CSI

Cisco SAFE Implementation

Student Guide

Version 2.0

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

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Table of Contents

1

Course Introduction

Overview

This lesson includes the following topics:

- Course objectives
- Course agenda
- Participant responsibilities
- General administration
- Graphic symbols
- Participant introductions
- Cisco security career certifications
- Lab topology overview

Course Objectives

This topic introduces the course and the course objectives.

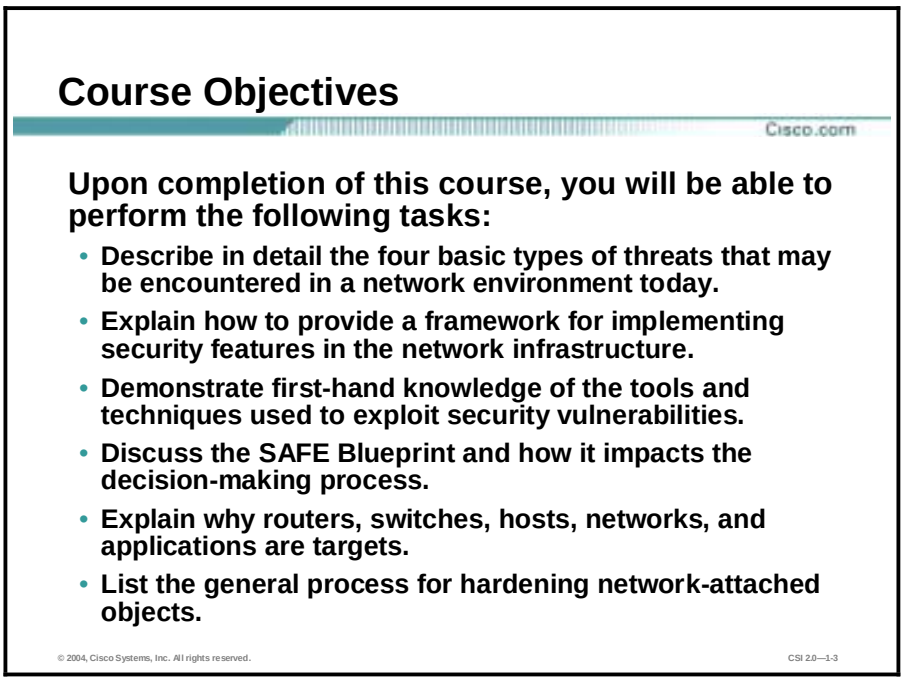

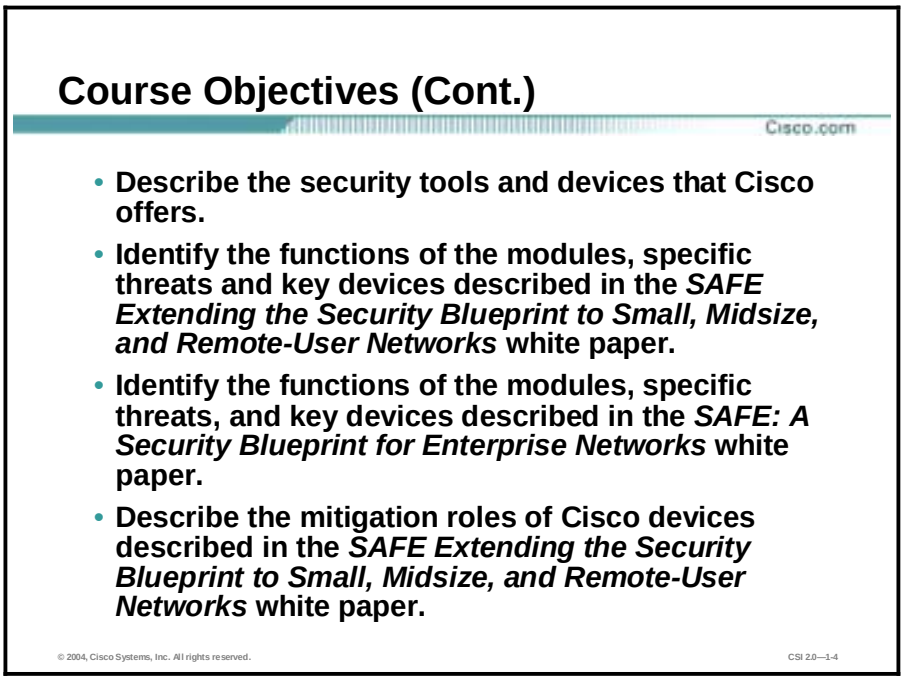

Course Objectives (Cont.)

Cisco.com

- Implement specific configurations to apply the mitigation roles described in the *SAFE Extending the Security Blueprint to Small, Midsize, and Remote-User Networks* white paper.
- Recommend alternative devices that can fulfill the same mitigation roles described in the *SAFE Extending the Security Blueprint to Small, Midsize, and Remote-User Networks* white paper.

Course Objectives (Cont.)

Cisco.com

- Discuss the technologies and blueprint involved in building a SAFE IP telephony network.
- Identify the functions of the modules, specific threats, and key devices described in the *SAFE: IP Telephony Security in Depth* white paper.
- Describe the mitigation roles of Cisco devices described in the *SAFE: IP Telephony Security in Depth* white paper.
- Discuss the technologies and blueprint involved in building a SAFE Wireless LAN.
- Identify the functions of the modules, specific threats, and key devices described in the *SAFE : Wireless LAN Security in Depth* white paper.
- Describe the mitigation roles of Cisco devices described in the *SAFE: Wireless LAN Security in Depth* white paper.

Course Agenda

Cisco.com

Day 1

- Lesson 1—Course Introduction
- Lesson 2—Security Fundamentals
- Lunch
- Lab—Vulnerabilities and Threats
- Lesson 3—SAFE Blueprint Overview
- Lesson 4—The Cisco Security Portfolio
- Lesson 5—SAFE Small Network Design

Day 2

- Lab—SAFE Small Network Design Implementation
- Lunch
- Lesson 6—SAFE Midsize Network Design
- Lab—SAFE Midsize Network Design Implementation

Course Agenda (Cont.)

Cisco.com

Day 3

- Lesson 7—SAFE Remote-User Network Implementation
- Lesson 8—SAFE Enterprise Network Design
- Lunch
- Lab—SAFE Remote-User Network Design Implementation

Day 4

- Lesson 9—SAFE IP Telephony Network Design
- Lunch
- Lesson 10—SAFE Wireless LAN Network Design
- Lab—SAFE Wireless LAN Network Design Implementation

Day 5

- Lab—Case Study (Optional)

Participant Responsibilities

Cisco.com

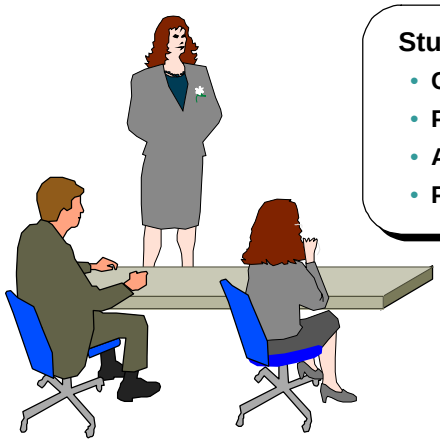

Student responsibilities

- Complete prerequisites
- Participate in lab exercises
- Ask questions
- Provide feedback

General Administration

Cisco.com

Class-related

- Sign-in sheet
- Length and times
- Break and lunch room locations
- Attire

Facilities-related

- Participant materials
- Site emergency procedures
- Restrooms
- Telephones/faxes

Graphic Symbols

Cisco.com

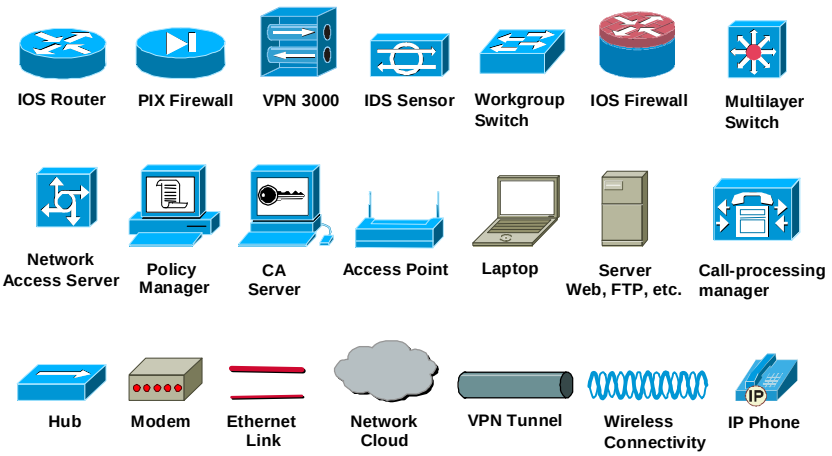

CSI 2.0—1-11

Participant Introductions

Cisco.com

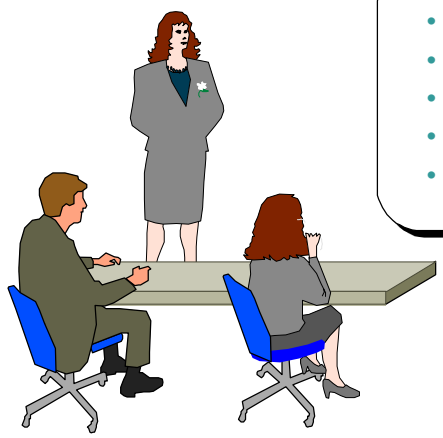

- Your name
- Your company
- Prerequisite skills
- Brief history
- Objective

CSI 2.0—1-12

Cisco Security Career Certifications

Cisco.com

Expand Your Professional Options —— and Advance Your Career

Cisco Certified Security Professional (CCSP) Certification

Professional-level recognition in designing and implementing Cisco security solutions

Required Exam	Recommended Training through Cisco Learning Partners
642-501	Securing Cisco IOS Networks
642-511	Cisco Secure Virtual Private Networks
642-531	Cisco Secure Intrusion Detection System
642-521	Cisco Secure PIX Firewall Advanced
642-541	Cisco SAFE Implementation

www.cisco.com/go/securitytraining

Cisco Security Career Certifications (Cont.)

Cisco.com

Enhance Your Cisco Certifications —— and Validate Your Areas of Expertise

Cisco Firewall, VPN, and IDS Specialists

Cisco Firewall Specialist

Required Exam	Recommended Training through Cisco Learning Partners
Pre-requisite: Valid CCNA certification	
642-501	Securing Cisco IOS Networks
642-521	Cisco Secure PIX Firewall Advanced

Cisco VPN Specialist

Required Exam	Recommended Training through Cisco Learning Partners
Pre-requisite: Valid CCNA certification	
642-501	Securing Cisco IOS Networks
642-511	Cisco Secure Virtual Private Networks

Cisco IDS Specialist

Required Exam	Recommended Training through Cisco Learning Partners
Pre-requisite: Valid CCNA certification	
642-501	Securing Cisco IOS Networks
642-531	Cisco Secure Intrusion Detection System

www.cisco.com/go/securitytraining

Lab Topology Overview

This topic describes the lab topology that is used in this course.

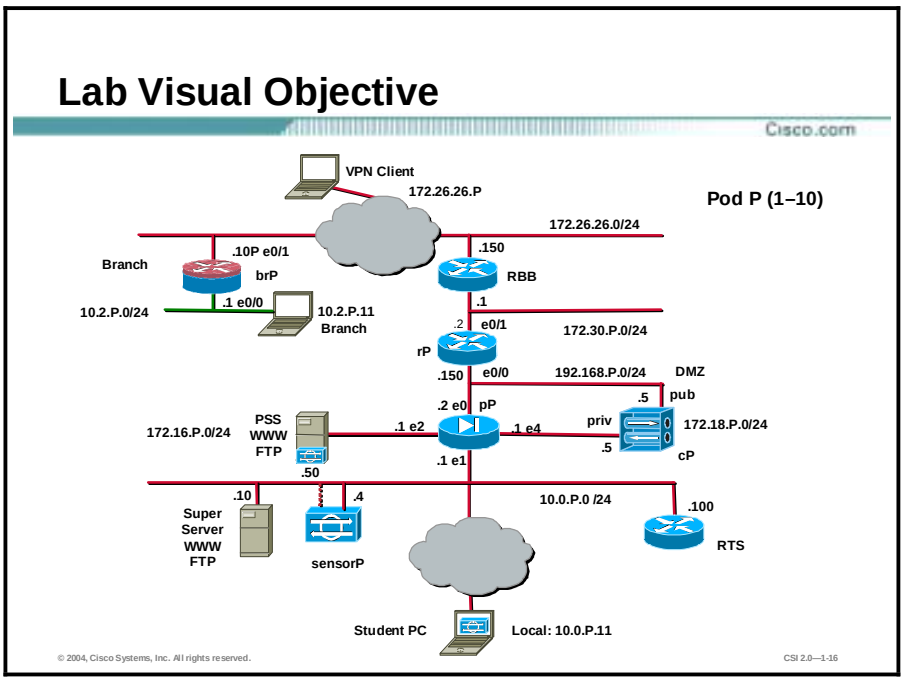

Each pair of students will be assigned a pod.

Note	The P in a command indicates your pod number.

Lab Visual Objective—SAFE Wireless LAN

Cisco.com

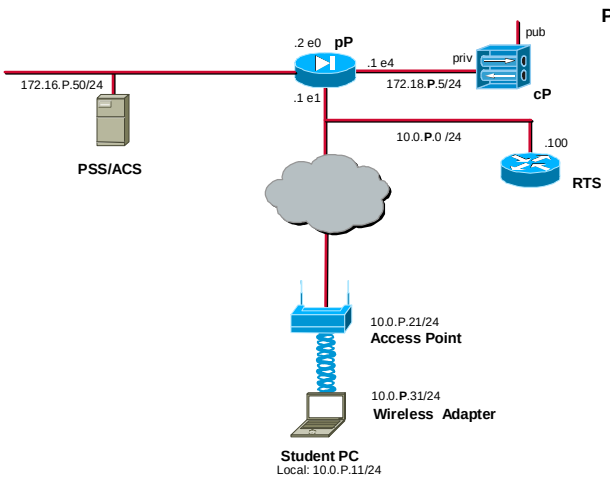

CSI 2.0—1-17

Lab Visual Objective—Case Study

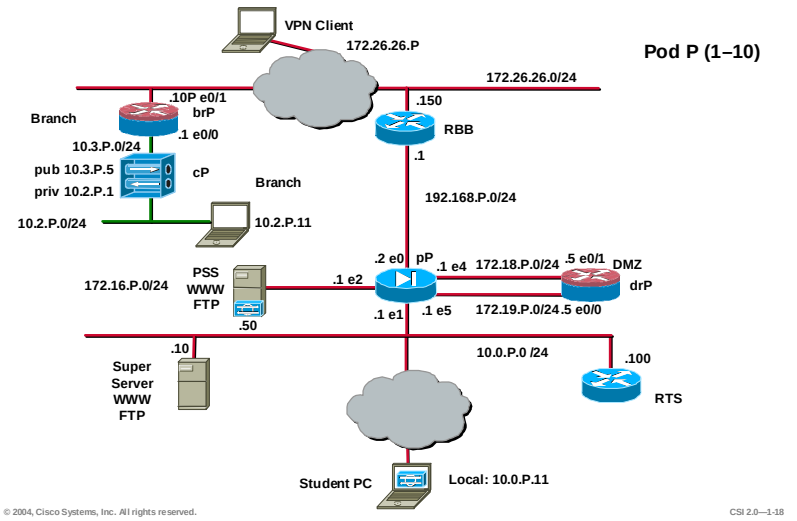

2

Security Fundamentals

Overview

This lesson describes security fundamentals. It includes the following topics:

- Objectives
- Need for network security
- Network security policy
- Primary network threats and attacks
- Reconnaissance attacks and mitigation
- Access attacks and mitigation
- Denial of service attacks and mitigation
- Worm, virus, and Trojan horse attacks and mitigation
- Management protocols and functions
- Summary

Objectives

This topic lists the lesson's objectives.

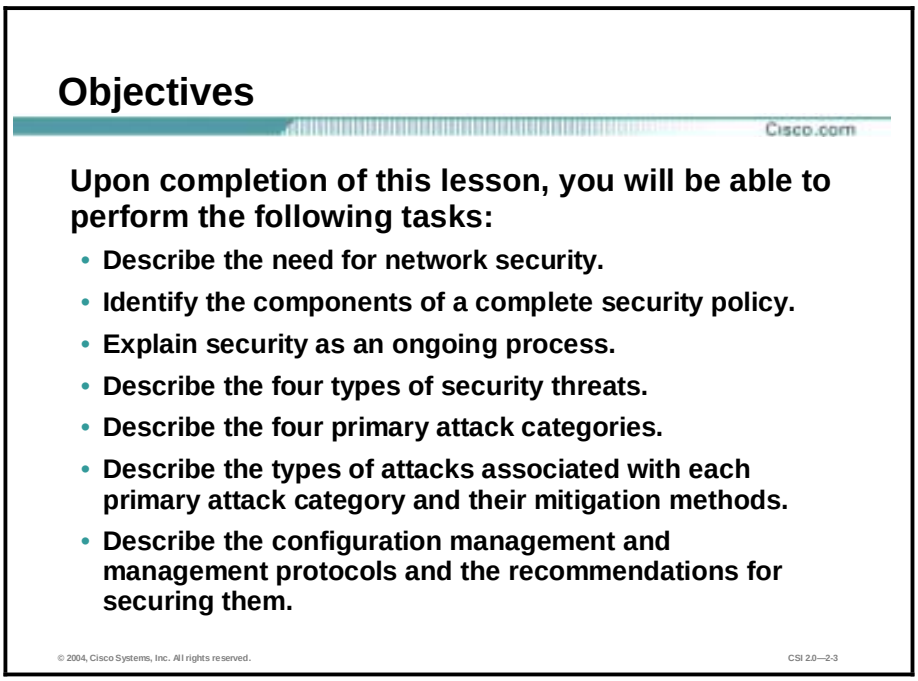

Need for Network Security

Over the past few years, Internet-enabled business, or e-business, has drastically improved companies' efficiency and revenue growth. E-business applications such as e-commerce, supply-chain management, and remote access enable companies to streamline processes, lower operating costs, and increase customer satisfaction. Such applications require mission-critical networks that accommodate voice, video, and data traffic, and these networks must be scalable to support increasing numbers of users and the need for greater capacity and performance. However, as networks enable more and more applications and are available to more and more users, they become ever more vulnerable to a wider range of security threats. To combat those threats and ensure that e-business transactions are not compromised, security technology must play a major role in today's networks.

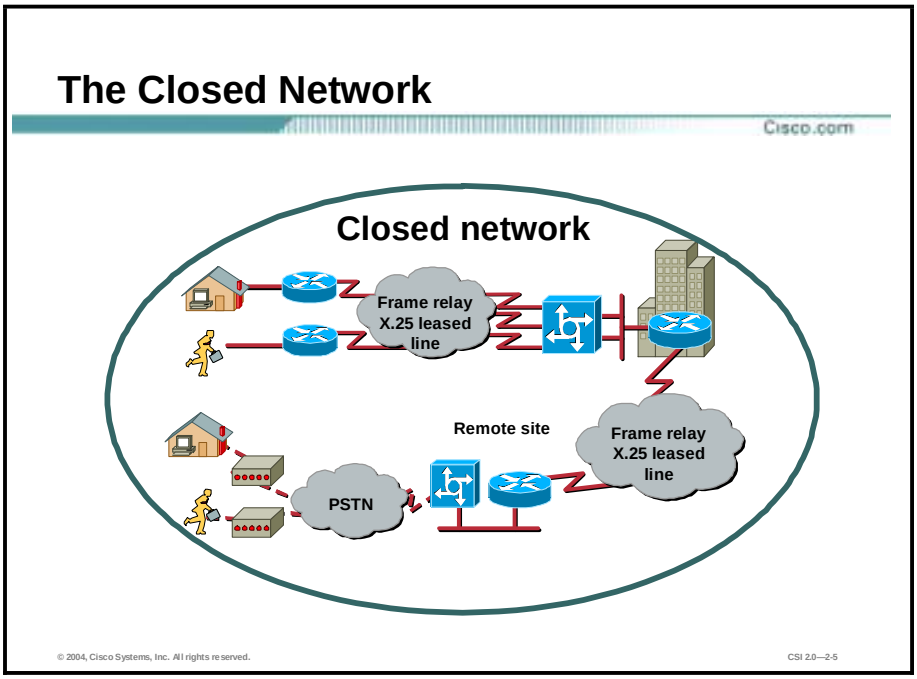

The closed network typically consists of a network designed and implemented in a corporate environment, and it provides connectivity only to known parties and sites without connecting to public networks. Networks were designed this way in the past and thought to be reasonably secure because there was no outside connectivity.

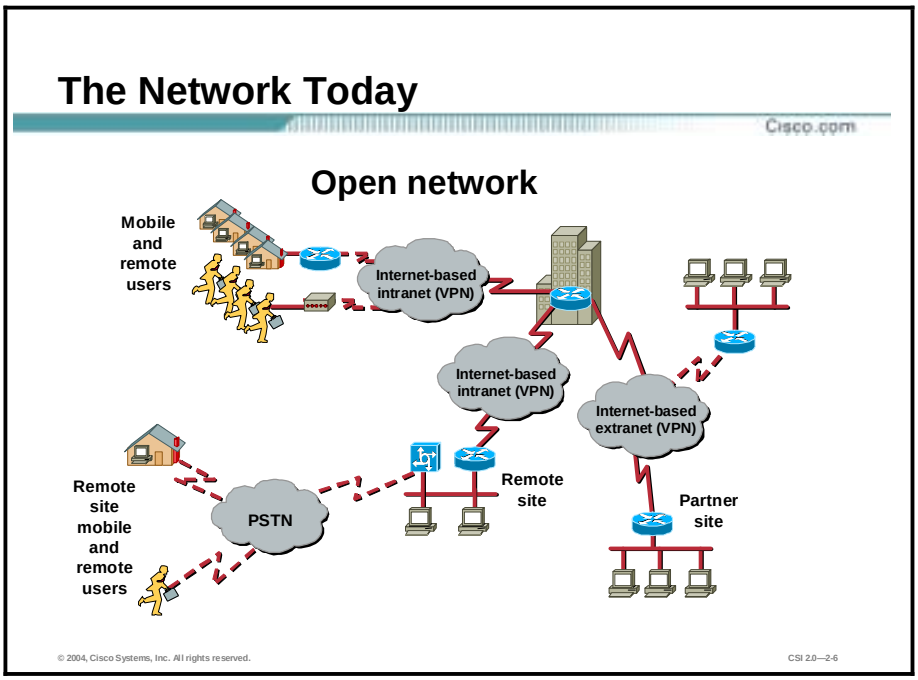

The networks of today are designed with availability to the Internet and public networks, which is a major requirement. Most of today's networks have several access points to other networks both public and private; therefore, securing these networks has become fundamentally important.

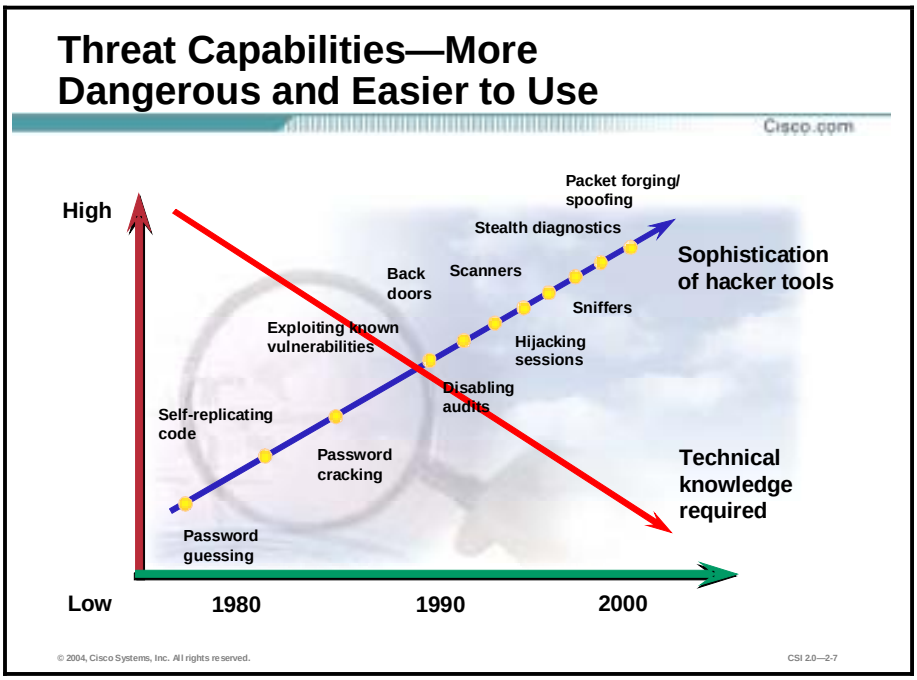

With the development of large open networks there has been a huge increase in security threats in the past 20 years. Not only have hackers discovered more vulnerabilities, but the tools used to hack a network have become simpler and the technical knowledge required has decreased. There are downloadable applications available that require little or no hacking knowledge to implement. There are also applications intended for troubleshooting a network that when used improperly can pose severe threats.

The Role of Security Is Changing

Cisco.com

As businesses become more open to supporting Internet-powered initiatives such as e-commerce, customer care, supply-chain management, and extranet collaboration, network security risks are also increasing.

CSI 2.0—2-8

Security has moved to the forefront of network management and implementation. It is necessary for the survival of many businesses to allow open access to network resources and ensure that the data and resources are as secure as possible.

Security is becoming more important because of the following:

- Required for e-business—The importance of e-business and the need for private data to traverse public networks has increased the need for network security.
- Required for communicating and doing business safely in potentially unsafe environments—Today's business environment requires communication with many public networks and systems, which produces the need for as much security as is possible.
- Networks require development and implementation of a corporate-wide security policy—Establishing a security policy should be the first step in migrating a network to a secure infrastructure.

