



Concepts

This chapter contains conceptual information that may be useful to Internet Service Providers or Network Administrators when configuring Cisco routers. To review some typical network scenarios, refer to [“Network Scenarios”](#) in Chapter 2. For information on specific configurations, refer to [Chapter 3, “Basic Router Configuration”](#) and [Chapter 4, “Advanced Router Configuration.”](#)

The following topics are included in this chapter:

- [Overview, page 1-1](#)
- [ADSL, page 1-3](#)
- [Network Protocols, page 1-3](#)
- [Routing Protocol Options, page 1-4](#)
- [PPP Authentication Protocols, page 1-5](#)
- [TACACS+, page 1-6](#)
- [Network Interfaces, page 1-6](#)
- [Dial Backup, page 1-8](#)
- [NAT, page 1-9](#)
- [Easy IP \(Phase 1\), page 1-9](#)
- [Easy IP \(Phase 2\), page 1-10](#)
- [VoIP, page 1-10](#)
- [QoS, page 1-11](#)
- [Access Lists, page 1-13](#)

Overview

The Cisco 826, 827, 828, 831, 836, and 837 and Cisco SOHO 76, 77, 78, 91, 96, and 97 routers are Cisco IOS-based members of the Cisco 800 router family with ATM/ADSL support. Depending on their feature set, the routers send data, voice, and video over high-speed ADSL lines to connect to the Internet or corporate intranets.

The data-only Cisco 826, 827, SOHO 76, and SOHO 77 have one 10Base-T Ethernet and one ADSL-over-ISDN or ADSL network port, respectively.

The data and voice Cisco 827-4V has four FXS/POTS ports in addition to the 10Base-T Ethernet and one ADSL network port, and it supports Voice over IP (VoIP). The four FXS/POTS ports will support loop-start functions for connecting to POTS devices up to 500 ft. The Cisco 827-4V router includes a digital signal processor (DSP) chip to support VoIP over ATM adaptation layer (AAL5) protocol.

AAL5 operates over the asymmetric digital subscriber line (ADSL) physical interface for both data and voice. The ADSL protocol supports EOC message sets defined in T1.413 DMT Issue 2 as limited by Digital Subscriber Line Access Multiplexers (DSLAMs). The ADSL controller and line interface unit are based on Alcatel chip sets.

The Cisco 828 router is a Cisco IOS-based member of the Cisco 800 router family with ATM/SHDSL support. The SOHO 78 router also supports ATM/SHDSL. The routers send data, voice, and video over high-speed G.SHDSL lines to connect to the Internet or corporate intranets.

Both the Cisco 828 router and the SOHO 78 router provide a 4-port Ethernet hub, in addition to the G.SHDSL port.

Both the Cisco 831 router and the SOHO 91 Ethernet-to-Ethernet routers can connect a corporate telecommuter or small office to an ISP over a broadband or Ethernet connection to corporate LANs or the Internet. The routers are capable of bridging and multiprotocol routing between LAN and WAN ports. The Cisco 831 router is a hardware encryption-capable router offering business-class features to small offices and enterprise telecommuters. The SOHO 91 router offers software encryption capability without hardware encryption.

The Cisco 836 and Cisco SOHO 96 routers are ADSL routers with an integrated switch. These routers provide a 4-port Ethernet switch for the LAN and an ADSL physical interface for the WAN compatibility. The Cisco 836 router is a hardware encryption-capable, Ethernet-to-ADSL router offering business-class features to small offices and enterprise telecommuters. The Cisco SOHO 96 router offers software encryption capability without hardware encryption. Both these routers provide an ISDN basic rate interface (BRI) S/T interface as a backup for the ADSL interface.

The Cisco 837 router and SOHO 97 are ADSL routers with an integrated switch. These routers provide a 4-port Ethernet switch for LAN and an ADSL physical interface for WAN compatibility. The Cisco 837 router is a hardware encryption-capable, Ethernet-to-ADSL router offering business-class features to small offices and enterprise telecommuters. The SOHO 97 router offers software encryption capability without hardware encryption.

