosts and routers use RSVP to deliver QoS requests to the routers along the paths of the data stream and to maintain router and host state to provide the requested service, usually bandwidth and latency. RSVP uses a mean data rate, the largest amount of data the router will keep in queue, and minimum QoS to determine bandwidth reservation.

WFQ or WRED acts as the workhorse for RSVP, setting up the packet classification and scheduling required for the reserved flows. Using WFQ, RSVP can deliver an Integrated Services Guaranteed Service. Using WRED, it can deliver a Controlled Load service. WFQ continues to provide its advantageous handling of non reserved traffic by expediting interactive traffic and fairly sharing the remaining bandwidth between high-bandwidth flows, and WRED provides its commensurate advantages for non-RSVP flow traffic. RSVP can be deployed in existing networks with a software upgrade.

Tag Switching Allows Flexible Traffic Engineering

Cisco's Tag Switching feature contains the mechanisms to interoperate with and take advantage of both RSVP and IP Precedence signaling. The Tag Switching header contains a three-bit field that can be used as a traffic prioritization signal. It can also be used to map particular flows and classes of traffic along engineered Tag Switching paths to obtain the required QoS through the Tag Switching portion of a network. The QoS capabilities provided by Tag Switching, along with its general operation, is covered in Tag Switching white papers and technical documentation.

Cisco's QoS Policy, Management, and Accounting Capabilities

Cisco IOS software provides technologies that enable policy control, management, and accounting of the QoS techniques described in this document. The following sections provide an overview of these technologies.

QoS Policy Control

The QoS policy control architecture is being developed as a key initial piece of the CiscoAssure policy networking initiative. This initiative leverages standards-based QoS policy control protocols and mechanisms to implement QoS policy from a single console interface. The CiscoAssure architecture is covered in detail in CiscoAssure specific documentation. The focus in the following sections is on packet-level services required in the infrastructure for QoS policy to be implemented.

At the infrastructure level, packet classification is a key capability for each policy technique that allows the appropriate packets traversing a network element or particular interface to be selected for QoS service. These packets can then be marked for the appropriate IP Precedence in some cases, or identified as an RSVP. Policy control also requires integration with underlying link layer network technologies, or non-IP protocols.

SNA ToS

SNA ToS in conjunction with Data Link Switching+ (DLSw+), allows mapping of traditional SNA Class-of-Service (CoS) into IP Differentiated service. This feature takes advantage of both QoS signaling and pieces of the architecture. DLSW+ opens four TCP sessions and maps each SNA ToS traffic into a different session. Each session is marked by IP Precedence. Cisco's congestion control technologies (custom queuing, priority queuing, and weighted fair queuing) acts on these sessions to provide a bandwidth guarantee or other improved handling across an intranet, as shown in Figure 14. This provides a migration path for traditional SNA customers onto an IP-based intranet, while preserving the performance characteristics expected of SNA.

Figure 14 SNA ToS

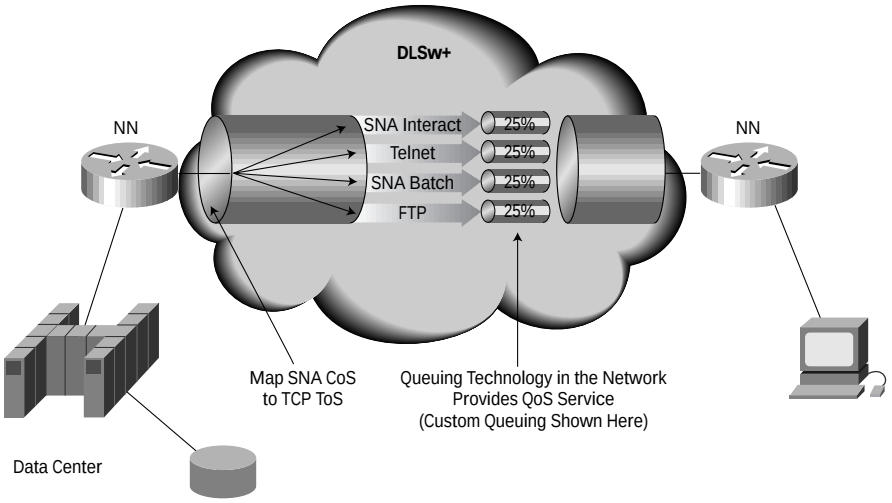

DLSW+ supports the following applications:

- LAN Network Manager (LNM)
- Native Service Point (NSP)
- Down Stream Physical Unit (DSPU)
- Advanced Peer-to-Peer Networking (APPN)
- Source-route-bridged FDDI LANs

Thus, traditional mainframe-based, mission-critical applications can take advantage of evolving IP intranets and extranets without sacrificing the QoS capabilities historically provided by SNA networking.

QoS Policy Setting with Policy-Based Routing (PBR)

Cisco IOS policy-based routing (PBR) allows you to classify traffic based on extended access list criteria, set IP Precedence bits, and even route to specific traffic-engineered paths that may be required to allow a specific QoS service through the network. By setting precedence levels on incoming traffic and using them in combination with the queuing tools described earlier in this paper, you can create differentiated service. These tools provide you with powerful, simple, and flexible options for implementing QoS policies in your network.

You can also set up PBR as a way to route packets based on configured policies. Some applications or traffic can benefit from QoS-specific routing-transferring stock records to a corporate office (for example-on a higher-bandwidth, higher-cost link for a short time), while transmitting routine application data such as e-mail over a lower-bandwidth, lower-cost link. PBR can be used to direct packets to take different paths than the path derived from the routing protocols. It provides a more flexible mechanism for routing packets, complementing the existing mechanisms provided by routing protocols.