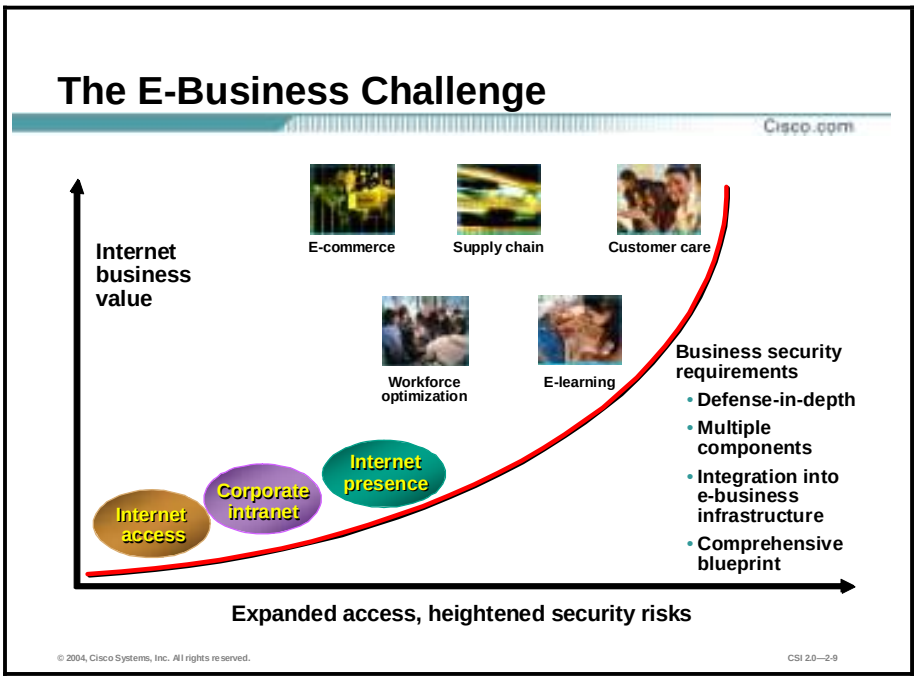

Security must be a fundamental component of any e-business strategy. As enterprise network managers open their networks to more users and applications, they also expose these networks to greater risk. The result has been an increase in business security requirements.

The Internet has radically shifted expectations of companies' abilities to build stronger relationships with customers, suppliers, partners, and employees. Driving companies to become more agile and competitive, e-business is giving birth to exciting new applications for e-commerce, supply-chain management, customer care, workforce optimization, and e-learning—applications that streamline and improve processes, speed up turnaround times, lower costs, and increase user satisfaction.

E-business requires mission-critical networks that accommodate ever-increasing constituencies and demands for greater capacity and performance. These networks also need to handle voice, video, and data traffic as networks converge into multiservice environments.

Legal and Governmental Policy Issues

Cisco.com

- **Many governments have formed cross-border task forces to deal with privacy issues.**
- **The outcome of international privacy efforts is expected to take several years to develop.**
- **National laws regarding privacy are expected to continue to evolve worldwide.**

As concerns about privacy increase, many governments have formed cross-border task forces to deal with privacy issues. International privacy efforts are expected to take several years to develop and even longer to implement globally. National laws regarding privacy are expected to continue to evolve worldwide.

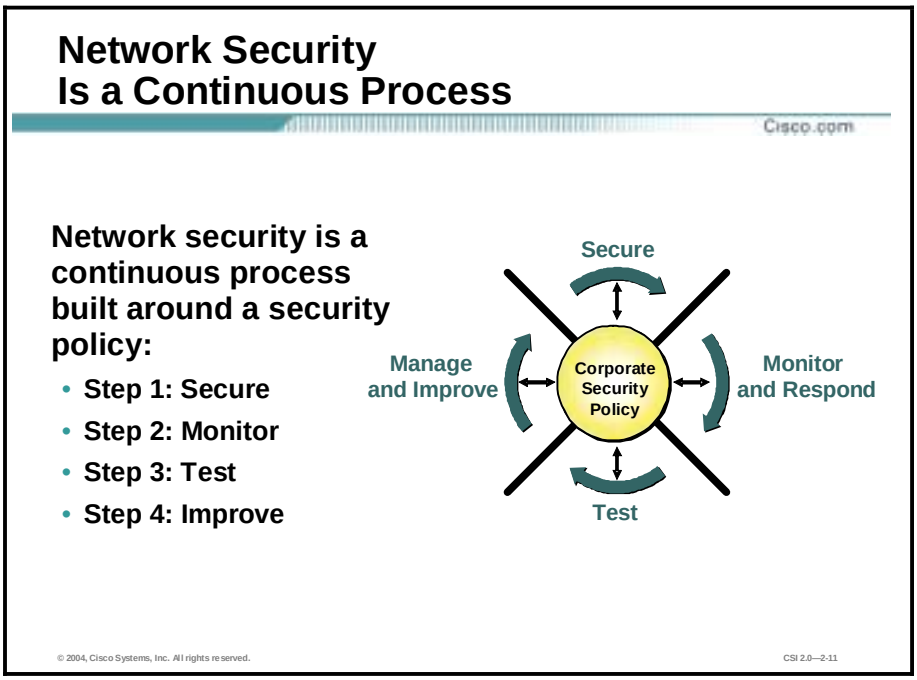

After setting appropriate policies, a company or organization must methodically consider security as part of normal network operations. This process could be as simple as configuring routers to not accept unauthorized addresses or services, or as complex as installing firewalls, intrusion detection systems (IDSs), centralized authentication servers, and encrypted virtual private networks (VPNs). Network security is a continuing process:

- Secure—The following are methods used to secure a network:
  - Authentication
  - Encryption
  - Firewalls
  - Vulnerability patching
- Monitor—To ensure that a network remains secure, it is important to monitor the state of security preparation. Network vulnerability scanners can proactively identify areas of weakness, and IDSs can monitor and respond to security events as they occur. Using security monitoring solutions, organizations can obtain unprecedented visibility into both the network data stream and the security posture of the network.
- Test—Testing security is as important as monitoring. Without testing the security solutions in place, it is impossible to know about existing or new attacks. The hacker community is an ever-changing environment. You can perform this testing yourself or outsource it to a third party such as the Cisco Security Posture Assessment (SPA) group.
- Improve—Monitoring and testing provides the data necessary to improve network security. Administrators and engineers should use the information from the monitor and test phases to make improvements to the security implementation as well as to adjust the security policy as vulnerabilities and risks are identified.

Network Security Policy

A security policy can be as simple as an acceptable use policy for network resources or it can be several hundred pages in length and detail every element of connectivity and associated policies.

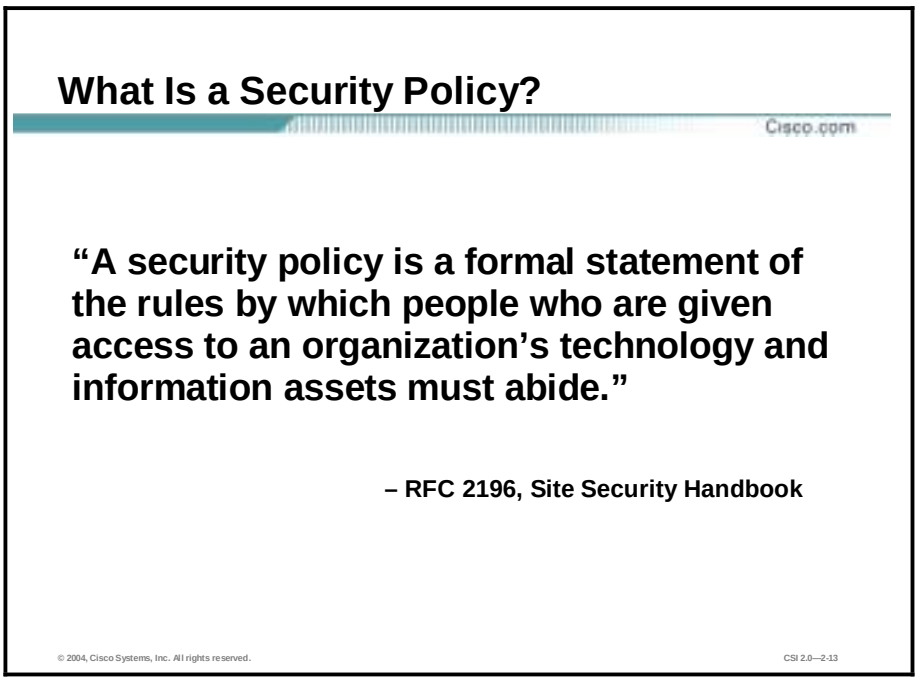

According to the Site Security Handbook (RFC 2196), "A security policy is a formal statement of the rules by which people who are given access to an organization's technology and information assets must abide." It further states, "A security policy is essentially a document summarizing how the corporation will use and protect its computing and network resources."

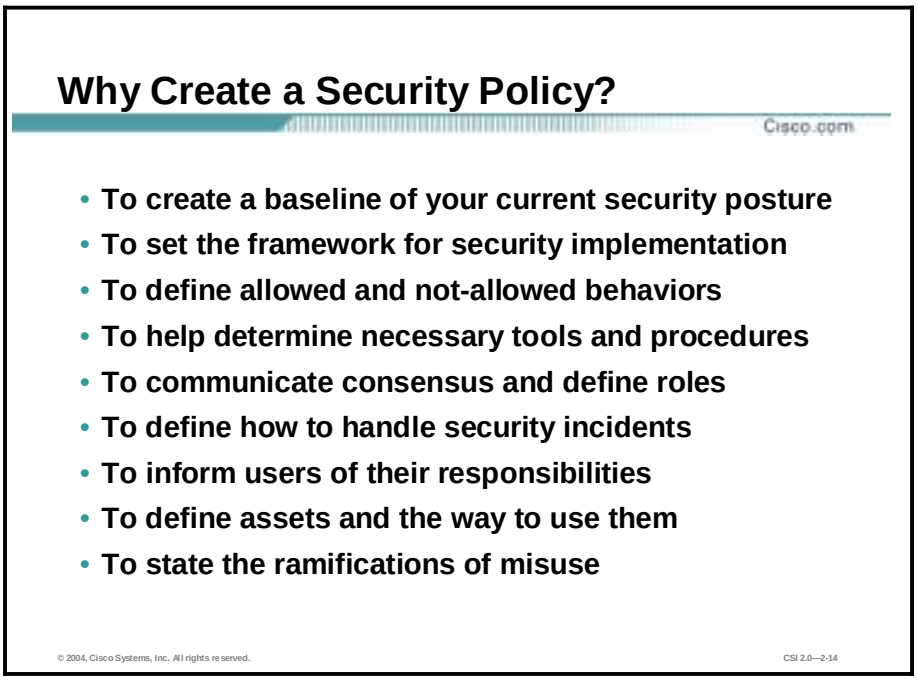

Security policies provide many benefits and are worth the time and effort needed to develop them. Developing a security policy:

- Provides a process for auditing existing network security.
- Provides a general security framework for implementing network security.
- Defines which behavior is and is not allowed.
- Helps determine which tools and procedures are needed for the organization.
- Helps communicate consensus among a group of key decision makers and define responsibilities of users and administrators.
- Defines a process for handling network security incidents.
- Enables global security implementation and enforcement. Computer security is now an enterprise-wide issue, and computing sites are expected to conform to the network security policy.
- Creates a basis for legal action if necessary.

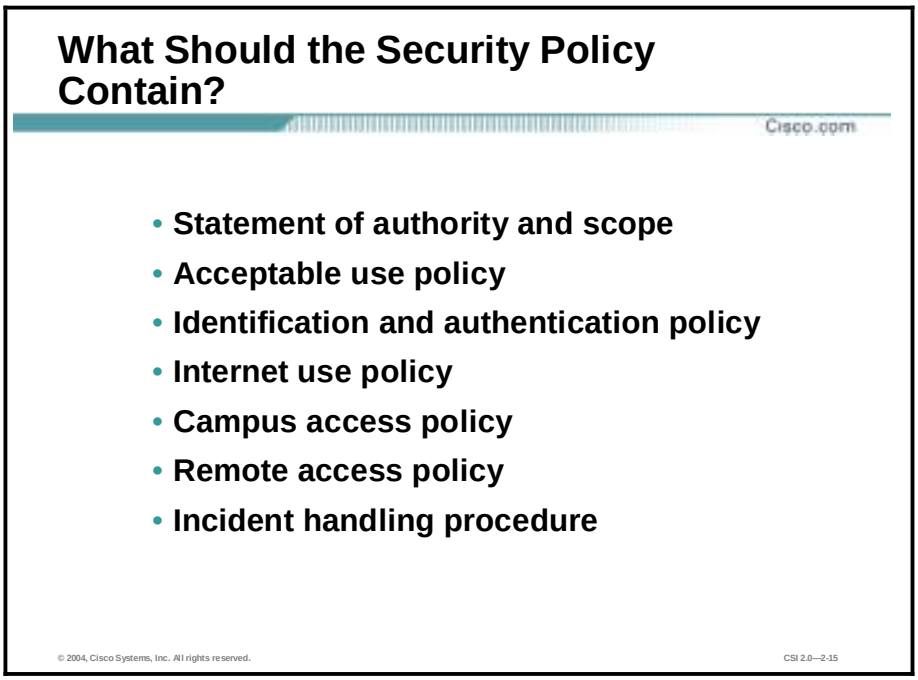

The following are some of the key policy components:

- Statement of authority and scope—This topic specifies who sponsors the security policy and what areas the policy covers.
- Acceptable use policy—This topic specifies what the company will and will not allow regarding its information infrastructure.
- Identification and authentication policy—This topic specifies what technologies, equipment, or combination of the two the company will use to ensure that only authorized individuals have access to its data.
- Internet access policy—This topic specifies what the company considers ethical and proper use of its Internet access capabilities.
- Campus access policy—This topic specifies how on-campus users will use the company's data infrastructure.
- Remote access policy—This topic specifies how remote users will access the company's data infrastructure.
- Incident handling procedure—This topic specifies how the company will create an incident response team and the procedures it will use during and after an incident.

Primary Network Threats and Attacks

This topic provides an overview of primary network threats and attacks.

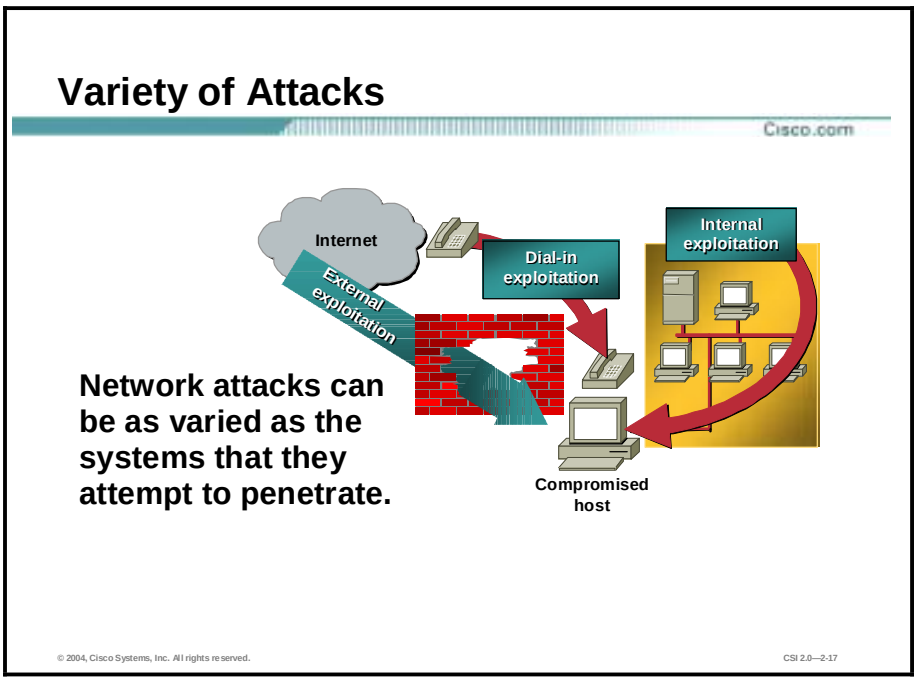

Without proper protection, any part of any network can be susceptible to attacks or unauthorized activity. Routers, switches, and hosts can all be violated by professional hackers, company competitors, or even internal employees. In fact, according to several studies, more than half of all network attacks are waged internally. The Computer Security Institute (CSI) in San Francisco, California, estimates that between 60 and 80 percent of network misuse comes from inside the enterprises where the misuse has taken place. To determine the best ways to protect against attacks, IT managers should understand the many types of attacks that can be instigated and the damage that these attacks can cause to e-business infrastructures.

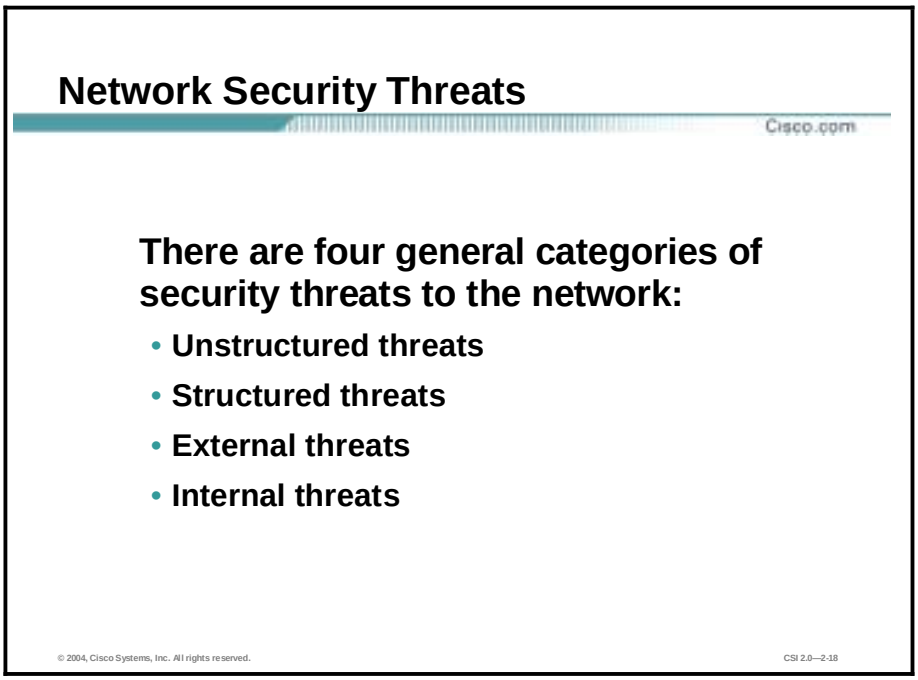

There are four general threats to network security:

- Unstructured threats—These threats primarily consist of random hackers using various common tools, such as malicious shell scripts, password crackers, credit card number generators, and dialer daemons. Although hackers in this category may have malicious intent, many are more interested in the intellectual challenge of cracking safeguards than in creating havoc.
- Structured threats—These threats are created by hackers who are more highly motivated and technically competent. Typically, such hackers act alone or in small groups to understand, develop, and use sophisticated hacking techniques to penetrate unsuspecting businesses. These groups are often involved in the major fraud and theft cases reported to law enforcement agencies. Occasionally, such hackers are hired by organized crime, industry competitors, or state-sponsored intelligence collection organizations.
- External threats—These threats consist of structured and unstructured threats originating from an external source. These threats may have malicious and destructive intent, or they may simply be errors that generate a threat.
- Internal threats—These threats typically involve disgruntled former or current employees. Although internal threats may seem more ominous than threats from external sources, security measures are available for reducing vulnerabilities to internal threats and responding when attacks occur.

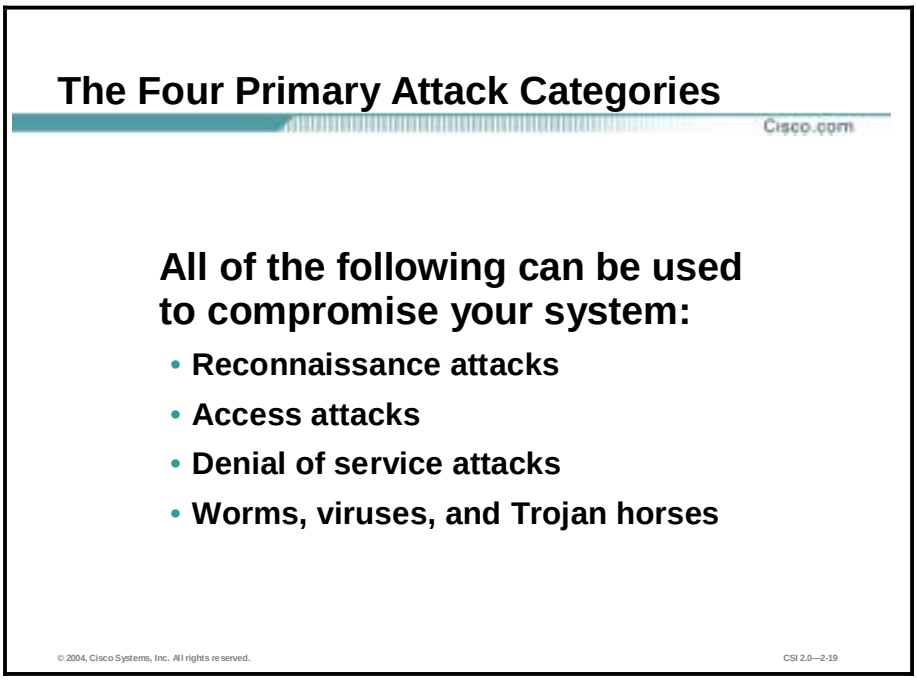

There are four types of network attacks:

- Reconnaissance attacks—An intruder attempts to discover and map systems, services, and vulnerabilities.
- Access attacks—An intruder attacks networks or systems to retrieve data, gain access, or escalate access privileges.
- Denial of service (DoS) attacks—An intruder attacks your network in a way that damages or corrupts your computer system or denies you and others access to your networks, systems, or services.
- Worms, viruses, and Trojan horses—Malicious software is inserted onto a host in order to damage a system, corrupt a system, replicate itself, or deny services or access to networks, systems, or services.

Reconnaissance Attacks and Mitigation

This topic describes reconnaissance attacks and their mitigation.

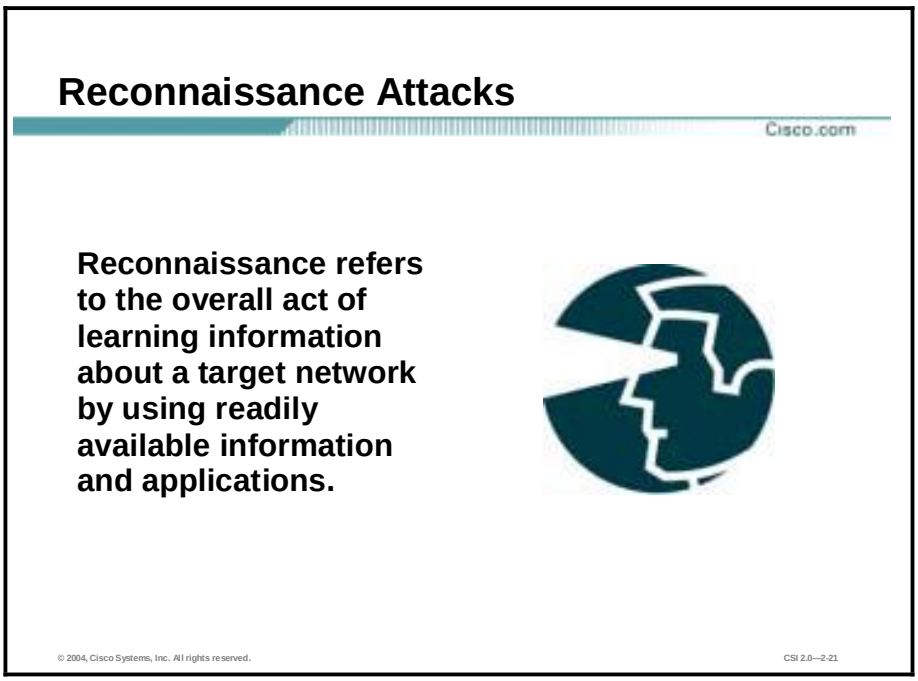

Reconnaissance is the unauthorized discovery and mapping of systems, services, or vulnerabilities. It is also known as information gathering and, in most cases, precedes an actual access or DoS attack. The malicious intruder typically conducts a ping sweep of the target network first to determine which IP addresses are alive. After this has been accomplished, the intruder determines which services or ports are active on the live IP addresses. From this information, the intruder queries the ports to determine the application type and version as well as the type and version of the operating system running on the target host.

Reconnaissance is somewhat analogous to a thief casing a neighborhood for vulnerable homes to break into, such as an unoccupied residence, a house with an easy-to-open door or window, and so on. In many cases the intruders go as far as “rattling the door handle,” not to go in immediately if it is opened, but to discover vulnerable services that they can exploit later when there is less likelihood that anyone is looking.

Reconnaissance attacks can consist of the following:

- Packet sniffers
- Port scans
- Ping sweeps
- Internet information queries

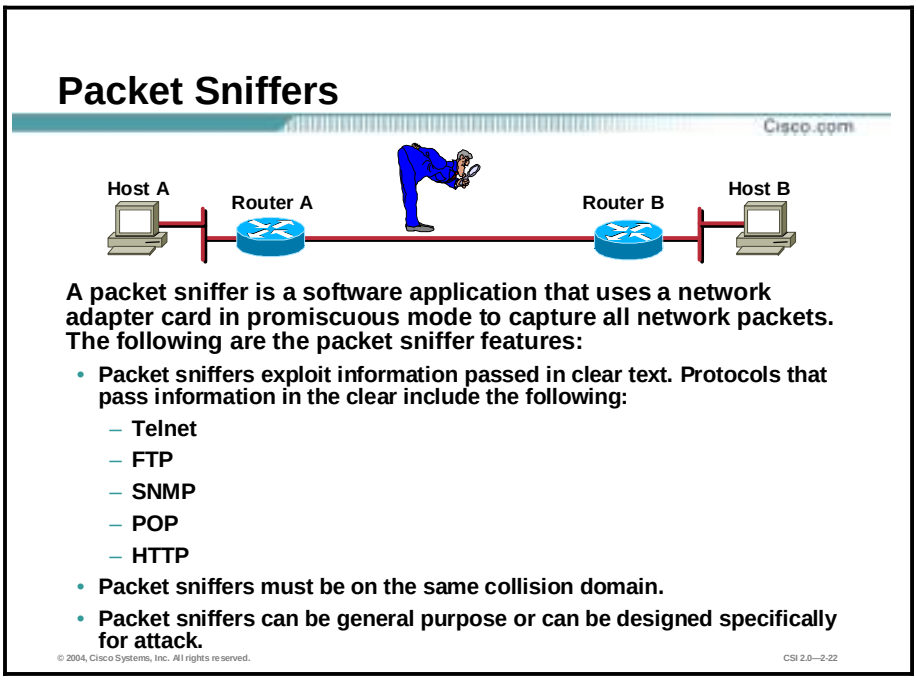

A packet sniffer is a software application that uses a network adapter card in promiscuous mode (a mode in which the network adapter card sends all packets received on the physical network wire to an application for processing) to capture all network packets that are sent across a LAN.

