



CCNP ISCW Quick Reference Sheets

Denise Donohue, CCIE No. 9566
Jay Swan

Chapter 1 Network Conceptual Models	3
Chapter 2 Providing SOHO/ Teleworker Connectivity.....	5
Chapter 3 Frame Mode MPLS.....	16
Chapter 4 IPsec	24
Chapter 5 Cisco Device Hardening	38
Chapter 6 Cisco IOS Threat Defenses.....	54



About the Authors

Denise Donohue, CCIE® No. 9566, is a manager of solutions engineering for ePlus Technology in Maryland. She is responsible for designing and implementing data and VoIP networks, supporting companies based in the National Capital region. Prior to this role, she was a systems engineer for the data consulting arm of SBC/AT&T. Denise was a Cisco® instructor and course director for Global Knowledge and did network consulting for many years. Her CCIE is in Routing and Switching.

Jay Swan is a senior network engineer for the Southern Ute Indian Tribe Growth Fund in Ignacio, CO. Prior to this position, he was a Cisco instructor and course director for Global Knowledge. Jay has also worked in IT in the higher education and service provider fields. He holds CCNP® and CCSP® certifications.

About the Technical Review

Rus Healy, CCIE No. 15025, works as a senior engineer for Annese & Associates, a Cisco partner in upstate New York. He also holds CCNP and CCDP® certifications. His other interests include bicycling, skiing, and camping with his family, as well as competitive amateur radio events.

CHAPTER 1

CCNP ISCW Quick Reference Sheets by Denise Donohue and Jay Swan

Chapter 1

Network Conceptual Models 4

Intelligent Information Network 4

Service-Oriented Network Architecture 4

Cisco Enterprise Architecture 5

Network Conceptual Models

Conceptual models allow network designers to move from looking at the network as a collection of devices to viewing it as a way to provide services to users, no matter where the users and the services are located. Cisco currently defines three models as the building blocks of a world-class enterprise network: the Intelligent Information Network, the Service-Oriented Network Architecture, and the Cisco Enterprise Architecture.

Intelligent Information Network

The Intelligent Information Network (IIN) seeks to create a holistic network that integrates with and enables your business processes. It allows centralized control and interoperation of distributed systems. This control and interoperability can provide increased network security and efficiency.

The IIN consists of three components:

- **Integrated transport**—Data, voice, and video all transported over a secure IP network.

- **Integrated services**—Shared/virtualized resources, such as storage and servers.
- **Integrated applications**—The network is application aware, enhancing the efficiency of applications. This component includes Application-Oriented Networking (AON), which offloads shared, common functions, such as logging and security, to the network.

Service-Oriented Network Architecture

An intelligent network is delivered using the Service-Oriented Network Architecture (SONA) framework. SONA sees a converged network as the connecting thread for all the portions of the network and the services provided. The network is application aware; that is, it contains the intelligence needed to tie all the various types of traffic together to deliver required services. SONA defines three layers:

- **Network Infrastructure**—IT resources such as servers, users, WANs, and office locations all connected and accessible to each other
- **Integrated Services**—Services such as voice, network management, mobility, security, and storage that are delivered using the network infrastructure
- **Application**—Business applications that function using the integrated services

In an enterprise, the campus and branch offices, teleworker access, and WAN access all fall under the Network Infrastructure layer of the SONA. This categorization allows workers in all types of locations to access the services and applications of the other layers. The Cisco Enterprise Architecture defines how each of these components should be designed and structured.

Cisco Enterprise Architecture

The Cisco Enterprise Architecture model divides the network into building blocks and gives best practices for the architecture of each one. The traditional three-layer model (Core, Distribution, and Access) is still around and can be integrated into the design of components of the Enterprise Architecture model.

Enterprise Architecture building blocks include the following:

- **Campus**—The enterprise core, or headquarters. The campus building block contains routing, switching, security, Voice over IP (VoIP), wireless, and so on.
- **Data center**—Server and application resources. Redundant data centers provide business continuity and allow load balancing.
- **Branch office**—Remote locations that contain services similar to the campus but are administered centrally rather than at each location.

- **Teleworker**—Either a small office, home office, or a mobile user. Extends data (and possibly voice) services to these users over a virtual private network (VPN) using broadband WAN access.
- **WAN**—Connects all the different blocks together. Converges voice, video, and data over an IP WAN that provides security, quality of service (QoS), and ubiquitous access.

WAN options between the campus and branch offices include traditional Layer 2 connections such as Frame Relay, ATM, and leased lines. Multiprotocol Label Switching (MPLS) can provide any-to-any connectivity between the sites and is highly scalable. IPsec VPNs across the Internet can also be used.

This short cut is concerned mainly with how the campus, branch, and small office, home office (SOHO)/teleworker portions of the network use the WAN to communicate with each other to provide network services to their users.

CHAPTER 2

CCNP ISCW Quick Reference Sheets by Denise Donohue and Jay Swan

Chapter 2

Providing SOHO/Teleworker Connectivity 7

Broadband Cable 7

Digital Subscriber Line 10

Configuring DSL CPE 13

Providing SOHO/Teleworker Connectivity

The traditional teleworker solution consists of a virtual private network (VPN) client on the user's computer connecting over the Internet to a VPN concentrator, firewall, or Cisco Adaptive Security Appliance (ASA) at the corporate site. This requires only a dialup or broadband Internet connection and a dialup or broadband modem. However, this approach has several shortcomings:

- Dialup does not provide the necessary bandwidth to take advantage of all the corporate services, such as Voice over IP (VoIP).
- There is no centralized control of the teleworker equipment, so security, virus protection, and so forth are left to the teleworker to implement.
- There is no control over quality of service (QoS) for advanced services.
- It is hard for corporate IT staff to support.

The Cisco Business-Ready Teleworker Solution addresses these issues with the traditional teleworker approach. It seeks to secure corporate data by using IPsec VPNs, allow corporate control of the connection components, and provide a scalable architecture as part of disaster planning. It consists of an always-on broadband connection, a corporate-owned and -managed router configured for VPN and QoS, IP phone, and (optionally) video equipment.

Two typical broadband connection types are digital subscriber line (DSL) and cable.

Broadband Cable

CATV, or Community Antenna Television, was originally developed to provide improved TV signals by sharing antennas and satellite dishes. It used coaxial cable to transport the TV signals to each subscriber. Current systems typically use a combination of fiber and coaxial. When both fiber and coaxial cables are used, the system is referred to as a *hybrid fiber-coaxial (HFC) network*.

Broadband cable uses frequency-division multiplexing (FDM) to deliver data over a radio frequency (RF) network. Cable provides relatively inexpensive high-speed Internet access that supports analog and digital video, voice, and data. Its downsides include possible bandwidth and security issues because it is a shared medium. The provider can increase bandwidth by using smaller service areas and more channels. Security can be addressed by the user within the cable modem or with an onsite router.

Cable Components

Five basic components comprise a cable TV/data system:

- **Antenna site**—Receives TV signals from antennas or satellite dishes.
- **Headend site**—Converts TV signals for distribution to end users and converts data for transport to and from end users. Similar to a telephone central office.
- **Transportation network**—Links the antenna site to the headend or the headend to the distribution network. Can use microwave, coaxial cable, or fiber cable.
- **Distribution network**—Carries signals between the end user and the transportation network. Consists of trunks and feeder cables. Backbone trunks are either fiber or coaxial cable. Feeder cables are usually coaxial and connect the distribution network to the subscriber drops.
- **Subscriber drops**—Connects the customer premises equipment (CPE) such as TV, set-top box, or cable modem to the distribution network. Uses coaxial cable.

Figure 2-1 shows more detail about how these components work together to deliver combined cable TV and data to an end user. Each acronym and component is further explained following the figure.

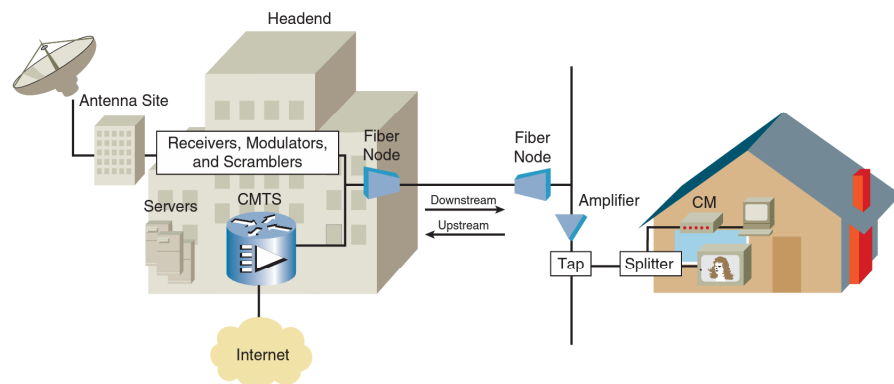


FIGURE 2-1 Components of a Cable System

In Figure 2-1, video signals are received at the antenna site and sent to the local headend site. Feeds from other headend sites might be used, too. Each TV channel has receivers, modulators, and scramblers. The signals are transferred via RF to a fiber transmitter. Data signals from the Internet or various servers (such as e-mail or content servers) are modulated by the Cable Modem Termination System (CMTS) router and sent as RF signals to the fiber transmitter. They are then translated to optical signals and sent downstream, toward the end users. A fiber node in the distribution network translates them back to RF signals and sends them over the coaxial cable. Signals are boosted at intervals by an amplifier. A tap divides the signal for sending to a particular subscriber's residence. At the residence, a splitter divides the signals into data and video. Video is sent to the TV or set-top box, and data is sent to a cable modem (CM). The cable modem demodulates the signal back to digital.

Cable Standards

Worldwide, three standards control cable TV systems:

- **National Television Standards Committee (NTSC)**—Analog television standard used in North America. Specifies a 6-MHz channel width.
- **Phase-Alternating Line (PAL)**—Color television standard used in most of the rest world. Specifies 6-, 7-, or 8-MHz channel widths.
- **Système Electronique Couleur avec Mèmoire (SECAM)**—Standard used in France and some Eastern European countries. Specifies an 8-MHz channel.

The standard for sending data over cable systems is Data-Over-Cable Service Interface Specifications (DOCSIS). This standard, developed by Cablelabs, defines physical and data link layer requirements for cable modems. Cablelabs also certifies cable modems and CMTS systems to work with the standard. At the physical layer, DOCSIS specifies channel widths and modulation methods. At the data link layer it specifies access methods, some QoS capabilities, and some security features.

Cable RF waves use different frequencies for upstream and downstream signaling. Downstream signals are allowed 810 MHz of bandwidth, in the 50- to 860-MHz range. This is then further subdivided into channels of 6-, 7-, or 8-MHz depending on the standard used. Upstream signals have only 37 MHz of bandwidth, in the 5- to 42-MHz range.

Provisioning the Cable Modem

Cable modems communicate with their CMTS across whatever physical networks connect them. The service provider must have the necessary auxiliary services, such as DHCP, TFTP, and Time of Day (TOD), available at its headend. When a cable modem boots, it registers with the CMTS and acquires its configuration using the following steps:

- Step 1.** The cable modem scans for a downstream channel to use for communication with the CMTS. Once it finds a channel, the CM locks it in.
- Step 2.** CMTS tells the CM the parameters to use for upstream messages.
- Step 3.** Communication is established at Layers 1 and 2 (physical and data link layers).
- Step 4.** The cable modem broadcasts for a DHCP server. It obtains an IP address, the address of the TFTP and TOD servers, and the name of the TFTP file to download.
- Step 5.** The cable modem downloads the DOCSIS configuration file from the TFTP server.
- Step 6.** The cable modem forwards the configuration file to the CMTS and attempts to register with it. If the configuration is valid, the modem is registered. The two devices negotiate QoS and security settings.
- Step 7.** The user device—either a PC or a router—requests an IP address, DNS server, and default gateway information from the cable provider.

Digital Subscriber Line

Voice does not use all the available bandwidth on a phone line—it uses frequencies only up to about 3 kHz. DSL was created to use the space between 3 kHz and 1 MHz to send data traffic over a telephone local loop. Thus, both voice and data can be sent simultaneously over the same connection (some variants of DSL use the entire spectrum, however, so no voice can be sent). DSL is a physical layer medium that extends between the subscriber's DSL modem and the provider's DSL access multiplexer (DSLAM.)

Types of DSL

Asymmetrical DSL has higher downstream (from the provider's central office [CO] to the subscriber) bandwidth than upstream (from the subscriber to the CO.) Symmetrical DSL has the same bandwidth both downstream and upstream. You will sometimes see these referred to as “asynchronous” and “synchronous” DSL.

The various types of DSL include the following:

- **ADSL**—Asymmetric DSL supports both voice and data. Downstream bandwidth goes up to 8 Mbps; upstream goes up to 1 Mbps. Two other versions, ADSL2 and ADSL2+, provide 24 Mbps downstream and 1.5 Mbps upstream. The maximum distance from the CO is 18,000 feet, or 5.46 km.
- **RADSL**—Rate-adaptive DSL changes the rate based on the local loop.

- **VDSL**—Very-high-rate DSL can be either symmetric or asymmetric and can carry voice along with data. Maximum symmetric bandwidth is 26 Mbps; maximum asymmetric is 52 Mbps downstream and 13 Mbps upstream. The maximum distance from the CO is 4,500 feet, or 1.37 km.
- **IDSL**—ISDN DSL carries only digital data (other forms of DSL send analog signals). It uses both ISDN B channels and the D channel, for a symmetric bandwidth of 144 kbps. The maximum distance for IDSL is 18,000 feet, or 5.46 km.
- **SDSL**—Symmetric DSL carries only data, with a maximum for both downstream and upstream of 768 kbps. The distance limitation is 22,000 feet, or 6.7 km. It is a proprietary technology that uses only one twisted pair of wires.
- **HDSL**—High-data-rate DSL uses two twisted pairs of wires to achieve a maximum symmetrical bandwidth of 2.048 Mbps. Its maximum distance from the CO is 12,000 feet, or 3.7 km. HDSL carries only data, no voice.
- **G.SHDSL**—Symmetric high-speed DSL has a symmetrical data rate of 2.3 Mbps and the longest maximum distance: 28,000 feet, or 8.52 km. It also carries only data, no voice.