CAR Manages Access Bandwidth Policy and Performs Policing

Similar in some ways to PBR, the Committed Access Rate (CAR) feature allows you to classify and police traffic on an incoming interface. It also allows specification of policies for handling traffic that exceeds a certain bandwidth allocation. CAR looks at traffic received on an interface, or a subset of that traffic selected by access list criteria, compares its rate to a configured token bucket, and then takes action based on the result (for example, drop or rewrite IP Precedence). CAR is currently available only in Cisco IOS 11.1cc, a special version designed initially for IP-only networks on the 7200, and in distributed mode on VIP processors on the 7500 (see “Cisco IOS QoS Capabilities Matrix” in the appendices for more details).

QoS Management

A variety of mechanisms described below, help control and manage QoS in the network. In addition, a number of the accounting tools described in the next section play a key role in proactively managing and designing QoS services in the network.

Netsys Network Connectivity and Performance Policies

Cisco Netsys Service-Level Management Suite 4.0 provides a policy-based service-level management solution that allows you to define, monitor, and assess network connectivity, security, and performance policies, and to troubleshoot problems quickly. The Cisco Netsys Service-Level Management Suite consists of three products: Cisco Netsys Connectivity Service Manager, Cisco Netsys Performance Service Manager, and Cisco Netsys LAN Service Manager. The Netsys Connectivity Tools, the first in a series of simulation-based planning and problem-solving products, assist network managers and analysts. They also assist with problem-solving, design, and planning activities focusing on network connectivity, route, and flow analysis.

The Netsys Performance Baseliner and Performance Solver developed by Netsys Technologies, Inc., are simulation-based network modeling tools that assist network managers, analysts, and designers with performance-related problem solving and planning functions.

Building upon the actual network configuration captured and modeled by the Netsys Connectivity Tools, the Performance Baseliner and Solver add network application traffic and performance analysis functions. With the Performance Tools, users can create a network baseline from configuration and performance data, then analyze the interactions between traffic flow, topology, routing parameters, and Cisco IOS features. They can also diagnose and solve operational problems, test scenarios, tune the network configuration for improved performance, and plan for incremental network changes.

QoS MIB Support

Cisco IOS provides QoS MIB support through a queuing MIB, and through support of standard MIBs for RSVP and the IETF Integrated Services MIB definitions.

QoS Accounting Tools

The Cisco IOS software offers NetFlow and Cisco Enterprise Accounting (CEA) to provide accounting capabilities.

Netflow Network Accounting and Statistics Application

NetFlow software identifies IP packet flows, performs efficient statistics collection, accelerates security filtering, and exports the collected statistics to downstream collectors, all while minimizing router performance impact. The Cisco NetFlow family also includes a set of management utilities, the FlowCollector and the FlowAnalyzer, and Cisco applications, Netsys and Cisco Enterprise Accounting, all designed to help your network operate more cost effectively through flexible network billing, planning, and monitoring.

Cisco is also working with a number of partners to deliver comprehensive solutions for NetFlow-based billing, planning and monitoring. NetFlow provides a key building block for delivering advanced, QoS-based services by providing comprehensive data collection and export.

Cisco Enterprise Accounting

Solutions that help identify, monitor, and control network bandwidth costs and proactively manage network usage are key to today's successful information systems (IS) organizations. Cisco Enterprise Accounting (CEA), a new family of network management software, delivers powerful, easy-to-use solutions that help manage and control your network costs. CEA is an indispensable tool for managers who want to gain a better understanding of their network costs and ensure that their networks perform at optimum levels.

A member of the Cisco network management family, CEA is designed to lower the overall cost of network ownership, detect equipment problems, and provide valuable information to IS staff. It allows an enterprise to make informed decisions about the costs of owning and operating a network.

Network QoS Application Examples

Cisco IOS provides QoS services for a wide range of applications, from mission-critical to new bandwidth-intensive multimedia. Cisco's goal for QoS service is simple to deliver the right end-to-end QoS solutions for every important application, including integrated voice, data, and video networking. Cisco has already delivered much of the technology required to achieve this vision.

As previously described, QoS service spans enterprise, Internet-service-provider, and small-business networks, and much of the technology to deliver QoS service is common between them. The examples below span all three segments.

QoS for Mission-Critical Applications: Oracle Sales Database

One of the primary business requirements for QoS service is to provide priority service for mission-critical traffic. The following example illustrates a possible application of Cisco IOS software QoS technology for this purpose.

In this example, a field sales force for a pharmaceutical company needs timely access to the sales database. This is an Oracle-based sales application primarily provided at sales office locations; see Figure 15. The sales application has been deployed for several years; however, the growth in traffic from Web applications on the corporate intranet has started to reduce application

performance on the WAN links because of increasing congestion. Corporate policy dictates that the sales application receive priority treatment, as this application directly impacts the corporate bottom line. It is also necessary to make sure that intranet users also get service from on the WAN, but they have a lower importance.