The Cisco 831, Cisco 836, Cisco 837, SOHO 91, SOHO 96, and SOHO 97 routers support switch functions which enable the routers to be connected as a 10/100 BASE-T device. These routers crossover functionality enable them to detect MDI/MDIX to any other PC or hub with a straight-through cable or crossover cable.

Table 1-1 summarizes what interface each Cisco model supports.

Table 1-1 Interface Supported in Each Cisco Router

Interface Supported	Cisco Router Model
Ethernet to Ethernet	831, SOHO 91
Ethernet to ADSL over ISDN	826, SOHO 76, 836, SOHO 96
Ethernet to ADSL over POTS	827, 827H, 827-4V, 837, SOHO 77, SOHO 77H, SOHO 97

ADSL

ADSL is a technology that allows both data and voice to transmit over the same line. It is a packet-based network technology that allows high-speed transmission over twisted-pair copper wire on the local loop (“last mile”) between a network service provider (NSP) central office and the customer site, or on local loops created either within a building or campus.

The benefit of ADSL over a serial or dial-up line is that it is always on and always connected, increasing bandwidth and lowering the costs compared with a dial-up or leased line. ADSL technology is asymmetric in that it allows more bandwidth from an NSP’s central office to the customer site than from the customer site to the central office. This asymmetry, combined with always-on access (which eliminates call setup), makes ADSL ideal for Internet and intranet surfing, video-on-demand, and remote LAN access.

SHDSL

SHDSL is a technology based on the G.SHDSL (G.991.2) standard that allows both data and voice to be transmitted over the same line. SHDSL is a packet-based network technology that allows high-speed transmission over twisted-pair copper wire between a network service provider (NSP) central office and a customer site, or on local loops created within either a building or a campus.

G.SHDSL devices can extend reach from central offices and remote terminals to approximately 26,000 feet, at symmetrical data rates from 72 kbps up to 2.3 Mbps. In addition, it is repeatable at lower speeds, which means there is virtually no limit to its reach.

SHDSL technology is symmetric in that it allows equal bandwidth between an NSP’s central office and a customer site. This symmetry, combined with always-on access (which eliminates call setup), makes SHDSL ideal for LAN access.

Network Protocols

Network protocols enable the network to pass data from its source to a specific destination over LAN or WAN links. Routing address tables are included in the network protocols to provide the best path for moving the data through the network.

IP

The best known Transmission Control Protocol/Internet Protocol (TCP/IP) at the internetwork layer is IP, which provides the basic packet delivery service for all TCP/IP networks. In addition to the physical node addresses, the IP protocol implements a system of logical host addresses called IP addresses. The IP addresses are used by the internetwork and higher layers to identify devices and to perform internetwork routing. The Address Resolution Protocol (ARP) enables IP to identify the physical address that matches a given IP address.

IP is used by all protocols in the layers above and below it to deliver data, which means that all TCP/IP data flows through IP when it is sent and received regardless of its final destination.

IP is a connectionless protocol, which means that IP does not exchange control information (called a *handshake*) to establish an end-to-end connection before transmitting data. In contrast, a connection-oriented protocol exchanges control information with the remote computer to verify that it

is ready to receive data before sending it. When the handshaking is successful, the computers have established a connection. IP relies on protocols in other layers to establish the connection if connection-oriented services are required.

IPX exchanges routing information using Routing Information Protocol (RIP), a dynamic distance-vector routing protocol. RIP is described in more detail in the following subsections.

Routing Protocol Options

Routing protocols include the following:

- Routing Information Protocol (RIP)
- Enhanced Interior Gateway Routing Protocol (Enhanced IGRP)

RIP and Enhanced IGRP protocols differ in several ways, as shown in [Table 1-2](#).

Table 1-2 *RIP and Enhanced IGRP Comparison*

Protocol	Ideal Topology	Metric	Routing Updates
RIP	Suited for topologies with 15 or fewer hops.	Hop count. Maximum hop count is 15. Best route is one with lowest hop count.	By default, every 30 seconds. You can reconfigure this value and also use triggered extensions to RIP.
Enhanced IGRP	Suited for large topologies with 16 or more hops to reach a destination.	Distance information. Based on a successor, which is a neighboring router that has a least-cost path to a destination that is guaranteed to not be part of a routing loop.	Hello packets sent every 5 seconds plus incremental updates sent when the state of a destination changes.