Several network applications distribute network packets in clear text; that is, the information sent across the network is not encrypted. Because the network packets are not encrypted, they can be processed and understood by any application that can pick them up off the network and process them.

A network protocol specifies how packets are identified and labeled, which enables a computer to determine whether a packet is intended for it. Because the specifications for network protocols, such as TCP/IP, are widely published, a third party can easily interpret the network packets and develop a packet sniffer. (The real threat today results from the numerous freeware and shareware packet sniffers that are available, which do not require the user to understand anything about the underlying protocols.)

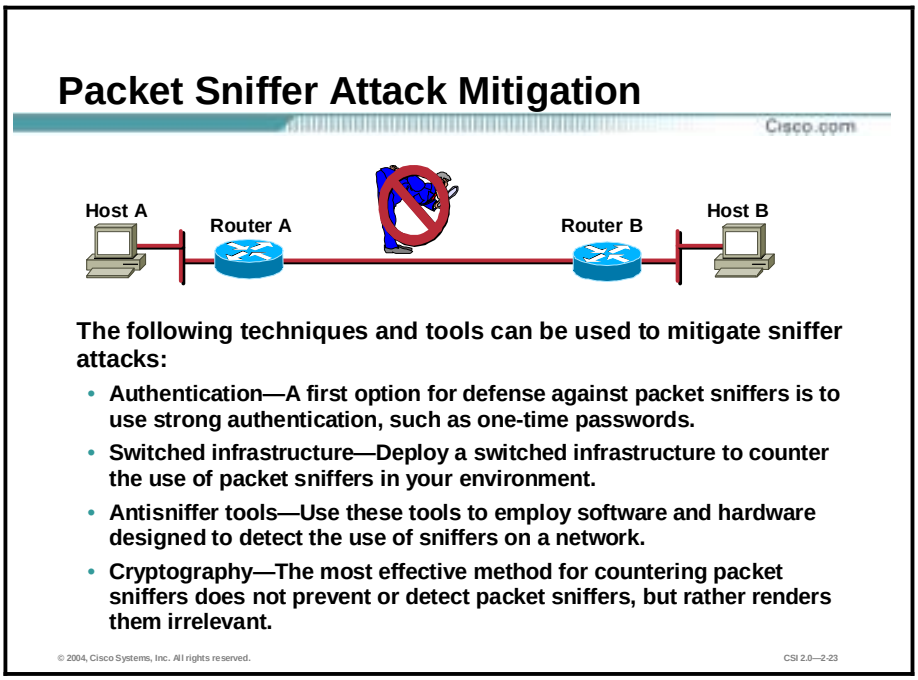

The following techniques and tools can be used to mitigate packet sniffer attacks:

- Authentication—Using strong authentication is a first option for defense against packet sniffers. Strong authentication can be broadly defined as a method of authenticating users that cannot easily be circumvented. A common example of strong authentication is one-time passwords (OTPs).

  An OTP is a type of two-factor authentication. Two-factor authentication involves using something you have combined with something you know. Automated teller machines (ATMs) use two-factor authentication. A customer needs both an ATM card and a personal identification number (PIN) to make transactions. With OTPs you need a PIN and your token card to authenticate to a device or software application. A token card is a hardware or software device that generates new, seemingly random, passwords at specified intervals (usually 60 seconds). A user combines that password with a PIN to create a unique password that works only for one instance of authentication. If a hacker learns that password by using a packet sniffer, the information is useless because the password has already expired. Note that this mitigation technique is effective only against a sniffer implementation that is designed to grab passwords. Sniffers deployed to learn sensitive information (such as e-mail messages) will still be effective.

- Switched infrastructure—This technique can be used to counter the use of packet sniffers in your network environment. For example, if an entire organization deploys switched Ethernet, hackers can gain access only to the traffic that flows on the specific port to which they connect. A switched infrastructure obviously does not eliminate the threat of packet sniffers, but it can greatly reduce their effectiveness.

- Antisniffer tools—Software and hardware designed to detect the use of sniffers on a network can be employed. Such software and hardware does not completely eliminate the threat, but like many network security tools, they are part of the overall system. These so-called antisniffers detect changes in the response time of hosts to determine whether the hosts are processing more traffic than their own. One such network security software tool, which is available from Security Software Technologies, is called AntiSniff.

- Cryptography—Rendering packet sniffers irrelevant is the most effective method for countering packet sniffers, even more effective than preventing or detecting packet sniffers. If a communication channel is cryptographically secure, the only data a packet sniffer will detect is cipher text (a seemingly random string of bits) and not the original message. The Cisco deployment of network-level cryptography is based on IPSec, which is a standard method for networking devices to communicate privately using IP. Other cryptographic protocols for network management include Secure Shell Protocol (SSH) and Secure Sockets Layer (SSL).

Port Scans and Ping Sweeps

Cisco.com

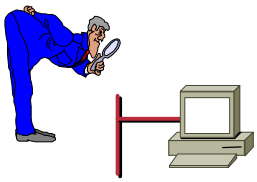

These attacks can attempt to:

- **Identify all services on the network**
- **Identify all hosts and devices on the network**
- **Identify the operating systems on the network**
- **Identify vulnerabilities on the network**

Port scans and ping sweeps are typically applications built to run various tests against a host or device in order to identify vulnerable services. The information is gathered by examining IP addressing and port or banner data from both TCP and UDP ports.

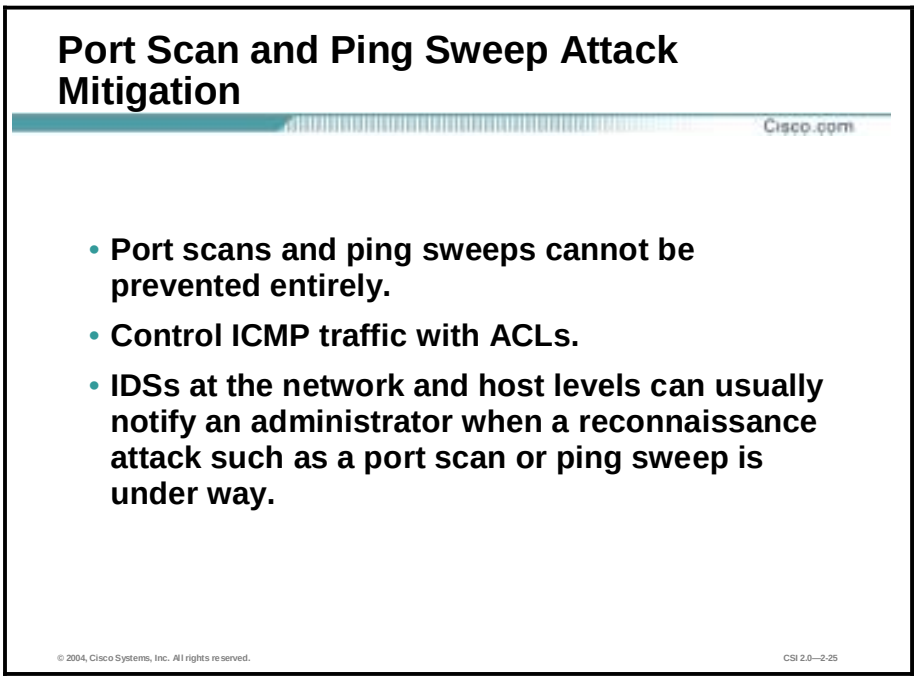

If ICMP echo and echo reply are turned off on edge routers, for example, ping sweeps can be stopped, but at the expense of network diagnostic data. However, port scans can easily be run without full ping sweeps; they simply take longer because they need to scan IP addresses that might not be live. IDSs at the network and host levels can usually notify an administrator when a reconnaissance attack is under way. This warning allows the administrator to better prepare for the coming attack or to notify the Internet service provider (ISP) that is hosting the system launching the reconnaissance probe.

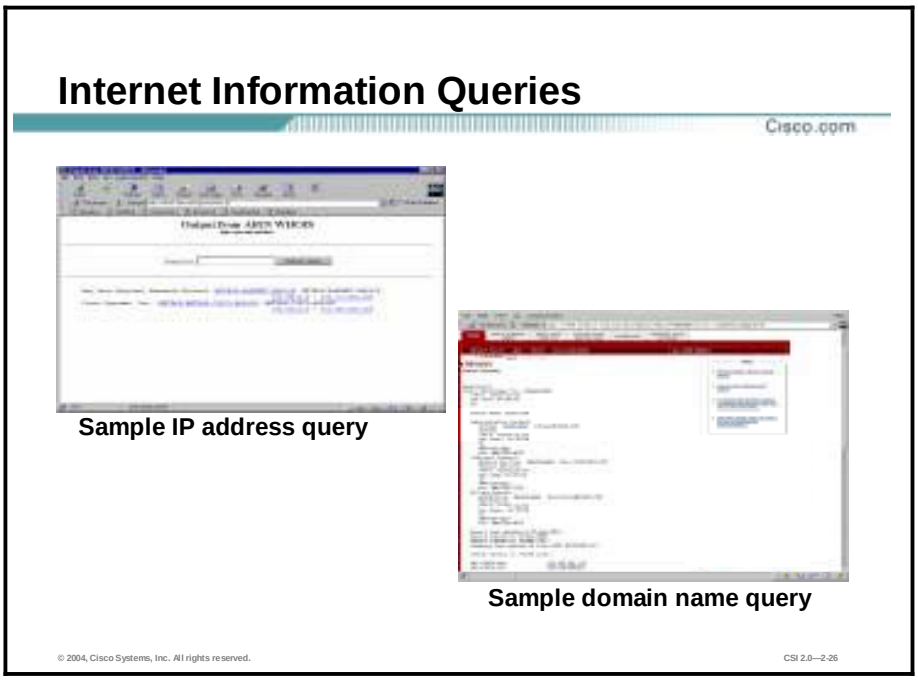

The figure demonstrates how existing Internet tools can be used for network reconnaissance (for example, an IP address query or a Domain Name System [DNS] query).

DNS queries can reveal such information as who owns a particular domain and what addresses have been assigned to that domain. Ping sweeps of the addresses revealed by the DNS queries can present a picture of the live hosts in a particular environment. After such a list is generated, port scanning tools can cycle through all well-known ports to provide a complete list of all services running on the hosts discovered by the ping sweep. Finally, the hackers can examine the characteristics of the applications that are running on the hosts. This step can lead to specific information that is useful when the hacker attempts to compromise that service.

IP address queries can reveal information such as who owns a particular IP address or range of addresses and what domain is associated with them.

Access Attacks and Mitigation

This topic describes specific access attacks and their mitigation.

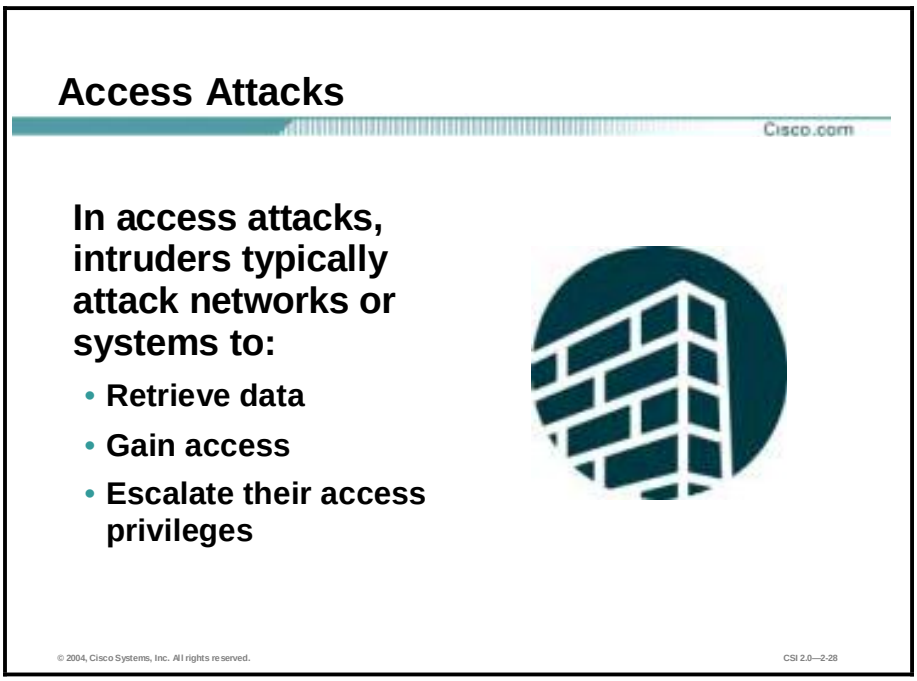

Access attacks exploit known vulnerabilities in authentication services, FTP services, and Web services to gain entry to Web accounts, confidential databases, and other sensitive information. Access attacks can consist of the following:

- Password attacks
- Trust exploitation
- Port redirection
- Man-in-the-middle attacks

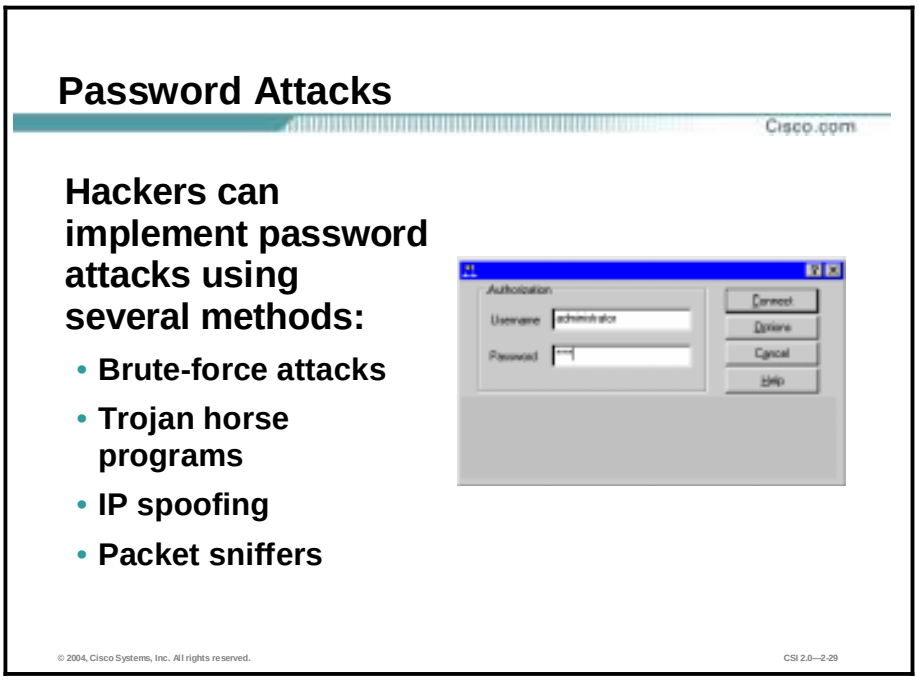

Password attacks can be implemented using several methods, including brute-force attacks, Trojan horse programs, IP spoofing, and packet sniffers. Although packet sniffers and IP spoofing can yield user accounts and passwords, password attacks usually refer to repeated attempts to identify a user account, password, or both. These repeated attempts are called brute-force attacks.

Often a brute-force attack is performed using a program that runs across the network and attempts to log in to a shared resource, such as a server. When an attacker gains access to a resource, he or she has the same access rights as the user whose account has been compromised. If this account has sufficient privileges, the attacker can create a back door for future access, without concern for any status and password changes to the compromised user account.

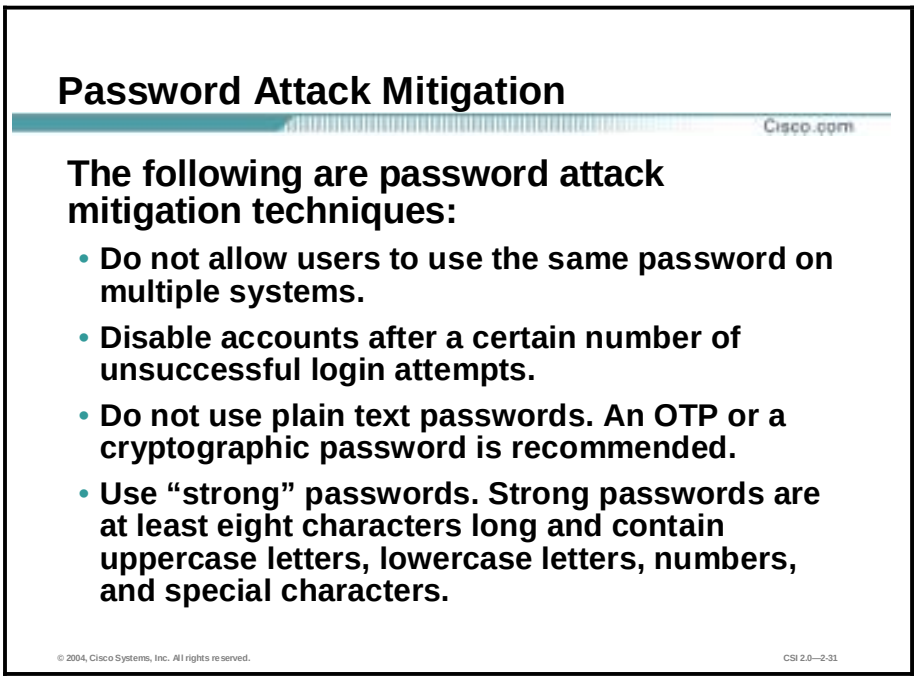

The following are password attack mitigation techniques:

- Do not allow users to have the same password on multiple systems—Most users will use the same password for each system they access, and often personal system passwords will be the same as well.
- Disable accounts after a specific number of unsuccessful logins—This practice helps to prevent continuous password attempts.
- Do not use plain-text passwords—Use of either an OTP or encrypted password is recommended.
- Use “strong” passwords—Many systems now provide strong password support and can restrict a user to the use of strong passwords only. Strong passwords are at least eight characters long and contain uppercase letters, lowercase letters, numbers, and special characters.

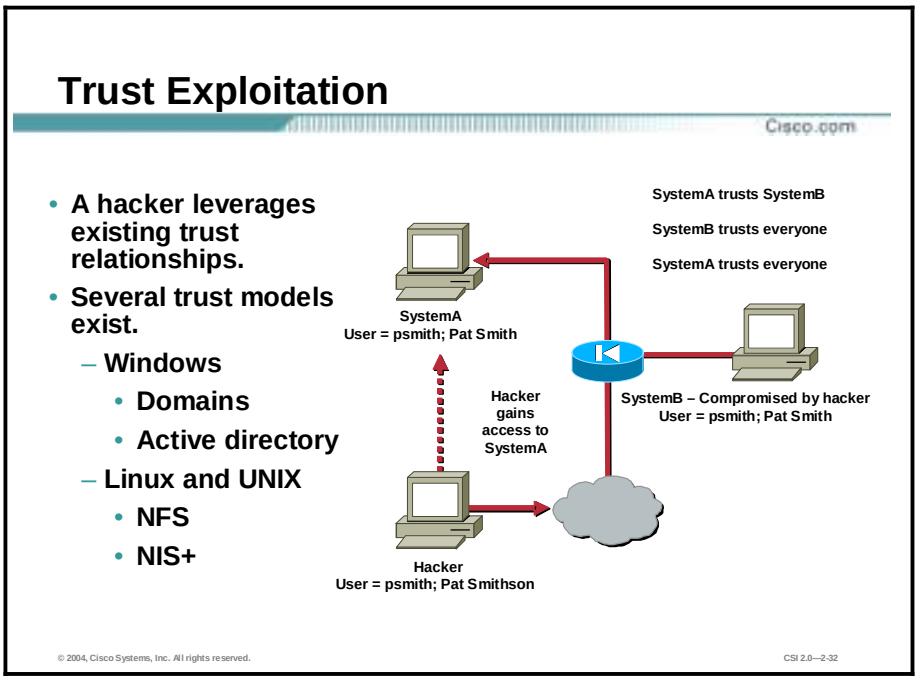

Although it is not an attack in itself, trust exploitation refers to an individual's taking advantage of a trust relationship within a network. The classic example is a perimeter network connection from a corporation. These network segments often house DNS, Simple Mail Transfer Protocol (SMTP), and HTTP servers. Because they all reside on the same segment, a compromise of one system can lead to the compromise of other systems if those other systems in turn trust systems attached to the same network. Another example is a system on the outside of a firewall that has a trust relationship with a system on the inside of a firewall. When the outside system is compromised, the attacker can leverage that trust relationship to attack the inside network.

Trust Exploitation Attack Mitigation

Cisco.com

- **Systems on the outside of a firewall should never be absolutely trusted by systems on the inside of a firewall.**
- **Such trust should be limited to specific protocols and should be validated by something other than an IP address where possible.**

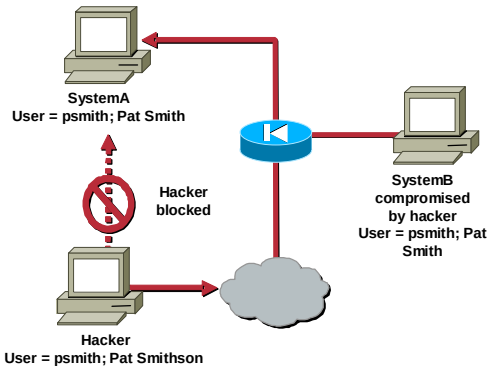

You can mitigate trust exploitation-based attacks through tight constraints on trust levels within a network. Systems on the outside of a firewall should never be absolutely trusted by systems on the inside of a firewall. Such trust should be limited to specific protocols and should be authenticated by something other than an IP address where possible.

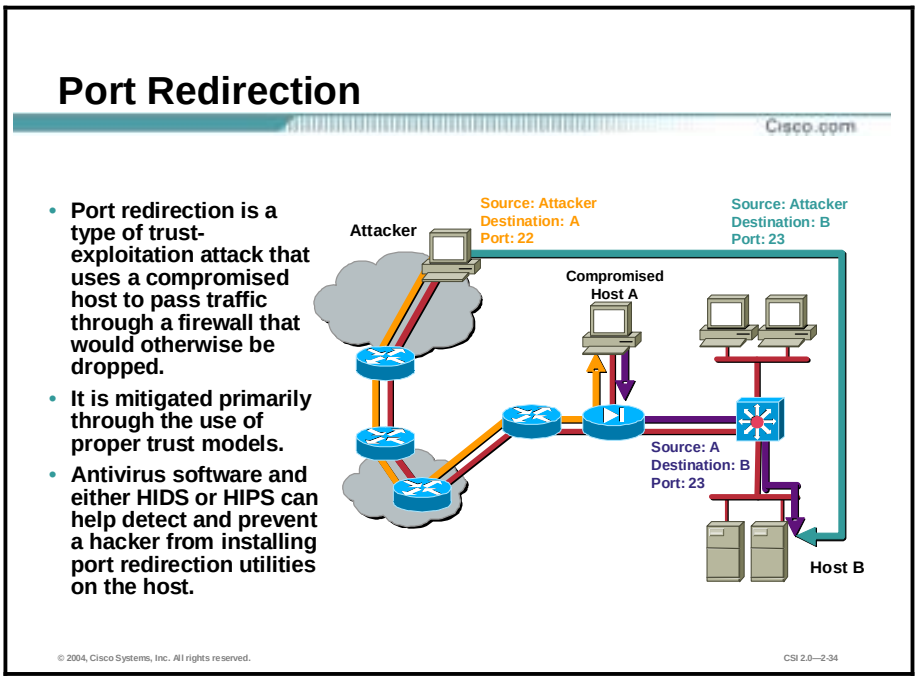

Port redirection attacks are a type of trust exploitation attack that uses a compromised host to pass traffic through a firewall that would otherwise be dropped. Consider a firewall with three interfaces and a host on each interface. The host on the outside can reach the host on the public services segment (commonly referred to as a Demilitarized Zone [DMZ]), but not the host on the inside. The host on the public services segment can reach the host on both the outside and the inside. If hackers were able to compromise the public services segment host, they could install software to redirect traffic from the outside host directly to the inside host. Though neither communication violates the rules implemented in the firewall, the outside host has now achieved connectivity to the inside host through the port redirection process on the public services host. An example of an application that can provide this type of access is netcat.

Port redirection can be mitigated primarily through the use of proper trust models, which are network specific (as mentioned earlier). Assuming a system under attack, a HIDS or HIPS can help detect a hacker and prevent installation of such utilities on a host.

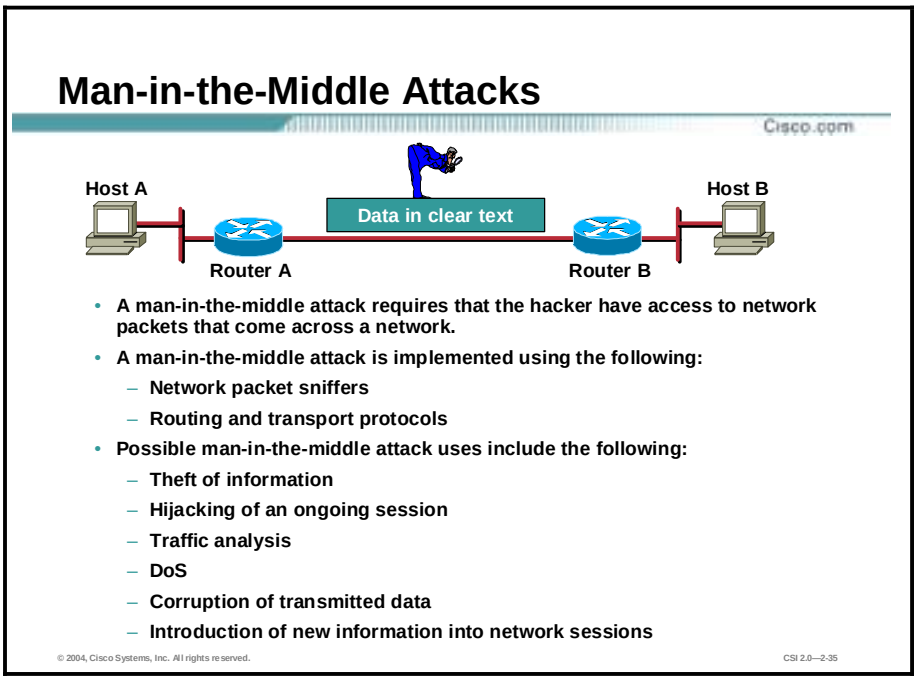

A man-in-the-middle attack requires that the attacker have access to network packets that come across the network. Such attacks are often implemented using network packet sniffers and routing and transport protocols. The possible uses of such attacks are theft of information, hijacking of an ongoing session to gain access to your internal network resources, traffic analysis to derive information about your network and its users, denial of service, corruption of transmitted data, and introduction of new information into network sessions.