Figure 15 Sales Database Application

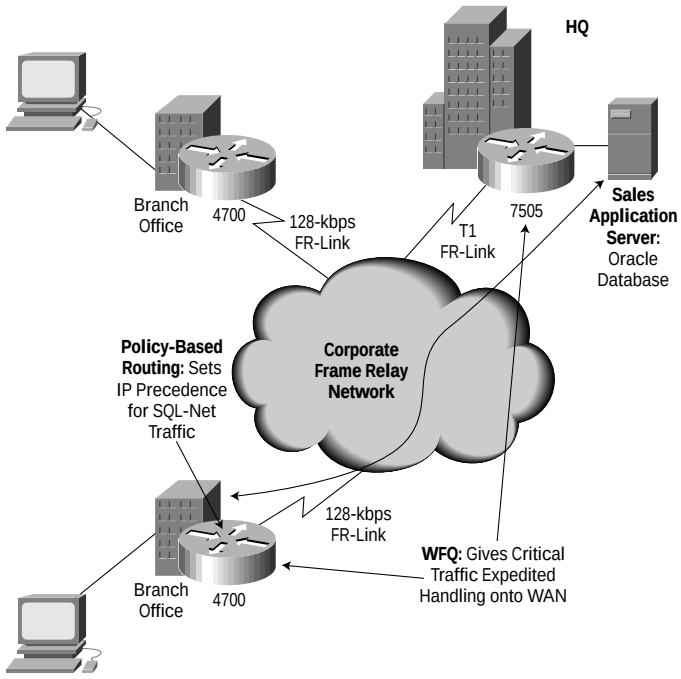

To implement this policy, the company uses WFQ with IP Precedence signaling to ensure that data to the sales application receives the network resource it needs, while ensuring that other WAN users also receive service, albeit slightly less timely if the sales application is being used heavily.

Figure 15 illustrates the branch router configured to recognize high-priority traffic destined for the sales application server over a Frame Relay network. At each branch, PBR is configured to recognize traffic destined for the sales application. At headquarters, the campus router has PBR configured to recognize the return traffic. All routers use WFQ to provide needed bandwidth for this traffic and allow it to reach its destination ahead of other traffic on the network.

QoS for Mission-Critical Applications: IBM SNA Order-Entry Application

Mainframe-based order-entry applications are widely used in many industries. One key requirement for many of these applications is timely, predicable user response time. Traditionally, end-to-end SNA services have been used to create such networks; however, the same level of priority service can now be provided for these traditional applications on IP-based networks.

Using SNA ToS with queuing features in Cisco IOS software, data center applications that have traditionally run on SNA can now be migrated to IP networks, while still preserving the reliability and service levels that are available through native SNA.

Figure 16 SNA ToS for an Intranet

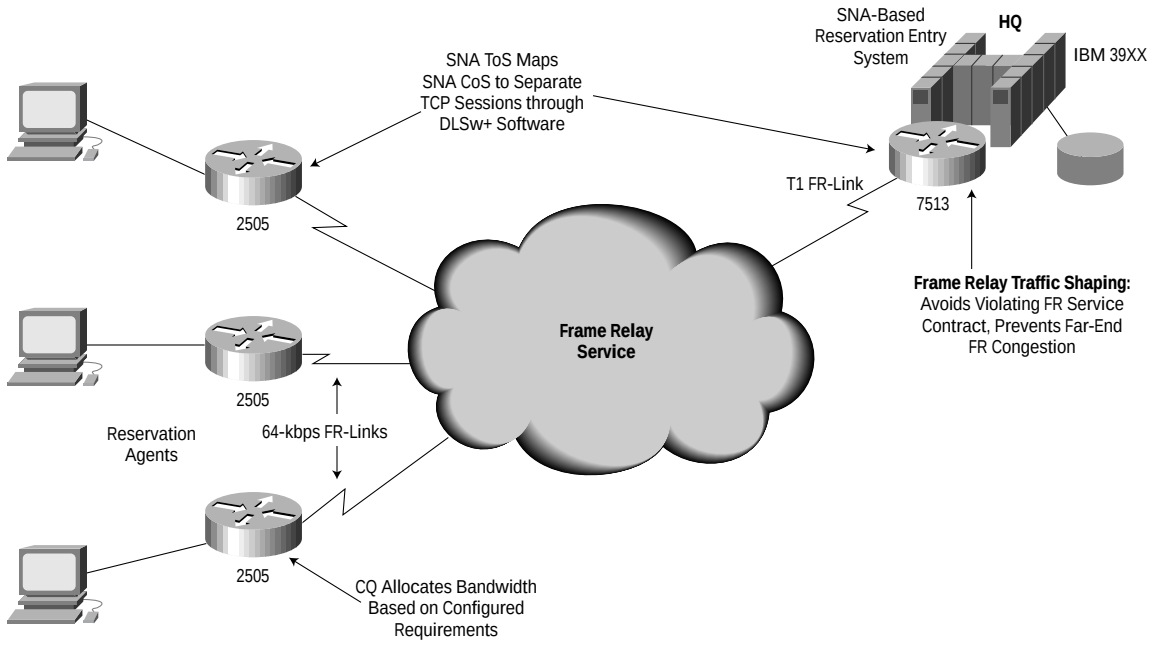

In Figure 16, CQ is used to ensure that appropriate network resources are allocated for the reservation system, which has transitioned from a leased-line network to a new IP-based network. The new network also allows other applications to share the same infrastructure without impacting the reservation system.

SNA ToS leveraging DLSw+ is configured on the Cisco routers, so different types of SNA traffic such as batch or interactive can be kept strictly separate through the IP network. CQ is configured to provided 60 percent of the link bandwidth for the important reservation traffic, leaving the remaining 40 percent for other uses. In addition, FRTS can be used at the mainframe site to constrain traffic on various DLCs to remote locations. In this way, far-end access links, (in this case 64-kbps links), are not overloaded by a burst from some other application transmitting on the T1 link from the mainframe site.