ADSL

ADSL is a popular residential service because it can carry both voice and data over one twisted pair of wires. This capability is accomplished by using either a splitter or a filter. A splitter takes the incoming analog signals and splits off the frequencies under 4 MHz to a voice line. It sends all other traffic to the DSL line. Splitters are more typically used at the CO than the subscriber premises because they require a technician to install them. A filter, or microfilter, requires no installation. It simply connects to the phone line on one end and the telephone on the other. It passively filters out any signals in the DSL range so that only voice reaches the telephone.

Figure 2-2 shows how ADSL components work together in a typical residential implementation. The telephone company's CO forwards both plain old telephone service (POTS) and DSL data traffic over the same line to the subscriber. The line enters at the network interface device (NIDS) and branches toward the telephone and the PC. A low-pass filter blocks everything but voice frequencies from reaching the phone. A DSL modem (or router with a DSL interface) forwards data to the PC. When the CO receives traffic from the subscriber, a splitter sends voice frequencies to the PSTN switch and DSL frequencies to the DSLAM. The DSLAM sends data traffic to a router for forwarding to the Internet.

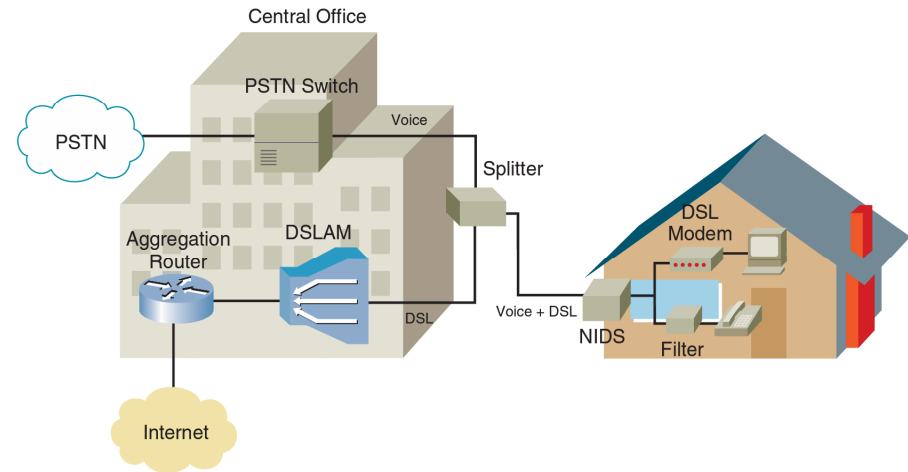


FIGURE 2-2 Components of an ADSL System

Carrierless Amplitude and Phase Line Coding

Carrierless Amplitude and Phase (CAP) is a DSL line-coding method that divides the bandwidth into three channels: one for voice, one for downstream data, and one for upstream data. Each type of traffic is carried within one frequency band, and so CAP is termed a *single-carrier modulation technique*. The bands are fairly wide. Voice uses 0 to 4 kHz, upstream traffic uses 25 to 160 kHz, and downstream uses 240 kHz to 1.5 MHz. CAP is simple to understand and implement but does not scale as well as Discrete Multi-Tone (DMT) modulation.

Discrete Multi-Tone Line Coding

Discrete Multi-Tone (DMT) is the most widely used method of ADSL line coding. It divides the DSL frequency band into 256 channels of 4 kHz each. Some channels are duplex and used for both downstream and upstream traffic. Others are used only for downstream. Channel quality is constantly monitored, and the channels used can be changed when conditions warrant. DMT is more complex than CAP but is also more flexible and scalable and can achieve higher speeds.

A version of DMT, G.Lite ADSL, uses half the number of channels as DMT.

Layer 2 over DSL

Recall that DSL is a Layer 1 (physical layer) technology. There are three methods of carrying data at Layer 2 over DSL:

- **Bridging**—Based on RFCs 1483 and 2684. Ethernet traffic is just bridged from the subscriber PCs, through the DSL modem and the DSLAM, to a provider router. Bridging is not as secure or scalable as other methods.
- **PPP over Ethernet (PPPoE)**—The most common Layer 2 method of carrying data over DSL. PPP traffic is encapsulated in Ethernet frames.
- **PPP over ATM (PPPoA)**—PPP packets are routed over ATM between the subscriber equipment and the provider.

PPPoE

When PPPoE is used, a PPP session is established, similar to when using dialup. Either Password Authentication Protocol (PAP) or Challenge Handshake Authentication Protocol (CHAP) authentication can be used. The provider's aggregation router and the subscriber's CPE establish a session between them. There are three ways to do this:

- The PPP client on a subscriber router with a DSL interface terminates both the DSL and the PPP sessions. The router can allow multiple users over the DSL with just one PPP login, by acting as a DHCP server and doing Network Address Translation (NAT) or Port Address Translation (PAT) for the subscriber users. The router obtains its outside IP address via PPP's IP Control Protocol (IPCP.)
- A DSL modem terminates the DSL session, and the PPP client on a CPE router terminates the PPP session. The router can act as a DHCP server and do NAT/PAT, to allow multiple internal users. It obtains its outside address via IPCP.
- A DSL modem terminates the DSL session, and a PPP client on the subscriber PC terminates the PPP session. Traffic is bridged from the PC to the aggregation router. This allows only a single DSL user. If multiple users are at the same residence, they must each have their own PPP login, and they each obtain an IP address via IPCP.

PPP was created to be used over a point-to-point connection, and Ethernet is inherently multipoint, so PPPoE uses a PPP server discovery process. After a server has been discovered, a virtual point-to-point link can be established, and the PPP session process can continue. The PPP server discovery stage has four steps:

- Step 1.** The PPP client sends a PPPoE Active Discovery Initiation (PADI) broadcast.
- Step 2.** Any PPP servers (aggregation routers) reply with a PPPoE Active Discovery Offer (PADO), sent as a unicast to the client's MAC address.
- Step 3.** The client replies to the server with a PPPoE Active Discovery Request (PADR).
- Step 4.** The server confirms the association with a PPPoE Active Discovery Session-confirmation (PADS) message.

When these steps have been completed, the normal PPP session negotiations proceed, and a session is established.

PPPoA

PPPoA requires a CPE router; traffic is routed from the subscriber PCs to the aggregation router—it cannot be bridged as with PPPoE. The PPP session is established between the CPE router and the aggregation router. Multiple users are supported if the CPE router is configured to do DHCP and NAT. Traffic between the CPE router and the aggregation router is encapsulated as ATM at Layer 2, rather than Ethernet. Therefore, the CPE router must have an ATM interface.

Configuring DSL CPE

When a CPE router is used as the PPP client, it must be configured. The configuration will differ depending on whether you are using PPPoE or PPPoA. PPPoE can be used with an Ethernet or an ATM interface on the router connecting to the DSL network. An Ethernet interface is used if the router connects to a DSL modem; an ATM interface is used if the router connects directly to the DSL network. A dialer interface must also be created and configured with PPP parameters.

Configuring PPPoE CPE

The following tasks must be completed to configure a CPE router with for PPPoE:

1. Configure the internal and external interfaces.
2. Configure a dialer interface.
3. Configure NAT/PAT.
4. Configure the router to act as a DHCP server.
5. Configure a default route.

First, configure the internal Ethernet interface with an IP address. It will be the default gateway for the users. Also, configure it as the inside interface for NAT.

CHAPTER 2

Do not put an IP address on the external Ethernet interface. Enable PPPoE on it, and assign it to a PPPoE client dialer pool. The final configuration on the two Ethernet interfaces should be similar to that shown in Example 2-1.

Example 2-1 Configuring Ethernet Interfaces for PPPoE

```
interface FastEthernet0/0
  description DSL interface
  no ip address
  pppoe enable
  pppoe-client dial-pool-number 1
!
interface FastEthernet0/1
  description Internal interface
  ip address 172.16.1.1 255.255.255.0
  ip nat inside
```

If the external interface is ATM, the configuration changes slightly. You must configure the ATM permanent virtual circuit (PVC) information and assign the interface to a PPPoE client dialer pool. Leave the DSL operating mode at its default to auto-detect the correct modulation, as shown in Example 2-2.

Example 2-2 Configuring an ATM Interface for PPPoE

```
interface ATM1/0
  description DSL interface
  no ip address
  dsl operating-mode auto
  pvc 1/100
  pppoe-client dial-pool-number 1
```

Second, configure a dialer interface and assign it to the same dialer pool as the Ethernet interface. Give it a PPP encapsulation and configure the PPP parameters on it. Make it the NAT outside interface. Limit the maximum transmission unit (MTU) size to 1492 bytes, to allow for the PPP and Ethernet headers. Because DSL is an always-on connection, a dialer list is not required. Your configuration should look similar to Example 2-3. Verify PPP operation with the **show ppp session** command. You can also debug PPP with the commands **debug ppp authentication** and **debug ppp negotiation**.

Example 2-3 Configuring a Dialer Interface for PPPoE

```
interface Dialer1
  ip address negotiated
  ip mtu 1492
  ip nat outside
  encapsulation ppp
  dialer pool 1
  ppp authentication chap
  ppp chap password dslpass
```

Next, configure the router to do NAT or PAT. NAT translates one internal address to one external one. PAT can translate multiple internal addresses to one external one. Most residential and SOHO subscribers use PAT. To configure it, identify the traffic that must be translated using an access list. Then tell the router to translate those IP addresses to the IP address of the dialer interface, and to “overload” that external IP address. The **overload** command causes the router to use PAT. Be

CHAPTER 2

sure to designate the inside and outside interfaces (see Example 2-1 for those commands.) PAT configuration is done in global command mode, and shown in Example 2-4.

Verify your NAT/PAT operation with the **show ip nat translations** command.

Example 2-4 Configuring NAT/PAT

```
access-list 100 permit ip 172.16.1.0 0.0.0.255 any
!
ip nat inside source list 100 interface Dialer1 overload
```

The next task is to configure the router to serve IP addresses to internal hosts. To set up basic DHCP, create a pool of addresses for assigning to clients, specify the clients' default gateway, and import the DNS information obtained from the DSL provider via PPP. Example 2-5 shows what this might look like. On the router, verify IP address assignment using the command **show ip dhcp binding**. On a Windows-based user computer, verify the IP address using the DOS command **ipconfig /all**.

Example 2-5 Configuring a Router as a DHCP Server

```
ip dhcp pool Users
import all
network 172.16.1.0 255.255.255.0
default-router 172.16.1.1
```

Finally, configure a static default route. It should point to the dialer interface rather than an IP address, as shown in Example 2-6.

Example 2-6 Configuring a Static Default Route

```
ip route 0.0.0.0 0.0.0.0 Dialer1
```

Configuring PPPoA CPE

Cisco routers support three types of PPPoA connections:

- Cisco proprietary PPPoA
- The Internet Engineering Task Force (IETF)'s Multiplex (MUX)-encapsulated PPPoA
- The IETF's Logical Link Control (LLC)-encapsulated PPPoA

Configuring PPPoA involves almost the same tasks as configuring PPPoE. You still must set up the internal Ethernet interface, a dialer interface, PAT, DHCP, and a static default route. The main difference is in the configuration of the external interface. Because it is ATM, you need to configure virtual path identifier (VPI) and virtual circuit identifier (VCI) information to match that of the provider. The type of ATM encapsulation must be specified, and the PPPoA enabled, and linked to the virtual dialer interface. A dialer pool is associated with PVC, as shown in Example 2-7.

Example 2-7 Configuring the PPPoA ATM Interface

```

interface ATM1/0
  description DSL interface
  no ip address
  dsl operating-mode auto
  pvc 1/100
    encapsulation aal5mux ppp dialer
    dialer pool-member 1

```

Troubleshooting ADSL

When troubleshooting ADSL problems, start with checking Layer 1 physical connectivity. If that checks out okay, progress to checking Layer 2 connectivity. Finally, check Layer 3.

Troubleshooting ADSL at Layer 1

Try the following procedures when troubleshooting DSL physical layer problems:

1. Check for the carrier detect light on the router's ATM interface. If it is off, use the **show interfaces atm** *interface_number* command to check the interface status. If the interface status is down, try swapping out the RJ-11 cable connecting to the wall jack. The middle pins are used with ADSL signaling. If that doesn't work, contact the provider to make sure that the DSL service has been started.

2. If the interface status shown in the **show interfaces atm** command is "administratively down," enable the interface with the command **no shutdown**.
3. If there is a carrier detect light, check communication with the DSLAM. The subscriber DSL modem should *train* to the DSLAM. This allows them to negotiate settings such as speed. To verify this, use the **show dsl interface atm** *interface_number* command and look for the Modem Status field. It should say "Showtime." There should also be a nonzero value in the Speed field.
4. Check the DSL modulation type. Verify with the provider that your router's chipset is supported and find out what operating mode it should be in to support the correct modulation. Set this on the router with the command **dsl operating-mode {auto | ansi-dmt | itu-dmt | splitterless}**. If the modulation type is unknown, use the **auto** option.

Troubleshooting ADSL at Layer 2

If the interface status shows that it is up, and the line protocol is up, move to troubleshooting Layer 2 issues. Try the following to look for PVC or PPP problems:

1. Use the command **ping-atm-interface-atm** *interface_number* *vpi* *vci* **seg-loopback** to check that your PVC is configured on the next-hop ATM switch, which is typically the DSLAM. This command

sends management traffic called Operation, Administration, and Maintenance (OAM) packets to the DSLAM. You should receive a normal ping response if the PVC is configured.

2. Debug the events occurring on the interface processor with the **debug atm events** command. This should show no output when everything is working well; when there are problems, however, it can show useful information such as the VPI/VCI number that the DSLAM expects. The ISCW course recommends beginning a continuous ping over the Internet (not over the internal network) to the router's IP address before giving this command.
3. Verify that the router is receiving data by using the **show interfaces atm interface_number** command. Look for packets input and output.
4. If the previous procedures show that everything is working, check for PPP problems. PPP should go through three phases: Link Control Protocol (LCP) negotiation, authentication, and Network Control Protocol (NCP) negotiation. The IP address is assigned by IPCP during the NCP phase. Use the commands **debug ppp negotiation** and **debug ppp authentication** to see whether there is a failure at any of these phases.

When debugging PPP, look first for a lack of response from the aggregation router. If data link parameters cannot be negotiated, LCP will not open. If the authentication parameters are incorrectly configured, CHAP authentication will fail. If IPCP fails, the IP parameters are likely configured incorrectly either on the CPE or on the aggregation router.