An example of a man-in-the-middle attack could be someone who is working for your ISP and who can gain access to all network packets transferred between your network and any other network.

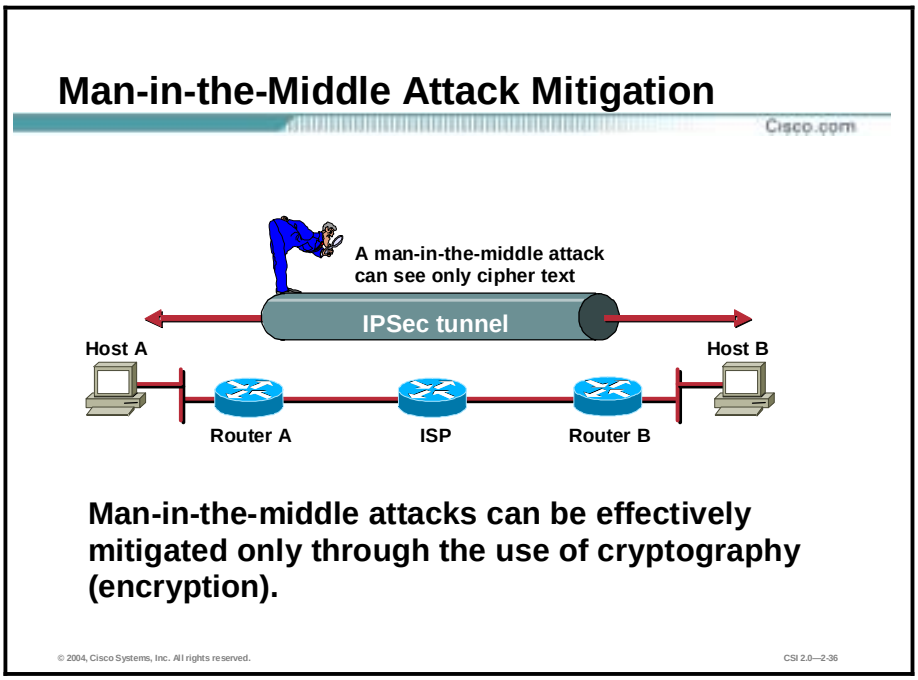

Man-in-the-middle attack mitigation is achieved, as shown in the figure, by encrypting traffic in an IPSec tunnel, which would allow the hacker to see only cipher text.

Denial of Service Attacks and Mitigation

This topic describes specific DoS attacks and their mitigation.

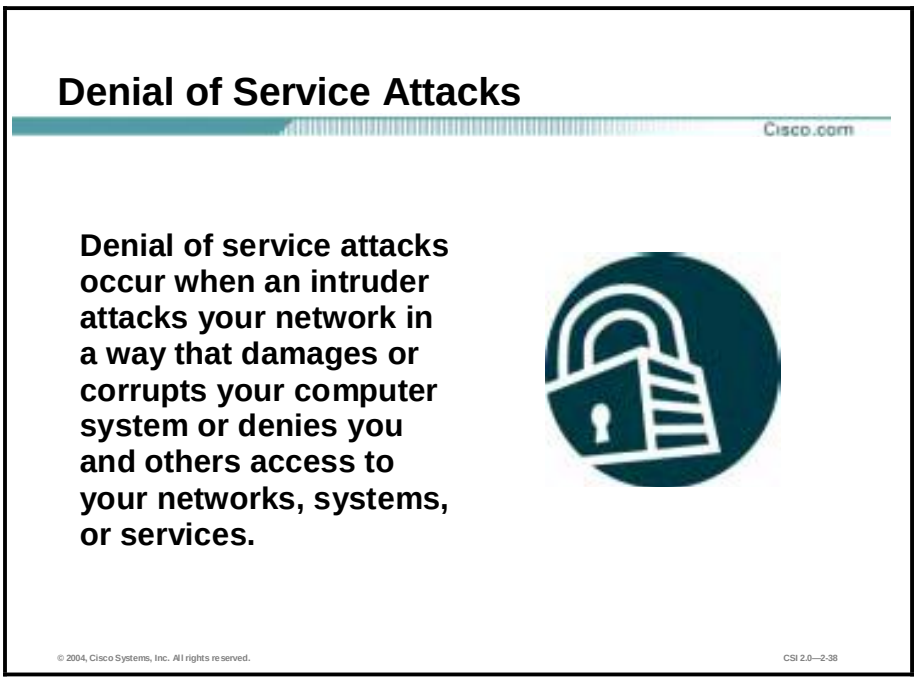

Certainly the most publicized form of attack, DoS attacks are also among the most difficult to completely eliminate. Even within the hacker community, DoS attacks are regarded as trivial and considered bad form because they require so little effort to execute. Still, because of their ease of implementation and potentially significant damage, DoS attacks deserve special attention from security administrators. If you are interested in learning more about DoS attacks, researching the methods employed by some of the better-known attacks can be useful. DoS attacks can consist of the following:

- IP spoofing
- Distributed denial of service (DDoS)

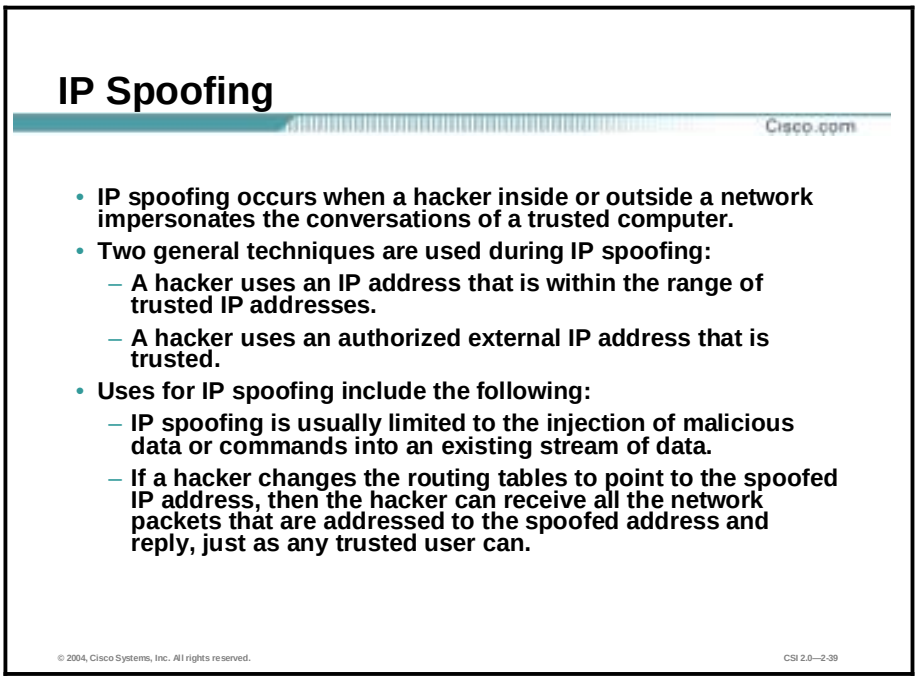

An IP spoofing attack occurs when an attacker outside your network pretends to be a trusted computer, either by using an IP address that is within the range of IP addresses for your network or by using an authorized external IP address that you trust and to which you wish to provide access to specified resources on your network.

Normally, an IP spoofing attack is limited to the injection of data or commands into an existing stream of data passed between a client and server application or a peer-to-peer network connection. To enable bidirectional communication, the attacker must change all routing tables to point to the spoofed IP address. Another approach the attacker could take is simply not to worry about receiving any response from the applications. For example, if an attacker is attempting to get a system to mail him or her a sensitive file, application responses are unimportant.

However, if an attacker manages to change the routing tables to point to the spoofed IP address, he or she can receive all the network packets that are addressed to the spoofed address and reply just as any trusted user can. Like packet sniffers, IP spoofing use is not restricted to people who are external to the network.

Although this use is not as common, IP spoofing can also provide access to user accounts and passwords, and it can also be used in other ways. For example, an attacker can emulate one of your internal users in ways that prove embarrassing for your organization; the attacker could send e-mail messages to business partners that appear to have originated from someone within your organization. Such attacks are easier when an attacker has a user account and password, but they are possible when simple spoofing attacks are combined with knowledge of messaging protocols.

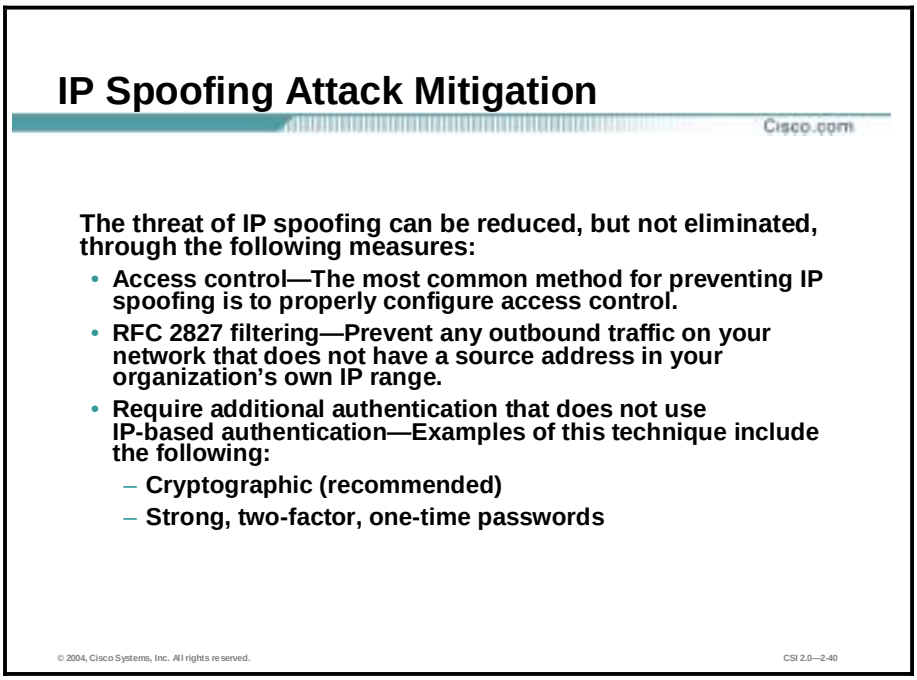

The threat of IP spoofing can be reduced, but not eliminated, through the following measures:

- Access control—The most common method for preventing IP spoofing is to properly configure access control. To reduce the effectiveness of IP spoofing, configure access control to deny any traffic from the external network that has a source address that should reside on the internal network. Note that this helps prevent spoofing attacks only if the internal addresses are the only trusted addresses. If some external addresses are trusted, this method is not effective.
- RFC 2827 filtering—You can prevent users of your network from spoofing other networks (and be a good Internet citizen at the same time) by preventing any outbound traffic on your network that does not have a source address in your organization's own IP range.

  This filtering denies any traffic that does not have the source address that was expected on a particular interface. For example, if an ISP is providing a connection to the IP address 15.1.1.0/24, the ISP could filter traffic so that only traffic sourced from address 15.1.1.0/24 can enter the ISP router from that interface. Note that unless all ISPs implement this type of filtering, its effectiveness is significantly reduced.
- Additional authentication—The most effective method for mitigating the threat of IP spoofing is the same as the most effective method for mitigating the threat of packet sniffers: namely, eliminating its effectiveness. IP spoofing can function correctly only when devices use IP address-based authentication; therefore, if you use additional authentication methods, IP spoofing attacks are irrelevant. Cryptographic authentication is the best form of additional authentication, but when that is not possible, strong two-factor authentication using OTPs can also be effective.

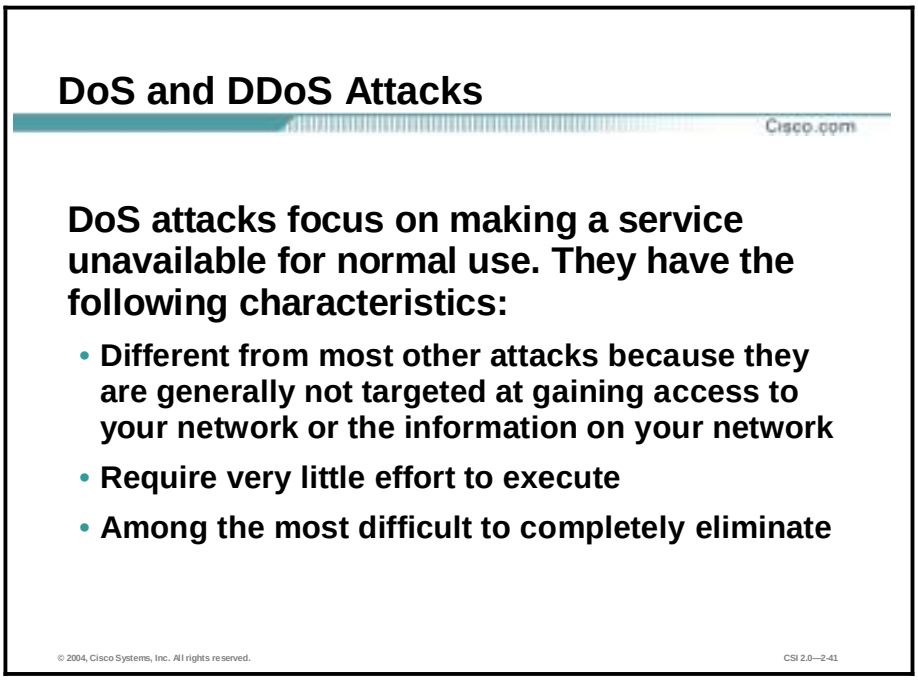

DoS attacks are different from most other attacks because they are not targeted at gaining access to your network or the information on your network. These attacks focus on making a service unavailable for normal use, which is typically accomplished by exhausting some resource limitation on the network or within an operating system or application. These attacks require little effort to execute because they typically take advantage of protocol weaknesses or because the attacks are carried out using traffic that would normally be allowed into a network. DoS attacks are among the most difficult to completely eliminate because of the way they use protocol weaknesses and “native” traffic to attack a network.

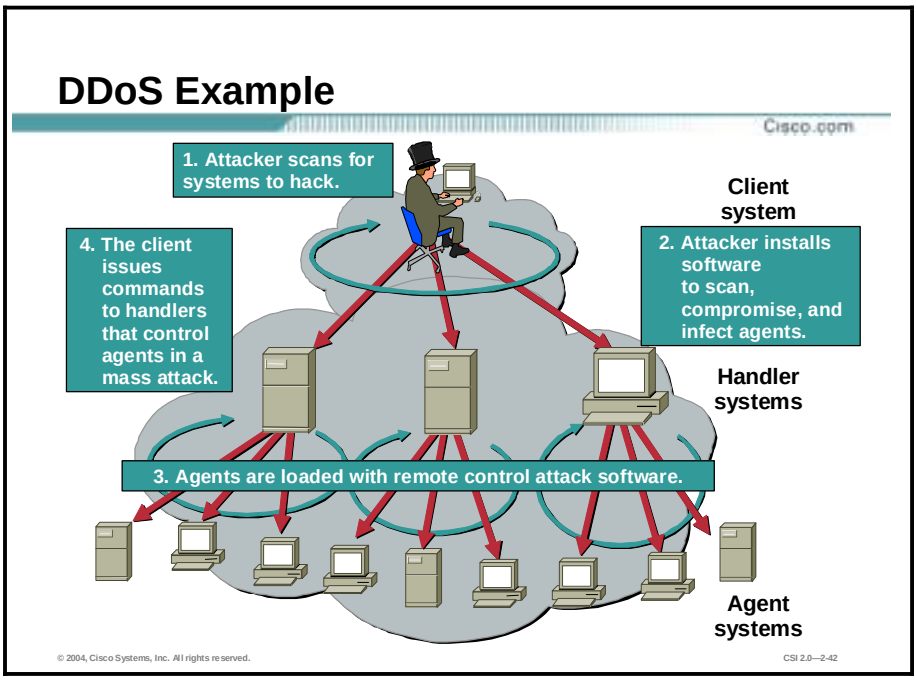

DDoS attacks are the “next generation” of DoS attacks on the Internet. This type of attack is not new—UDP and TCP SYN flooding, Internet Control Message Protocol (ICMP) echo request floods, and ICMP directed broadcasts (also known as smurf attacks) are similar—but the scope certainly is new. Victims of DDoS attacks experience packet flooding from many different sources, possibly spoofed IP source addresses, that bring their network connectivity to a grinding halt. In the past, the typical DoS attack involved a single attacker’s attempt to flood a target host with packets. With DDoS tools, an attacker can conduct the same attack using thousands of systems.

In the figure, the hacker uses a terminal to scan for systems to hack. When the handler systems are accessed, the hacker then installs software on them to scan for, compromise, and infect agent systems. When the agent systems are accessed, the hacker then loads remote control attack software to carry out the DoS attack.

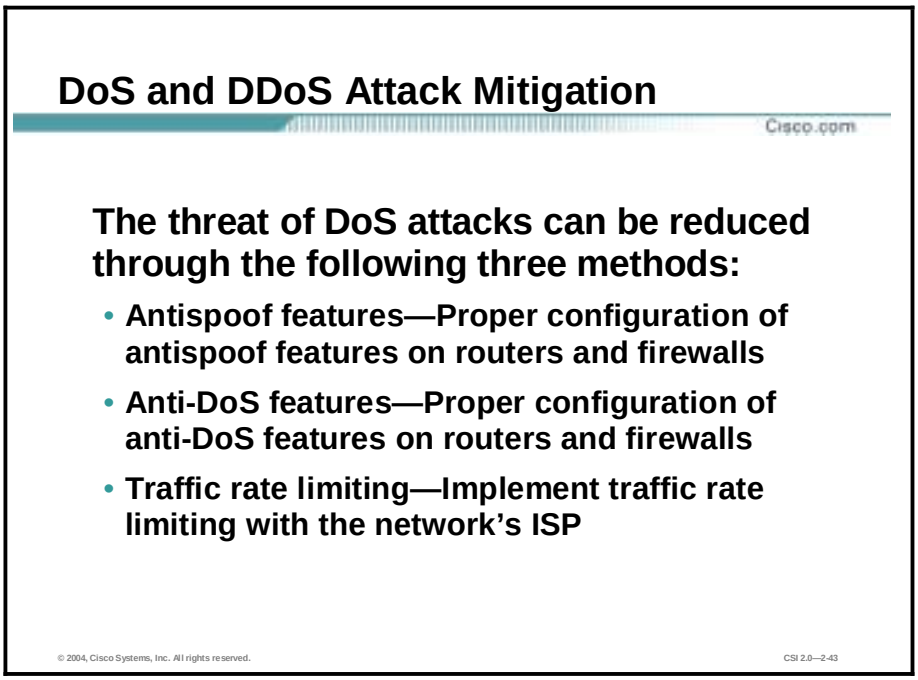

When they involve specific network server applications, such as an HTTP server or an FTP server, these attacks can focus on acquiring and keeping open all the available connections supported by that server, effectively locking out valid users of the server or service. DoS attacks can also be implemented using common Internet protocols, such as TCP and ICMP. While most DoS attacks exploit a weakness in the overall architecture of the system being attacked rather than a software bug or security hole, some attacks compromise the performance of your network by flooding the network with undesired, and often useless, network packets and by providing false information about the status of network resources.

The threat of DoS attacks can be reduced through the following three methods:

- Antispoof features—Proper configuration of antispoof features on your routers and firewalls can reduce your risk. This configuration includes RFC 2827 filtering at a minimum. If hackers cannot mask their identities, they might not attack.
- Anti-DoS features—Proper configuration of anti-DoS features on routers and firewalls can help limit the effectiveness of an attack. These features often involve limits on the amount of half-open connections that a system allows at any given time.
- Traffic rate limiting—An organization can implement traffic rate limiting with its ISP. This type of filtering limits the amount of nonessential traffic that crosses network segments at a certain rate. A common example is to limit the amount of ICMP traffic allowed into a network because it is used only for diagnostic purposes. ICMP-based DDoS attacks are common.

Worm, Virus, and Trojan Horse Attacks and Mitigation

This topic describes worm, virus, and Trojan horse attacks and their mitigation.

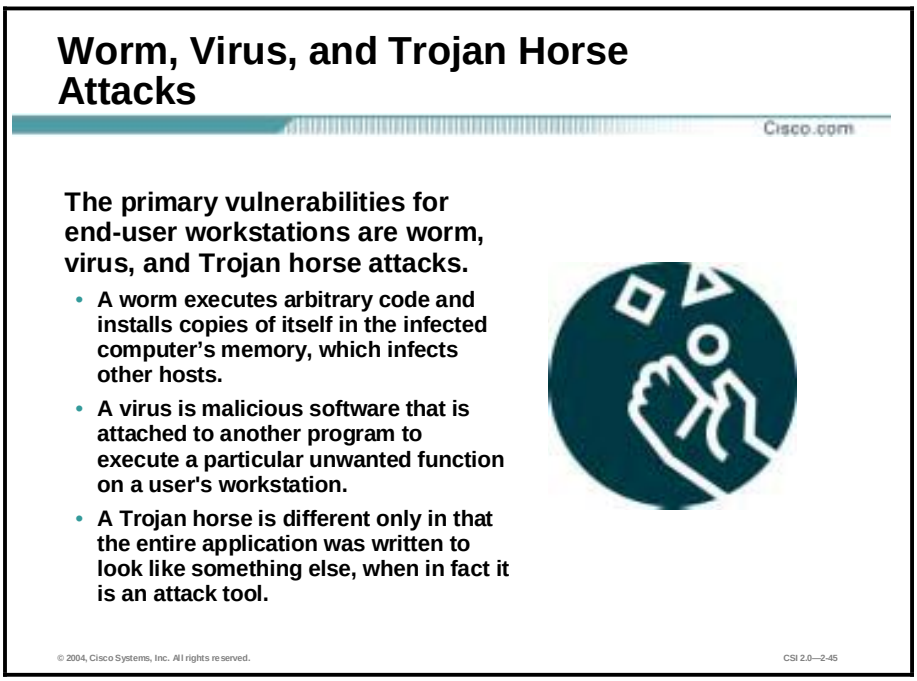

The primary vulnerabilities for end-user workstations are worm, virus, and Trojan horse attacks.

- A worm executes arbitrary code and installs copies of itself in the infected computer's memory, which infects other hosts.
- A virus is malicious software that is attached to another program to execute a particular unwanted function on a user's workstation.
- A Trojan horse is different only in that the entire application was written to look like something else, when in fact it is an attack tool.

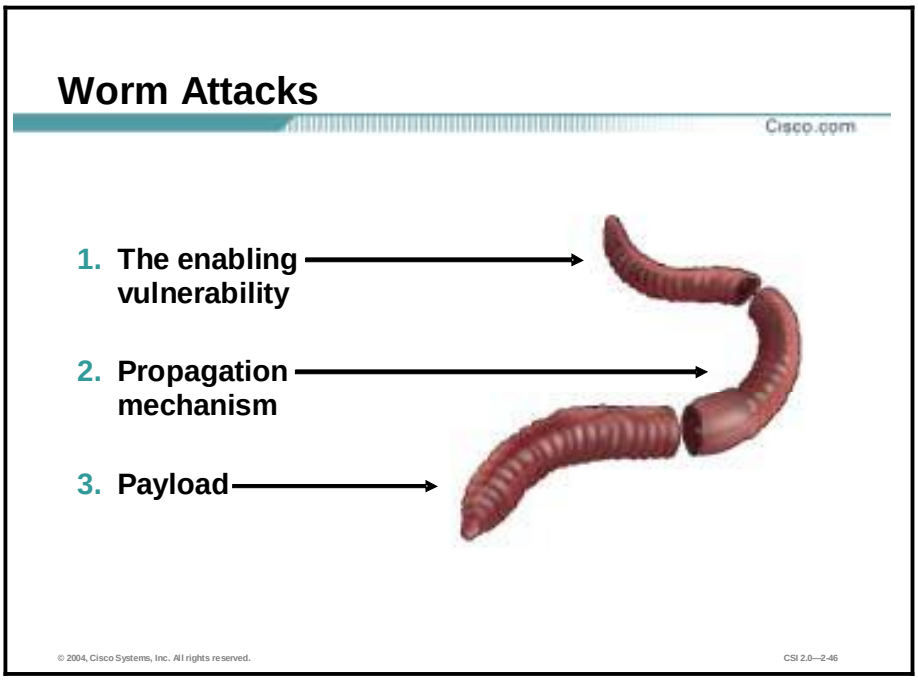

The anatomy of a worm attack is as follows:

- The enabling vulnerability—A worm installs itself using an exploit vector on a vulnerable system.
- Propagation mechanism—After gaining access to devices, a worm replicates and selects new targets.
- Payload—Once the device is infected with a worm, the attacker has access to the host—often as a privileged user. Attackers could use a local exploit to escalate their privilege level to administrator.

Typically, worms are self-contained programs that attack a system and try to exploit a vulnerability in the target. Upon successful exploitation of the vulnerability, the worm copies its program from the attacking host to the newly exploited system to begin the cycle again. A virus normally requires a vector to carry the virus code from one system to another. The vector can be a word-processing document, an e-mail message, or an executable program. The key element that distinguishes a computer worm from a computer virus is that human interaction is required to facilitate the spread of a virus.