QoS for Packetized Voice

One of the most promising uses for IP networks is to allow sharing of voice traffic with the traditional data and LAN-to-LAN traffic. Typically, this can help reduce transmission costs by reducing the number of network connections, sharing existing connections and infrastructure, and so on.

Cisco has a wide range of voice networking products and technologies, including a number of Voice on IP (VoIP) solutions. To provide the required voice quality, however, QoS capability must be added to the traditional data-only network. Cisco IOS Software QoS features give VoIP traffic the service it needs, while providing the traditional data traffic with the service it needs as well.

Figure 17 QoS VoIP Solution

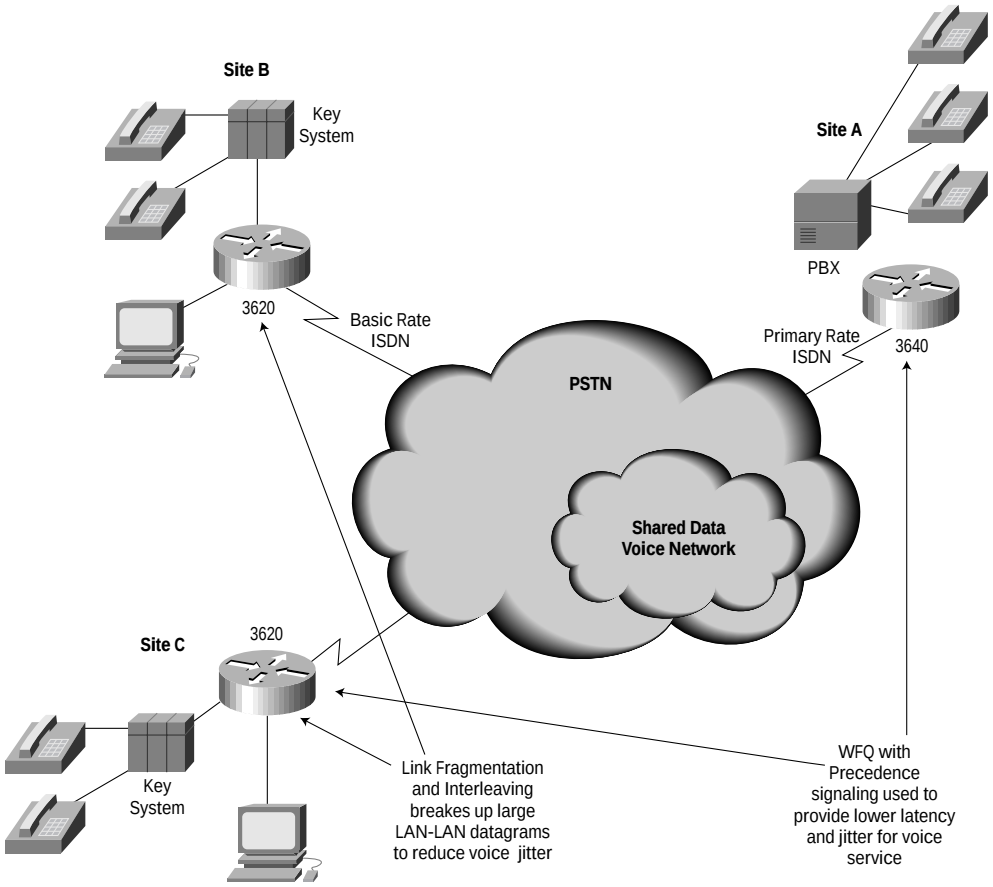

Figure 17 shows a business that has chosen to reduce some of its voice costs by combining voice traffic onto its existing IP network. Voice traffic at each office is digitized on voice modules on 3600 processors. This traffic is then routed via H.323 Gatekeeper, which also requests specific QoS for the voice traffic. In this case, IP Precedence is set to High for the voice traffic. WFQ is enabled on all the router interfaces for this network. WFQ automatically expedites the forwarding of high precedence voice traffic out each interface, reducing delay and jitter for this traffic.

Since the IP network was originally handling LAN-to-LAN traffic, many datagrams traversing the network are large 1500 byte packets. On slow links (below T1/E1 speeds), voice packets may be forced to wait behind one of these large packets, adding tens or even hundreds of milliseconds to the delay. LFI is used in conjunction with WFQ to break up these “jumbograms” and interleave the voice traffic to reduce this delay as well as jitter.

QoS for Streaming Video

One of the most significant challenges for IP-based networks, which have traditionally provided only best-effort service, has been to provide some type of service guarantees for different types of traffic. This has been a particular challenge for streaming video applications that often require a significant amount of reserved bandwidth to be useful.

In the network shown below, RSVP is used in conjunction with ATM PVCs to provide guaranteed bandwidth to a mesh of locations. RSVP is configured from within Cisco IOS to provide paths from the router networks, at the edges, and through the ATM core. Simulation traffic then uses these guaranteed paths to meet the constraints of geographically distributed real-time simulation. Video-enabled machines at the various sites also use this network to do live video conferencing.