CHAPTER 3

CCNP ISCW Quick Reference Sheets by Denise Donohue and Jay Swan

Chapter 3**Frame Mode MPLS 19**

Cisco Express Forwarding.....19

MPLS Routers.....20

MPLS Labels21

Configuring Frame Mode MPLS.....23

MPLS VPNs.....24

Frame Mode MPLS

Multiprotocol Label Switching (MPLS) is a technology that provides any-to-any connectivity between remote sites, using only a single WAN connection per site. The “magic” to MPLS is done in the WAN provider’s network where traffic is switched from hop to hop, rather than routed. MPLS tunnels each company’s traffic through the provider’s network, providing an extra level of separation and security. Some large companies set up their own private MPLS networks, but most rely on the provider. Most companies just connect their WAN edge routers to a provider MPLS edge router and have no MPLS configuration on their own routers at all.

The destination IP network usually determines the path between sites, but MPLS allows other considerations to influence the path, such as the following:

- Virtual private network (VPN) destination
- Quality of service (QoS) settings
- Source address
- Outbound interface
- Layer 2 circuit

Having path-selection options other than IP route enables MPLS to support non-IP protocols. MPLS uses labels that are read and acted upon at Layer 2 to indicate the selected path. MPLS routers remove the received label and insert a new label before the packet is forwarded to the next hop.

Cisco Express Forwarding

MPLS relies on Cisco Express Forwarding (CEF) to indicate the next hop for a packet to use. Cisco routers can use three types of packet switching:

- **Process switching**—The CPU must be interrupted and a route table lookup done for every packet. This is the slowest type of switching.
- **Fast switching**—A route table lookup is done only for the first packet in a flow. The next-hop information, including the Layer 2 header, is cached and used for the remainder of the packets in the flow. Faster than process switching.
- **CEF switching**—The router builds tables of next-hop and Layer 2 information before any traffic is received. This is the fastest type of switching.

CEF takes information from the IP routing table and builds its own table, the Forwarding Information Base (FIB). Because the CEF table is based on the routing table, any route changes are immediately reflected in it. CEF also builds an adjacency table, which contains the Layer 2 header for each next-hop neighbor. When a packet needs to be forwarded through the router, CEF can usually do all the processing in hardware, making it extremely fast. With MPLS, an extra field with label information is added to the FIB.

MPLS Routers

MPLS defines two roles for routers. A *Label Switch Router (LSR)* has all its interfaces within the MPLS network and does its path selection primarily based on labels. An *Edge LSR* has some interfaces in the MPLS network and some in a normal IP network, and so does some routing and some label switching. An LSR is sometimes referred to as a *provider (P) router*, and the edge LSR as a *provider edge (PE) router*.

LSRs function at two planes, the control plane and the data plane. The control plane handles routing protocols and a label-exchange protocol called Label Distribution Protocol (LDP). It contains the routing table and the Label Information Base (LIB). The data plane contains the CEF FIB and adjacency table and the MPLS Label FIB (LFIB); it forwards traffic based on those. Figure 3-1 shows the functions at each plane.

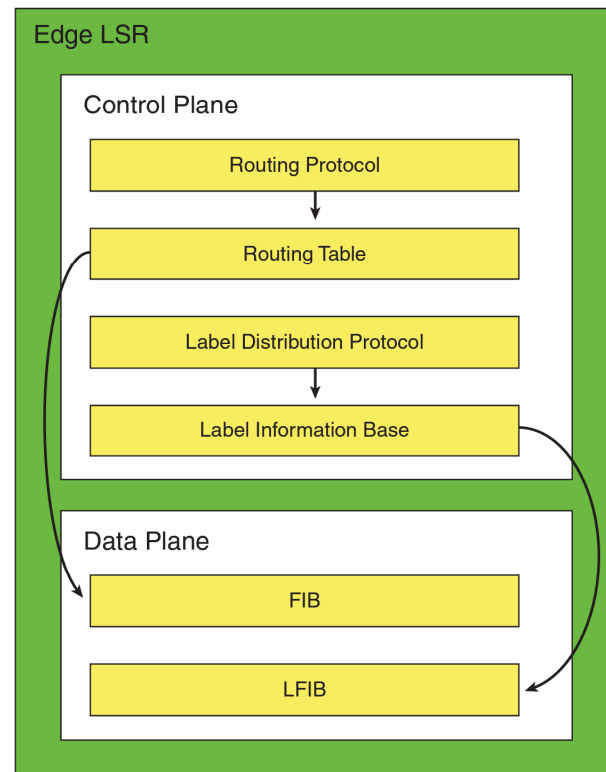


FIGURE 3-1 Functions of the Control and Data Planes on an LSR

Figure 3-1 shows an edge LSR, which does both routing and label switching. Four forwarding scenarios could occur in this router:

- An unlabeled IP packet is received and is forwarded unlabeled to a next hop in the IP network.
- An unlabeled IP packet is received, a label is imposed, and it is forwarded to a next hop in the MPLS network.
- A labeled packet is received, the label is swapped, and it is forwarded to a next hop in the MPLS network.
- A labeled packet is received, the label is removed, and it is sent as an unlabeled, regular IP packet.

MPLS Labels

MPLS has two modes: Cell mode and Frame mode. Cell mode is used with ATM, and the virtual path identifier / virtual circuit identifier (VPI/VCI) values are used as the MPLS label. Frame mode is used with any Layer 2 protocol that uses frames and inserts a 32-bit label between the Layer 2 and Layer 3 headers. MPLS routers can do one of three things with a label. They can *impose* a label, which means they insert it into the header. They can *swap* a label, which means they remove one label and replace it with another. Or they can *pop* a label, which means they remove the label. Figure 3-2 shows the structure of an MPLS label.

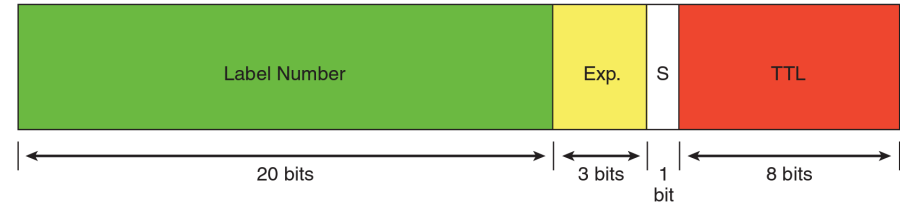


FIGURE 3-2 MPLS Label

The label's Number field is 20 bits, the 3 Experimental bits are used to indicate IP precedence, the 1-bit Bottom of the Stack field indicates whether this is the last label, and the last 8-bit field indicates a Time to Live value. Multiple labels can be imposed, such as when using an MPLS VPN, MPLS traffic engineering, or a combination of the two. Each LSR only acts based on the top label.

Special Ethertypes are used in an Ethernet header to indicate that an MPLS label follows, including

- **0x8847**—Indicates a unicast labeled IP packet
- **0x8848**—Indicates a multicast labeled IP packet

Label Distribution and Label Tables

Routers send label information to each other using LDP, but they must first have a route to a network before creating a label for it. The process has four steps:

- Step 1.** Routing protocols distribute route information. The routing table and CEF FIB are built based on this information.
- Step 2.** The LSR assigns a locally significant label to each destination network. This is recorded in the LIB table. Label values 0 through 15 are reserved for special use.
- Step 3.** Each LSR sends the local label for each network to its neighbors via LDP. This is done asynchronously—a router does not wait to receive a label from its downstream router before advertising its own label. Labels are advertised to every neighbor, even the one chosen as the next hop for that network.
- Step 4.** The LSR records label information received from its neighbors in its LIB, FIB, and LFIB.

Each forwarding table is used as follows:

- The LIB lists each IP network, the local label for that network, and any labels received from neighbors for that network. It is in the control plane.
- The FIB is used to forward unlabeled IP packets. It is in the data plane.
- The LFIB lists each label, what label to swap it for, and the next-hop neighbor. It is in the data plane.

Penultimate Hop Popping

The LSR directly connected to a destination network (typically the PE router) is referred to as the *ultimate hop* for that network. The router right before it is the *penultimate hop* for that network. The directly connected LSR can advertise a label value of 3 for that destination network. Label 3 tells the neighbor router to pop the label before forwarding the packet—referred to as *penultimate hop popping* (PHP). This is recorded in the neighbor's FIBs as a null label.

PHP saves time and work for the PE LSR. Without it, the PE router would have to do an LFIB lookup, remove the top label, and then do a FIB lookup. With PHP, it only needs to do the FIB lookup, which helps optimize MPLS performance.

Figure 3-3 puts all these concepts together to show how MPLS routers use labels and each of the tables to forward traffic.

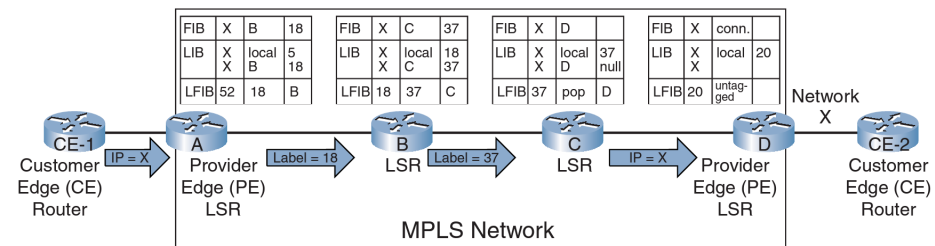


FIGURE 3-3 MPLS Packet Forwarding

In Figure 3-3, router CE-1 forwards a packet destined for network X, which is directly connected to LSR D. LSR A receives it as an IP packet, does a FIB lookup, and imposes the label of 18. This is the value that LSR B advertised to it for network X. It forwards the labeled packet to LSR B.

LSR B receives the labeled packet, does an LFIB lookup, and swaps the labels. It forwards the packet to LSR C with a label of 37, the label value advertised by C for network X.

LSR C receives the labeled packet and does an LFIB lookup. It sees that it must pop the label before forwarding the packet to the next hop, router D. When the label is popped, there are no other labels, so the packet is forwarded to LSR D as an IP packet.

The packet is addressed to the WAN interface of CE-2, so LSR D forwards it as an IP packet to CE-2.

Configuring Frame Mode MPLS

Three steps are required to configure a router to run MPLS:

- Step 1.** Enable CEF.
- Step 2.** Enable MPLS on the desired interfaces.
- Step 3.** Optionally adjust the interface maximum transmission units (MTU).

This following subsections examine each of these steps more closely.

Enabling CEF

CEF can be enabled either for all interfaces or on individual interfaces only. If it is not already enabled by default, use the global command **ip cef [distributed]** to enable it on all interface. The **distributed** keyword is used to enable distributed CEF on line cards or processors capable of running it.

To enable CEF on an individual interface, use the **ip route-cache cef** interface command. Note that CEF is not supported on logical interfaces, such as loopbacks.

Verify CEF operation with any of several variables of the **show ip cef** command. The **show ip cef detail** option is good to start with because it shows a summary of CEF routes and a list of all adjacencies.

Enabling MPLS

MPLS is globally enabled by default on Cisco routers. If it has been disabled, enable it with the global command **mpls ip**. It must then be enabled on each interface that will participate in label switching; enable it with the interface command **mpls ip**.

LDP is the default label distribution protocol in Cisco IOS Release 12.4(3) and later. There is an older Cisco prestandard version called

Tag Distribution Protocol (TDP.) You might need to enable it if you are connecting to a Cisco router that does not support LDP; if you need to specify the protocol, use the interface command **mpls label protocol** [**tdp** | **ldp** | **both**]. Both protocols can run on the same router, and even on the same interface. In some Cisco IOS versions, MPLS commands show in the running configuration as “tag-switching” commands.

Increasing the MTU Size

The MPLS Frame mode tag adds 4 bytes to every frame. You might need to increase the interface MTU size to accommodate this, to prevent packets from being fragmented. The MTU is automatically increased on WAN interfaces, but must be manually configured on LAN interfaces.

Ethernet uses a default MTU of 1500 bytes. If you are using an MPLS implementation that uses just one label, increase it to 1504 bytes. MPLS VPNs and MPLS traffic engineering (TE) use two labels, so you must increase the MTU size to 1508 bytes if you are using either of these. Increase the MTU to 1512 bytes when using MPLS VPNs with TE, to accommodate three headers.

To manually set MTU size, use the **mpls mtu bytes** command in interface configuration mode. You may also need to enable jumbo frame support on the connecting switch.

MPLS VPNs

There are two basic types of VPNs: overlay and peer to peer. In an overlay VPN, the service provider sets up the connections. Frame Relay permanent virtual circuits (PVC) are an example of an overlay VPN. The service provider does not participate in the customer’s routing when using an overlay VPN. In a peer-to-peer VPN, the service provider transports the customer’s routes across its network. Only one circuit per customer site is required, but the service provider is required to have knowledge of each customer’s routes. Customers may be required to re-IP address their networks, depending on whether the provider uses a dedicated or a shared PE router. In addition, there is no separation of customer routes.

MPLS VPNs provide the advantages of both types and minimize their drawbacks. They provide the following:

- The service provider participates in customer routing, thus providing optimum paths through the provider network.
- Each customer’s routes are kept separate from other customers’ routes.
- Overlapping IP addresses are permitted, so customers do not have to renumber.

MPLS VPNs use a two-label stack. In a traditional VPN, the IP header is hidden by a tunnel IP header. In an MPLS VPN, the label identifying the interface to the customer router is hidden by a label identifying the PE router connected to that customer. MPLS switching through the

provider network is based on the top label until it reaches the edge (or *egress*) router. The top label is popped, and the egress router reads the second label to learn where to send that traffic. PHP can be used with MPLS VPNs. If so, the PE router can use the second label to identify the VPN customer and do a route lookup based on it. Otherwise, the PE router must do two lookups.

Handling Customer Routes

MPLS PE routers use a separate virtual routing instance for each customer, called a Virtual Routing and Forwarding (VRF) table. Each customer router advertises its routes to its PE router. C routers can use a standard routing protocol to advertise their routes. The PE router looks like any other neighbor to the C router. Because VRFs must be configured on the PE router, the routing protocol needs to support them. EIGRP, OSPF, RIPv2, BGP, and static routing support VRFs.