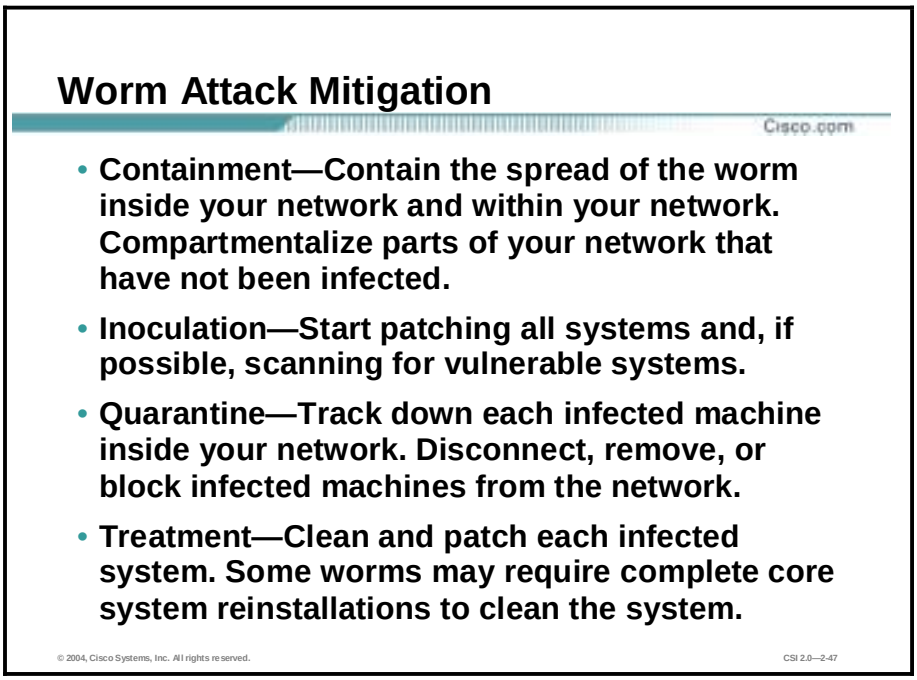

Worm attack mitigation requires diligence on the part of system and network administration staff. Coordination between system administration, network engineering, and security operations personnel is critical in responding effectively to a worm incident. The following are the recommended steps for worm attack mitigation:

- Containment
- Inoculation
- Quarantine
- Treatment

Typical incident response methodologies can be subdivided into six major categories. The following categories are based on the network service provider security (NSP-SEC) incident response methodology:

- Preparation—Acquire the resources to respond.
- Identification—Identify the worm.
- Classification—Classify the type of worm.
- Traceback—Trace the worm back to its origin.
- Reaction—Isolate and repair the affected systems.
- Post mortem—Document and analyze the process used for the future.

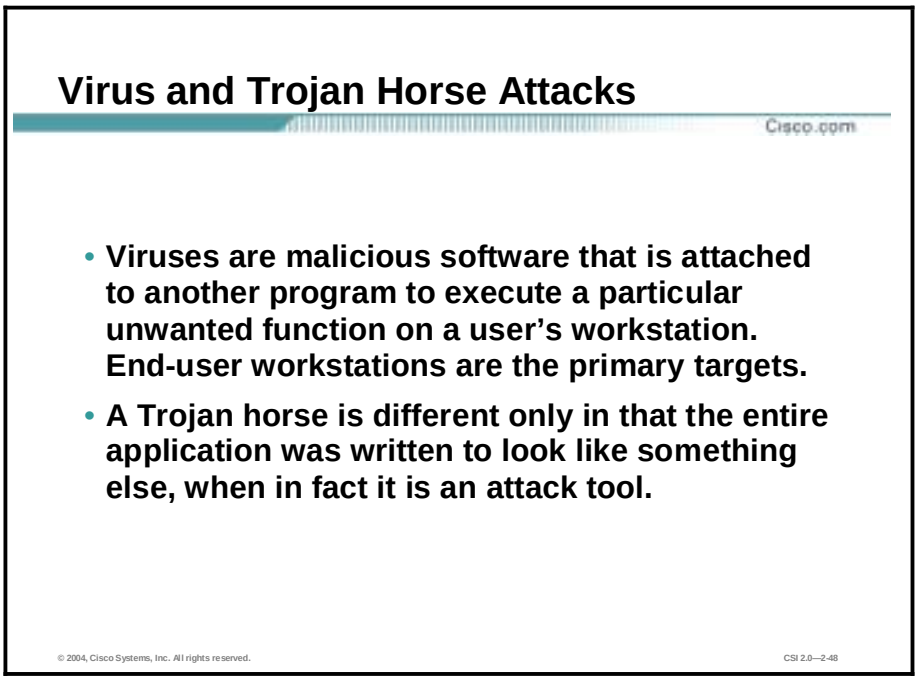

The primary vulnerabilities for end-user workstations are viruses and Trojan horse attacks. Viruses are malicious software that is attached to another program to execute a particular unwanted function on a user's workstation. An example of a virus is a program that is attached to command.com (the primary interpreter for Windows systems) that deletes certain files and infects any other versions of command.com that it can find.

A Trojan horse is different only in that the entire application was written to look like something else, when in fact it is an attack tool. An example of a Trojan horse is a software application that runs a simple game on the user's workstation. While the user is occupied with the game, the Trojan horse mails a copy of itself to every user in the user's address book. The other users receive the game and then play it, thus spreading the Trojan horse.

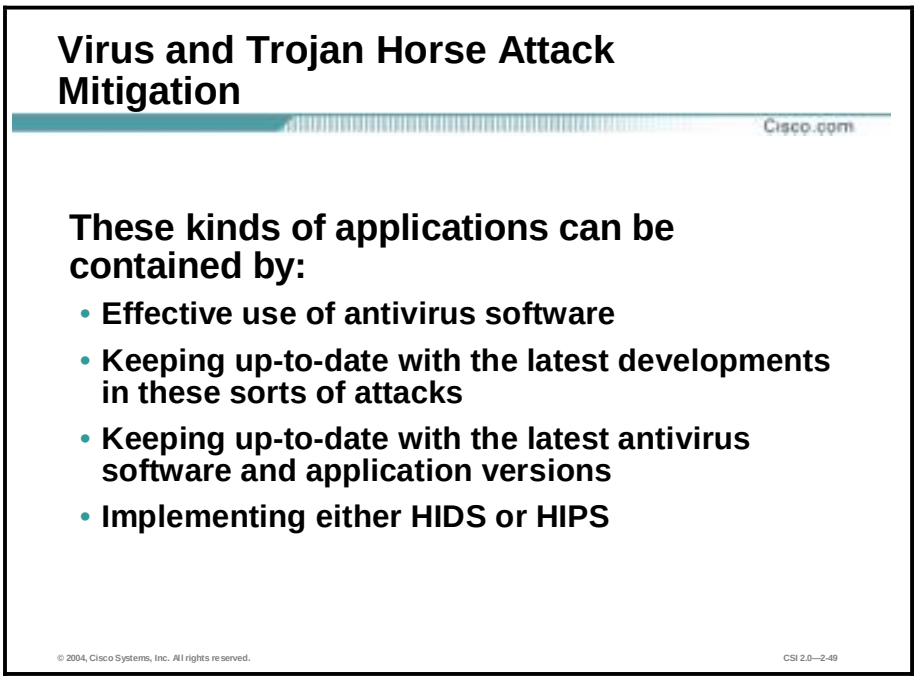

These kinds of applications can be contained through the effective use of antivirus software at the user level and potentially at the network level. Antivirus software can detect most viruses and many Trojan horse applications and prevent them from spreading in the network. Keeping up-to-date with the latest developments in these sorts of attacks can also lead to a more effective posture against these attacks. As new virus or Trojan applications are released, enterprises need to keep up-to-date with the latest antivirus software and application versions. Implementing HIDS or HIPS also helps detect and prevent a hacker from installing viruses and Trojans.

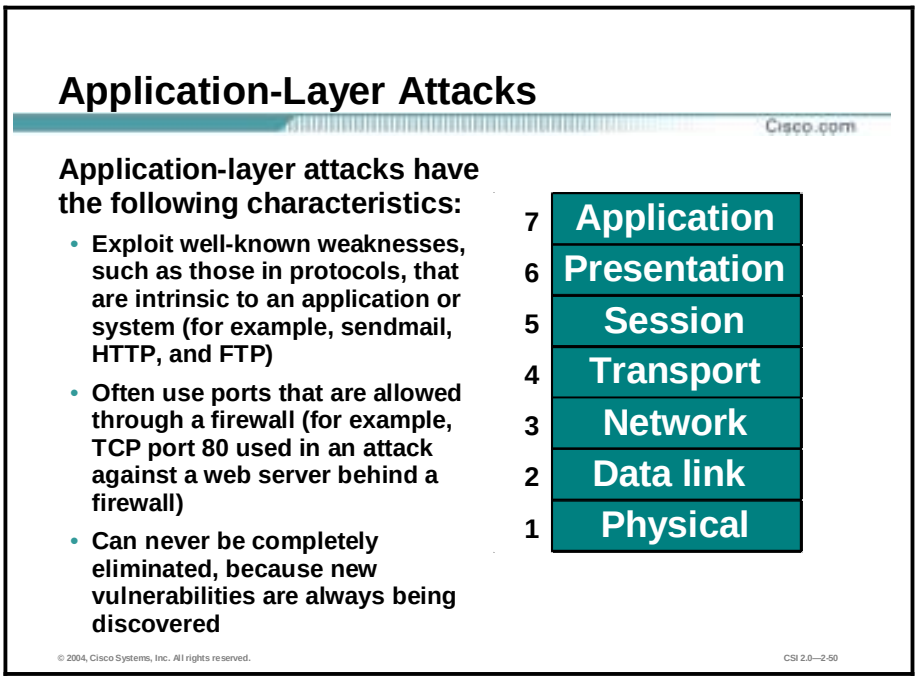

Application-layer attacks can be implemented using several different methods:

- One of the most common methods is exploiting well-known weaknesses in software commonly found on servers, such as sendmail, PostScript, and FTP. By exploiting these weaknesses, attackers can gain access to a computer with the permissions of the account running the application, which is usually a privileged, system-level account.
- Trojan horse program attacks are implemented using programs that an attacker substitutes for common programs. These programs may provide all the functionality that the normal program provides, but also include other features that are known to the attacker, such as monitoring login attempts to capture user account and password information. These programs can capture sensitive information and distribute it back to the attacker. They can also modify application functionality, such as applying a blind carbon copy to all e-mail messages so that the attacker can read all of your organization's e-mail.

  One of the oldest forms of application-layer attacks is a Trojan horse program that displays a screen, banner, or prompt that the user believes is the valid login sequence. The program then captures the information that the user enters and stores or e-mails it to the attacker. Next, the program either forwards the information to the normal login process (normally impossible on modern systems) or simply sends an expected error to the user (for example, Bad Username/Password Combination), exits, and starts the normal login sequence. The user, believing that he or she has incorrectly entered the password (a common mistake experienced by everyone), re-enters the information and is allowed access.
- One of the newest forms of application-layer attacks exploits the openness of several new technologies: the HTML specification, web browser functionality, and HTTP. These attacks, which include Java applets and ActiveX controls, involve passing harmful programs across the network and loading them through a user's browser.

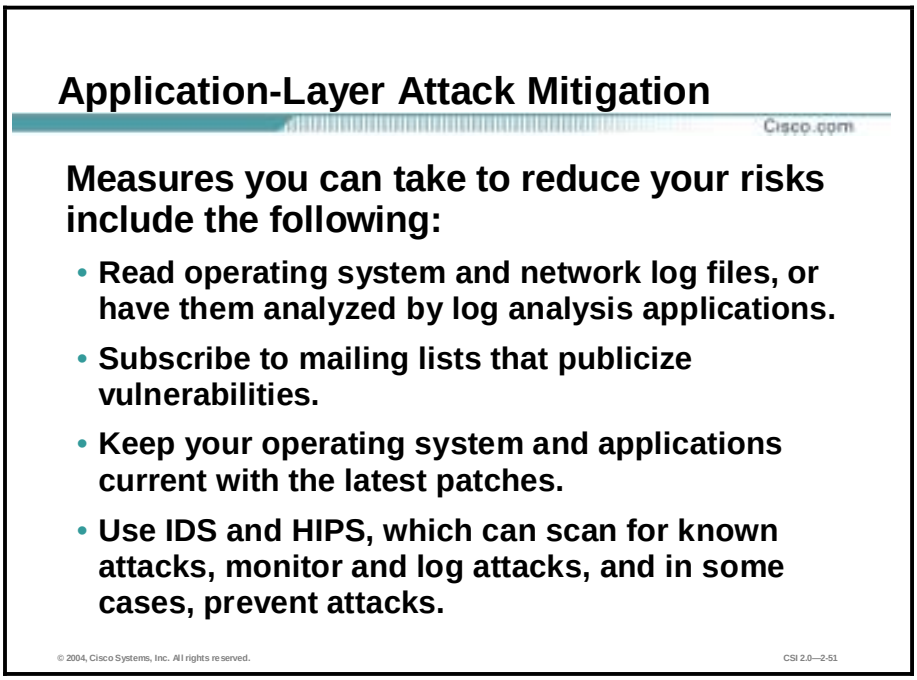

The following are some measures you can take to reduce your risks for application-layer attacks:

- Read operating system and network log files or have them analyzed—It is important to review all logs and take action accordingly.
- Subscribe to mailing lists that publicize vulnerabilities—Most application and operating system vulnerabilities are published on the Web by various sources.
- Keep your operating system and applications current with the latest patches—Always test patches and fixes in a nonproduction environment. This practice prevents downtime and keeps errors from being generated unnecessarily.
- Use IDS and HIPS to scan for known attacks, monitor and log attacks, and in some cases, prevent attacks—The use of IDS and HIPS can be essential to identifying security threats and mitigating some of those threats. In most cases, it can be done automatically.

Management Protocols and Functions

The protocols used to manage your network can in themselves be a source of vulnerability. This topic examines common management protocols and how they can be exploited.

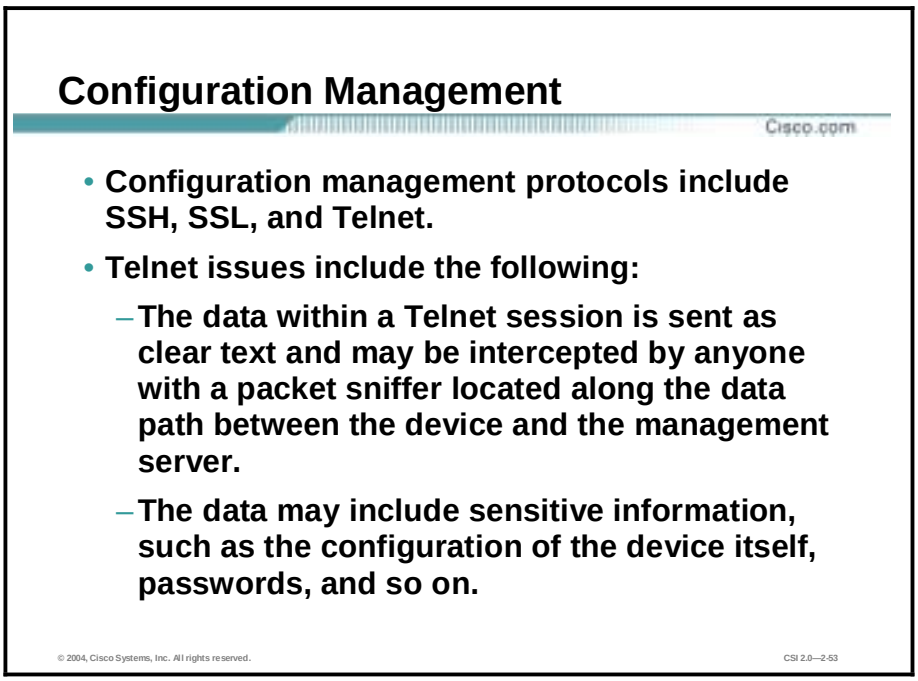

If the managed device does not support any of the recommended protocols, such as SSH and SSL, Telnet may have to be used (although this protocol is not highly recommended). The network administrator should recognize that the data within a Telnet session is sent as clear text and may be intercepted by anyone with a packet sniffer located along the data path between the managed device and the management server. The clear text may include important information, such as the configuration of the device itself, passwords, and other sensitive data.

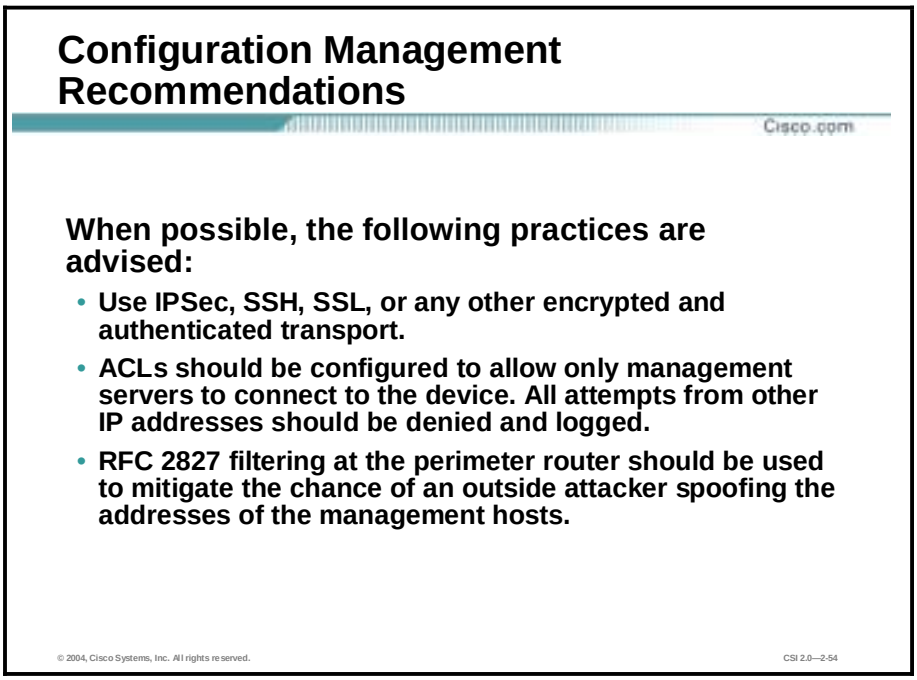

Regardless of whether SSH, SSL, or Telnet is used for remote access to the managed device, access control lists (ACLs) should be configured to allow only management servers to connect to the device. All attempts from other IP addresses should be denied and logged. RFC 2827 filtering at the ingress router should also be implemented to reduce the chance of an attacker from outside the network spoofing the addresses of the management hosts.

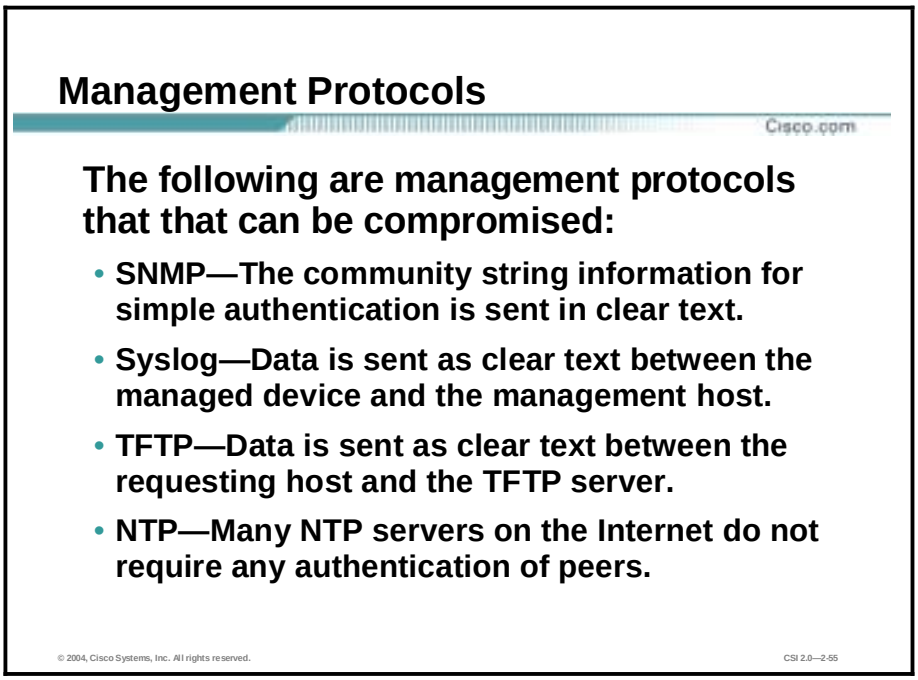

Simple Network Management Protocol (SNMP) is a network management protocol that can be used to retrieve information from a network device (commonly referred to as read-only access) or to remotely configure parameters on the device (commonly referred to as read-write access). SNMP uses passwords, called community strings, within each message as a very simple form of security. Unfortunately, most implementations of SNMP on networking devices today send the community string in clear text along with the message. Therefore, SNMP messages may be intercepted by anyone with a packet sniffer located along the data path between the device and the management server, and the community string may be compromised.

Syslog, which is information generated by a device that has been configured for logging, is sent as clear text between the managed device and the management host. Syslog has no packet-level integrity checking to ensure that the packet contents have not been altered in transit. An attacker may alter Syslog data in order to confuse a network administrator during an attack.

Trivial File Transfer Protocol (TFTP) is used for transferring configuration or system files across the network. TFTP uses UDP for the data stream between the requesting host and the TFTP server.

As with other management protocols that send data in clear text, the network administrator should recognize that the data within a TFTP session might be intercepted by anyone with a packet sniffer located along the data path between the device and the management server. Where possible, TFTP traffic should be encrypted within an IPSec tunnel in order to reduce the chance of its being intercepted.

Network Time Protocol (NTP) is used to synchronize the clocks of various devices across a network. Synchronization of the clocks within a network is critical for digital certificates and for correct interpretation of events within Syslog data.

A secure method of providing clocking for the network is for network administrators to implement their own master clocks for private networks synchronized to Coordinated Universal Time (UTC) via satellite or radio. However, clock sources are available for synchronization via

the Internet, for network administrators who do not wish to implement their own master clocks because of cost or other reasons.

An attacker could attempt a DoS attack on a network by sending bogus NTP data across the Internet in an attempt to change the clocks on network devices in such a manner that digital certificates are considered invalid. Further, an attacker could attempt to confuse a network administrator during an attack by disrupting the clocks on network devices. This scenario would make it difficult for the network administrator to determine the order of Syslog events on multiple devices.

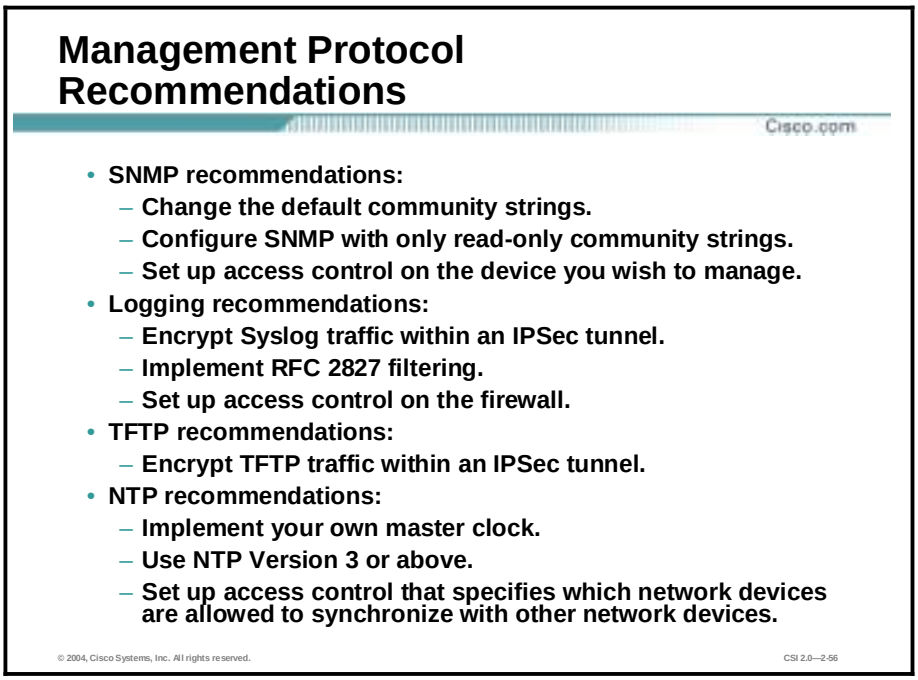

The following are SNMP recommendations:

- Change easy-to-guess default community strings.
- Configure SNMP with only read-only community strings.
- Set up access control on the device you wish to manage via SNMP to allow access by only the appropriate management hosts.

When possible, the following practices are advised:

- Encrypt Syslog traffic within an IPSec tunnel.
- When allowing Syslog access from devices on the outside of a firewall, you should implement RFC 2827 filtering at the perimeter router.
- ACLs should also be implemented on the firewall in order to allow Syslog data from only the managed devices themselves to reach the management hosts.
- When possible, TFTP traffic should be encrypted within an IPSec tunnel in order to reduce the chance of its being intercepted.

The following are NTP recommendations:

- Implement your own master clock for private network synchronization.
- Use NTP Version 3 or above because these versions support a cryptographic authentication mechanism between peers.
- Use ACLs that specify which network devices are allowed to synchronize with other network devices.

Summary

This topic summarizes what you learned in this lesson.