Figure 18

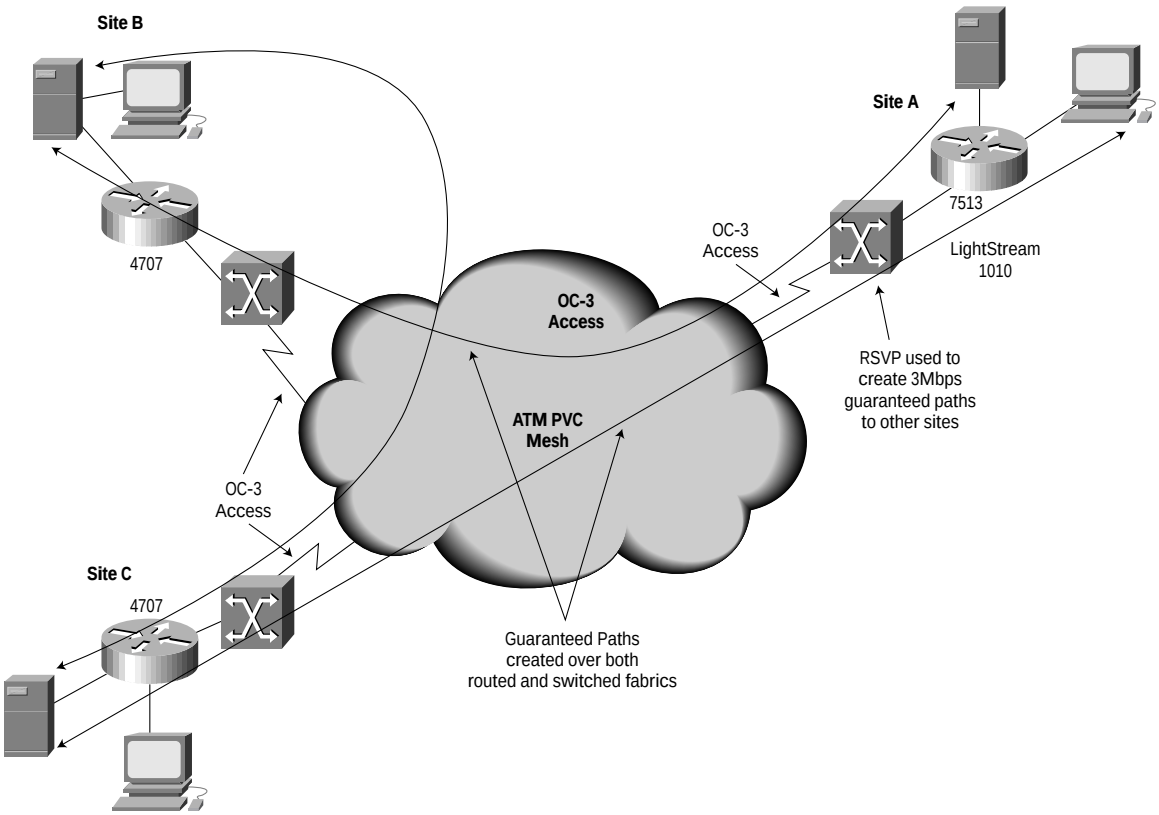

In this instance, OC-3 ATM links are configured with multiple 3 Mbps PVCs connecting to various remote sites. RSVP ensures that QoS from this PVC is extended to the appropriate application across the local routed network. In the future, Cisco IOS will extend this RSVP capability to dynamically set up ATM SVCs. This will reduce configuration complexity and add a great degree of automatic configuration.

QoS Looking Forward

In a continued evolution towards end-to-end QoS services, Cisco is expanding QoS interworking to operate more seamlessly across heterogeneous link layer technologies, and working very closely with our host platform partners to ensure interoperation between networks and end systems.

ATM QoS Interworking

ATM is a pervasive backbone technology, both in large-scale service provider networks and in campus LAN backbones. Thus, one major area for QoS interworking is at the IP-to-ATM boundary.

IP-to-ATM QoS Interworking for Differentiated Service

Cisco IOS software is being upgraded to ensure that IP Differentiated services, using IP Precedence signaling, are automatically preserved across ATM PVCs in backbone networks. These enhancements will take advantage of new capabilities in the next generation of ATM interface cards and will provide contiguous classes of service across both routed and ATM infrastructures.

RSVP-to-ATM Interworking for Guaranteed Service

Capabilities are also being added so that RSVP signaling can set up ATM SVCs of the appropriate parameters to dynamically reserve bandwidth across routed networks with ATM backbones.

Switched LAN QoS Interworking

Campus networks are increasingly migrating to switched LAN technology, and increasingly the underlying technology is based on Ethernet, whether traditional switched 10 Mpbs, Fast-Ethernet or Gig-Ethernet. Although these networks are often lightly loaded, QoS will be an increasingly important capability, especially on heavily loaded uplinks and other key congestion points. QoS will become even more important as new networked multimedia applications get deployed for normal business activities.

LAN QoS Interworking for Differentiated Service

Cisco IOS software will support IP Differentiated services by providing a mapping from IP Precedence signaling to the IEEE's 802.1p frame prioritization standard. This will allow differentiated services to be mapped seamlessly across the data-link layer technology on the campus and onto the WAN network to provide end-end QoS services.

LAN QoS Interworking for Guaranteed Service

Cisco IOS software will also support the IETF's Subnet Bandwidth Manager (SBM), currently an Internet Draft. SBM will extend RSVP capabilities to campus switched networks by enabling campus switches with native RSVP capabilities. This will allow RSVP reservations to be extended from the host, through the campus network, and over the WAN to provide true end-to-end reservations for applications that require reserved bandwidth.

Expanded Host/Client Support for RSVP

The major host and server providers in the industry have developed client-side RSVP support for their platforms: Microsoft for NT version 5.0, Sun for Solaris, and Hewlett Packard for HP-UX. Cisco has performed interoperability testing with these implementations, and Cisco networks support host-initiated RSVP requests from QoS-aware applications. In addition, Microsoft is shipping RSVP as part of the Winsock API, and future releases of applications such as Microsoft's NetShow and NetMeeting will be RSVP enabled.

