
# Concepts

This chapter contains conceptual information that may be useful to Internet service providers and network administrators when configuring the Cisco 806 router. To review some typical network scenarios, read Chapter 2, "Network Scenarios." For information on specific configurations, read Chapter 3, "Feature-by-Feature Router Configurations."

The following topics are included in this chapter:

- Cisco 806 Router Overview
- Network Protocols
- Routing Protocol Options
- Policy-Based Routing
- IP Multicasting
- Point-to-Point Protocol over Ethernet
- Security Features
- Ethernet
- Network Address Translation
- Internet Protocol Control Protocol
- Dynamic Host Configuration Protocol Client and Server
- NetMeeting
- Network Time Protocol Server
- Service Assurance Agent

# Cisco 806 Router Overview

The Cisco 806 router is a Cisco IOS-based member of the Cisco 800 router family. The router is a fixed-configuration IP router with security features to provide a secure Ethernet gateway for users in small offices, branch offices and home offices using broadband access to the Internet. It is designed to work with digital subscriber line (DSL), cable, or long-reach Ethernet (LRE) modems, or with an Ethernet switch serving a multitenant unit.

Among the features that the Cisco 806 router supports are IP Security (IPSec), NetMeeting, and IP multicasting. Users in remote locations using a public network can exchange data, access corporate intranets, participate in video conferences using their web browsers, and access corporate multicast material such as distance learning courses and videotaped presentations.

The Cisco 806 router has four 10BaseT Ethernet ports that function as a hub; this router also has one 10BaseT Ethernet wide area network (WAN) port.

Cisco 806 router Flash memory includes Webflash, which is a 2-Megabyte partition separate from system Flash. Webflash is only used by the Cisco Router Web Setup software.

# Network Protocols

Network protocols enable the network to pass data from its source to a specific destination over LAN or WAN links. The Cisco 806 router supports the Internet Protocol (IP). To enable the transport of other protocols over IP, the Cisco 806 router supports Generic routing encapsulation (GRE) tunneling protocol.

## 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 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, before sending data, a connection-oriented protocol first exchanges control information with the remote computer to verify that it is ready to receive data. When the handshaking is successful, the computers have established a connection. IP relies on the upper layer protocol TCP to establish the connection if connection-oriented services are required.

## GRE Tunneling Protocol

GRE is a tunneling protocol developed by Cisco that can encapsulate a wide variety of protocol packet types inside IP tunnels, creating a virtual point-to-point link to Cisco routers at remote points over an IP network. By connecting multiprotocol subnetworks in a single-protocol backbone environment, IP tunneling using GRE allows network expansion across a single-protocol backbone environment. GRE tunneling is commonly used in conjunction with IPSec.

# Routing Protocol Options

Supported 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 Table 1-1 shows.

*Table 1-1 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. | Distance information. 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. | 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 Release 12.0(1)T documentation set.

# Enhanced IGRP

Enhanced IGRP is an advanced Cisco-proprietary distance-vector and link-state routing protocol. Enhanced IGRP uses a more sophisticated metric 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 multiprotocol network environments, thus minimizing the size of the routing tables and the amount of routing information.

# Policy-Based Routing

Policy-based routing (PBR) can be used when administrative issues dictate that traffic be routed through specific paths. By using policy-based routing, customers can implement policies that selectively cause packets to take different paths.

PBR also provides a mechanism for marking packets so that certain kinds of traffic receive differentiated, preferential service when used in combination with queuing techniques enabled through the Cisco IOS software. These queuing techniques are extremely powerful, simple, and flexible tools for network managers who implement routing policies in their networks.

PBR provides a mechanism for expressing and implementing forwarding/routing of data packets, based on the policies defined by the network administrators. It provides a more flexible mechanism for routing packets than the existing mechanism provided by routing protocols that use the destination of the packet to route it.

Routers forward packets to the destination addresses, based on information from static routes or dynamic routing protocols such as RIP, Open Shortest Path First (OSPF), or Enhanced IGRP. Instead of routing by the destination address, policy-based routing allows network administrators to determine and implement routing policies to allow or deny paths based on the following:

- Identity of a particular end system
- Application

- Protocol
- Size of packets

# IP Multicasting

IP multicasting is a bandwidth-conserving technology that reduces traffic by simultaneously delivering a single stream of information to thousands of corporate recipients and homes. Applications that take advantage of multicasting include video conferencing, corporate communications, distance learning, and distribution of software, stock quotes, and news.

Without multicasting, these applications must be run by two inefficient schemes—unicasting and broadcasting. In unicasting, one copy of data is sent to each receiver. Although unicasting is a simple mechanism for one-to-one communication, it demands too much bandwidth from the network for one-to-many communication. In broadcasting, a single copy of data is sent to every user in the network, circumventing the bandwidth problem. However, it is not suitable if only few receivers requested the data.

IP multicasting solves the inherent bottlenecks created when you need information transferred from a single sender to multiple recipients. By sending only one copy of the information to the network and letting the network intelligently replicate the packet only where it was requested, you conserve bandwidth and network resources on both the sending end and the receiving end of a transmission.

# Point-to-Point Protocol over Ethernet

The Point-to-Point Protocol (PPP) encapsulates network layer protocol information over point-to-point links. Point-to-Point Protocol over Ethernet (PPPoE) allows a PPP session to be initiated on a simple bridging Ethernet-connected client.

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)

PPPoE 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 in which a remote office Cisco 806 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 in which a remote office Cisco 806 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 anytime 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 more secure than PAP.