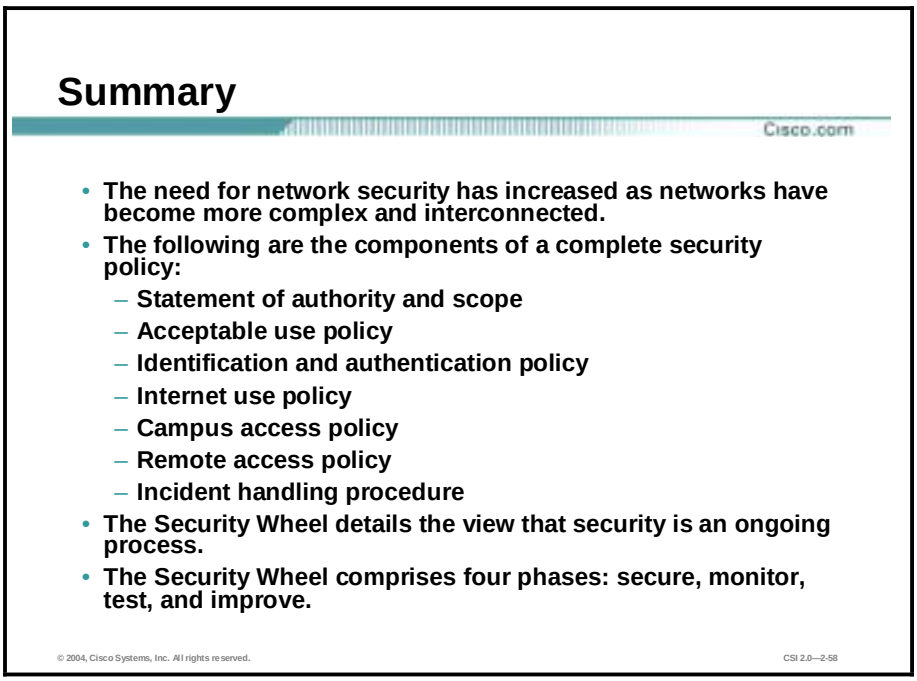

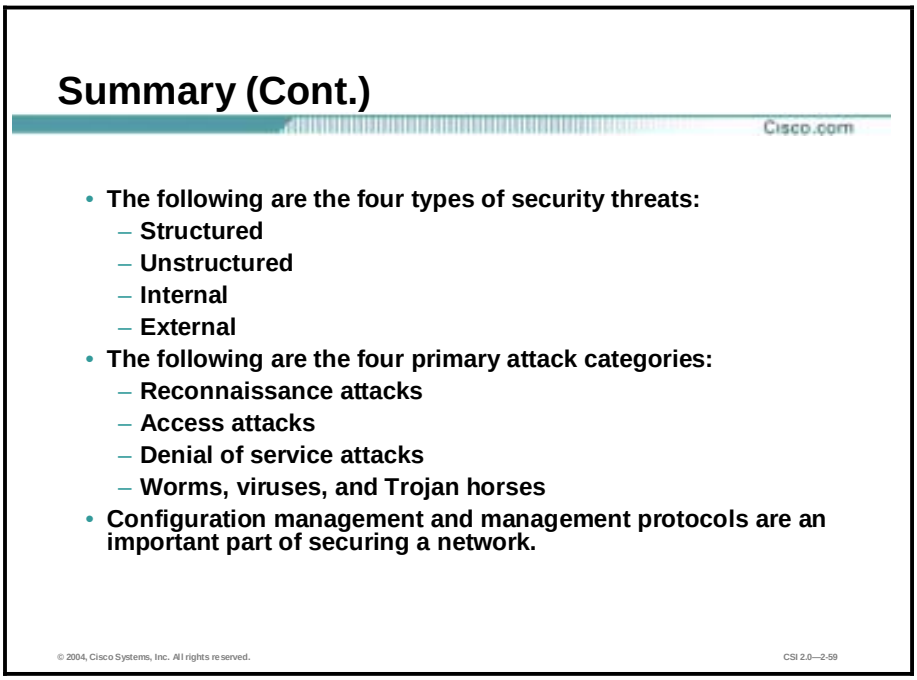

3

SAFE Blueprint Overview

Overview

This lesson introduces and gives an overview of the SAFE Blueprint. Cisco has significantly enhanced the SAFE Blueprint and extended network security and virtual private network (VPN) options to small branch offices, teleworkers, small-to-enterprise networks, IP telephony networks, and wireless networks. It includes the following topics:

- Objectives
- SAFE Blueprint overview
- Design fundamentals
- SAFE axioms
- Summary

Objectives

This topic lists the lesson's objectives.

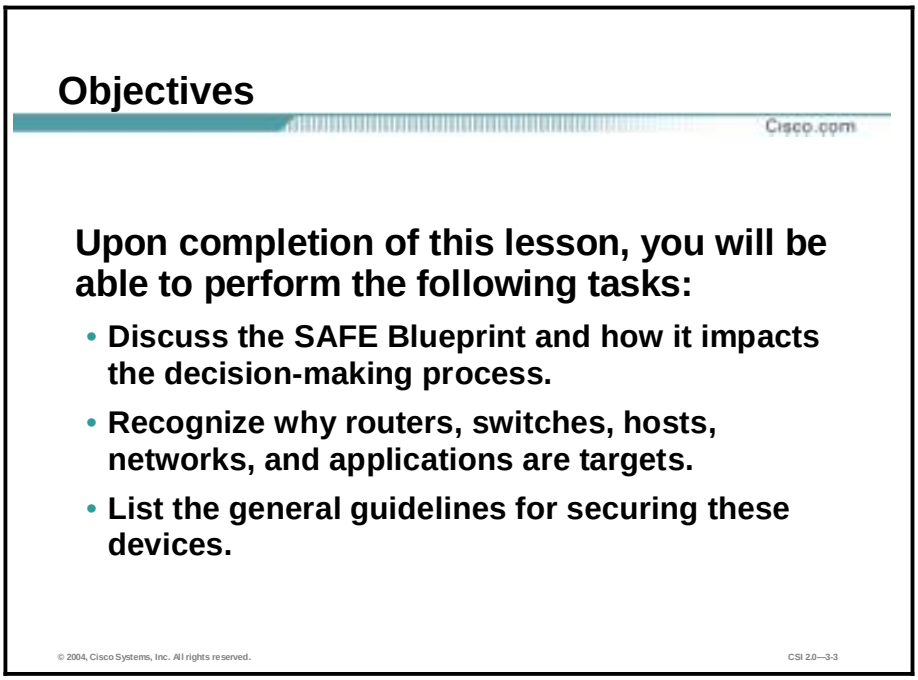

SAFE Blueprint Overview

SAFE emulates as closely as possible the functional requirements of today's networks. This topic provides an architectural overview of the SAFE white papers.

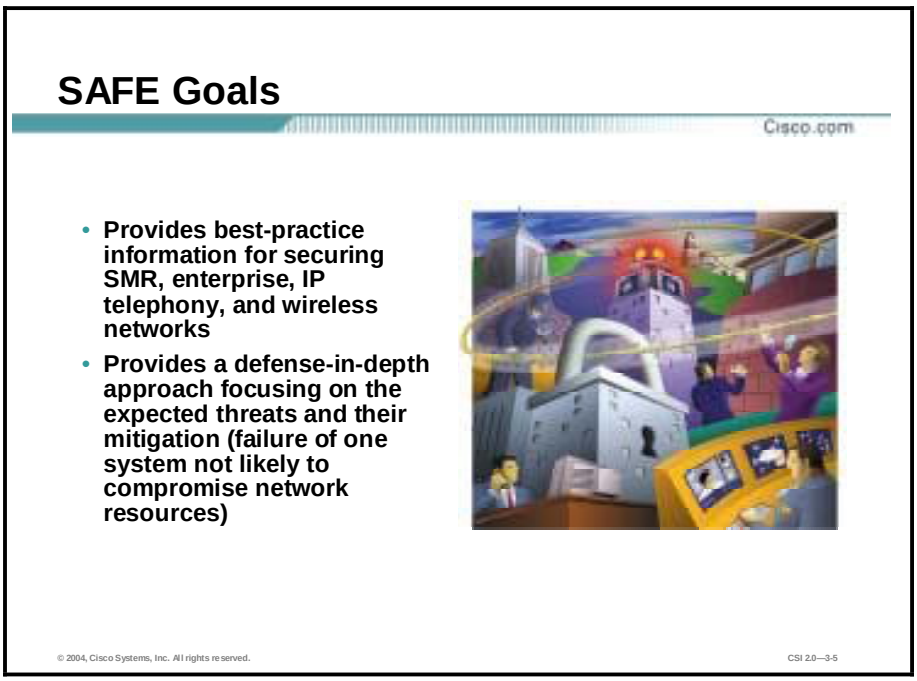

SAFE Blueprint serves as a guide to network designers considering the security requirements of their network. SAFE takes a defense-in-depth approach to network security design. This type of design focuses on the expected threats and their methods of mitigation, rather than on locations of firewalls, intrusion detection systems, and other mitigation technologies. This strategy results in a layered approach to security where the failure of one security system is not likely to lead to the compromise of network resources. SAFE is based on Cisco products and those of Cisco partners.

SAFE Blueprint focuses heavily on threats encountered in networks today. Network designers who understand these threats can better decide where and how to deploy mitigation technologies. Without a full understanding of the threats involved in network security, deployments tend to be incorrectly configured, are too focused on security devices, or lack threat response options. By taking the threat-mitigation approach, SAFE Blueprint should provide network designers with information for making sound network security choices.

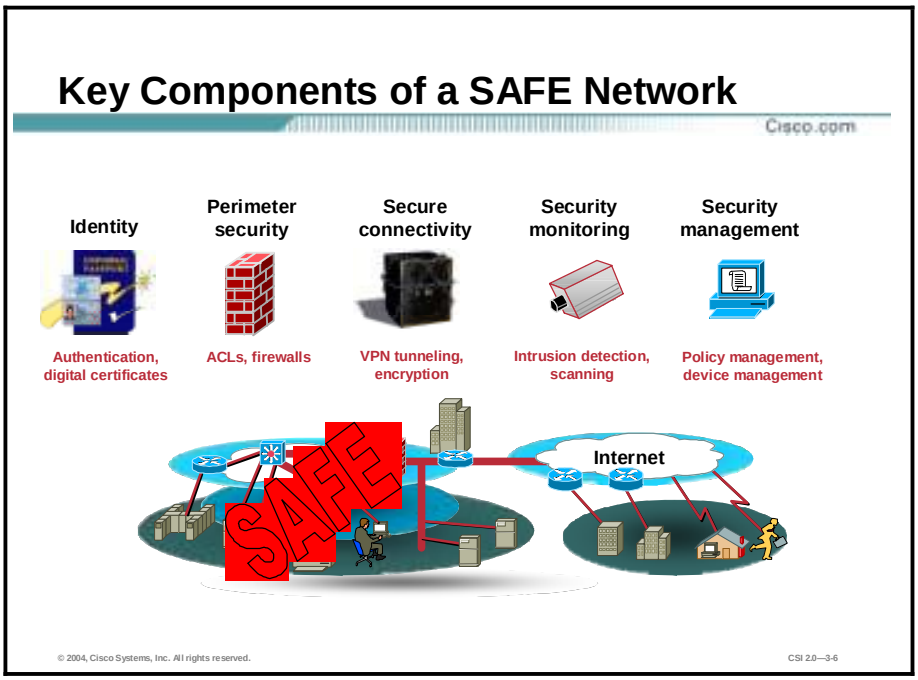

The key components of a SAFE network are fundamental to the success of an implementation. These key components are broken down as follows:

- Identity—Authentication and digital certificates
- Perimeter security—Access control lists (ACLs) and firewalls
- Secure connectivity—VPN tunneling and encryption
- Security monitoring—Intrusion detection and scanning
- Security management—Policy management, device management, and directory services

The SAFE Blueprint identifies these components as fundamental in protecting all networks including small, midsize, remote access, enterprise, IP telephony, and wireless networks.

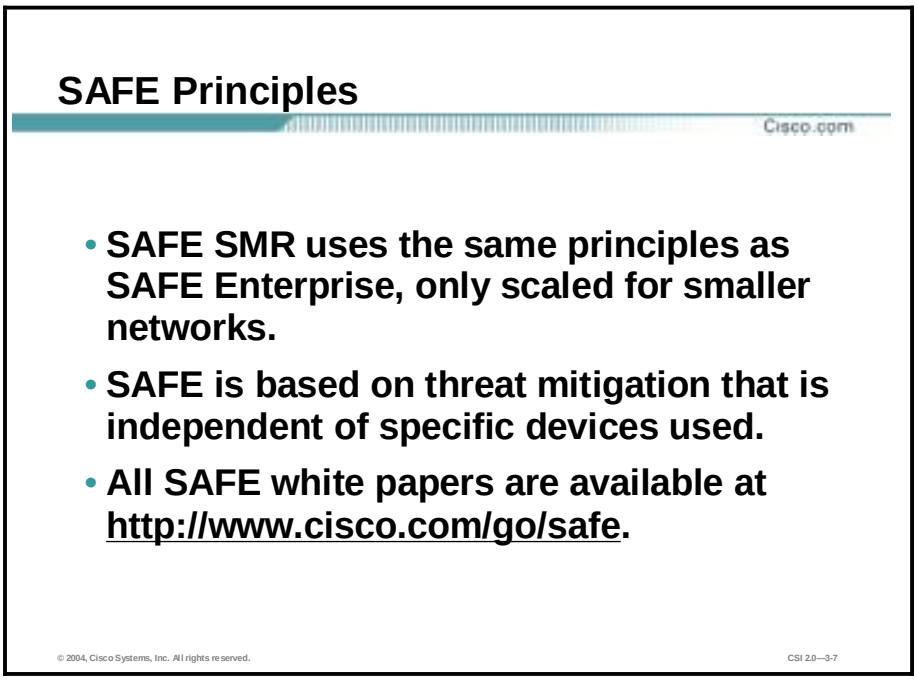

SAFE Extending the Security Blueprint to Small, Midsize, and Remote-User Networks (SAFE SMR) principles were developed to take the principles of SAFE Enterprise and size them appropriately for smaller networks. This includes branches of larger enterprises as well as standalone, small-to-midsize security deployments. It also includes information on remote user networks, such as teleworkers and mobile workers. The principles are not necessarily device specific. The design considerations used for this course are based on Cisco products and those of its partners.

For further information on SAFE white papers, go to http://www.cisco.com/go/safe.

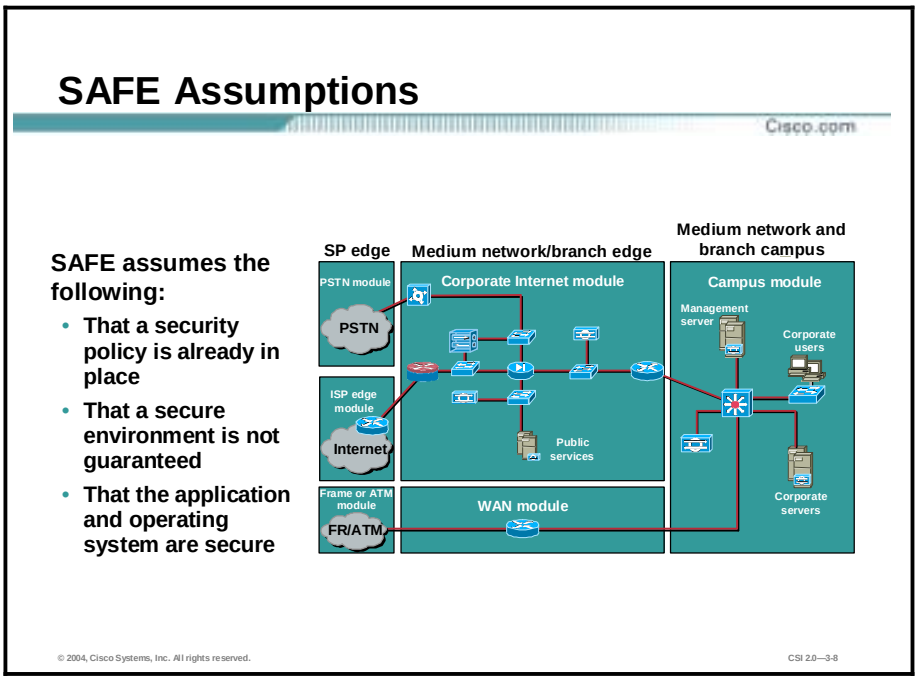

SAFE makes the following assumptions:

- The security policy is already in place—Deploying security technologies without an associated policy is not recommended.
- SAFE does not guarantee a secure environment—Following the guidelines in this course or the SAFE Blueprint does not guarantee a secure environment, nor does it guarantee that you will prevent all penetrations. However, you can achieve reasonable security by doing the following:
  - Establishing a good security policy
  - Staying up-to-date on the latest developments in the hacker and security communities
  - Maintaining and monitoring all systems with sound system-administration practices
  - Following the guidelines of this course
- Application and operating system vulnerabilities are not comprehensively covered—Proper application and operating system monitoring and maintenance is understood as one of the fundamentals of network security and is therefore not covered in depth in this course.

Design Fundamentals

Implementation decisions vary, depending on the network functionality required. This topic covers the SAFE design objectives that guide the decision-making process.

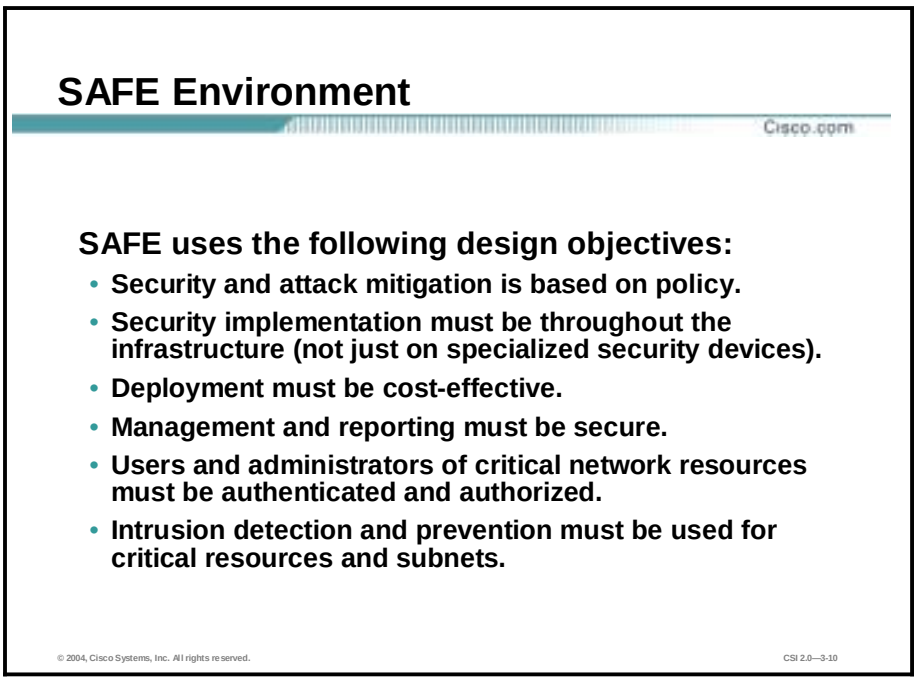

SAFE uses the following objectives, which are based on the SAFE Blueprint:

- Security and attack mitigation is based on policy—A properly implemented security policy without proper security practices can be less effective at mitigating the threat to enterprise resources than a comprehensive security product implementation without an associated policy. Cisco assumes that a security policy has been developed and implemented appropriately.
- Security implementation must be throughout the infrastructure—It is important to understand that network security extends far beyond a simple perimeter. It is necessary to take an overall approach to network security, including all types of threats.
- Deployment must be cost-effective—At many points in the network design process, you need to choose between using integrated functionality in a network device and using a specialized functional appliance. Integrated functionality is a major consideration in the implementation of a SAFE SMR network for cost-effectiveness.
- Management and reporting must be secure—It is recommended that management of devices inside the “private” network use Out-of-Band Management whenever necessary. Other circumstances such as location, budget, and so on affect this decision, as well as when devices outside the network require management and reporting. In these cases In-Band Management may be necessary.
- Users and administrators of critical network resources must be authenticated and authorized—It’s always necessary to ensure that users and administrators are accessing network resources with appropriate authentication and authorization such as digital certificates, Terminal Access Controller Access Control System Plus (TACACS+), and key exchange.

- Intrusion detection and prevention must be used for critical resources and subnets—Deployment of intrusion detection and prevention is necessary to mitigate many of the expected threats discussed in this course.

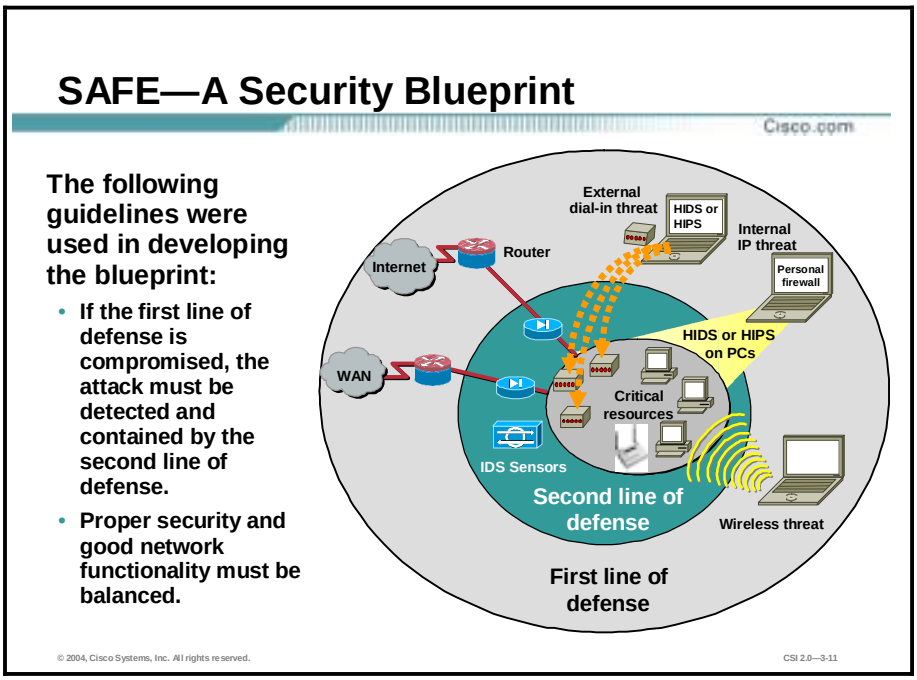

First and foremost, SAFE is a security blueprint. A SAFE Blueprint must prevent most attacks from successfully affecting valuable network resources. However, while being secure, the network must continue to provide critical services that users expect.

The following guidelines are used when developing the blueprint:

- If the first line of defense is compromised, the attack must be detected and contained by the second line of defense.
- Proper security and good network functionality must be balanced.

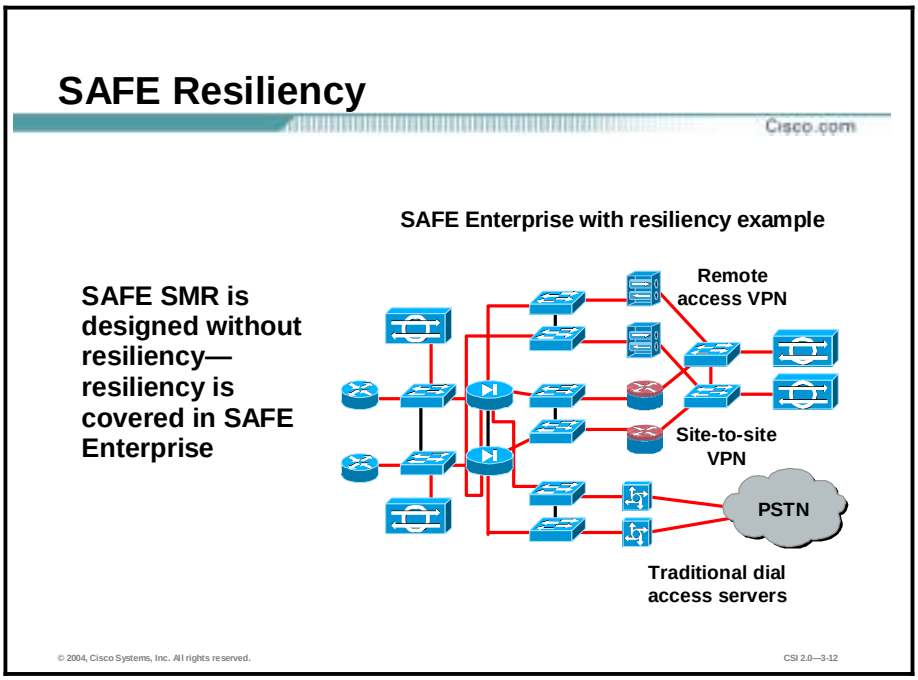

Unlike SAFE Enterprise, SAFE SMR is designed without resiliency. The approach taken for SAFE SMR is to provide a security architecture without resilient and redundant practices for both cost savings and ease of integration. Those interested in designing secure networks in a resilient environment should concentrate on the SAFE Enterprise method of design.

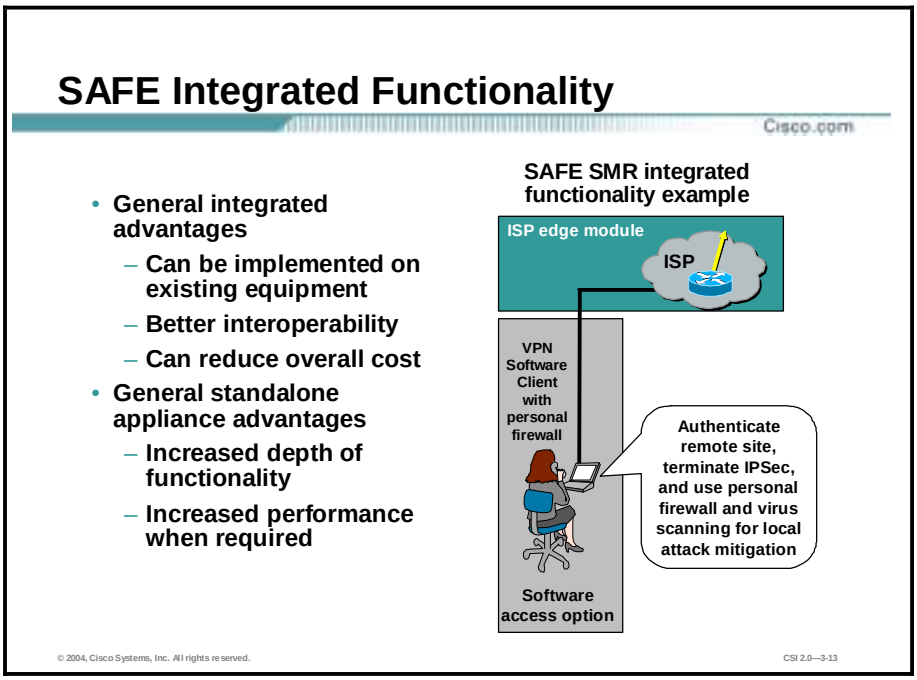

The advantages to integrated functionality are as follows:

- Can be implemented on existing equipment—Many devices, such as routers and firewalls, can provide multiple functions, including routing and packet filtering.
- Offers better interoperability—A router running a Cisco IOS firewall is less likely to introduce problems into the network than two separate devices.
- Can reduce overall cost—It is less expensive to integrate functionality into a single device rather than purchasing two separate devices.

The advantages to standalone appliances are as follows:

- Increased depth of functionality—A standalone appliance can provide functionality that is not available in an integrated product.
- Increased performance when required—A standalone appliance can provide bandwidth and throughput advantages to the network.

Throughout the SAFE Blueprint, both integrated systems and appliances are used. When the example design requirements used for the development of the architecture did not dictate a specific choice, the developers of SAFE SMR opted to go with integrated functionality in order to reduce the overall cost of the solution.