QoS Policy Networking

Cisco has lead the IETF effort to standardize the first QoS policy protocol, currently in Internet Draft status and called Common Open Policy Service (COPS). Cisco IOS software will add COPS support as this work moves toward standardization. In addition, as part of the CiscoAssure initiative, QoS policy will be coordinated with security policy, name services, and so on, in coordinated network policy administration services.

Appendix 1: QoS Features Matrices

Table 1 Cisco IOS QoS Capabilities Matrix

Feature	Function	Classification by:	QoS Signaling	Protocol Support	Interfaces Supported	VC / SubInterface Support
FIFO Queuing[1]	Congestion Management	—	—	Multiprotocol	All	All
Priority Queuing (PQ)	Congestion Management	• Protocol • Source interface	—	Multiprotocol	Output – Ethernet, FR, ATM, SMDS, Serial	Per VC or Subinterface in 11.2
Custom Queuing (CQ)	Congestion Management	• Protocol • Source interface	—	Multiprotocol	Output – Ethernet, FR, ATM, SMDS, Serial	Per VC or Subinterface in 11.2
Weighted Fair Queuing (WFQ)[2] **and DWFQ**[3]	Congestion Management	• Flow	IP Precedence, RSVP	Multiprotocol	Output – Ethernet, FR, ATM, SMDS, Serial	Per VC or Subinterface in 11.2
	—	• Flow or Class	IP Precedence	IP	VIP w/ SONET – Output	Per Interface in 11.1CC
Weighted Random Early Detection (WRED) and DWRED[3]	Congestion Avoidance	• Class	IP Precedence, RSVP	Multiprotocol, but built for TCP/IP	Output – Ethernet, FR, ATM, Serial	—
	—	• Class	IP Precedence	IP	Output – VIP w/ SONET	Per Interface in 11.1CC
Generic Traffic Shaping (GTS)[4,5]	Congestion Management/Shaping	Extended Access List[6]	IP Precedence, RSVP	Multiprotocol	Output – Ethernet, FR, ATM, Serial	Per Subinterface, but not per VC or DLCI
Frame Relay Traffic Shaping (FRTS)[7]	Congestion Management/Shaping	Extended Access List[6]	—	Multiprotocol	Output – Frame Relay	Per DLCI, PQ and CQ only; Not on Subinterface
Link Fragmentation and Interleaving (LFI)	Link Efficiency	Packet/Frame Size	—	IP, Multilink PPP	Output – (requires WFQ)	—
RTP Header Compression (RTP-HC)[5]	Link Efficiency	Multimedia/RTP Traffic	—	Ignores Non-RTP Packets	FR, HDLC or PPP	—
SNA Type of Service[5]	Classification/Policy Setting	SNA CoS	IP Precedence	IP -DLSw+	Output–Ethernet, FR, ATM, Serial	—
Policy Based Routing (PBR)	Classification / Policy Setting	Extended Access List[6]	Sets IP Precedence or ToS	IP	Input/Output Any Interface, Output, Input	Per Subinterface.
Committed Access Rate (CAR)	Classification/ Congestion Management/Policy Setting	Extended Access List[6]	Sets IP Precedence	IP	Input/Output VIP w/ SONET–Output, Input	
BGP Policy Propagation	Policy	IP Precedence	Sets IP Prec. in Reverse Flow Direction	IP, Requires BGP Connectivity	All	—

1. FIFO is the default queuing mode when an interface does not support fancy queuing techniques (e.g. X.25, GRE Tunnel, LAPB, or PPP Compressed). FIFO can be configured using the no-fair-queue command.

2. For interfaces that support fancy queuing, WFQ is on by default on WAN interfaces at E1 or lower speeds, including PPP, HDLC, FR, SMDS, and so on, and is inactive on LANs and high-speed WAN lines.

3. Distributed WFQ and WRED are performance-enhanced versions that run on VIP processors on the 7500 platform.

4. GTS supports subinterfaces on FR links and software interfaces on LAN or PPP links. It does not support point-multipoint FR DLCIs, but does support multipoint links in general. GTS works with any output queuing.

5. Uses CQ, PQ, WFQ, WRED, or FIFO as output queuing method.

6. Extended Access Lists allow a more flexible way to characterize and classify traffic than was previously available for PQ and CQ (for example, by source and destination address, source and destination port, and so on).

7. FRTS supports multipoint FR subinterfaces and can work with FIFO, PQ, or CQ, but not WFQ.

Table 2 IOS QoS Availability Matrix.