# Security Features

This section discusses the security features available on the Cisco 806 router. It discusses the following features:

- IPSec
- Access Lists
- Remote Authentication Dial-In User Service
- Terminal Access Controller Access Control System Plus

## IPSec

IPSec is a framework of open standards for ensuring secure private communications over IP networks. Based on standards developed by the Internet Engineering Task Force (IETF), IPSec ensures confidentiality, integrity, and authenticity of data communications across a public IP network. IPSec provides a necessary component of a standards-based, flexible solution for deploying a network-wide security policy.

Cisco's implementation of IPSec uses the Data Encryption Standard (DES) and triple DES.

# 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 acknowledgement (ACK) or reset (RST) bits are set. (Set ACK or RST bits indicate that the packet is not the first in a session and that the packet therefore belongs to an established session.) This filter criterion would be part of an access list applied permanently to an interface.

# Remote Authentication Dial-In User Service

Remote Authentication Dial-In User Service (RADIUS) is a distributed client/server system that secures networks against unauthorized access. In the Cisco implementation, RADIUS clients run on Cisco routers and send authentication requests to a central RADIUS server that contains all user authentication and network service access information.

RADIUS is a fully open protocol, distributed in source code format, that can be modified to work with any security system available on the market.

Cisco supports RADIUS under its AAA security paradigm. RADIUS can be used with other AAA security protocols, such as Terminal Access Controller Access Control System Plus (TACACS+), Kerberos, or local username lookup.

Use RADIUS in the following network environments that require access security:

- Networks with multiple-vendor access servers, each supporting RADIUS—For example, access servers from several vendors use a single RADIUS server-based security database. In an IP-based network with multiple vendors' access servers, dial-in users are authenticated through a RADIUS server that has been customized to work with the Kerberos security system.
- Networks already using RADIUS—You can add a Cisco router with RADIUS to the network. This might be the first step when you make a transition to a TACACS+ server.

- Networks in which a user must access only a single service—Using RADIUS, you can control user access to a single host, to a single utility such as Telnet, or to a single protocol such as PPP. For example, when a user logs in, RADIUS indentifies this user as having authorization to run PPP using IP address 10.2.3.4, and the defined access list is then started.
- Networks that require resource accounting—You can use RADIUS accounting independent of RADIUS authentication or authorization. The RADIUS accounting functions allow data to be sent at the start and end of services, indicating the amount of resources (such as time, packets, bytes, and so on) used during the session. An Internet service provider (ISP) might use a freeware-based version of RADIUS access control and accounting software to meet special security and billing needs.

RADIUS is not suitable in the following network security situations:

- Multiprotocol access environments—RADIUS does not support the following protocols:
  - AppleTalk Remote Access Protocol (ARAP)
  - NetBIOS Frame Protocol Control Protocol (NBFCP)
  - NetWare Asynchronous Services Interface (NASI)
  - X.25 packet assembler/disassembler (PAD) connections.
- Router-to-router situations—RADIUS does not provide two-way authentication. RADIUS can be used to authenticate from one router to a non-Cisco router if the non-Cisco router requires RADIUS authentication.
- Networks using a variety of services—RADIUS generally binds a user to one service model.

# Terminal Access Controller Access Control System Plus

Cisco 806 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.

# Ethernet

The Cisco 806 router supports two Ethernet network interfaces. 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.

# Network Address Translation

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 numerical 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.

NAT supports H.323 signaling for the Netmeeting application.

# Internet Protocol Control Protocol

NAT and PPP/Internet Protocol Control Protocol (IPCP) enable a Cisco router to automatically negotiate its own registered WAN interface IP address from a central server and make it possible for all remote hosts to access the Internet using this single registered IP address. Because existing port-level multiplexed NAT functionality within the Cisco IOS software is used, IP addresses on the remote LAN are invisible to the Internet.

With PPP/IPCP, the Cisco 806 router automatically negotiates a globally unique (registered) IP address for the dialer interface from the ISP router.

# Dynamic Host Configuration Protocol Client and Server

Dynamic Host Configuration Protocol (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 manually assign an IP address to each client.

DHCP configures the router to forward User Datagram Protocol (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 NAT is used, DHCP relay cannot be used on the Cisco 806 router. The built-in DHCP server should be used instead.

# NetMeeting

Microsoft NetMeeting is a conferencing tool that enables individuals to hold meetings with each other from their computers. Individuals using computers with the NetMeeting software and an Internet or LAN connection can hold meetings remotely, using such tools as a shared whiteboard, application and document sharing, file transfers, and (with the necessary hardware) audio and video conferencing.

The Cisco 806 router supports H.323 signaling for NetMeeting. Refer to the Cisco IOS documentation set for specific NetMeeting support information.

# Network Time Protocol Server

The Network Time Protocol (NTP) provides a synchronized time base for networked routers, servers, and other devices. It also coordinates the time of network events, which aids in understanding and troubleshooting the time sequence of network events. For example, call records for specific users can be

correlated within one millisecond. Using information from an NTP server, you can compare time logs from different networks, which is essential for tracking security incidents, analyzing faults, and troubleshooting.

# Service Assurance Agent

The Service Assurance Agent (SA Agent) is both an enhancement to and a new name for the Response Time Reporter (RTR) feature that was introduced in Cisco IOS Release 11.2. The feature allows you to monitor network performance by measuring key Service Level Agreement (SLA) metrics such as response time, network resources, availability, jitter, connect time, packet loss and application performance. With Cisco IOS Release 12.0(5)T, the SA Agent provides new capabilities that enable you to

- Monitor the Domain Name Server, DHCP Server, and data-link switching (DLSw) peer stack and tunnel performance. Thresholds can be used to trigger additional collection of time delay statistics.
- Monitor network one-way delay variance (jitter) and packet loss.
- Monitor web server response time.