RIP

RIP is an associated protocol for IP, and is widely used for routing Internet protocol traffic. RIP is a distance-vector routing protocol, which means that it uses distance (hop count) as its metric for route selection. *Hop count* is the number of routers that a packet must traverse to reach its destination. For example, if a particular route has a hop count of 2, then a packet must traverse two routers to reach its destination.

By default, RIP routing updates are broadcast every 30 seconds. You can reconfigure the interval at which the routing updates are broadcast. You can also configure triggered extensions to RIP so that routing updates are sent only when the routing database is updated. For more information on triggered extensions to RIP, refer to the Cisco IOS 12.0(1)T documentation set. For information on accessing the documentation, see the [“Obtaining Documentation”](#) in “About This Guide.”

Enhanced IGRP

Enhanced IGRP is an advanced Cisco proprietary distance-vector and link state routing protocol, which means it uses a metric more sophisticated than distance (hop count) for route selection. Enhanced IGRP uses a metric based on a successor, which is a neighboring router that has a least-cost path to a destination that is guaranteed not to be part of a routing loop. If a successor for a particular destination does not exist but neighbors advertise the destination, the router must recompute a route.

Each router running Enhanced IGRP sends hello packets every 5 seconds to inform neighboring routers that it is functioning. If a particular router does not send a hello packet within a prescribed period, Enhanced IGRP assumes that the state of a destination has changed and sends an incremental update.

Because Enhanced IGRP supports IP, you can use one routing protocol for multi-protocol network environments, minimizing the size of the routing tables and the amount of routing information.

PPP Authentication Protocols

The Point-to-Point Protocol (PPP) encapsulates network layer protocol information over point-to-point links.

PPP originally emerged as an encapsulation protocol for transporting IP traffic over point-to-point links. PPP also established a standard for the assignment and management of IP addresses, asynchronous (start/stop) and bit-oriented synchronous encapsulation, network protocol multiplexing, link configuration, link quality testing, error detection, and option negotiation for such capabilities as network-layer address negotiation and data-compression negotiation. PPP supports these functions by providing an extensible Link Control Protocol (LCP) and a family of Network Control Protocols (NCPs) to negotiate optional configuration parameters and facilities.

The current implementation of PPP supports two security authentication protocols to authenticate a PPP session:

- Password Authentication Protocol (PAP)
- Challenge Handshake Authentication Protocol (CHAP)

PPP with PAP or CHAP authentication is often used to inform the central site which remote routers are connected to it.

PAP

PAP uses a two-way handshake to verify the passwords between routers. To illustrate how PAP works, imagine a network topology where a remote office Cisco 827 router is connected to a corporate office Cisco 3600 router. After the PPP link is established, the remote office router repeatedly sends a configured username and password until the corporate office router accepts the authentication.

PAP has the following characteristics:

- The password portion of the authentication is sent across the link in clear text (not scrambled or encrypted).
- PAP provides no protection from playback or repeated trial-and-error attacks.
- The remote office router controls the frequency and timing of the authentication attempts.

CHAP

CHAP uses a three-way handshake to verify passwords. To illustrate how CHAP works, imagine a network topology where a remote office Cisco 827 router is connected to a corporate office Cisco 3600 router.

After the PPP link is established, the corporate office router sends a challenge message to the remote office router. The remote office router responds with a variable value. The corporate office router checks the response against its own calculation of the value. If the values match, the corporate office router accepts the authentication. The authentication process can be repeated any time after the link is established.

CHAP has the following characteristics:

- The authentication process uses a variable challenge value rather than a password.
- CHAP protects against playback attack through the use of the variable challenge value, which is unique and unpredictable. Repeated challenges limit the time of exposure to any single attack.
- The corporate office router controls the frequency and timing of the authentication attempts.

**Note**

Cisco recommends using CHAP because it is the more secure of the two protocols.