C routes are then advertised via Multiprotocol BGP (MP-BGP) to other PE routers participating in that VRF. BGP runs only between the edge routers; internal P routers use an Interior Gateway Protocol (IGP) such as OSPF or EIGRP to tell them how to reach the PE routers. P routers have no knowledge of customer routes. Thus, PE routers do the following types of routing:

- IGP, BGP, or static routing with its customer routers to exchange IPv4 routes
- MP-BGP with its peer PE routers to exchange VPNv4 routes

- IGP with its neighboring P routers to exchange core network routes

Route Distinguishers

To support customers with overlapping IP address space, MPLS providers use route distinguishers (RD). An RD is a 64-bit prefix added to each customer's IP address to make it globally unique. The resulting 96-bit IP address is called a *VPNv4 address*.

RDs are used to propagate routes across provider networks in the following way:

- Step 1.** A C router advertises its normal IPv4 networks to the PE router.
- Step 2.** The PE router prepends the RD to the C networks to create VPNv4 addresses.
- Step 3.** The PE router advertises the VPNv4 addresses to its MP-BGP peers.
- Step 4.** Other PE routers strip the RD from the network address and advertise the route to C routers in the same VRF.

NOTE

The RD is used only within the MPLS network and only to create a globally unique address. C routers never see the RD and are not aware of their VPNv4 address.

Route Targets

Sometimes, customer sites need to participate in more than one MPLS VPN. Export and import route targets (RT) are attributes attached to BGP routes to indicate which VPNs the route belongs to. This allows the creation of complex topologies with overlapping VPNs:

- **Export RT**—Attached to routes when they are imported into the VRF database to identify the VPNs to which the route belongs. These routes are then advertised to other PE routers.
- **Import RT**—Used by the PE router receiving the routes to identify which VRFs should receive the routes. VRFs with an import RT matching the route's RT will import the route. These networks are then installed in the VRF table and advertised to the appropriate customer routers.

CHAPTER 4

CCNP ISCW Quick Reference Sheets by Denise Donohue and Jay Swan

Chapter 4

IPsec 28

IPsec Headers.....28

Authentication Methods.....29

Encryption Methods.....29

Key Management.....30

Establishing an IPsec VPN.....31

Configuring a Site-to-Site VPN
Using Cisco IOS Commands.....31

Configuring an IPsec Transform Set32

Configuring a Site-to-Site VPN Using SDM34

Monitoring and Troubleshooting IPsec VPNs.....35

Using GRE with IPsec35

Configuring a GRE Tunnel Using Cisco IOS Commands.....36

High-Availability VPNs.....37

Cisco Easy VPN39

Configuring the Cisco VPN Client41

IPsec

IP Security, or IPsec, is a set of rules for securing data communications across a public, untrusted network such as the Internet. It provides the following:

- Data confidentiality by encrypting portions of a packet
- Data integrity by ensuring the packet has not been altered in transit
- Data source authentication to ensure that the data originated with a trusted source
- Anti-replay protection to ensure that packets are not copied and sent

IPsec standards do not specify exactly how packets should be encrypted or authenticated; it relies on other protocols to accomplish those functions. For encryption, it can use Data Encryption Standard (DES), Triple Data Encryption Standard (3DES), and Advanced Encryption Standard (AES). For authentication, it can use Hash-based Message Authentication Codes (HMAC). An HMAC combines a hash function such as Message Digest 5 (MD5) and Secure Hash Algorithm 1 (SHA-1) with a shared secret key. MD5 uses a 128-bit hash, whereas SHA-1 uses a 160-bit hash. Only 96 bits of the SHA-1 hash are used with IPsec, however.

IPsec Headers

IPsec defines two types of headers: Authentication Header and Encapsulating Security Payload.

Authentication Header

Authentication Header (AH) is IP protocol number 51. It authenticates the packet, including the IP header, but does not encrypt the packet payload. AH works by creating an MD5 or SHA-1 hash from the IP header (except any changeable fields such as Time to Live) and the packet payload. It sends this hash in an AH header after the Layer 3 IP header. The receiving host also creates a hash value from the IP header and original payload and compares the two hashes. If they match, the packet was unchanged during transit. A shared key is used to create the hashes, so a match also serves to authenticate the source of the packet. AH is rarely used without ESP.

Encapsulating Security Payload

Encapsulating Security Payload (ESP), IP protocol number 50, encrypts packet payloads and can optionally authenticate and do integrity checks by using it with AH. It adds a header and a trailer to the packet. When used with AH, the packet is encrypted first and then put through the hash mechanism.

IPsec Modes

IPsec can operate in either Transport mode or Tunnel mode. The headers differ based on the mode used:

- Transport mode IPsec uses the original IP header. The data payload can be encrypted, and the packet can be authenticated from the ESP header back. Transport mode is often used with generic routing encapsulation (GRE) tunnels, because GRE hides the original IP address.
- Tunnel mode IPsec replaces the original IP header with a tunnel header. The ESP header is placed after the new header, before the original one. The original IP header can be encrypted along with the data payload, and the packet can be authenticated from the ESP header back. Tunnel mode adds about 20 bytes to the packet.

Figure 4-1 shows the packet headers in the two IPsec modes.

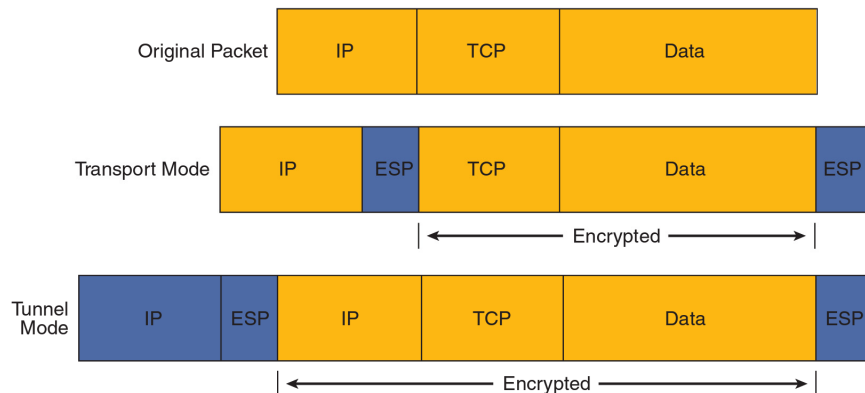


FIGURE 4-1 Transport Mode Versus Tunnel Mode IPsec

Tunnel mode ESP can cause problems when used with Network Address Translation (NAT). The original TCP or UDP header is encrypted and hidden, so there are no Layer 4 port numbers for NAT to use. *NAT Traversal* detects the existence of a NAT device and adds a UDP header after the tunnel IP header. NAT can then use the port number in that UDP header.

Authentication Methods

Several authentication methods are supported with IPsec virtual private networks (VPN):

- Username and password
- A one-time password
- Biometric features, such as fingerprint
- Preshared key values
- Digital certificates

Encryption Methods

IPsec encryption uses key values to encrypt and decrypt data. Keys can be either symmetric or asymmetric. Symmetric keys use the same value to both encrypt and decrypt the data. These include DES, 3DES, and AES. Asymmetric keys use one value to encrypt the data and another one to decrypt it. Diffie-Hellman and RSA use asymmetric keys.

NOTE

RSA is not an acronym—it is the initials of the last names of the algorithm’s inventors: Ron Rivest, Adi Shamir, and Len Adleman.

Symmetric Key Algorithms

DES uses a 56-bit key and can be broken fairly easily. It is a block cipher—it encrypts 64-bit blocks of data at a time.

3DES is also a block cipher, but it encrypts each block, decrypts it, and then encrypts it again. A 56-bit key is used each time, thus equaling a key length of 168 bits. It is more secure than DES but also requires more processing power.

AES is a stronger block cipher encryption method than DES or 3DES. It uses a 128-bit data block and a key length of 128, 192, or 256 bits. AES has been approved for use with government classified data.

Asymmetric Key Algorithm

RSA uses asymmetric keys and can be used for signing messages and encrypting them. A public key is used to encrypt or sign the data. It can only be decrypted with a private key held by the receiver. RSA is slower than symmetrical key algorithms, but more secure if a large enough key is used. A key length of 2048 bits is recommended.

Diffie-Hellman Key Exchange

The Diffie-Hellman protocol solves the problem of exchanging keys over an insecure network. Each device creates a public key and a private key. They exchange their public keys in the open, unencrypted. They each then use the other device’s public key and their own private key to generate a shared secret key that each will use.

Key Management

The Public Key Infrastructure (PKI) manages encryption and identity information such as public keys and certificates. It consists of the following components:

- Peer devices that need to communicate securely.
- Digital certificates that validate the peer’s identity and transmit their public key.
- Certificate authorities (CA), also known as *trustpoints*, that grant, manage, and revoke certificates. This could be a third-party CA or an internal one. Cisco has a Cisco IOS Certificate Server.
- Optional registration authorities (RA) that handle certificate enrollment requests.
- A way to distribute Certificate Revocation Lists (CRL), such as HTTP or Lightweight Directory Access Protocol (LDAP).

PKI credentials, such as RSA keys and digital certificates, can be stored in a router's nonvolatile random-access memory (NVRAM). They can also be stored in USB eTokens on routers that support them.

Establishing an IPsec VPN

When IPsec establishes a VPN between two peer hosts, it sets up a security association (SA) between them. SAs are unidirectional, so each bidirectional data session requires two. The Internet Security Association and Key Management Protocol (ISAKMP) defines how SAs are created and deleted. There are five basic steps:

- Step 1.** Interesting traffic arrives at the router—"Interesting" traffic is that which should be sent over the VPN. This is specified by a crypto access list. Any traffic not identified as "interesting" is sent in the clear, unprotected.
- Step 2.** Internet Key Exchange (IKE) Phase 1—Negotiates the algorithms and hashes to use, authenticates the peers, and sets up an ISAKMP SA. Has two modes: Main and Aggressive. Main mode uses three exchanges during Phase 1. Aggressive mode sends all the information in one exchange. The proposed settings are contained in transform sets, which list the proposed encryption algorithm, authentication algorithm, key length, and mode. Multiple transform sets can be specified, but both peers must have at least one matching transform set; otherwise, the session is torn down.

- Step 3.** IKE Phase 2—Uses the secure communication channel created in Phase 1 to set up the SAs for ESP/AH, negotiating the SA parameters and settings to be used to protect the data transmitted. Periodically renegotiates the SAs. SAs have lifetimes that can be measured in either amount of data transferred or length of time. May do an additional Diffie-Hellman key exchange during Phase 2.
- Step 4.** Data is transferred along the VPN between the two peers. It is encrypted by one peer and decrypted by the other, according to the transform sets negotiated.
- Step 5.** Tunnel termination—the IPsec session drops either because of direct termination or timeout.

Configuring a Site-to-Site VPN Using Cisco IOS Commands

Configuring a site-to-site IPsec VPN using Cisco IOS commands requires six steps, as follows:

- Step 1.** Configure the ISAKMP policy.
- Step 2.** Configure the IPsec transform set or sets.
- Step 3.** Configure a crypto access control list (ACL).
- Step 4.** Configure a crypto map.
- Step 5.** Apply the crypto map to the outgoing interface.
- Step 6.** Optionally configure and apply an ACL that permits only IPsec or IKE traffic.

Configuring an ISAKMP Policy

To configure an ISAKMP policy, first create the policy, and then give the parameters. These parameters might include such things as type of encryption, type of hash, type of authentication, SA lifetime, and Diffie-Hellman group. The following example shows an ISAKMP policy configuration, along with the options available with each parameter. Options will vary based on Cisco IOS version:

```
IPSEC_RTR(config)#crypto isakmp policy ?
<1-10000> Priority of protection suite
IPSEC_RTR(config)#crypto isakmp policy 1
!
IPSEC_RTR(config-isakmp)#encryption ?
3des Three key triple DES
aes AES - Advanced Encryption Standard.
des DES - Data Encryption Standard (56 bit keys).
IPSEC_RTR(config-isakmp)#encryption 3des
!
IPSEC_RTR(config-isakmp)#hash ?
md5 Message Digest 5
sha Secure Hash Standard
IPSEC_RTR(config-isakmp)#hash sha
!
IPSEC_RTR(config-isakmp)#authentication ?
pre-share Pre-Shared Key
rsa-encr Rivest-Shamir-Adleman Encryption
rsa-sig Rivest-Shamir-Adleman Signature
IPSEC_RTR(config-isakmp)#authentication pre-share
```

```
!
IPSEC_RTR(config-isakmp)#group ?
1 Diffie-Hellman group 1
2 Diffie-Hellman group 2
5 Diffie-Hellman group 5
IPSEC_RTR(config-isakmp)#group 2

IPSEC_RTR(config-isakmp)#lifetime ?
<60-86400> lifetime in seconds
IPSEC_RTR(config-isakmp)#lifetime 300
```

Configuring an IPsec Transform Set

An IPsec transform set defines how VPN data will be protected by specifying the IPsec protocols that will be used. You can specify up to four transforms and the algorithm to use with each. You can also configure either Tunnel or Transport mode (Tunnel is default). Transforms include combinations of the following:

- AH with either MD5 or SHA-1
- ESP encryption using DES, 3DES, AES, or others
- ESP authentication using MD5 or SHA-1
- Compression using the Lempel-Ziv-Stac (LZS) algorithm

The following example shows a transform set with ESP encryption and authentication. Note that these commands are all given as part of the same command:

```
IPSEC_RTR#conf t
Enter configuration commands, one per line. End with CNTL/Z.
IPSEC_RTR(config)#crypto ipsec transform-set TRANSFORM1 esp-aes
192 esp-md5-hmac
```

Configuring a Crypto ACL

You use a crypto ACL to identify traffic that should be protected by the IPsec VPN. Any traffic permitted in the ACL will be sent over the VPN. Traffic denied by the ACL will not be dropped—it will just be sent normally.

The following example shows a crypto ACL that permits traffic from two internal networks—172.16.1.0 and 172.16.4.0—if it is bound to the server network of 10.6.3.0.

NOTE

When configuring the crypto ACL on the router at the other end of the tunnel, be sure to reverse the source and destination IP addresses.

```
IPSEC_RTR(config)#access-list 172 permit ip 172.16.1.0 0.0.0.255
10.6.3.0 0.0.0.255
IPSEC_RTR(config)#access-list 172 permit ip 172.16.4.0 0.0.0.255
10.6.3.0 0.0.0.255
```

Configuring a Crypto Map

A crypto map pulls together the transform sets and crypto ACLs and associates them with a remote peer. A sequence number can be used when configuring a crypto map. Multiple crypto maps with the same name but different sequence numbers form a crypto map set. Traffic is evaluated against each crypto map depending on its sequence number to see whether it should be protected. This permits more complex and granular traffic filtering.