Feature	Cisco IOS Version and Switching Mode[1]						Platform Support	Rule of Thumb Max Aggregate Thruput on 7200
	10.3	11.0	11.1	11.2	11.1CC	11.3	—	—
FIFO Queuing[2,3]	P,F,N,O	P, F, N, O	P, F, N, O	P, F, N, O	P, F, N, O, dCEF	P, F, N, O	All Cisco IOS Platforms	Interface Speed
Priority Queuing (PQ)	P	P, F, N, O	P, F, N, O	P, F, N, O	P, F, N, O	P, F, N, O	All Cisco IOS Platforms	10 Mbps
Custom Queuing (CQ)	P	P, F, N, O	P, F, N, O	P, F, N, O	P, F, N, O	P, F, N, O	All Cisco IOS Platforms	10 Mbps
Weighted Fair Queuing (WFQ)[4]	—	P, F, N, O	P, F, N, O	P, F, N, O	P, F, N, O	P, F, N, O	All Cisco IOS Platforms	10 Mbps
DWFQ[5]	—	—	—	—	dCEF	—	7500 VIP	T3-VIP2/40 OC3-VIP2/504
Weighted Random Early Detection (WRED)	—	—	—	P, F, N, O	—	P, F, N, O	All Cisco IOS Platforms	20 Mbps
DWRED[5]	—	—	—	—	dCEF	—	7500 VIP	T3-VIP2/40 OC3-VIP2/50
Generic Traffic Shaping (GTS)[6,7]	—	—	—	P, F	—	P, F	All Cisco IOS Platforms	10 Mbps
Frame Relay Traffic Shaping (FRTS)[8]	—	—	—	P, F 11.2(9)	—	P, F	All Cisco IOS Platforms	10 Mbps
Link Fragmentation and Interleaving (LFI)	—	—	—	—	—	P	All Cisco IOS Platforms	Intended for Sub- T1 speeds
RTP Header Compression (RTP-HC)[7]	—	—	—	P	—	P	All Cisco IOS Platforms	10 Mbps
SNA Type of Service[7]	—	—	—	—	—	P, F	All Cisco IOS Platforms	—
Policy Based Routing (PBR)	—	—	—	P, F11.2(9)	P	P, F	All Cisco IOS Platforms	10 Mbps
Committed Access Rate (CAR)	—	—	—	—	dCEF, F, CEF	—	7500 VIP or RSP, 7200	T3/E3 per VIP2-40 OC3 per VIP2-50
BGP Policy Propagation	—	—	—	—	dCEF, CEF	—	7500 VIP or RSP, 7200	N/A

1. Switching Mode: P=Process, F=Fast, N=NetFlow, O=Optimum, CEF=Cisco Express Forwarding, d=Distributed (VIP), dCEF=Distributed CEF.

2. FIFO is the default queuing mode when an interface does not support fancy queuing techniques (e.g. X.25, GRE Tunnel, LAPB, or PPP Compressed). FIFO can be configured using the no-fair-queue command.

3. Fancy queuing modes are not supported for X.25 in 11.2.

4. For interfaces that support fancy queuing, WFQ is on by default on WAN interfaces at E1 or lower speeds, including PPP, HDLC, FR, SMDS, and so on, and is inactive on LANs and high-speed WAN lines.

5. Distributed WFQ and WRED are performance-enhanced versions that run on VIP processors on the 7500 platform.

6. GTS supports subinterfaces on FR links and software interfaces on LAN or PPP links. It does not support point-multipoint FR DLCIs, but does support multipoint links in general. GTS works with any output queuing.

7. Uses CQ, PQ, WFQ, WRED, or FIFO as output queuing method.

8. FRTS supports multipoint FR subinterfaces and can work with FIFO, PQ, or CQ, but not WFQ.

Appendix 2: QoS Terminology

Quality of Service Definitions

The term QoS is an umbrella for all related technology and offerings in this area. Thus, all of the definitions that follow below are subsets of QoS. In the past, QoS has sometimes had more specific connotations for particular technologies, such as ATM, but the term is now used more broadly to refer to a network's ability to provide better service to selected network traffic for various technologies, including IP routed networks, Frame-Relay, ATM, Ethernet, and 802.1 networks, SONET, and so on.

End-to-End (or Edge-to-Edge) QoS Service Levels

Service levels refer to the actual QoS capabilities, meaning the ability of a network to deliver service needed by a specific network application from end-to-end. This can also include edge-to-edge, as in the case of a network that connects other networks rather than hosts or end systems, (the typical service provider network, for example), with some level of control over bandwidth, jitter, delay, and loss, provided by the network.

QoS Strictness

The "strictness" of the QoS service describes how tightly the service can be bound by specific bandwidth, delay, jitter, and loss characteristics. For example, the delay, loss, and jitter characteristics can be offered to within tight tolerances on a terrestrial TDM circuit, or for an ATM Variable Bit Rate Real-Time (VBR-rt), or Constant Bit Rate (CBR) service; whereas they are much harder to bound on a typical Internet IP connection. Essentially, QoS can provide three levels of strictness from end-to-end or edge-to-edge: best effort, differentiated, and guaranteed.

Best-Effort Service

Also known as lack of QoS, best-effort service is basic connectivity with no guarantees. The Internet today is a good example of best-effort service. Although best effort service is the lack of QoS, it provides us with a reference point on the nonstrict end of the spectrum (see Strictness of QoS Service Level Guarantee, above). Also, best effort is suitable for a wide range of networked applications such as general file transfers or e-mail.

Differentiated Service (also called "Soft" QoS)

Some traffic is treated better than the rest (faster handling, more bandwidth on average, lower loss rate on average). This is a statistical preference, not a hard and fast guarantee. With proper engineering, including edge policing, differentiated service can provide expedited handling appropriate for a wide class of applications, including lower delay for mission-critical applications, packet voice applications, and so on. Typically, differentiated service is associated with a course level of traffic classification (see Packet Classification), which means that traffic gets grouped or aggregated into a small number of classes, with each class receiving a particular QoS in the network. However, this does not have to be the case, as classification granularity is an independent issue.

The Differentiated Services (DiffServ) working group in the IETF is working on specific standards and definitions for services that fall under the broad umbrella of Differentiated or Soft QoS as defined above. This effort is largely focused on the use of the ToS field in IPv4 header as a QoS signaling mechanism, and it aims to provide definitions appropriate for aggregated flows for any level of aggregation. At least two services have initially been defined under this effort: The "Assured Service" and the "Preferred Service," each with slightly different definitions of service that from a technical perspective might be called "engineered best effort."