TACACS+

Cisco 800-series routers support the Terminal Access Controller Access Control System Plus (TACACS+) protocol through Telnet. TACACS+ is a Cisco proprietary authentication protocol that provides remote access authentication and related network security services, such as event logging. User passwords are administered in a central database rather than in individual routers. TACACS+ also provides support for separate modular authentication, authorization, and accounting (AAA) facilities that are configured at individual routers.

Network Interfaces

This section describes the network interface protocols that Cisco 800-series routers support. The following network interface protocols are supported:

- Ethernet
- ATM

Ethernet

Ethernet is a baseband LAN protocol that transports data and voice packets to the WAN interface using carrier sense multiple access collision detect (CSMA/CD). The term is now often used to refer to all CSMA/CD LANs. Ethernet was designed to serve in networks with sporadic, occasionally heavy traffic requirements, and the IEEE 802.3 specification was developed in 1980 based on the original Ethernet technology.

Under the Ethernet CSMA/CD media-access process, any host on a CSMA/CD LAN can access the network at any time. Before sending data, CSMA/CD hosts listen for traffic on the network. A host wanting to send data waits until it detects no traffic before it transmits. Ethernet allows any host on the network to transmit whenever the network is quiet. A collision occurs when two hosts listen for traffic, hear none, and then transmit simultaneously. In this situation, both transmissions are damaged, and the hosts must retransmit at some later time. Algorithms determine when the colliding hosts should retransmit.

ATM

Asynchronous Transfer Mode (ATM) is a high-speed, multiplexing and switching protocol that supports multiple traffic types including voice, data, video, and imaging.

ATM is composed of fixed-length cells that switch and multiplex all information for the network. An ATM connection is simply used to transfer bits of information to a destination router or host. The ATM network is considered a LAN with high bandwidth availability. Unlike a LAN, which is connectionless, ATM requires certain features to provide a LAN environment to the users.

Each ATM node must establish a separate connection to every node in the ATM network that it needs to communicate with. All such connections are established through a permanent virtual circuit (PVC).

PVC

A PVC is a connection between remote hosts and routers. A PVC is established for each ATM end node with which the router communicates. The characteristics of the PVC that are established when it is created are set by the ATM adaptation layer (AAL) and the encapsulation type. An AAL defines the conversion of user information into cells. An AAL segments upper-layer information into cells at the transmitter and reassembles the cells at the receiver.

Cisco routers support the AAL5 format, which provides a streamlined data transport service that functions with less overhead and affords better error detection and correction capabilities than AAL3/4. AAL5 is typically associated with variable bit rate (VBR) traffic and unspecified bit rate traffic (UBR). Cisco 800-series routers also support AAL1 and 2 formats.

ATM encapsulation is the wrapping of data in a particular protocol header. The type of router you are connecting to the router determines the type of ATM PVC encapsulation types.

The routers support the following encapsulation types for ATM PVCs:

- LLC/SNAP (RFC 1483)
- VC-MUX (RFC 1483)
- PPP (RFC 2364)

Each PVC is considered a complete and separate link to a destination node. Users can encapsulate data as needed across the connection. The ATM network disregards the contents of the data. The only requirement is that data be sent to the router's ATM subsystem in a manner that follows the specific AAL format.

Dialer Interface

A dialer interface assigns PPP features (such as authentication and IP address assignment method) to a PVC. Dialer interfaces are used when configuring PPP over ATM.

Dialer interfaces can be configured independently of any physical interface and applied dynamically as needed.

Dial Backup

Dial backup provides protection against WAN downtime by allowing user to configure a backup modem line connection. The following can be used to bring up the dial backup feature in the Cisco IOS software:

- [Backup Interface](#)
- [Floating Static Routes](#)
- [Dialer Watch](#)

Backup Interface

A backup interface is an interface that stays idle until certain circumstances occur, such as WAN downtime, at which point it is activated. The backup interface can be a physical interface such as Basic Rate Interface (BRI), or an assigned backup dialer interface to be used in a dialer pool. While the primary line is up, the backup interface is placed in standby mode. In standby mode, the backup interface is effectively shut down until it is enabled. Any route associated with the backup interface does not appear in the routing table.