Integrated functionality is often attractive because you can implement it on existing equipment, or because the features can interoperate with the rest of the device to provide a better functional solution. Appliances are often used when the depth of functionality required is very advanced or when performance needs require the use of specialized hardware.

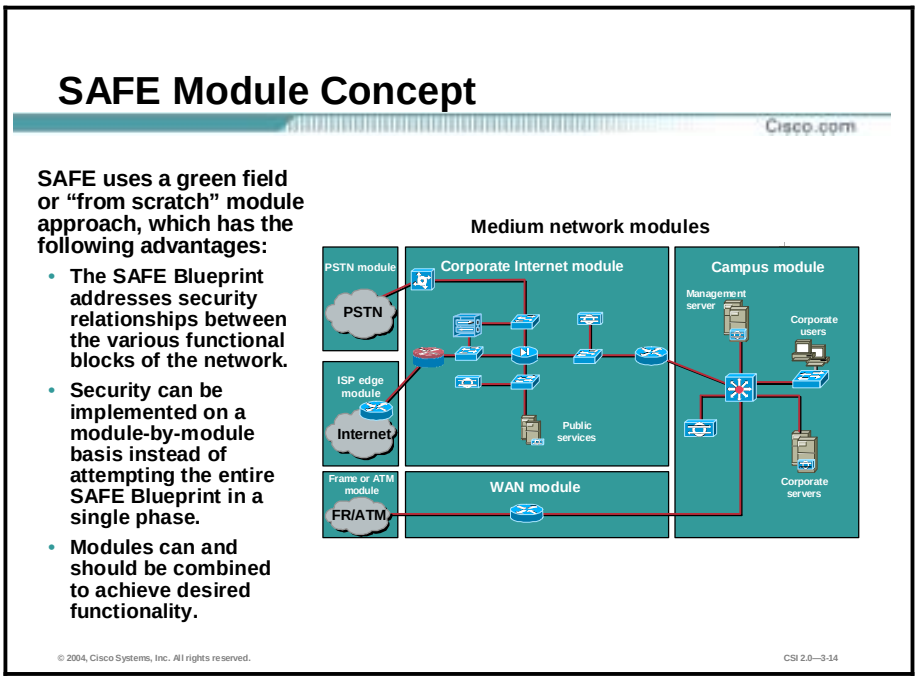

Although it is true that most networks cannot be easily dissected into clear-cut modules, the green field or "from scratch" modular approach provides a guide for implementing different security functions throughout the network. Engineers are not expected to design their networks to be identical to the SAFE implementation, but rather to use a combination of the modules described and integrate them into the existing network. The advantages to this approach are as follows:

- The architecture addresses security relationships between the various functional blocks of the network.
- Security can be implemented on a module-by-module basis instead of attempting the entire architecture in a single phase.
- Modules can and should be combined to achieve desired functionality.

The diagram in the figure is an example of a SAFE midsize network and its respective modules, which include the campus module, corporate Internet module, and the various edge modules.

SAFE Axioms

This topic covers the axioms used by SAFE.

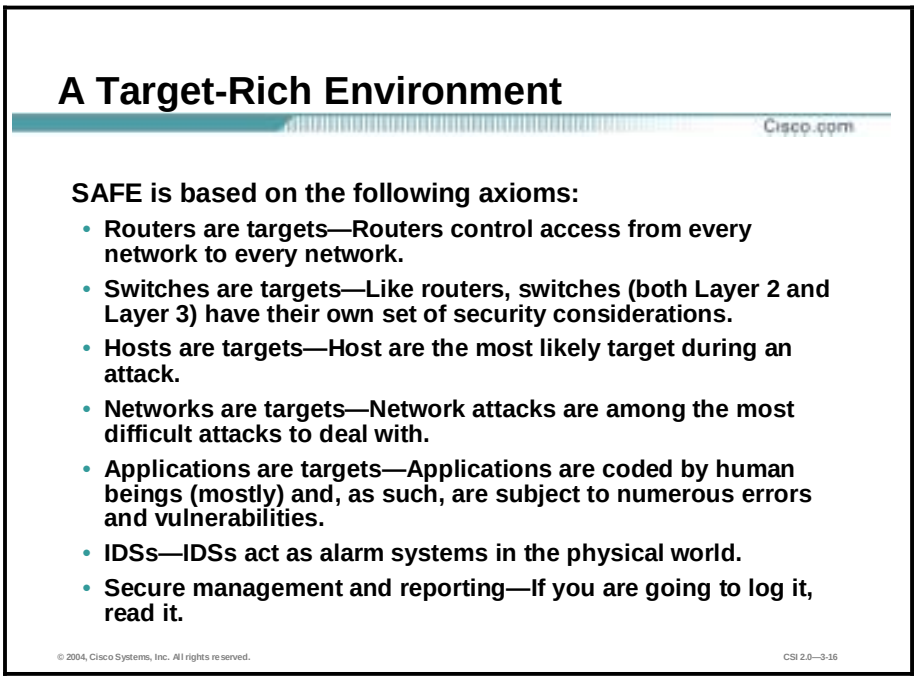

The following are axioms for identifying appliances and applications that are primary network targets:

- Routers are targets—Router security is a critical element in any security deployment.
- Switches are targets—Unlike routers, not as much information is available about the security risks in switches and what can be done to mitigate those risks. Most of the risks associated with routers are applicable to switches as well.
- Hosts are targets—A host presents some of the most difficult challenges from a security perspective and is the most likely target during an attack. There are numerous hardware platforms, operating systems, and applications, all of which have updates, patches, and fixes available at different times.
- Networks are targets—The attacks on networks are the most difficult challenges because, typically, they take advantage of an intrinsic characteristic in the way your network operates.
- Applications are targets—Attacks on applications can be benign or malignant. It is the malignant attacks that require the most attention.
- Intrusion detection systems (IDSs) act as alarms—When an IDS detects something that it considers an attack, it can either take corrective action itself or notify a management system for actions by the administrator.
- Secure management and reporting is important—Logging and reading information from many devices can be challenging. It is important to be able to identify priority events. Then you can take action for those events and deal with them appropriately.

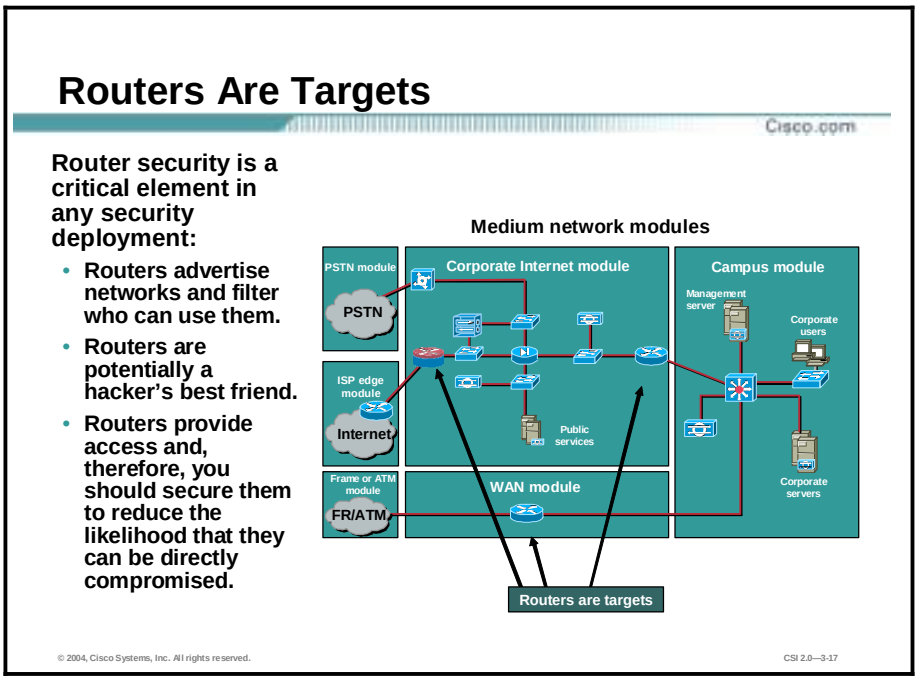

Routers control access from network to network. They advertise networks and filter what can use them, and they are potentially a hacker's best friend. Because of this, router security is a critical element in any security deployment. It is important for security professionals to be completely up-to-date on current router documentation and possible threats to routers. The following URL provides the most current Cisco documentation: http:/www.cisco.com/warp/public/707/21.html.

The figure indicates the location of the routers in a SAFE SMR network example.

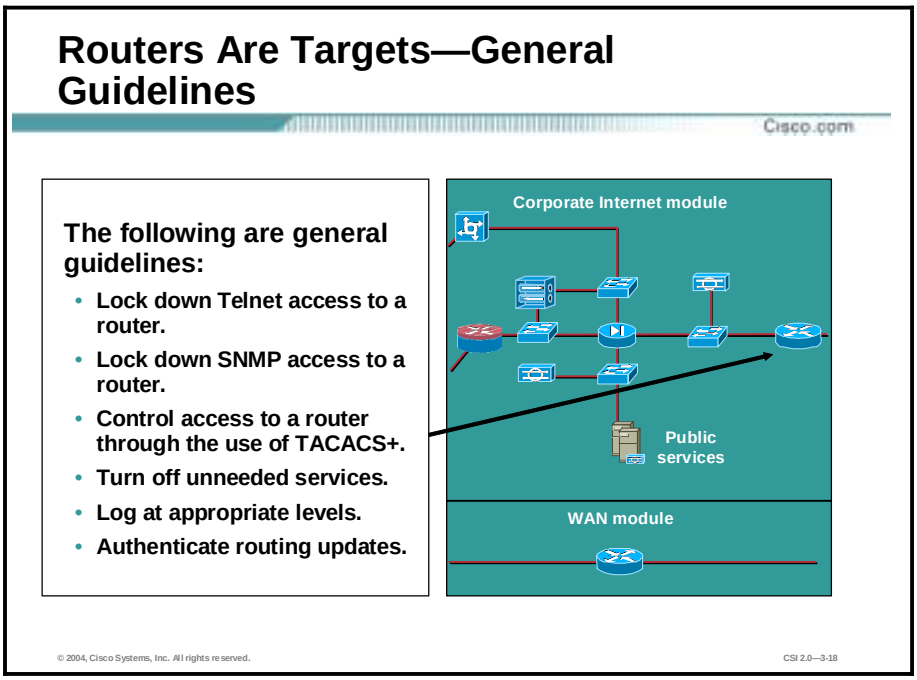

The following guidelines should be followed when securing routers:

- Lock down Telnet access to a router—Interactive Telnet access is available not only on the standard Telnet TCP port (port 23), but on a variety of higher-numbered ports as well. All interactive access mechanisms use the Cisco IOS teletype (TTY) abstraction (in other words, they all involve sessions on “lines” of one sort or another). Local asynchronous terminals and dialup modems use standard lines, known as TTYs. Remote network connections, regardless of the protocol, use virtual teletypes (VTYs). The best way to protect a system is to make certain that appropriate controls are applied on all lines, including both VTY lines and TTY lines.
- Lock down Simple Network Management Protocol (SNMP) access to a router—SNMP is widely used for router monitoring, and frequently for router configuration changes as well. Unfortunately, version 1 of SNMP, which is the most commonly used, uses a very weak authentication scheme based on a community string, which is a fixed password transmitted over the network without encryption. If at all possible, use SNMP version 2, which supports a Message Digest 5 (MD5)-based digest authentication scheme, and allows for restricted access to various management data. If you must use SNMP version 1, you should be careful to choose unobvious community strings (not, for example, “public” or “private”). If at all possible, you should avoid using the same community strings for all network devices; use a different string or strings for each device, or at least for each area of the network. Do not make a read-only string the same as a read-write string. If possible, periodic SNMP version 1 polling should be done with a read-only community string; read-write strings should be used only for actual write operations.
- Control access to a router through the use of Terminal Access Controller Access Control System Plus (TACACS+)—TACACS+ is a protocol providing detailed accounting information and flexible administrative control over authentication and authorization processes to control unauthorized access. TACACS+ is facilitated through authentication, authorization, and accounting (AAA).

- Turn off unneeded services—As a general rule, any unnecessary service should be disabled in any router that is reachable from a potentially hostile network. The following services are sometimes useful, but should be disabled if they are not actively being used:
  - Finger
  - Network Time Protocol (NTP)
  - Cisco Discovery Protocol (CDP)
- Log at appropriate levels—It is necessary to log information on the router (for example, access, fault, and warning logs).
- Authenticate routing updates—If you are using a dynamic routing protocol that supports authentication, it is a good idea to enable that authentication. This prevents some malicious attacks on the routing infrastructure, and can also help to prevent damage caused by misconfigured rogue devices on the network.

The figure identifies the location of the router in a SAFE SMR network example.

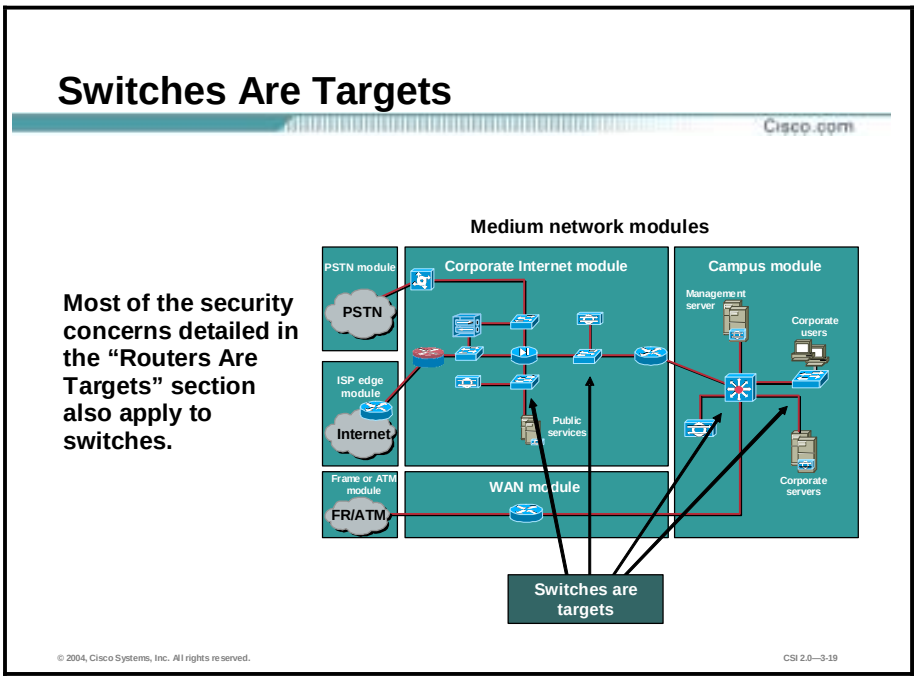

Like routers, switches (both Layer 2 and Layer 3) have their own set of security considerations. Unlike routers, not as much information is available about the security risks in switches and what can be done to mitigate those risks. Most of the security techniques detailed in the topic “Routers Are Targets” apply to switches.

The figure identifies the location of the switches in a SAFE SMR network example.

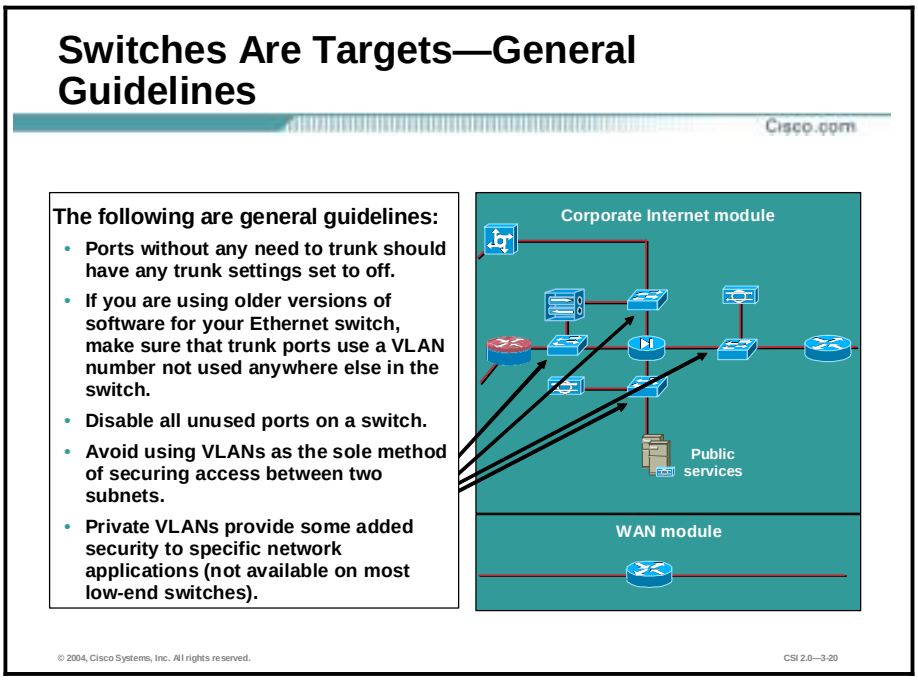

The following guidelines are in addition to router-specific guidelines and apply to both Layer 2 and Layer 3 switches:

- Ports without any need to trunk should have trunk settings set to off—This prevents a host from becoming a trunk port and receiving all traffic that would normally reside on a trunk port.
- If you are using older versions of software for your Ethernet switch, make sure that trunk ports use a VLAN number not used anywhere else in the switch—This prevents packets tagged with the same VLAN as the trunk port from reaching another VLAN without crossing a Layer 3 device.
- Disable all unused ports on a switch—This prevents hackers from plugging in to unused ports and communicating with the rest of the network.
- Avoid using VLANs as the sole method of securing access between two subnets—The capability for human error, combined with the understanding that VLANs and VLAN-tagging protocols were not designed with security in mind, makes their use in sensitive environments inadvisable.
- Private VLANs provide some added security to specific network applications (not available on most low-end switches)—They work by limiting which ports within a VLAN can communicate with ports in the same VLAN.

The figure identifies the location of the switches in the SAFE SMR network example.

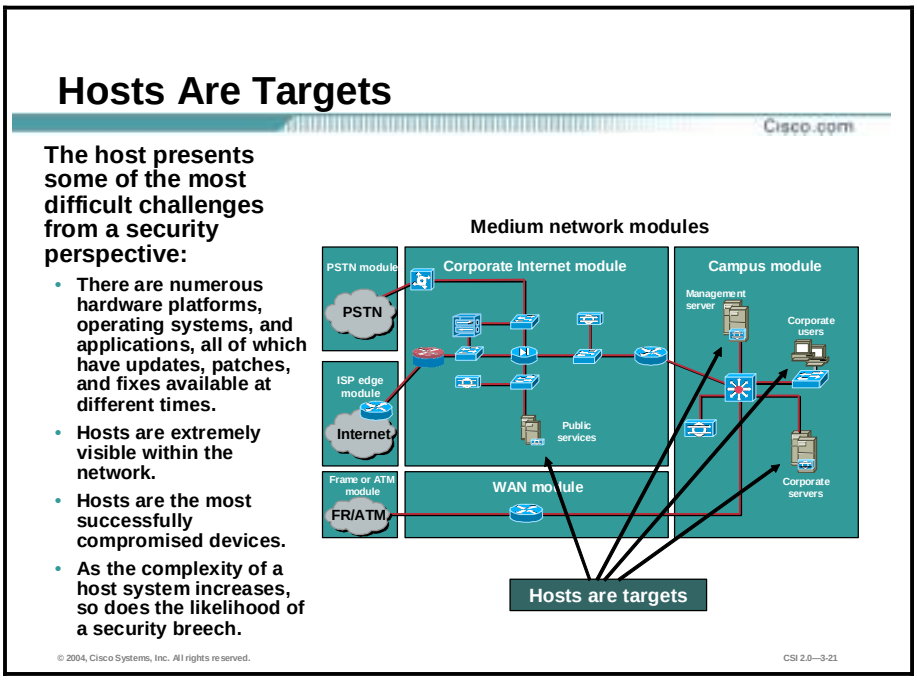

Being the most likely target during an attack, the host presents some of the most difficult challenges from a security perspective. There are numerous hardware platforms, operating systems, and applications, all of which have updates, patches, and fixes available at different times.

Because hosts provide the application services to other hosts that request them, they are extremely visible within the network. For example, many people have visited www.whitehouse.gov, which is a host, but few have attempted to access s2-0.whitehouseisp.net, which is a router. Because of this visibility, hosts are the most frequently attacked devices in any network intrusion attempt.

In part because of the aforementioned security challenges, hosts are also the most successfully compromised devices. For example, a given web server on the Internet might run a hardware platform from one vendor, a network card from another, an operating system from still another vendor, and a web server that is either open source or from yet another vendor. Additionally, the same web server might run applications that are freely distributed over the Internet, and might communicate with a database server that starts the variations all over again. That is not to say that the security vulnerabilities are specifically caused by the multisource nature of all of this, but rather that as the complexity of a system increases, so does the likelihood of a failure.

The figure identifies the location of the hosts in the SAFE SMR network example.

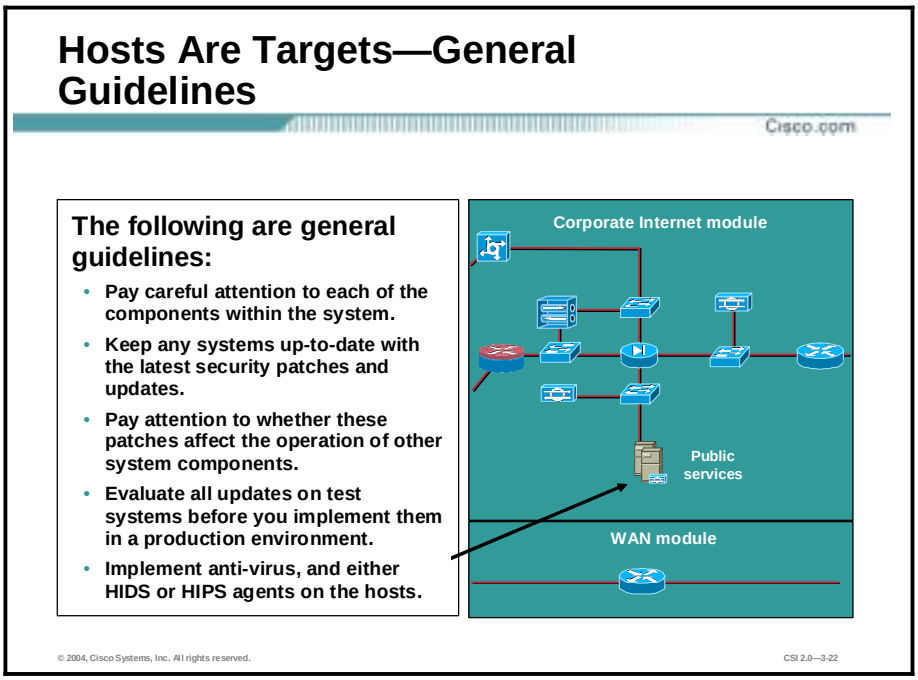

The following are guidelines that can be a major factor in maintaining a secure environment for hosts:

- Pay careful attention to each of the components within the system—These components include hardware and software.
- Keep any systems up-to-date with the latest patches, fixes, and updates—Patches and fixes are being created constantly for software and hardware. It is a good practice to include this rule in your organization's change management policy.
- Pay attention to how these patches affect the operation of other system components—Read release notes and updates on all changes prior to implementation.
- Evaluate all updates on test systems before you implement them in a production environment—This practice ensures that changes are successful and also inhibits possible effects on other components.
- Implement anti-virus, and either HIDS or HIPS agents on the hosts—Because of the specificity of their role, HIDS or HIPS are often better at preventing specific attacks on an end point.

The figure identifies the location of the hosts in the SAFE SMR network example.

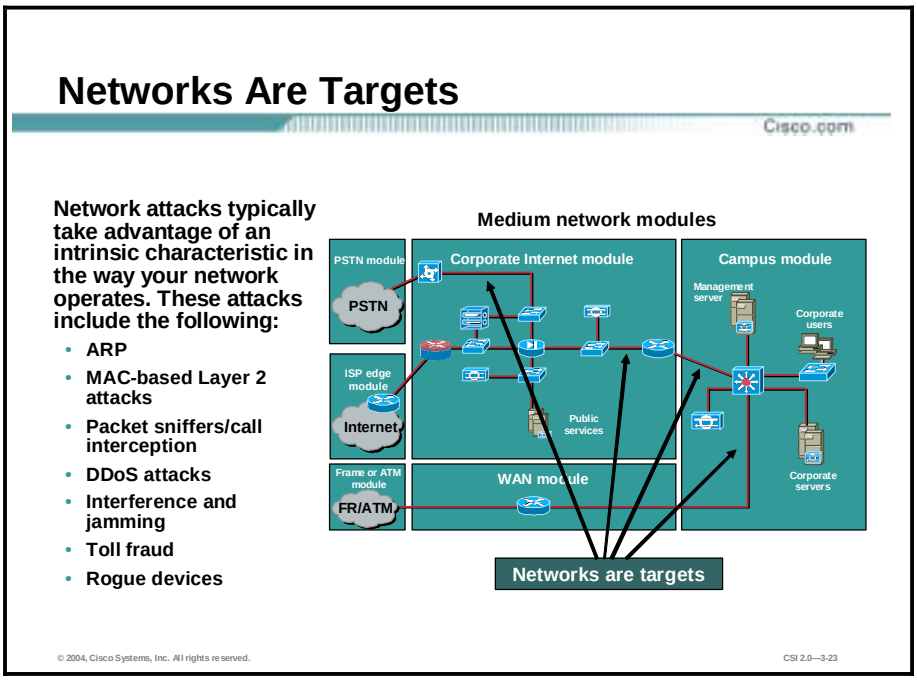

Network attacks are among the most difficult attacks to deal with because they typically take advantage of an intrinsic characteristic in the way your network operates. These attacks include:

- Address Resolution Protocol (ARP)—Identifies network component addresses for further attack.
- Media Access Control (MAC)-based Layer 2 attack—Identifies the MAC layer address for further attack.
- Packet sniffer—A software application that uses a network adapter card in promiscuous mode to capture all network packets that are sent across a particular collision domain. Call interception poses a threat to IP telephony networks.
- Distributed denial of service (DDoS) attack—Causes multiple machines to simultaneously send spurious data to an IP address. The following are common forms of DDoS attacks:
  - Internet Control Message Protocol (ICMP) floods
  - TCP SYN floods
  - UDP floods
- Interference and jamming—It is easy to interfere with wireless communications. A simple jamming transmitter can make communications impossible.
- Toll fraud—Allows unknown phones to be configured.
- Rogue devices—With access to the local switched segment, the hacker might be able to insert a phone into the voice segment with a spoofed MAC address, assume the target phone's identity, and intercept a call. Rogue access points in a wireless network can give unauthorized access to the enterprise.