The following example shows a crypto map that links the transform set and ACL configured in previous examples:

```
IPSEC_RTR(config)#crypto map TO_SERVERS 10 ipsec-isakmp
% NOTE: This new crypto map will remain disabled until a peer
and a valid access list have been configured.
IPSEC_RTR(config-crypto-map)#set peer 10.1.1.1
IPSEC_RTR(config-crypto-map)#match address 172
IPSEC_RTR(config-crypto-map)#set transform-set TRANSFORM1
```

Applying the Crypto Map to an Interface

After the crypto map has been configured, it must be applied to an interface for it to take effect. It is applied at the *outgoing* interface—the one that VPN traffic will use to reach the other end of the VPN. You might need to use a static route or otherwise adjust your routing to force traffic bound for the VPN destination networks to use the correct outgoing interface.

CHAPTER 4

The following example shows the crypto map TO_SERVERS applied to interface serial 0/0/0. Note that the router replies with a message that ISAKMP is now enabled:

```
IPSEC_RTR(config)#int s0/0/0
IPSEC_RTR(config-if)#crypto map TO_SERVERS
IPSEC_RTR(config-if)#
01:19:16: %CRYPTO-6-ISA_KMP_ON_OFF: ISAKMP is ON
```

Configuring an Optional Interface Access List

You might want to have an interface ACL on the VPN interface. Typically, you would permit only IPsec-related traffic, and perhaps routing protocol traffic, in and out that interface. Keep in mind the following port numbers when configuring the ACL:

- ESP is IP protocol 50.
- AH is IP protocol 51.
- IKE uses UDP port 500.
- NAT Traversal uses UDP port 4500.

The source and destination addresses should be the IP addresses of the outgoing VPN interfaces. The following example shows an ACL that permits IPsec traffic between two hosts:

```
IPSEC_RTR(config)#access-list 101 permit ahp host 10.1.1.2 host
10.1.1.1
```

```
IPSEC_RTR(config)#access-list 101 permit esp host 10.1.1.2 host
10.1.1.1
IPSEC_RTR(config)#access-list 101 permit udp host 10.1.1.2 eq
isakmp host 10.1.1.2
IPSEC_RTR(config)#access-list 101 permit udp host 10.1.1.2 host
10.1.1.2 eq isakmp
!
IPSEC_RTR(config)#interface s 0/0/0
IPSEC_RTR(config-if)#ip address 10.1.1.2 255.255.255.252
IPSEC_RTR(config-if)#ip access-group 101 out
```

Configuring a Site-to-Site VPN Using SDM

Cisco Security Device Manager (SDM) provides a graphical user interface (GUI) for configuring and monitoring your routers. It has wizards to simplify common tasks and is designed to allow small to medium-size businesses to easily deploy their routers. It comes with 800 to 3800 series routers at no extra charge. You can use the SDM to configure site-to-site VPNs, among other things.

The VPN Wizard comes with two IKE policies and an IPsec transform set. It also has a way for you to enter information manually and edit configurations created by the wizard.

The Quick Setup Wizard requires just one screen. On it, you enter the following information:

- Outgoing interface
- Peer IP address
- Authentication information
- What traffic to encrypt

SDM then shows you a recap of the configuration. You click the **Finish** button to apply it.

You can also use the Step-by-step Setup Wizard, which leads you through each of the tasks separately, letting you have more control over the settings. It also lets you review the completed configuration before applying it.

When the configuration is complete for one side of the tunnel, you can click the **Generate Mirror** button to generate a configuration for the router on the other side of the tunnel. When both routers are configured, the **Test Tunnel** button lets you verify that it is working. There is also a tab to monitor the VPN tunnel status.

Monitoring and Troubleshooting IPsec VPNs

Some useful Cisco IOS commands for monitoring your IPsec VPNs include the following:

- **show crypto isakmp sa**—This command shows all the IKE SAs currently active on the router. Look for a status of QM_IDLE to verify that the SA is active.
- **show crypto ipsec sa**—This command shows the parameters used by each SA and shows traffic flow. Look for the count of packets being encrypted and decrypted, to verify the VPNs operation.

To troubleshoot VPN problems, first verify IP connectivity. If that exists, review your configuration one more time. If the configuration looks correct on both peers, you can view detailed information about the IKE negotiations by using the command **debug crypto isakmp**.

Using GRE with IPsec

GRE is a tunneling protocol that can support multiple Layer 3 protocols, such as IP, IPX, and AppleTalk. It also allows the use of multicast routing protocols across the tunnel. It adds a 20-byte IP header and a 4-byte GRE header, hiding the existing packet headers. The GRE header contains a Flag field, and a Protocol Type field to identify the Layer 3 protocol being transported. It may optionally contain a tunnel checksum, tunnel key, and tunnel sequence number. GRE does not encrypt traffic or use any strong security measures to protect the traffic.

GRE can be used along with IPsec to provide data source authentication, data confidentiality, and assurance of data integrity. GRE over IPsec tunnels are typically configured in a hub-and-spoke topology over an untrusted WAN to minimize the number of tunnels that each router must maintain.

Figure 4-2 shows how the GRE and IPsec headers work together.

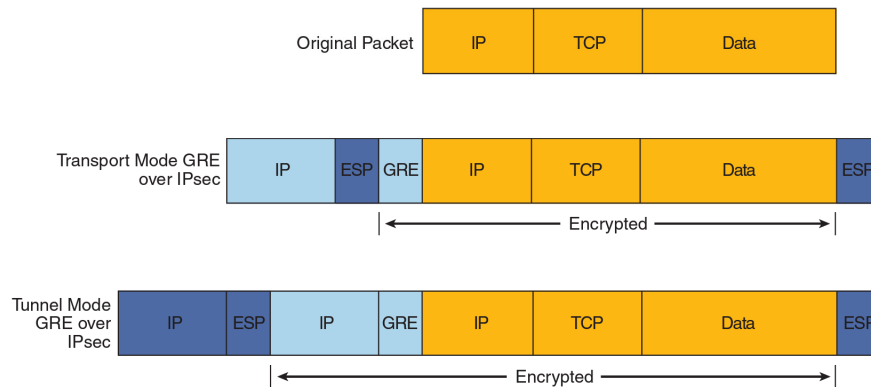


FIGURE 4-2 GRE over IPsec Headers

Configuring a GRE Tunnel Using Cisco IOS Commands

To configure GRE using Cisco IOS commands, you must first configure a logical tunnel interface. GRE commands are then given under that interface. You must specify a source and destination for the tunnel; the source is a local outgoing interface. You may also give the tunnel inter-

face an IP address and specify the Tunnel mode. GRE is the default mode.

The following example shows a tunnel interface configured for GRE. The **mode** command is shown only as a reference—because it is the default, it would not normally appear in the configuration:

```
interface Tunnel1
 ip address 172.16.5.2 255.255.255.0
 tunnel source Serial0/0
 tunnel destination 10.1.1.1
 tunnel mode gre ip
```

Configuring a GRE over IPsec Tunnel Using the SDM

The Site-to-Site VPN Wizard has an option for configuring a GRE over IPsec VPN. It leads you through creating the tunnel interface and specifying the parameters such as tunnel source and destination and interface IP address. The wizard also leads you through creating an optional backup GRE tunnel in case the primary one goes down.

Next, the IPsec parameters are presented. The wizard walks you through creating an IKE policy and transform set. You can either use the ones included with SDM or create new ones through the wizard.

You then have the option of configuring either static routing or the dynamic routing protocols Enhanced Interior Gateway Routing Protocol (EIGRP) or Open Shortest Path First (OSPF).

As before, the wizard shows a summary of the VPN configuration so that you can review it before applying it by clicking *Finish*. You can test the tunnel and monitor its status through SDM. Use the same **show crypto** and **debug** commands shown earlier to verify and troubleshoot the VPN. In addition, the **show interfaces** command shows the status of the logical tunnel interface as the physical outgoing interface.

High-Availability VPNs

Four typical types of failures that affect a VPN, and ways to mitigate them, are as follows:

- **Failure of an access link**—Use multiple links to mitigate an access link failure.
- **Failure of a remote IPsec peer**—Use multiple peers to mitigate a remote peer failure.
- **Failure of a VPN device**—Use multiple devices in critical locations to mitigate a device failure.
- **Failure someplace along the Internet path** Provide multiple independent paths to mitigate a path failure.

Routers can use a routing protocol or IPsec's Dead Peer Detection (DPD) to detect a failure across a VPN. Hot Standby Routing Protocol (HSRP), Virtual Router Redundancy Protocol (VRRP), and Gateway Load Balancing Protocol (GLBP) can detect local failures.

Detecting a Failure Using DPD

DPD sends periodic keepalives to its remote peer (similar to an older proprietary Cisco IOS method called IKE Keepalives). Periodic keepalives help routers quickly detect the failure of a remote peer, but they also mean more packets to encrypt and decrypt. DPD has an on-demand option. If the router needs to send traffic and has not received anything from the peer recently, it sends a DPD message to verify the peer's status. No messages are sent if there is no traffic.

You can configure multiple peers in a crypto map. Then, if DPD discovers that the primary peer (designated by the **default** keyword) is down, IPsec removes any SAs associated with that peer. It can then fail over to the backup peer listed in the crypto map. Configure the router to use DPD with the command **crypto isakmp keepalive seconds [retries] [periodic | on-demand]**. The **periodic** option causes the router to use keepalives, whereas the **on-demand** option causes it to use on-demand DPD. If neither option is specified, the IPsec peers negotiate the type of DPD.

Detecting a Failure Using HSRP

HSRP, VRRP, and GLBP are all protocols that allow multiple routers to share a single IP address. They are typically used for default gateway redundancy on a local LAN. You can also use these protocols on a WAN-facing interface. They use the shared, or virtual, IP address as the peer address for IPsec. Then, if the physical router fails, one of the

other routers in the group takes over. The IPsec tunnel drops and is then reestablished using the same peer IP address but a new physical router.

HSRP defines an *active* and a *standby* router, which form a *standby group*. The active router answers traffic to the virtual IP address. If it fails, the standby router takes over. HSRP routers share a virtual MAC address, too. Determining the return path to a remote site can be a problem when using HSRP at the headend. Either additionally configure HSRP on the internal interfaces of the routers, or use Reverse Path Injection (RRI). RRI injects remote networks into the internal routing protocol and is enabled with the command **reverse-route** in a crypto map.

Using IPsec Stateful Failover

DPD, routing protocols, and HSRP provide stateless failover—the tunnel drops and must be re-created if a peer fails. Stateful failover maintains SA information between an active and a failover device. It depends on two protocols:

- **HSRP**—Tracks both the inside and outside interfaces. If either goes down, it removes the router from active status, and ownership of the IPsec and IKE SAs passes to the standby router.
- **Stateful Switchover (SSO)**—Synchronizes IPsec state information between the active and standby HSRP routers using Stream Control Transmission Protocol (SCTP) and Inter-Process Communication (IPC) protocol.

For stateful failover to work, you must configure HSRP, use the virtual IP address in the IPsec peering, and configure SSO. The following example shows this configuration on one router. The configuration must be duplicated exactly on the other router; both routers must use the same Cisco IOS version and must be connected via a switch or hub:

Crypto map configuration

```
crypto dynamic-map VPN1 10
  set peer 10.1.1.1
  set transform-set TRANSFORM1
  match address 172
  reverse-route
!
crypto map HAVPN 10 ipsec-isakmp dynamic VPN1
!
```

HSRP Configuration—must be done on both inside and outside interfaces

```
interface FastEthernet0/0
  ip address 10.3.7.3 255.255.255.0
  standby 1 ip 10.3.7.1
  standby 1 preempt
  standby 1 name IPSEC1
  standby 1 track FastEthernet0/1
  crypto map HAVPN redundancy IPSEC1 stateful
!
```

SSO Configuration

```
redundancy inter-device
  scheme standby IPSEC1
!
```

IPC Configuration

```

ipc zone default
  association 1
  protocol sctp
  local-port 5555
  local-ip 10.3.7.3
  retransmit-timeout 300 10000
  path-retransmit 10
  assoc-retransmit 10
  remote-port 5555
  remote-ip 10.1.1.1

```

Using an IPsec Tunnel as a Backup WAN Link

You can configure an IPsec VPN tunnel over the Internet as a backup to your primary WAN link. Switchover can use either an IGP or HSRP. Ensure that the primary WAN link is used when it is available by either adjusting the routing protocol metric or using floating static routes.

Cisco Easy VPN

Easy VPN allows a server to push down VPN configuration to a client. It is a way to create site-to-site VPNs without manually configuring each remote router. Therefore, it is good for remote sites without technical support. It can also be used with software clients for remote users.

Cisco Easy VPN dynamically handles the following items:

- Negotiating VPN tunnel parameters
- Establishing the VPN tunnel based on those parameters
- NAT, PAT, or ACL configuration
- User authentication
- Managing encryption and decryption keys
- Authenticating, encrypting, and decrypting traffic

Cisco Easy VPN has two components: a server and a remote client. The *Easy VPN Server* can be a Cisco router, PIX firewall, or Cisco VPN concentrator. It contains security policies and pushes those to remote clients. The *Easy VPN Remote* can be a Cisco router, PIX or ASA firewall, or a hardware client or a software client. It contacts the server and receives policies from it to establish the tunnel.

Establishing an Easy VPN IPsec Session

The steps in establishing an IPsec session differ slightly when using Easy VPN. They are as follows:

- Step 1.** The remote client contacts the server and begins IKE Phase 1. If preshared keys are used, the client initiates Aggressive mode. If digital certificates are used, the client initiates Main mode.

- Step 2.** The remote client attempts to establish an ISAKMP SA with the server. It sends proposals with various combinations of hashes, authentication types, and Diffie-Hellman groups.
- Step 3.** The server finds a match to one of the client's proposals, accepts it, and establishes an ISAKMP SA. The device is now authenticated.
- Step 4.** When using Extended Authentication (XAuth), the server issues a username and password challenge. It checks the remote client's response against a RADIUS or TACACS+ authentication, authorization, and accounting (AAA) server, or uses tokens. The user is now authenticated.
- Step 5.** The server pushes configuration parameters to the device. At a minimum, this must include an IP address.
- Step 6.** RRI creates a static route to the remote client.
- Step 7.** IPsec Quick mode is used to negotiate an IPsec SA. When this is complete, the VPN is established.