Because the backup interface command is dependent on the router's identifying that an interface is physically down, it is commonly used to back up ISDN BRI connections and async lines and leased lines. The interfaces to such connections go down when the primary line fails, and the backup interface quickly identifies such failures.

Floating Static Routes

Floating static routes are static routes that have an administrative distance greater than the administrative distance of dynamic routes. Administrative distances can be configured on a static route so that the static route is less desirable than a dynamic route. In this manner, the static route is not used when the dynamic route is available. However, if the dynamic route is lost, the static route can take over, and the traffic can be sent through this alternate route. If this alternate route uses a Dial-on-Demand Routing (DDR) interface, then that interface can be used as a backup feature.

Dialer Watch

Dialer watch is a backup feature that integrates dial backup with routing capabilities. Dialer watch provides reliable connectivity without having to define traffic of interest to trigger outgoing calls at the central router. Hence, dialer watch can be considered regular DDR with no requirement for traffic of interest. By configuring a set of watched routes that define the primary interface, you are able to monitor and track the status of the primary interface as watched routes are added and deleted.

When a watched route is deleted, dialer watch checks for at least one valid route for any of the IP addresses or networks being watched. If there is no valid route, the primary line is considered down and unusable. If there is a valid route for at least one of the watched IP networks defined and the route is pointing to an interface other than the backup interface configured for dialer watch, the primary link is considered up and dialer watch does not initiate the backup link.

NAT

Network address translation (NAT) provides a mechanism for a privately addressed network to access registered networks, such as the Internet, without requiring a registered subnet address. This mechanism eliminates the need for host renumbering and allows the same IP address range to be used in multiple intranets.

NAT is configured on the router at the border of an *inside network* (a network that uses nonregistered IP addresses) and an *outside network* (a network that uses a globally unique IP address; in this case, the Internet). NAT translates the inside local addresses (the nonregistered IP addresses assigned to hosts on the inside network) into globally unique IP addresses before sending packets to the outside network.

With NAT, the inside network continues to use its existing private or obsolete addresses. These addresses are converted into legal addresses before packets are forwarded onto the outside network. The translation function is compatible with standard routing; the feature is required only on the router connecting the inside network to the outside domain.

Translations can be static or dynamic. A static address translation establishes a one-to-one mapping between the inside network and the outside domain. Dynamic address translations are defined by describing the local addresses to be translated and the pool of addresses from which to allocate outside addresses. Allocation occurs in numeric order and multiple pools of contiguous address blocks can be defined.

NAT eliminates the need to readdress all hosts that require external access, saving time and money. It also conserves addresses through application port-level multiplexing. With NAT, internal hosts can share a single registered IP address for all external communications. In this type of configuration, relatively few external addresses are required to support many internal hosts, thus conserving IP addresses.

Because the addressing scheme on the inside network may conflict with registered addresses already assigned within the Internet, NAT can support a separate address pool for overlapping networks and translate as appropriate.

Easy IP (Phase 1)

The Easy IP (Phase 1) feature combines Network Address Translation (NAT) and PPP/Internet Protocol Control Protocol (IPCP). This feature enables a Cisco router to automatically negotiate its own registered WAN interface IP address from a central server and to enable all remote hosts to access the Internet using this single registered IP address. Because Easy IP (Phase 1) uses existing port-level multiplexed NAT functionality within the Cisco IOS software, IP addresses on the remote LAN are invisible to the Internet.

The Easy IP (Phase 1) feature combines NAT and PPP/IPCP. With NAT, the router translates the nonregistered IP addresses used by the LAN devices into the globally unique IP address used by the dialer interface. The ability of multiple LAN devices to use the same globally unique IP address is known as *overloading*. NAT is configured on the router at the border of an inside network (a network that uses nonregistered IP addresses) and an outside network (a network that uses a globally unique IP address; in this case, the Internet).

With PPP/IPCP, the Cisco routers automatically negotiate a globally unique (registered) IP address for the dialer interface from the ISP router.