Technologies that can provide differentiated service (by our previous definition) for portions of the end-to-end connection include:

- IP-WRED, WFQ, combined with IP Precedence signaling or PQ on a single link
- ATM-Unspecified Bit Rate (UBR) and Available Bit Rate (ABR), especially if no Minimum Cell Rate (MCR) can be specified in the implementation
- Frame prioritization in campus switches in conjunction with 802.1p signaling

Guaranteed Service (also called “Hard” QoS or “reserved bandwidth”)

Guaranteed service is an absolute reservation of network resources, typically bandwidth, which implies reservation of buffer space along with the appropriate queuing disciplines, and so on, to ensure that specific traffic gets a specific service level. Bandwidth is typically used as a proxy for the other QoS attributes (delay, jitter, and so on), as the widest audience most easily understands it. Typically, guaranteed service is associated with a fine level of traffic classification (see Packet Classification below), often down to the level of individual traffic flows, which means that particular flows have network resources reserved for them so that required guarantees can be met. However, this does not have to be the case, as aggregated flows may receive guaranteed service.

The Integrated Services (IntServ) working group in the IETF has developed specific standards and definitions for services that fall under an umbrella definition of Guaranteed or Hard QoS provided above. This effort attempted to provide an architecturally sound way to specify flows in the Internet with varying requirements. The primary services in use today are Controlled Load Service and Guaranteed Service, each having precise definitions in the context of this work. RSVP was developed as a QoS signaling mechanism to provide these types of flow-based services.

Technologies that can provide guaranteed service for portions of the end-to-end connection include:

- IP-WFQ combined with RSVP signaling or CQ on a single link.
- Ethernet-SBM (when used with a compliant switch)
- ATM-VBR and CBR
- FR-CIR

Packet Classification

Packet Classification organizes a packet into a group useful for QoS or other handling such as security or firewalling, on the network. Classification may be done over a range of granularities, from groups of aggregated flows to individual flows or even subflows.

Typically, classification is done in a way similar to defining access lists, that is, based on some contents of the packet header. In this case, a packet may be classified by information in the L2, L3, or L4 headers (source/destination addresses, port numbers, subarea address, applications, user, as well as various Layer 2 attributes, such as Data Link Connection Identifier (DLCI) or ATM Virtual Path Connection (VPC), and so on. Classification can also be done based on information within the packet payload, such as is done with H.323 Proxy service, or CBAC Firewall functionality. Classifications can broad for aggregated flows such as “traffic destined for a subnetwork X,” or as narrow as a single flow or even subflow.

Policing and Shaping

Policing means limiting the amount of traffic flowing into or out of a particular interface to achieve a specific policy goal. Policing typically refers to actions taken by the network to monitor and control traffic to protect network resources such as bandwidth against unintended or malicious behavior.

Traffic shaping may be used to achieve policing goals as described herein, or to do congestion management as described below.

These activities are typically done at the edge of the network, where it is used to control the amount of traffic entering the network from a particular ingress point (an administratively separate network or a host system).

Congestion Management (including Scheduling and Queuing)

Typically, this means dealing with congestion at multiplexing points in the network, including how to order or schedule traffic, leaving a congested interface to provide appropriate QoS for a flow or set of aggregated flows. Typically, some type of queuing mechanism is used for congestion management. Traffic shaping can also be considered a congestion management mechanism, depending on its particular application in the network.

Congestion Avoidance (also called Drop/Flow Control)

Congestion avoidance is the action a network takes to avoid circumstances in which flows or aggregated flows no longer receive their associated service levels due to excessive traffic loads at points in the network. This action could be achieved through various means, including constructive application of drop policy to provide implicit feedback to host systems to reduce network traffic during congestion.

QoS Policy

QoS policy is a set of actions a network takes to configure and signal for a particular QoS service to be provided to a particular traffic classification.

QoS Signaling

QoS signaling is the means to deliver a QoS service requirement across the network. Either in-band signaling (for example, IP Precedence or 802.1p) or out-of-band signaling (RSVP) is used to indicate that a particular QoS service is desired for a particular traffic classification. IP Precedence and RSVP are the two most useful signaling mechanisms going forward because they both take advantage of the end-to-end nature of Layer 3 protocol and the growing ubiquity of IP as the network protocol of choice.

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe s.a.r.l.
Parc Evolic, Batiment L1/L2
16 Avenue du Quebec
Villebon, BP 706
91961 Courtaboeuf Cedex
France
http://www-europe.cisco.com
Tel: 33 1 6918 61 00
Fax: 33 1 6928 83 26

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Headquarters
Nihon Cisco Systems K.K.
Fuji Building, 9th Floor
3-2-3 Marunouchi
Chiyoda-ku, Tokyo 100
Japan
http://www.cisco.com
Tel: 81 3 5219 6250
Fax: 81 3 5219 6001

Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the Cisco Connection Online Web site at http://www.cisco.com.

Argentina · Australia · Austria · Belgium · Brazil · Canada · Chile · China (PRC) · Colombia · Costa Rica · Czech Republic · Denmark England · France · Germany · Greece · Hungary · India · Indonesia · Ireland · Israel · Italy · Japan · Korea · Luxembourg · Malaysia Mexico · The Netherlands · New Zealand · Norway · Peru · Philippines · Poland · Portugal · Russia · Saudi Arabia · Scotland · Singapore

5/98 SP