Using SDM to Configure the Easy VPN Server

You can use SDM to create and configure an Easy VPN Server. Before beginning, decide on the IKE authentication method, the user authentication method, and the IP addressing scheme you will use. In addition, set up AAA servers or certification authority information, any needed DNS resolution for the Easy VPN servers, and Network Time Protocol

if needed for key exchange. The tasks to create an Easy VPN server include the following:

- Create a privileged user.
- Configure an enable secret password.
- Enable the router to use the AAA server's database.
- Use the SDM's Easy VPN Wizard to configure the following:
 - The tunnel interface
 - IKE policies
 - RADIUS or TACACS+ policy lookup
 - User authentication
 - Local group policies
 - IPsec transform set

SDM enables you to test the configuration. In addition to **crypto show** commands mentioned earlier in this chapter, the following Cisco IOS commands help in verifying and troubleshooting your configuration:

- **debug crypto isakmp**—Shows IKE messages
- **debug aaa authentication**—Shows user authentication messages
- **debug aaa authorization**—Shows messages relating to group policy configuration
- **debug radius**—Shows RADIUS messages

Configuring the Cisco VPN Client

The Cisco VPN Client runs on a computer and communicates with an Easy VPN Server to create a VPN between the user and the server. Installing and configuring the Cisco VPN Client on a remote user's computer requires the following steps:

- Step 1.** Install the Cisco VPN Client software on the computer.
- Step 2.** Create at least one client connection entry. This includes configuring the IP address of the server and preshared keys or certificates. You can configure multiple connection entries, but only one can be used at a time.
- Step 3.** Configure authentication parameters such as group authentication, mutual group authentication, or certificate authentication.
- Step 4.** Configure transparent tunneling, which allows the tunnel to work through a router or firewall doing NAT or PAT. Access to local LAN resources, and the use of IPsec over UDP or TCP, can also be enabled here.
- Step 5.** Optionally add backup VPN servers.
- Step 6.** Configure the Internet connection. This can be either through a home LAN, or through dialup using Microsoft Dial-Up Networking or a local ISP.

Chapter 5**Cisco Device Hardening 43**

Mitigating Network Attacks.....43

Disabling Unused Cisco Router
Network Services and Interfaces45Securing Cisco Router Installations
and Administrative Access.....47

Privilege Levels50

Configuring Banner Messages.....50

Role-Based CLI.....50

Cisco IOS Resilient Configuration.....51

Mitigating Threats and Attacks with Access Lists51

Securing Management and Reporting Features53

Configuring AAA on Cisco Routers.....56

Cisco Device Hardening

Because of their important role in packet forwarding, Cisco routers make attractive targets for network attacks. This chapter reviews how to harden Cisco devices against the most common types of attacks.

Mitigating Network Attacks

Before starting to harden Cisco devices against attack, you need to understand the types of attacks that are in common use today.

Cisco Self-Defending Network

The Cisco Self-Defending Network strategy consists of three interrelated components:

- **Secure connectivity**—Virtual private network (VPN) solutions, including VPN-enabled routers, VPN concentrators, and VPN-enabled firewalls
- **Threat defense**—Cisco IOS-based and appliance-based firewalls
- **Trust and identity**—Network Access Control (NAC), Cisco Secure Access Control Server (ACS), and 802.1x

Types of Network Attacks

There are three major categories of network attacks:

- **Reconnaissance attacks**—In reconnaissance attacks, the intruder attempts to gain information about a network, typically in preparation for a more aggressive attack later. Methods used in reconnaissance attacks include ping sweeps, installation of packet sniffers to gather passwords and other sensitive information, port scans to discover vulnerable services, and Internet information queries (Domain Name System [DNS] records, Internet Assigned Numbers Authority [IANA] records, search engine queries, and so on).
- **Access attacks**—In access attacks, the intruder attempts to gain unauthorized access to a network. Methods used in access attacks include password cracking or guessing, trust exploitation, port redirection, man-in-the-middle (MitM) attacks, buffer overflows, and attacks against network applications
- **Denial-of-service (DoS) attacks**—In DoS attacks, the attacker tries to deny legitimate users access to a network resource. This might involve destruction of a compromised network system (such as erasing hard drives or operating system files) or just flooding the resource with more traffic than it can process. A DoS attack that is launched from a large number of hosts simultaneously is called a distributed DoS (DDoS) attack. These types of attacks are usually the most difficult to mitigate.

Mitigating Reconnaissance Attacks

You can mitigate reconnaissance attacks in several ways:

- **Firewall and intrusion prevention system (IPS)**—A firewall (either Cisco IOS based or appliance based) is an effective way to stop ping sweeps, port scans, and other network probes. An IPS can detect and sometimes take countermeasures against these probes.
- **Authentication**—Strong authentication is an effective way to defeat password sniffers. Use of two-factor authentication such as token cards makes it extremely difficult for an attacker to gather passwords with a packet sniffer because the password hashes expire continually.
- **Cryptography**—Even with strong authentication, an attacker with a packet sniffer could still gather other sensitive information on the network. Encrypting traffic with standards-based encryption protocols prevents this.
- **Antisniffer tools**—Several manufacturers offer tools designed to detect the presence of packet sniffers on a network.
- **Switched infrastructure**—By isolating collision domains to individual ports, switches make it more difficult for packet sniffers to find sensitive information. Advanced switch security tools such as DHCP inspection and dynamic Address Resolution Protocol (ARP) inspection add to this functionality.

Mitigating Access Attacks

You can mitigate access attacks as follows:

- **Strong password security**—A surprising number of access attacks are carried out through simple password-guessing or brute-force dictionary attacks against passwords. The use of encrypted or hashed authentication protocols (for instance, Secure Shell [SSH] for terminal access, TACACS+ for authentication, authorization, and accounting [AAA]) along with a strong password policy (requiring different passwords on different systems, locking out accounts after a string of unsuccessful attempts, and complex password requirements) greatly reduce the probability of password access attacks.
- **Principle of minimum trust**—Systems should not trust one another unnecessarily. A common trust exploitation attack occurs when an inside network host trusts a device in the demilitarized zone (DMZ). If an attacker is able to compromise the DMZ system, the DMZ system can be used as a stepping-stone to access and compromise the trusted internal system. Secure network designs take this into account by ensuring that inside systems do not trust DMZ systems unconditionally.
- **Cryptography**—The MitM attack, in which an attacker inserts himself between two trusted hosts and impersonates both to gather sensitive information, can be thwarted only by using cryptography in the communications channel between the trusted hosts.

Mitigating Denial-of-Service Attacks

DoS attacks are difficult to stop. Companies with a high-profile Internet presence should plan in advance their responses to potential DoS attacks. Historically, many DoS attacks were sourced from spoofed source addresses. These types of attacks can be thwarted through use of antispoofing access lists on border routers and firewalls. Today, however, many DoS attacks are carried out by distributed networks of real hosts that have been compromised for the purpose of building attack networks. Mitigating a large DoS or DDoS attack typically requires careful diagnostics, planning, and cooperation from Internet service providers (ISP).

Disabling Unused Cisco Router Network Services and Interfaces

In any network security strategy, you need to identify and deactivate services that are on by default but not actually used. This section discusses deactivating these services on Cisco routers.

Unused Router Interfaces

Attacks can exploit active, unused router interfaces to gain access to a router or to gather information about it. Disable unused router interfaces with the **shutdown** command.

Vulnerable Router Services

The following services have the potential to be exploited by attackers under certain conditions. If they are not required on a particular router, ensure that they are disabled. Remember, however, that many of these services perform important functions in some networks; they should only be disabled after considering the potential drawbacks of doing so:

- **BOOTP server**—Enabled by default. Disable with the **no ip bootp server** command.
- **Cisco Discovery Protocol (CDP)**—Enabled by default for most interface types. Disable on interfaces where not needed with the **no cdp enable** command.
- **Configuration auto-loading**—Disabled by default.
- **FTP/TFTP servers**—Disabled by default.
- **Network Time Protocol (NTP)**—Disabled by default, but necessary for many security features.
- **Packet assembler/disassembler (PAD) Service**—Enabled by default. Disable with the **no service pad** command.
- **TCP and UDP small servers**—For example, Echo, Chargen, Discard, Daytime. Disabled by default.
- **Maintenance Operations Protocol (MOP) service**—Enabled for some Ethernet interfaces by default.

- **Simple Network Management Protocol (SNMP)**—Disabled by default, but widely used.
- **HTTP**—Enabled by default. Disable with the **no ip http server** global command if not needed.
- **DNS**—Disabled by default.
- **Internet Control Message Protocol (ICMP) redirects**—Enabled by default. Disable with the **no ip redirects** command if not needed.
- **IP source routing**—Enabled by default. Disable with the **no ip source-route** command.
- **Finger service**—Disabled by default.
- **ICMP unreachable**—Enabled by default. Disable with the **no ip unreachable** command if not needed.
- **ICMP mask reply**—Disabled by default.
- **TCP keepalives**—Disabled by default. Enable with the **service tcp-keepalives** command.
- **Proxy ARP**—Enabled by default. Disable with the **no ip proxy-arp** command if unneeded.
- **IP directed broadcasts**—Disabled by default.

Hardening with AutoSecure

AutoSecure is a feature found in Cisco IOS Release 12.3T and later that automates many of the tasks involved in hardening a router. It can be operated in either Interactive mode or in Noninteractive mode. Interactive mode prompts the user with questions regarding security features such as enabling and disabling services. Noninteractive mode automatically hardens the router according to Cisco-recommended guidelines.

AutoSecure can selectively lock down the router with the following features:

- Management plane features, including disabling unneeded services
- Forwarding plane features, such as Cisco Express Forwarding (CEF) and basic access control lists (ACLs)
- Cisco IOS Firewall services
- Login and password security
- NTP
- SSH
- TCP Intercept

Configuring AutoSecure

In Cisco IOS Release 12.3(8)T and later, AutoSecure retains the pre-lockdown configuration in Flash under the name *pre_autosec.cfg*. If AutoSecure configuration fails, you can revert to the pre-lockdown configuration with the command **configure replace flash:pre_autosec.cfg**.

AutoSecure is initiated from the command-line interface (CLI) with the **auto secure** command. The **management**, **forwarding**, **ntp**, **ssh**, **fire-wall**, and **tcp-intercept** arguments activate the features described in the preceding list. By default, the AutoSecure script interactively prompts the user with questions about each feature that it configures. To switch to Noninteractive mode, add the **no-interact** keyword to the **auto secure** command.

Security Device Manager

The Security Device Manager (SDM) is a web-based security configuration tool that runs on Cisco routers. You can use it to lock down a router in a way similar to the AutoSecure script. You can also use it to audit an existing configuration for compliance with Cisco security recommendations.

Securing Cisco Router Installations and Administrative Access

Password-Creation Rules

Cisco router passwords are subject to the following restrictions:

- 1 to 25 characters in length
- Can include any alphanumeric characters, symbols, and spaces
- Cannot have a number as the first character
- Leading spaces ignored, but subsequent spaces (including trailing spaces) not ignored

Types of Router Passwords

Many different types of passwords are used for Cisco IOS routers. The most common ones are described here:

- **Enable secret**—The enable secret controls access to privileged EXEC mode on the router. The password is stored in a nonreversible one-way Message Digest 5 (MD5) hash. If the enable

secret is present in the configuration, it overrides the enable password. To configure the enable secret password, use the **enable secret** *password* command.

- **Enable password**—The enable password controls access to privileged EXEC mode on the router if the **enable secret** command is not present. The enable password is stored in clear text in the configuration, unless the **service password-encryption** command is present. To configure the enable password, use the **enable password** *password* command.
- **Line passwords**—Access to a router's tty lines can be controlled either with AAA or with individual passwords applied to the lines. AAA configuration is discussed later. To configure individual passwords on a TTY line, use the **password** *password* command in line configuration mode. Line passwords are stored in clear text in the configuration, unless the **service password-encryption** command is present. tty lines include the console port, vty lines for Telnet/SSH access, the AUX port, as well as regular tty lines. The **login** command must also be present in the line configuration for password prompts to be displayed.

Password-Length Enforcement

You can globally set a minimum length for all router passwords with the **security passwords min-length** *length* command.

Password Encryption

Many types of passwords are stored in clear text in the configuration, by default. To prevent casual discovery of these passwords, use the **service password-encryption** command. Software to decrypt passwords that are ciphered with this command is widely available on the Internet, so it should not be relied on to protect highly sensitive passwords.

Enhanced Username Password Security

When configuring a local username/password database on a router, use the **username** *username* **secret** *password* command. This protects the password with a MD5 hash rather than plain text or weak Type 7 encryption.

Password Example

Example 5-1 shows a configuration that incorporates all the password features previously discussed.

CHAPTER 5

Example 5-1 Configuring Router Passwords

```

service password-encryption

security passwords min-length 6
enable secret 5 $1$jS1$QbQmCuRx0Uf0GiMkxbHk0
enable password 7 053B0A2F70427A5A0111

username admin secret 5 $1$zltV$9tNP0xX4ehfQ0pKceNG6A/

line con 0
  password 7 053B0A2F70427A5A0111
  login
  stopbits 1
line aux 0
  password 7 01230A240A05325C3958
  login
  stopbits 1
line vty 0 4
  password 7 046B07265E2F781D110D
  login

```

Securing ROMMON

By default, Cisco routers allow a user connected to the console port to execute a keyboard break sequence to enter the ROM Monitor (ROMMON). From ROMMON, it is possible to perform a password override sequence to reset the enable secret while retaining the configu-

ration file. To prevent unauthorized users with physical access to the router from doing this, you can use the **no service password-recovery** command. After this command has been configured, it is impossible to reset any router passwords without completely erasing the configuration.

Rate-Limiting Authentication Attempts

Cisco IOS commands offer several ways to rate-limit authentication attempts:

- The **security authentication failure rate threshold-rate log** command enables you to set a number of failures after which a 15-second delay is imposed and a syslog message triggered.
- The **login block-for seconds attempts tries within seconds** command enables you to block login attempts for *seconds* if the number of login attempts exceeds *tries* within *seconds*. You can exclude a list of addresses from blocking by configuring the **login quiet-mode access-class {acl-name | acl-number}** command.
- The **login delay seconds** command enforces a minimum delay of *seconds* between successive login attempts. This helps mitigate dictionary attacks against the router.