Easy IP (Phase 2)

The Easy IP (Phase 2) feature combines Dynamic Host Configuration Protocol (DHCP) server and relay. DHCP is a client-server protocol that enables devices on an IP network (the DHCP clients) to request configuration information from a DHCP server. DHCP allocates network addresses from a central pool on an as-needed basis. DHCP is useful for assigning IP addresses to hosts connected to the network temporarily or for sharing a limited pool of IP addresses among a group of hosts that do not need permanent IP addresses.

DHCP frees you from having to assign an IP address to each client manually.

DHCP configures the router to forward UDP broadcasts, including IP address requests, from DHCP clients. DHCP allows for increased automation and fewer network administration problems by:

- Eliminating the need for the manual configuration of individual computers, printers, and shared file systems
- Preventing the simultaneous use of the same IP address by two clients
- Allowing configuration from a central site

**Note**

When using NAT, DHCP relay cannot be used on the Cisco 800-series routers. The built-in DHCP server should be used instead.

VoIP

The Cisco 827-4V router is a voice-and-data-capable router that provides Voice-over-IP (VoIP) functionality and can carry voice traffic (such as telephone calls and faxes) over an IP network.

Cisco voice support is implemented using voice packet technology. There are two primary applications for VoIP:

- It provides a central-site telephony termination facility for VoIP traffic from multiple voice-equipped remote office facilities.
- It provides a PSTN gateway for Internet telephone traffic. VoIP used as a PSTN gateway leverages the standardized use of H.323-based Internet telephone client applications.

In VoIP, the digital signal processor (DSP) segments the voice signal into frames and stores them in voice packets. These voice packets are transported by using IP in compliance with H.323 signaling standards.

H.323

H.323 is an International Telecommunication Union (ITU-T) standard that describes packet-based video, audio, and data conferencing. H.323 is an umbrella standard that describes the architecture of the conferencing system and refers to a set of other standards (H.245, H.225.0, and Q.931) to describe its actual protocol.

Cisco H.323 Version 2 support upgrades Cisco IOS software to comply with the mandatory requirements and several of the optional features of the version 2 specification. This upgrade enhances the existing VoIP gateway and the Multimedia Conference Manager (gatekeeper and proxy). A *gateway* allows H.323 terminals to communicate with non-H.323 terminals by converting protocols, and it is an endpoint on the LAN that provides real-time, two-way communications between H.323 terminals on the LAN and other ITU-T terminals in the WAN or to another H.323 gateway.

The *gatekeeper* maintains a registry of devices in the multimedia network. The devices register with the gatekeeper at startup and request admission to a call from the gatekeeper. The gatekeeper is an H.323 entity on the LAN that provides address translation and control access to the LAN for H.323 terminals and gateways. The gatekeeper may provide other services to the H.323 terminals and gateways, such as bandwidth management and locating gateways.

Voice Dial Peers

Dial peers enable outgoing calls from a particular telephony device. All of the voice technologies use dial peers to define the characteristics associated with a *call leg*.

A call leg is a discrete segment of a call connection that lies between two points in the connection. It is important to remember that these terms are defined from the *router* perspective. An inbound call leg means that an incoming call comes *to* the router. An outbound call leg means that an outgoing call is placed *from* the router. Dial peers are used for both inbound and outbound call legs.

For inbound call legs, a dial peer might be associated with the calling number or the voice-port number. Outbound call legs always have a dial peer associated with them. The destination pattern is used to identify the outbound dial peer. The call is associated with the outbound dial peer at setup time.

There are two kinds of dial peers that need to be configured for each voice implementation:

- POTS—(also known as “plain old telephone service” or “basic telephone service”) dial peer associates a physical voice port with a local telephone device. The key commands in your configuration are the **port** and **destination-pattern** commands. The **destination-pattern** command defines the telephone number associated with the POTS dial peer. The **port** command associates the POTS dial peer with a specific logical dial interface, normally the voice port connecting your router to the local POTS network.
- VoIP—dial peer associates a telephone number with an IP address. The key commands in your configuration are the **destination-pattern** command and the **session target** command. The **destination-pattern** command defines the telephone number associated with the VoIP dial peer. The **session target** command specifies a destination IP address for the VoIP dial peer. In addition, you can use VoIP dial peers to define characteristics such as IP precedence, additional QoS parameters, and codec.

QoS

This section describes Quality of Service (QoS) parameters, including the following:

- [IP Precedence](#)
- [PPP Fragmentation and Interleaving](#)
- [CBWFQ](#)
- [RSVP](#)
- [Low Latency Queuing](#)

QoS refers to the capability of a network to provide better service to selected network traffic over various technologies, including ATM, Ethernet and IEEE 802.1 networks, and 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 for future business applications in campus, WAN, and service provider networks.

QoS must be configured throughout your network, not just on your router running VoIP, to improve voice network performance. Not all QoS techniques are appropriate for all network routers. Edge routers and backbone routers in your network do not necessarily perform the same operations; the QoS tasks they perform might differ as well. To configure your IP network for real-time voice traffic, you need to consider the functions of both edge and backbone routers in your network.

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, an Internet service provider, or an enterprise network.

IP Precedence

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 expedite handling.

Features such as policy-based routing and committed access rate (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 with the signaling used optionally. IP Precedence enables service classes to be established using existing network queuing mechanisms (such as CBWFQ), with no changes to existing applications or complicated network requirements.

PPP Fragmentation and Interleaving

With multiclass multilink PPP interleaving, large packets can be multilink-encapsulated and fragmented into smaller packets to satisfy the delay requirements of real-time voice traffic; small real-time packets, which are not multilink encapsulated, are transmitted between fragments of the large packets. The interleaving feature also provides a special transmit queue for the smaller, delay-sensitive packets, enabling them to be transmitted earlier than other flows. Interleaving provides the delay bounds for delay-sensitive voice packets on a slow link that is used for other best-effort traffic.

In general, multilink PPP with interleaving is used in conjunction with CBWFQ and RSVP or IP Precedence to ensure voice packet delivery. Use multilink PPP with interleaving and CBWFQ to define how data is managed; use Resource Reservation Protocol (RSVP) or IP precedence to give priority to voice packets.

CBWFQ

In general, class-based weighted fair queuing (CBWFQ) is used in conjunction with multilink PPP and interleaving and RSVP or IP Precedence to ensure voice packet delivery. CBWFQ is used with multilink PPP to define how data is managed; RSVP or IP Precedence is used to give priority to voice packets.

There are two levels of queuing; ATM queues and Cisco IOS queues. CBWFQ is applied to Cisco IOS queues. A first-in-first-out (fifo) Cisco IOS queue is automatically created when a PVC is created. If you use CBWFQ to create classes and attach them to a PVC, a queue is created for each class.

CBWFQ ensures that queues have sufficient bandwidth and that traffic gets predictable service.

Low-volume traffic streams are preferred; high-volume traffic streams share the remaining capacity, obtaining equal or proportional bandwidth.

RSVP

RSVP enables routers to reserve enough bandwidth on an interface to ensure reliability and quality performance. RSVP allows end systems to request a particular QoS from the network. Real-time voice traffic requires network consistency. Without consistent QoS, real-time traffic can experience jitter, insufficient bandwidth, delay variations, or information loss. RSVP works in conjunction with current queuing mechanisms. It is up to the interface queuing mechanism (such as CBWFQ) to implement the reservation.

RSVP works well on PPP, HDLC, and similar serial-line interfaces. It does not work well on multi-access LANs. RSVP can be equated to a dynamic access list for packet flows.

You should configure RSVP to ensure QoS if the following conditions describe your network:

- Small-scale voice network implementation
- Links slower than 2 Mbps
- Links with high utilization
- Need for the best possible voice quality

Low Latency Queuing

Low latency queuing (LLQ) provides a low-latency strict priority transmit queue for real-time traffic. Strict priority queuing allows delay-sensitive data to be dequeued and sent first (before packets in other queues are dequeued), giving delay-sensitive data preferential treatment over other traffic.

Access Lists

With basic standard and static extended access lists, you can approximate session filtering by using the established keyword with the **permit** command. The established keyword filters TCP packets based on whether the ACK or RST bits are set. (Set ACK or RST bits indicate that the packet is not the first in the session and the packet therefore belongs to an established session.) This filter criterion would be part of an access list applied permanently to an interface.