Setting Timeouts

CLI sessions stay logged in for 10 minutes by default. You can change this with the **exec-timeout** *minutes seconds* command, in line configuration mode.

Privilege Levels

Cisco routers enable you to define up to 16 privilege levels with different command sets assigned to each. Normal user-level privileges are level 0. Standard privileged mode (access to all commands) is level 15. Example 5-2 shows how to assign the **traceroute** command to level 2, with a separate enable secret of cat. After applying this configuration, users at level 0 can no longer execute traceroutes.

Example 5-2 Configuring Privilege Levels

```
R2(config)#privilege exec level 2 traceroute
R2(config)#enable secret level 2 cat
```

Configuring Banner Messages

Example 5-3 shows how to configure a banner message that displays every time a user logs in to the router CLI. Note that the character used to begin and end the message is the percentage sign (%) in the example, but it can be any character that does not appear in the text of the message. The other type of banner commonly used in Cisco IOS

configurations is the “message-of-the-day” banner, which is configured with the **banner motd** command.

Example 5-3 Configuring a Login Banner

```
R2(config)#banner login %
Enter TEXT message. End with the character %.
You are entering the Test Network. Unauthorized access is prohibited.
All activity is logged.
%
R2(config)#
```

Role-Based CLI

Role-based CLI is a relatively new method of limiting access to CLI commands in a much more flexible way than privilege levels. Cisco IOS commands are grouped into “views,” which can then be assigned to users (or interfaces) in a variety of ways. The “root view” has access to all commands and can be used to create up to 15 additional views. Views can be offloaded to AAA servers for even more flexibility. “Superviews” enables you to group together multiple views, allowing you to assign multiple views to users with less configuration complexity. Example 5-4 shows a basic role-based CLI configuration similar to Example 5-2.

Example 5-4 Configuring Role-Based CLI

```

R2(config)#aaa new-model
R2(config)#exit
R2#enable view
Password:[enter level 15 password]
*Dec 16 19:44:39.411: %PARSER-6-VIEW_SWITCH: successfully set to
view 'root'.
R2#conf t
R2(config)#parser view TRACEROUTE_VIEW
*Dec 16 19:45:16.403: %PARSER-6-VIEW_CREATED: view
'TRACEROUTE_VIEW' successfully created.
R2(config-view)#password 5 cat
R2(config-view)#commands exec include traceroute
R2(config-view)#exit

```

You can verify role-based CLI configuration with command **show parser view [all]** command.

Cisco IOS Resilient Configuration

The Cisco IOS Resilient Configuration feature allows for faster recovery in situations in which an attacker has compromised a router and erased its Cisco IOS image / configuration file. This feature is available only on platforms with PCMCIA ATA Flash drives. When enabled, this feature saves nonerasable copies of the running Cisco IOS image and

running configuration to the Flash drive. If the configuration / Cisco IOS image is then erased, the secure backup copy of the Cisco IOS image can be booted from ROMMON. After the boot, you can restore the secure backup copy of the configuration from the Flash drive.

Mitigating Threats and Attacks with Access Lists

ACL Review

Standard access lists (standard ACLs) allow filtering based on source IP address only. Extended access lists (extended ACLs) allow filtering based on source or destination address and most other fields in the IP packet header (Layer 4 protocol type, source/destination port number, IP options, Differentiated Services Code Point [DSCP] values, fragmentation parameters, and so on).

Access lists can be either numbered or named. Numbered standard ACLs have numbers from 1 to 99 or 1300 to 1999. Numbered extended ACLs have numbers from 100 to 199 or 2000 to 2699.

You can apply access lists either inbound or outbound on an interface. Inbound ACLs affect traffic moving toward the interface. Outbound ACLs affect traffic leaving the interface.

Access lists are used extensively in router security configurations for permitting or denying access to services, mitigating address spoofing, mitigating various attack types, and more.

Mitigating Spoofed Addresses (Inbound)

You can use access lists to prevent packets with spoofed source addresses from entering your network. When configuring inbound antispoof ACLs, you should deny packets from, at a minimum, the following:

- Any internal address space
- Internal loopback addresses
- RFC 1918 reserved addresses
- Multicast addresses

Mitigating Spoofed Addresses (Outbound)

In addition to dropping inbound packets with spoofed source addresses, you should also configure ACLs to prevent packets from leaving your network with spoofed source addresses. No packets should leave your network that do not have source addresses inside your network.

Mitigating SYN Attacks

One common type of network attack is the half-open SYN attack. This is a DoS attack in which the attacker sends a large quantity of TCP SYN messages to a host without ever completing the three-way TCP handshake. This attack can result in the depletion of memory resources on the host. The most flexible way to mitigate this attack is to use the Cisco IOS Firewall feature set. The following subsections identify two other ways to mitigate half-open SYN attacks.

Using the **established** Keyword in ACLs

The **established** keyword in a TCP-based ACL entry permits only packets that have the TCP ACK bit set to pass the ACL entry. Example 5-5 demonstrates this.

Example 5-5 Using the **established** Keyword

```
R2(config)#access-list 150 permit tcp any any established
R2(config)#interface serial 1/0/0
R2(config-if)#ip access-group 150 in
```

Using TCP Intercept

The TCP Intercept feature permits half-open SYN connections only within configurable thresholds. Half-open SYN connections outside these thresholds are dropped. Example 5-6 demonstrates this.

Example 5-5 Using TCP Intercept

```
R2(config)#ip tcp intercept list 150
R2(config)#access-list 150 permit tcp any 10.1.1.0 255.255.255.0
R2(config)#interface serial 2/0/0
R2(config-if)#ip access-group 150 in
```

ACL Caveats

Remember the following caveats when configuring ACLs:

- **Implicit deny any**—All ACLs have an implicit **deny any** statement at the end. It is not displayed in the configuration. Any traffic not explicitly permitted is implicitly denied.
- **Evaluation order**—ACLs are evaluated from the top down, in order. Be sure not to place a statement at the top of the ACL that negates a later statement. Place the most specific statements at the top of the ACL.
- **ACL direction**—Inbound ACLs affect packets that are moving toward the interface. Outbound ACLs affect packets that are moving away from the interface. It can be easy to confuse these, especially on VLAN interfaces.

Securing Management and Reporting Features

In addition to securing the router itself, it is also important to secure the traffic used to manage the device and collect statistical information from it.

Types of Management Traffic

In-band management traffic flows inside the production network and is intermixed with production traffic. Although common in most networks, the risk of in-band management is that an attacker who compromises a system on the production network could interfere with management traffic, capture sensitive information from management packets, or mount further attacks against network management protocols. With in-band management, you should use encrypted protocols such as IPsec, SSH, or Secure Sockets Layer (SSL) rather than clear text protocols such as Telnet.

Out-of-band management traffic flows on an independent, purpose-built network and is kept totally separate from production traffic. Completely out-of-band management networks are most common in large networks.

With both types of management, you should ensure that the managed devices have synchronized clocks and that configuration archives and change logs are available.

Configuring Secure Shell

Traditionally, network administrators have used Telnet to manage routers and switches. The problem with Telnet, of course, is that it sends all traffic in clear text, allowing attackers with access to the network to sniff passwords and other sensitive information. You should use the encrypted SSH protocol to manage network devices wherever possible. To configure SSH, complete the following steps:

- Step 1.** Configure the domain name.
- Step 2.** Generate RSA keys.
- Step 3.** Optionally configure an SSH timeout interval and retry count.
- Step 4.** Disable Telnet.
- Step 5.** Enable SSH.

Example 5-6 demonstrates SSH configuration.

Example 5-6 Configuring SSH

```
R2(config)#ip domain-name test.com
R2(config)#crypto key generate rsa general-keys modulus 1024
The name for the keys will be: R2.test.com
% The key modulus size is 1024 bits
% Generating 1024 bit RSA keys, keys will be non-
  exportable...[OK]
*Dec 18 19:16:26.275: %SSH-5-ENABLED: SSH 1.99 has been enabled
R2(config)#ip ssh time-out 10
```

```
R2(config)#ip ssh authentication-retries 3
R2(config-line)#line vty 0 4
R2(config-line)#transport input none
R2(config-line)#transport input ssh
```

Configuring Syslog

Syslog is the most common way to archive router events for security monitoring. A router acts as a syslog client that sends messages to a syslog server. Syslog severities range from 0 (critical) to 7 (debugging). Example 5-7 demonstrates how to configure syslog.

Example 5-7 Configuring Syslog

```
R1(config)#logging 10.1.2.1
R1(config)#logging trap informational
R1(config)#logging source-interface loopback 0
R1(config)#logging on
```

Simple Network Management Protocol

SNMP is widely used to gather information from network nodes. SNMP is an application layer protocol that runs on top of TCP/IP, typically on UDP port 161. SNMP is most often used in read-only mode, in which information is read from the node, but no changes can be made. SNMP also supports read-write mode, which allows changes to be made to the node's configuration. SNMP information can be read

passively or sent based on triggered events. When used passively, a network management host reads the SNMP Management Information Bases (MIB) on the router to gather information on a periodic basis. The router can also send event-triggered SNMP traps to a network management host when a particular event occurs.

SNMP exists in Versions 1, 2, and 3. SNMPv1 and SNMPv2 lack strong security mechanisms. Read-only or read-write access is controlled via a community string, which is sent across the network in clear text. Using SNMP read-write with clear text community strings is particularly dangerous.

SNMPv3 supports strong security by enabling the use of MD5 or SHA hashed authentication and DES encryption with SNMP messages.

Example 5-8 shows a basic SNMPv3 configuration allowing a management host to read MIBs on the router.

Example 5-8 Configuring SNMPv3

```
R2(config)#snmp-server group SNMP_GROUP v3 auth
R2(config)#snmp-server user SNMP_USER SNMP_GROUP v3 auth md5
my_password
```

Network Time Protocol

NTP is used to synchronize device clocks in the network. Clock synchronization is important for correlating syslog messages and other

security features such as certificate-based encryption, routing protocol authentication key expiration, time-based ACLs, and more. NTP runs over UDP port 123. Time is tracked internally using universal coordinated time (UTC). You can configure a time zone on the router to display the correct local time. Cisco routers allow you to configure NTP to act as either a peer association or a server association. In a peer association, the local system is able to either synchronize to the other system, or the other system can synchronize to it. In a server association, the local system can *only* synchronize to the remote system.

Because time synchronization is a security-related feature, it is wise to configure a router to authenticate NTP information coming from a peer or server. This prevents an attacker from spoofing NTP packets to corrupt the system clock. For added security, you can use an ACL to restrict the IP address(es) with which the router can synchronize time.

Example 5-9 demonstrates configuration of an authenticated NTP server with an NTP ACL. In this example, the router is only allowed to synchronize with a server at 10.1.1.1 that shares the MD5 hashed key value my_ntp_key.

Example 5-9 Configuring Authenticated NTP

```
R2(config)#ntp authenticate
R2(config)#ntp authentication-key 1 md5 my_ntp_key
R2(config)#ntp trusted-key 1
R2(config)#ntp server 10.1.1.1 key 1
R2(config)#access-list 1 permit host 10.1.1.1
R2(config)#ntp access-group peer 1
```

Configuring AAA on Cisco Routers

AAA stands for authentication, authorization, and accounting:

- **Authentication**—Proves the identity of the person logging in to the router via a username and password, token card, biometrics, and so forth
- **Authorization**—Decides what actions the user is allowed to perform, such as accessing privileged mode or running certain commands
- **Accounting**—Maintains records of what the user did during the session, such as login/logout times and commands executed

AAA Services

Cisco routers support AAA either through local databases (using the **username/password** command) or through external security servers. External security servers can use one of two protocols:

- **TACACS+**—Runs over TCP port 49. Includes authentication and encryption of messages between the client and server.
- **RADIUS**—Widely supported, standardized in RFC 2865. Cisco allows the use of proprietary TACACS+ attributes via a vendor-specific attribute (VSA). Runs over UDP. Does not encrypt entire message; passwords are sent as an MD5 hash, but the rest of the message is sent in clear text.

Router Access Modes

You can use AAA in either character mode or packet mode. Character mode is used when logging in to the CLI on the router via a vty or tty line, the AUX port, or the console port. Packet mode is used when authenticating a user on a dialup or serial interface (for example, a PPP-authenticated ISDN dialup session).

Configuring AAA

Example 5-10 shows how to configure communications with the AAA security server using TACACS+. In this example, the TACACS+ server is located at 10.2.2.2 and is configured to use a single TCP socket for all connections, rather than a separate socket for each. This saves processing resources on both the router and server. The server must also be configured with the key T@C_key1.

Example 5-10 Configuring a TACACS+ Server

```
R2(config)#aaa new-model
R2(config)#tacacs-server host 10.2.2.2 single-connection
R2(config)#tacacs-server key T@C_key1
```

Example 5-11 shows how to configure communications with the AAA security server using RADIUS:

Example 5-11 Configuring a RADIUS Server

```
R2(config)#aaa new-model
R2(config)#radius-server host 10.3.3.3
R2(config)#radius-server key R@D_key1
```

Configuring CLI Authentication on a Cisco Router

Example 5-12 shows how to configure character mode AAA to authenticate a CLI session on a router's console and vty ports. In this example, a user on the console port will be authenticated using the AAA list called CUSTOM_LIST. The user will be authenticated against the TACACS+ server if it is available. If it is unavailable, the enable secret or enable password will be accepted instead. A user on one of the vty lines, on the other hand, will be authenticated using the default list. The default list authenticates first against the TACACS+ server (if it is available). If the server is unavailable, the vty user will be authenticated against the local username/password database. If a security server is the sole authentication method, you could get locked out of the router in the event that the security server is unavailable. For this reason, it is important to use either local authentication or enable password authentication as a fallback method.

Example 5-12 Configuring AAA CLI Authentication

```
R2(config)#aaa new-model
R2(config)#aaa authentication login default group tacacs+ local
R2(config)#aaa authentication login CUSTOM_LIST group tacacs+
enable
R2(config)#line con 0
R2(config-line)#login authentication CUSTOM_LIST
R2(config-line)#exit
R2(config)#line vty 0 4
R2(config-line)#login authentication default
```

Configuring Authorization

Example 5-13 shows how to configure basic command authorization using a TACACS+ security server. In this example, the user must be authorized to use the EXEC shell, and access to commands at privilege levels 1 and 15 must also be authorized. Authorization is performed against the TACACS+ server if it is available, and against the local username/password database if it is unavailable.

Example 5-13 Configuring AAA Authorization

```
R2(config)#aaa new-model
R2(config)#aaa authorization exec default group tacacs+ local
R2(config)#aaa authorization commands 1 default group tacacs+
local
R2(config)#aaa authorization commands 15 default group
tacacs+ local
```

Configuring Accounting

Example 5-14 shows how to configure basic accounting using a TACACS+ security server. In this example, accounting of EXEC commands is sent to the TACACS+ server.

Example 5-14 Configuring AAA Accounting

```
R2(config)#aaa new-model  
R2(config)#aaa accounting exec default start-stop group tacacs+
```

Troubleshooting AAA

The following commands are the most useful for troubleshooting AAA:

- **debug aaa authentication**
- **debug aaa authorization**
- **debug aaa accounting**

Chapter 6
Cisco IOS Threat Defenses 60

DMZ Design Review60

Firewall Technologies60

Cisco IOS Firewall.....61

Configuring Cisco IOS Firewalls62

Introducing Cisco IOS IPS64

Cisco IOS Threat Defenses

The router-hardening techniques discussed in Chapter 5, “Cisco Device Hardening,” help to protect the router against many types of infrastructure attacks. The Cisco IOS Firewall feature set enables you to integrate a stateful firewall and an intrusion prevention system (IPS) to protect end stations located behind the router.

DMZ Design Review

A demilitarized zone (DMZ) is an intermediate network between an organization’s “inside” network and the “outside” world. Most organizations use a DMZ to host their Internet-accessible devices, such as web servers or mail servers. Some type of security system (for example, stateful firewall, filtering router, application layer gateway) filters packets traveling between the outside world and the systems in the DMZ, and between the DMZ and the inside network. Depending on the design, there can be one filtering device that performs both functions or two separate devices.

Traffic initiated from the outside world should be filtered so that all traffic to nonessential services is dropped. If possible, the systems in the DMZ should not be allowed to initiate conversations with systems on the inside; all communications between the inside and the DMZ should be initiated from the inside. This reduces the probability of a trust exploitation attack in the event that an attacker compromises a DMZ system.

Firewall Technologies

A variety of firewall technologies exist:

- **Packet filtering**—A packet filter (such as an access list on a router) permits or denies packets based on information in the Layer 3 or Layer 4 packet headers.
- **Application layer gateway**—An application layer gateway (ALG) is a piece of software that intercepts application layer requests between the endpoints of a network conversation. The ALG typically passes requests from a client to a server and vice versa after inspecting the application layer packets to ensure that they pass configured security criteria. In some circumstances, the ALG may change the contents of packets moving in either direction.
- **Stateful packet filtering**—A stateful packet filter combines aspects of a packet filter and an ALG. The attributes of each communications session are maintained in a state table. Only

packets whose attributes match the rules of the state table are permitted to pass. For example, an HTTP response packet would typically only be allowed to pass if it is a response to a query packet that was previously permitted by the firewall. Modern stateful packet filters are also capable of tracking complex information about application sessions. For example, a Voice over IP (VoIP)-aware stateful firewall would typically be able to “know” that it should dynamically open UDP ports for a Real-time Transport Protocol (RTP) session based on information gained from examining the call setup traffic that takes place inside various call-control protocols.

Cisco IOS Firewall

The Cisco IOS Firewall is a stateful packet filter that is built in to Cisco IOS security images. Some of its features include the ability to dynamically alter router access control lists (ACL) to permit return traffic for sessions originated on the inside, the ability to track TCP sequence numbers and permit only expected TCP traffic, and the ability to mitigate some types of TCP-based and IP fragmentation-based denial-of-service (DoS) attacks.

TCP Handling in the Cisco IOS Firewall

When a router running the Cisco IOS Firewall detects an outbound TCP packet, it tracks the source and destination IP addresses, source

and destination TCP port numbers, the TCP flags, and the SYN/ACK numbers associated with the session. Only inbound packets whose packet headers match the expected parameters for a legitimate response to the session are permitted.

UDP Handling in the Cisco IOS Firewall

Because UDP packets do not have the same kind of state information as TCP packets (that is, there are no TCP flags or SYN/ACK numbers), UDP return packets are permitted based on matching source/destination IP addresses and port numbers and a configurable timeout interval. If the UDP return packet arrives outside the timeout window, or with unexpected packet headers, it is dropped.

Alerts and Audit Trails

The Cisco IOS Firewall can trigger, based on configurable parameters, syslog alerts and log audit information about firewall sessions to a syslog server.

Cisco IOS Authentication Proxy

The Cisco IOS Firewall can authenticate HTTP, HTTPS, Telnet, and FTP sessions against local username/password databases or against TACACS+ or RADIUS security servers. Therefore, an administrator can define specific access policies for each user rather than generic policies for entire subnets or interfaces.

Configuring Cisco IOS Firewalls

To configure the Cisco IOS Firewall, follow these five steps:

- Step 1.** Define external and internal interfaces.
- Step 2.** Configure access lists on the interfaces.
- Step 3.** Define inspection rules.
- Step 4.** Apply inspection rules to interfaces.
- Step 5.** Test and verify the configuration.

Defining External and Internal Interfaces

The external interface is the one connected to the “outside” network, whereas the internal interface is the one connected to the inside, protected network. For example, the external address might be connected to an Internet service provider (ISP), whereas the internal interface might be connected to your corporate LAN. Traffic arriving on the external interface is considered less trusted than traffic arriving on the internal interface. The most common type of firewall configuration is to allow outside traffic to pass the external interface only if it is a response to a legitimate session that was originated from the inside.

Configuring Access Lists on the Interfaces

Consider the following guidelines when configuring ACLs in association with the Cisco IOS Firewall:

- Extended ACLs (as opposed to standard ACLs) are required if you want to dynamically allow return traffic for sessions originated from the inside.
- Consider implementing antispoofing ACLs, as discussed in Chapter 5.
- If you want to enable application layer inspection for a protocol that is permitted through the firewall, that protocol must also be permitted by the relevant extended ACLs. For example, if you want to perform H.323 inspection, your extended ACLs must permit H.323.

Defining Inspection Rules

Inspection rules determine which application layer protocols are inspected at the firewall interface. Typically, only one inspection rule is defined, and all the protocols you want to inspect are added to it. The exception to this scenario is where you want to inspect different protocols in different directions. You define inspection rules with the **ip inspect** command. Example 6-1 demonstrates how to configure basic protocol inspection.

CHAPTER 6

Example 6-1 Basic Protocol Inspection

```
R2(config)#ip inspect name FW tcp alert on audit-trail on
timeout 300
R2(config)#ip inspect name FW ftp alert on audit-trail on
timeout 300
R2(config)#ip inspect name FW h323 alert on audit-trail on
timeout 300
R2(config)#ip inspect name FW udp alert on audit-trail on
timeout 300
```

Example 6-1 defines an inspection rule named FW with four protocols that will be inspected: generic TCP, FTP, H.323, and generic UDP. When a packet initiated from the inside interface exits the router, the inspection rule allows replies to that session to pass through the external interface's ACL, provided that the reply packet does not violate any parameters of the protocol. The **alert** and **audit-trail** keywords configure syslog alerting and auditing for the protocol. The **timeout** keyword sets the period (in seconds) after which the dynamic "hole" in the external ACL will be closed if there is no activity.

The **alert** and **audit-trail** keywords produce log messages only if the global commands **ip inspect audit-trail** and **no ip inspect alert-off** are also configured. The following output shows sample representative sample messages that the router sends when the alert and audit-trail features are active:

```
%FW-6-SESS_AUDIT_TRAIL: tcp session initiator (10.1.1.2:5590)
sent 22 bytes — responder (10.1.1.3:23) sent 88 bytes
%FW-4-ALERT_ON: getting aggressive, count (550/500) current
1-min rate: 250
%FW-4-ALERT_OFF: calming down, count (0/400) current 1-min
rate: 0
```

Applying Inspection Rules to Interfaces

After you have created an inspection rule, you must apply it to an interface. The most common configuration is to have the inspection rule applied inbound on the inside interface. This configuration allows the router to dynamically create holes in the ACLs applied to other interfaces that allow replies to sessions initiated by hosts on the inside interface. Example 6-2 demonstrate this.

Example 6-2 Applying Inspection Rule Inbound on Inside Interface

```
interface FastEthernet0/0
description inside interface
ip address 10.1.1.1 255.255.255.0
ip inspect FW in
```

Verifying Inspection

The following commands are useful for verifying Cisco IOS Firewall configurations:

- **show ip inspect** *inspection-name*
- **show ip inspect config**
- **show ip inspect interfaces**
- **show ip inspect session [detail]**
- **show ip inspect statistics**

- **show ip inspect all**
- **debug ip inspect detail**
- **debug ip inspect events**
- **debug ip inspect** *protocol*

Introducing Cisco IOS IPS

The Cisco IOS Intrusion Prevention System can help detect and mitigate attacks against routers and hosts. This section reviews the components and configuration of the IOS IPS.

Defining IDS/IPS Terms

The following terms are important to an understanding of intrusion detection systems (IDS) and intrusion prevention systems (IPS):

- **Intrusion detection system**—An IDS is a device that listens passively to network traffic and produces alerts when suspicious activity is detected. An IDS is often located outside the traffic forwarding path and monitors traffic that is copied to a Switched Port Analyzer (SPAN) port on a switch.
- **Intrusion prevention system**—An IPS is a device that not only alerts on suspicious activity, but that can also be configured to actively block it. An IPS is typically located in the forwarding path. The Cisco IOS IPS is a feature offered in Cisco IOS security images that allows the router to detect and respond to possible network attacks.
- **Signature-based approach**—Signature-based IDS/IPS devices detect possible attacks by matching preconfigured patterns (that is, “signatures”) in network traffic.
- **Policy-based approach**—Policy-based IDS/IPS devices detect attacks based on thresholds or other policies, such as a number of half-open TCP SYN sessions.
- **Anomaly-based approach**—Anomaly-based IDS/IPS devices profile network traffic and build up a set of patterns that is considered “normal.” Traffic that falls outside of normal parameters triggers alerts or other actions.
- **Honeypot approach**—“Honeypots” are systems that are deliberately left vulnerable to network attacks so that security researchers can analyze an attack methodology. The network design must prevent a compromised honeypot from ever having access to legitimate systems.
- **Host-based IDS/IPS**—A host-based IDS/IPS (HIDS/HIPS) resides on end-system hosts. It is typically written to prevent attacks against a particular operating system, such as the installation of unauthorized software.
- **Network-based IDS/IPS**—A network-based IDS/IPS (NIDS/NIPS) resides on the transport network. It may be a passive IDS located on a switch SPAN port or an active IPS colocated on a firewall or router.

Cisco IOS IPS Signatures

A Cisco router running the IPS module comes with 100 attack signatures preloaded in the Cisco IOS Software. Many additional signatures can be loaded by installing Signature Definition Files (SDF) on the router.

Cisco IOS IPS Alarms

When an attack signature is detected, the router can take any of the following configurable actions:

- Send an alarm
- Drop the packet
- Reset the TCP connection
- Block the source IP address of the packet for a configurable amount of time
- Block the connection for a configurable amount of time

Configuring Cisco IOS IPS

Example 6-3 demonstrates how to configure the most common Cisco IOS IPS features.

Example 6-3 Configuring Cisco IOS IPS

```
ip ips sdf location flash:sig.sdf
ip ips signature 1107 0 disable
ip ips signature 6190 0 list 199
ip ips name MY_IPS list 100
!
interface serial 1/0
 ip ips MY_IPS in
!
access-list 100 deny ip host 10.1.1.1 any
access-list 100 permit ip any any
!
access-list 199 deny ip host 172.16.1.1 any
access-list 199 permit ip any any
```

The commands in Example 6-3 function as follows:

- **ip ips sdf location**—Specifies the location of the signature definition file.
- **ip ips signature 1107 0 disable**—Disables signature 1107, subsignature 0.
- **ip ips signature 6190 0 list 199**—Specifies that signature 6190, subsignature 0 will be filtered against access list 199. Packets matching a **deny** statement in the ACL bypass the IPS engine, whereas packets matching a permit statement are scanned with the IPS engine.

- **ip ips name MY_IPS list 100**—Creates an IPS rule named MY_IPS and filters it against access list 100. Packets matching a **deny** statement in the ACL bypass the IPS engine, whereas packets matching a permit statement are scanned with the IPS engine.
- **ip ips MY_IPS in**—Specifies that packets inbound to the interface are scanned with the IPS rule MY_IPS.

CCNP ISCW Quick Reference Sheets

Denise Donohue, CCIE No. 9566

Jay Swan

Copyright© 2007 Cisco Systems, Inc.

Published by:
Cisco Press
800 East 96th Street
Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

First Digital Edition February 2007

ISBN10: 1-58705-314-4

ISBN13: 978-1-58705-314-6

Warning and Disclaimer

This Short Cut is designed to provide information about the CCNP ISCW exam. Every effort has been made to make this Short Cut as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The authors, Cisco Press, and Cisco Systems, Inc. shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this Short Cut or from the use of the discs or programs that may accompany it.

The opinions expressed in this Short Cut belong to the author and are not necessarily those of Cisco Systems, Inc.

Trademark Acknowledgments

All terms mentioned in this Short Cut that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc. cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Feedback Information

At Cisco Press, our goal is to create Short Cuts of the highest quality and value. Each Short Cut is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this Short Cut, or otherwise alter it to better suit your needs, you can contact us through email at feedback@ciscopress.com. Please make sure to include the Short Cut title and ISBN in your message.

We greatly appreciate your assistance.

Corporate and Government Sales

Cisco Press offers excellent discounts on this Short Cut when ordered in quantity for bulk purchases or special sales.

For more information please contact: U.S. Corporate and Government Sales 1-800-382-3419
corpsales@pearsontechgroup.com

For sales outside the U.S. please contact: International Sales international@pearsoned.com



Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 528-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
166 Robinson Road
#29-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

©2007 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, IQ Expertise, the IQ logo, IQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0609R)