The figure identifies the location of the networks in the SAFE SMR network example.

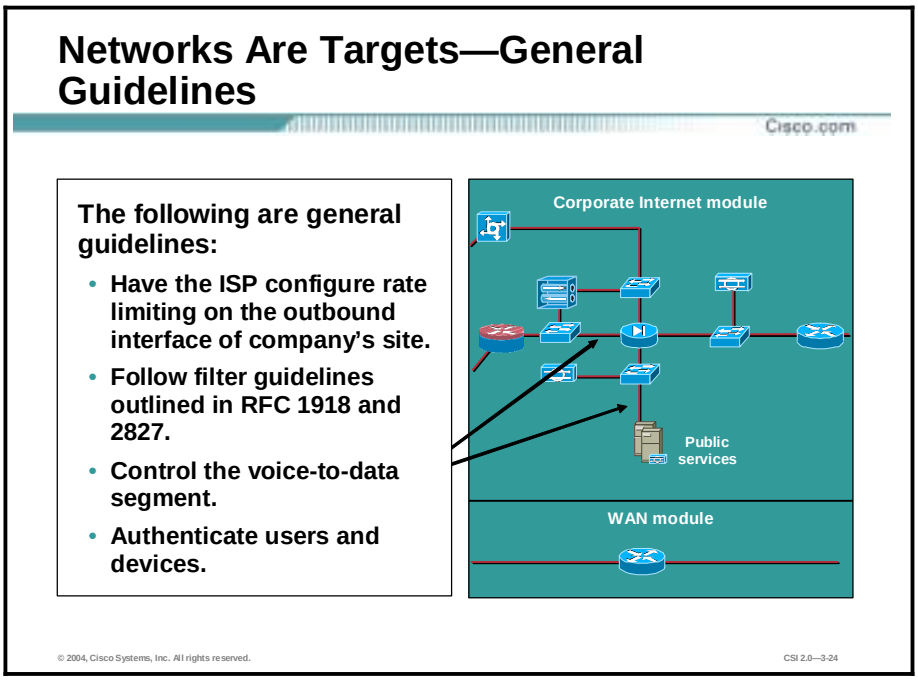

The following are guidelines for practices that can be a major factor in maintaining a secure environment for networks:

- Committed access rate (CAR) and TCP setup controls—CAR at the Internet service provider (ISP) edge and TCP setup controls at the firewall limit exposure to denial of service (DoS) attacks.
- RFC 1918—This RFC provides background on the allocation of IP addresses for private internets. It also provides implementation guidelines for companies that want to implement IP but do not want full connectivity to the Internet.
- RFC 2827—This RFC provides an effective and straightforward method for using ingress traffic filtering to prohibit DoS attacks that use forged IP addresses from being propagated from behind an ISP's aggregation point. If ISPs worldwide were to collectively implement the guidelines in RFC 2827, source address spoofing would be greatly diminished. Although this strategy does not prevent DDoS attacks, it does prevent such attacks from masking their source, which makes tracing back to the attacking networks much easier. Ask your ISP about which DDoS mitigation options they make available to their customers.
- Data and voice segmentation—Segmentation of data and voice traffic is key. IP-based telephony networking provides a means of providing telephony over the existing IP data network. However, for reasons including QoS, scalability, manageability, and security, IP telephony devices and IP data devices should be deployed on two logically disparate segments. Segmenting IP voice from the traditional IP data network greatly increases your ability to mitigate attacks and allows use of the same access, core, and distribution layers. Although the segments should be disparate, it is by no means recommended that you deploy two IP infrastructures. Technologies such as VLANs, access control, and stateful firewall provide the Layer 3 segmentation necessary to keep the voice and data segments separate at the access layer.
- Authentication of users and devices—Locking down the switched ports, segments, and services in the network will provide attack mitigation for rogue devices. As in any IP

network, there is value in limiting the capabilities of a rogue device plugged into the network. Best practices for data and voice networks include the following:

- Disabling unused ports
- Deploying a switched environment for a scalable IP phone deployment
- Statically assigning IP addresses to known MAC addresses
- Configuring the call-processing manager to deny unknown PC-based IP phone (soft Phone) access
- Using a utility such as arpwatch to monitor the MAC addresses in your voice segment
- Filtering in all segments to limit devices in unknown segments from connecting to the call-processing manager

The figure identifies the location of the networks in the SAFE SMR network example.

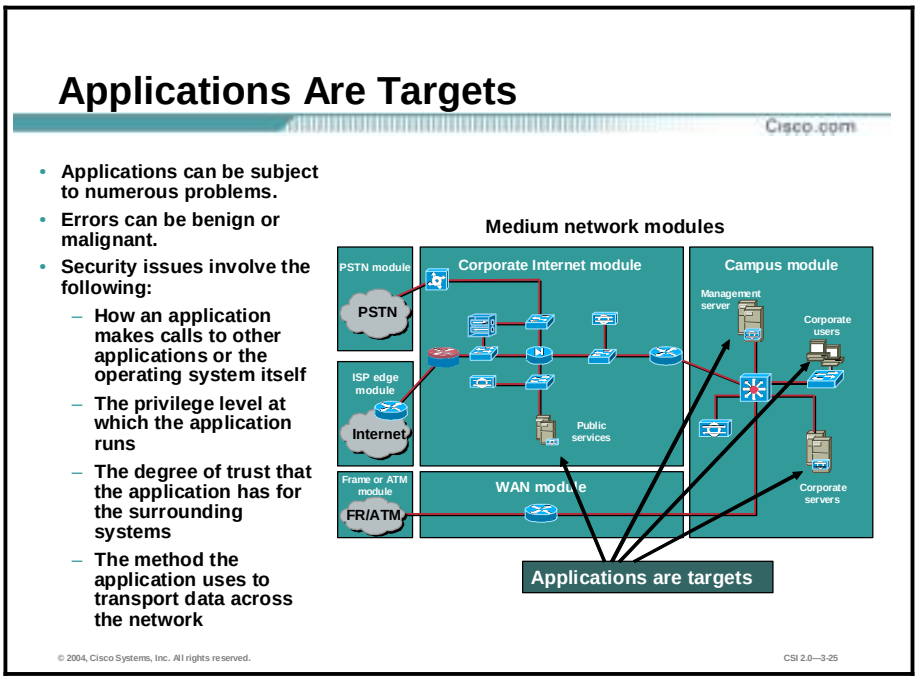

Applications are subject to numerous problems that can be benign or malignant. Security issues involve the following:

- How an application makes calls to other applications or the operating system itself
- The privilege level at which the application runs
- The degree of trust that the application has for the surrounding systems
- The method the application uses to transport data across the network

Applications are usually coded by human beings and, as such, are subject to numerous errors. These errors can be benign (for example, an error that causes your document to print incorrectly), or malignant (for example, an error that makes the credit card numbers on your database server available via anonymous FTP). It is the malignant problems that need careful attention.

The figure identifies the location of the applications in the SAFE SMR network example.

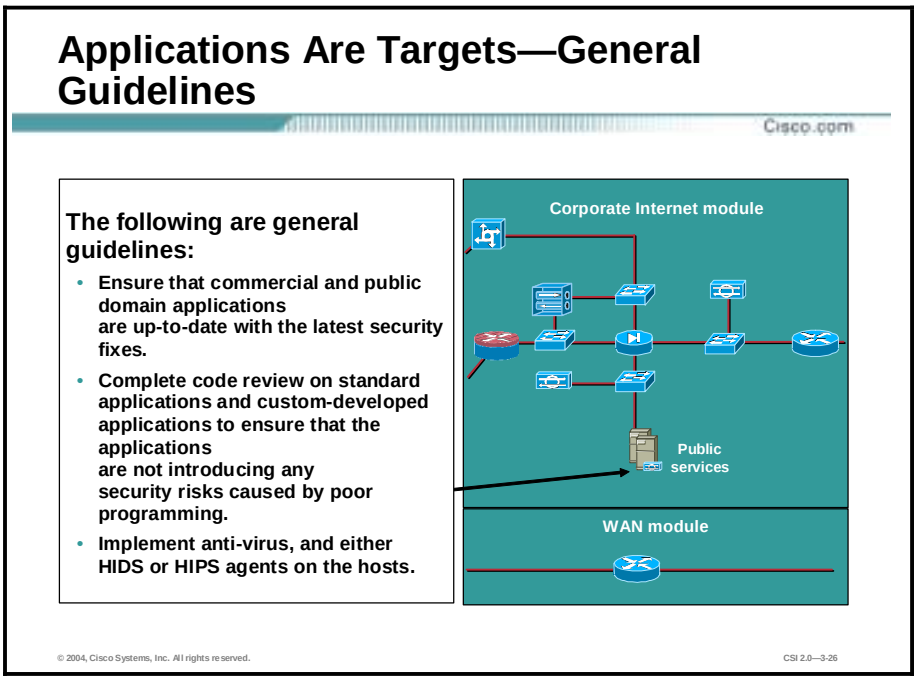

Care needs to be taken to ensure that commercial and public domain applications are up-to-date with the latest security fixes. Public domain applications, as well as custom-developed applications, also require code review to ensure that the applications are not introducing any security risks caused by poor programming. IDSs can help mitigate some of the attacks launched against applications and other functions within the network. Implement anti-virus and either HIDS or HIPS agents on the hosts for granular protection.

When changing any system or application, you must review release notes and documentation, follow your organization's change management policies, and test in a non-production environment prior to implementation.

The figure identifies the location of the applications in the SAFE SMR network example.

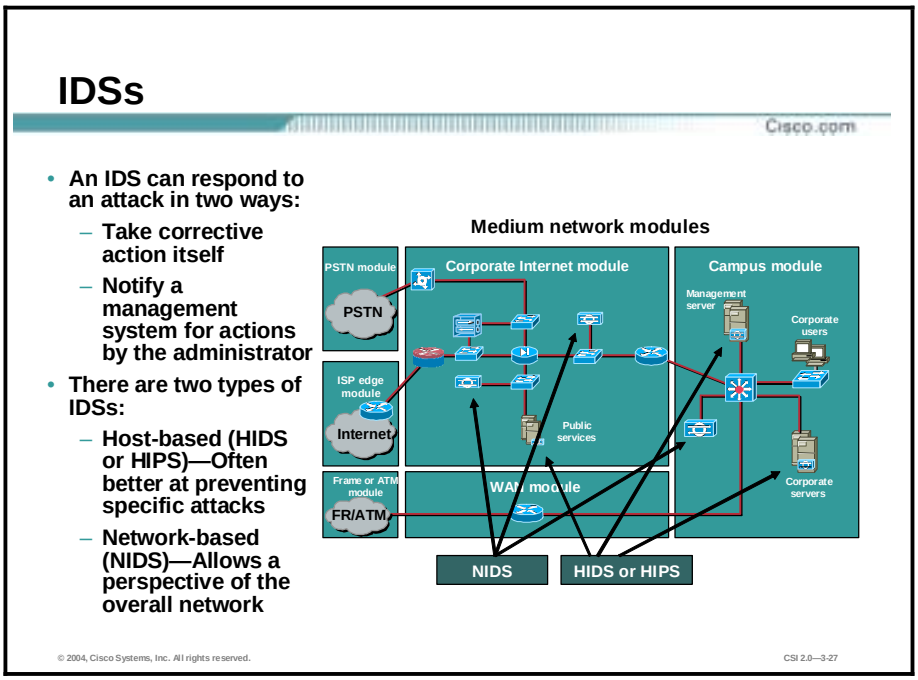

An IDS acts like an alarm system in the physical world. When an IDS detects something that it considers an attack, it can either take corrective action itself or notify a management system for actions by the administrator. Some systems are more or less equipped to respond to and prevent such an attack.

A host-based intrusion detection system (HIDS) works by intercepting operating system and application calls on an individual host. It can also operate by after-the-fact analysis of local log files. The former approach allows better attack prevention, whereas the latter approach dictates a more passive attack-response role. Because of the specificity of their role, either HIDS or HIPS are often better at preventing specific attacks than network-based IDSs (NIDSs), which usually issue only an alert upon discovery of an attack.

However, that specificity causes a loss of perspective of the overall network. This is where NIDS excels. Ideally, a combination of the two systems is recommended—HIDS or HIPS on critical hosts and NIDS looking over the whole network—for a complete IDS.

Several factors need to be considered when choosing between the types of IDSs to implement:

- Budget
- Number of devices needing to be monitored
- Topology of the network
- Number of personnel required to respond to attacks

The figure identifies the location of the IDSs in the SAFE SMR network example.

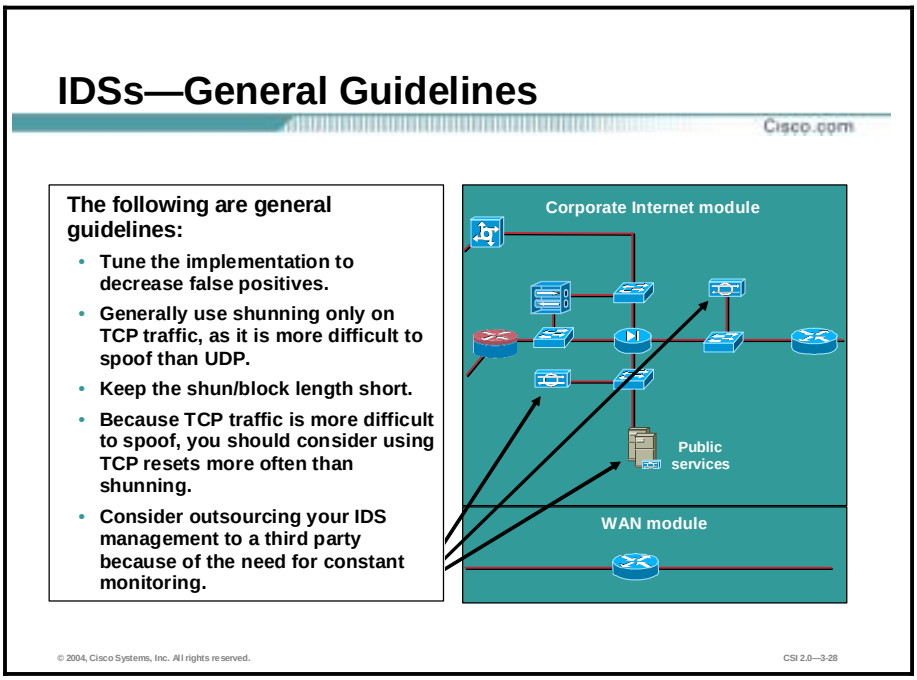

The following IDS guidelines can aid an administrator in preventing attacks:

- Tune the implementation to decrease false positives—False positives are alarms caused by legitimate traffic or activity. False negatives are attacks that the IDS fails to see. When you tune the IDS, you can configure it more specifically as to its threat-mitigation role.
- Generally use shunning only on TCP traffic, as it is more difficult to spoof than UDP—Shunning is the use of access control filters and should be carefully implemented.
- Keep the shun length short—Keeping the shun length short eliminates blocking traffic from a valid address that has been spoofed previously.
- Because TCP traffic is more difficult to spoof, you should consider using TCP resets more often than shunning—TCP resets operate only on TCP traffic and terminate an active attack by sending a TCP reset to both the attacker and the attacked host.
- Consider outsourcing your IDS management to a third party because of the need for constant monitoring—IT staff are often overworked (particularly in smaller organizations).

The figure identifies the location of the IDSs in the SAFE SMR network example.

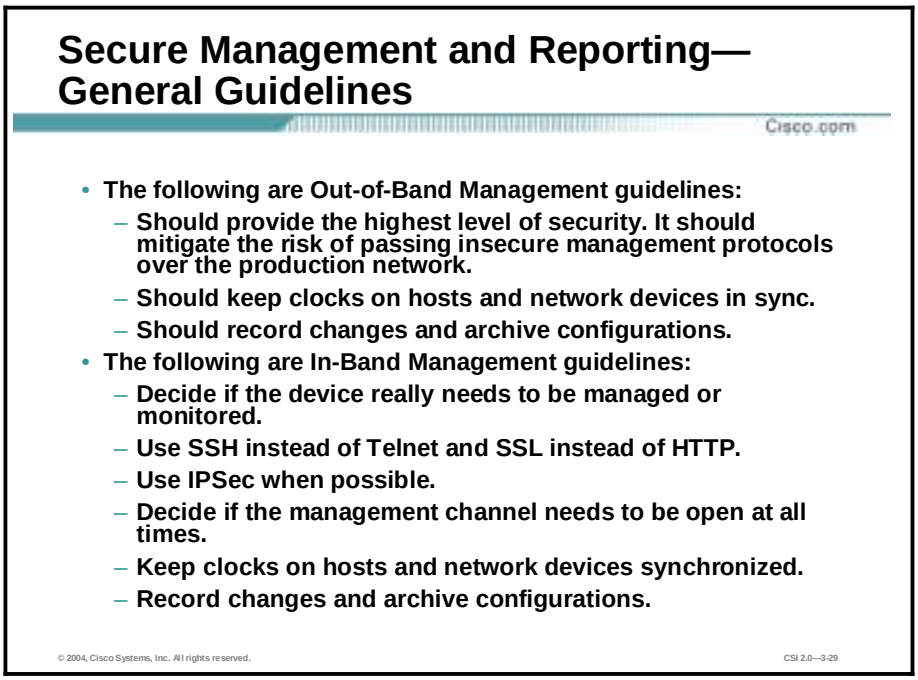

The following are out-of-band secure management guidelines for the architecture:

- It should provide the highest level of security, mitigating the risk of passing insecure management protocols over the production network.
- It should keep clocks on hosts and network devices synchronized.
- It should record changes and archive configurations.

The following are in-band secure management guidelines:

- Decide if the device really needs to be managed or monitored.
- Use Secure Shell (SSH) or Secure Sockets Layer (SSL) instead of Telnet.
- Use IPSec when possible, especially if the management protocol does not offer encryption.
- Decide if the management channel needs to be open at all times.
- Keep clocks on hosts and network devices synchronized.
- Record changes and archive configurations.

Even though Out-of-Band Management is recommended for devices in SAFE Enterprise, SAFE SMR recommends In-Band Management because the goal is cost-effective security deployment.

In the SAFE SMR architecture, management traffic flows in-band in all cases, and is made as secure as possible using tunneling protocols and secure variants to insecure management protocols. For example, SSH is used whenever possible instead of Telnet. With management traffic flowing in-band across the production network, it becomes very important to closely follow the axioms mentioned earlier in the lesson.

To ensure that log messages are time-synchronized to one another, clocks on hosts and network devices must be synchronized. For devices that support it, NTP provides a way to ensure that

accurate time is kept on all devices. When dealing with attacks, seconds matter because it is important to identify the order in which a specified attack occurred.

NTP is used to synchronize the clocks of various devices across a network. Synchronization of the clocks within a network is critical for digital certificates, and for correct interpretation of events within Syslog data. A secure method of providing clocking for the network is for the network administrator to implement their own master clock. The private network should then be synchronized to Universal Coordinated Time (UTC) via satellite or radio. However, clock sources are available which synchronize via the Internet if the network administrator does not wish to implement their own master clock because of costs or other reasons.

An attacker could attempt a DoS attack on a network by sending bogus NTP data across the Internet in an attempt to change the clocks on network devices in such a manner that digital certificates are considered invalid. Further, an attacker could attempt to confuse a network administrator during an attack by disrupting the clocks on network devices. This scenario would make it difficult for the network administrator to determine the order of Syslog events on multiple devices.

NTP version 3 and higher supports a cryptographic authentication mechanism between peers. The use of the authentication mechanism, as well as ACLs that specify which network devices are allowed to synchronize with other network devices, is recommended to help mitigate against such a scenario.

The network administrator should weigh the cost benefits of pulling the clock from the Internet with the possible risk of doing so and allowing it through the firewall. Many NTP servers on the Internet do not require any authentication of peers. Therefore, the network administrator must trust that the clock itself is reliable, valid, and secure. NTP uses UDP port 123.

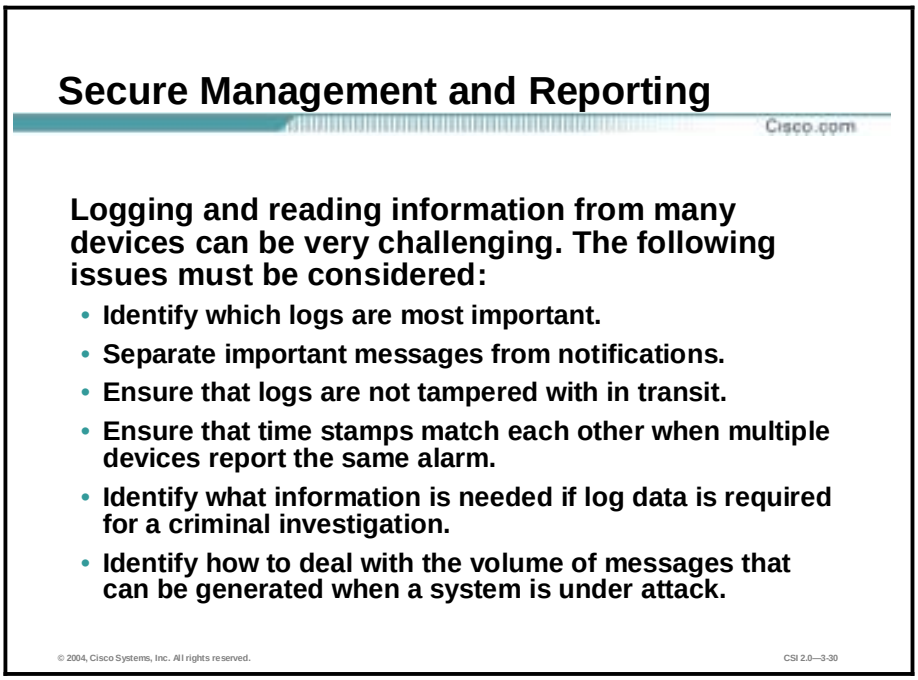

Logging and reading information from many devices can be very challenging. The following issues must be considered:

- Identify which logs are most important.
- Separate important messages from notifications.
- Ensure that logs are not tampered with in transit.
- Ensure that time stamps match each other when multiple devices report the same alarm.
- Identify what information is needed if log data is required for a criminal investigation.
- Identify how to deal with the volume of messages that can be generated when a system is under attack.

Each of these issues is company-specific and requires the input of management as well as the network and security teams to identify the priorities of reporting and monitoring. The implemented security policy should also play a large role in answering these questions.

From a reporting standpoint, most networking devices can send Syslog data that can be invaluable when troubleshooting network problems or security threats. You can send this data to your Syslog analysis host from any device whose logs you wish to view. This data can be viewed in real time or on-demand, and in scheduled reports. Depending on the device involved, you can choose various logging levels to ensure that the correct amount of data is sent to the logging device. You also need to flag device log data within the analysis software to permit granular viewing and reporting. For example, during an attack the log data provided by Layer 2 switches might not be as interesting as the data provided by the IDS.

To ensure that log messages are time-synchronized to one another, clocks on hosts and network devices must be synchronized. For devices that support it, NTP provides a way to ensure that accurate time is kept on all devices. When dealing with attacks, seconds matter because it is important to identify the order in which a specified attack occurred.

Configuration change management is another issue related to secure management. When a network is under attack, it is important to know the state of critical network devices and when the last known modifications occurred. Creating a plan for change management should be a part of your comprehensive security policy, but, at a minimum, you should record changes using authentication systems on the devices, and archive configurations via FTP or TFTP.

Summary

This topic summarizes the information you learned in this lesson.

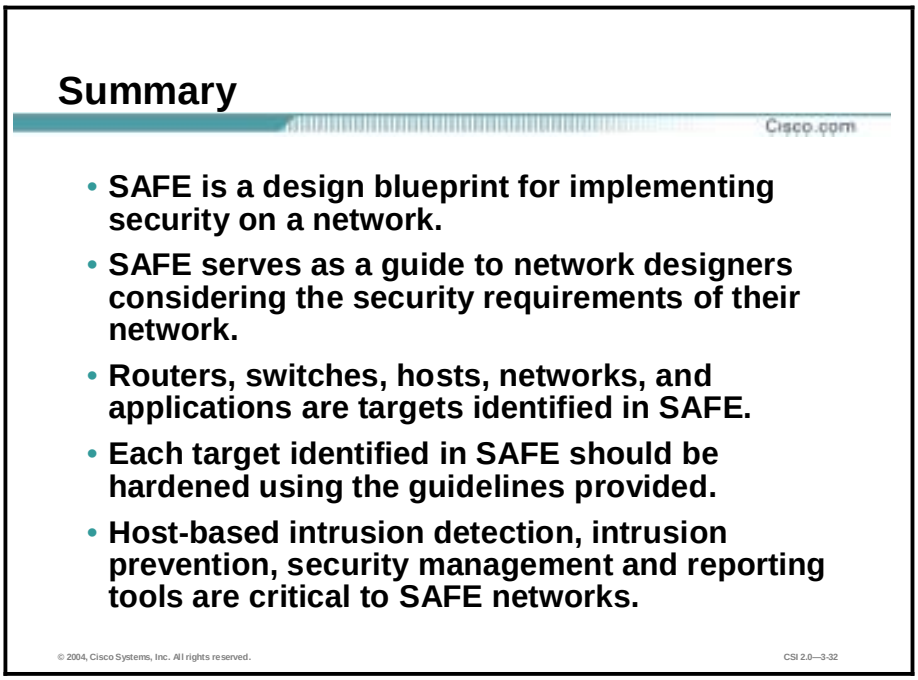

4

The Cisco Security Portfolio

Overview

This lesson introduces and gives an overview of a Cisco security portfolio. It includes the following topics:

- Objectives
- Cisco security portfolio overview
- Secure connectivity—VPN solutions
- Secure connectivity—The VPN 3000 Concentrator series
- Secure connectivity—Cisco VPN-optimized routers
- Perimeter security firewalls—Cisco PIX Firewall and Cisco IOS Firewall
- Intrusion protection—IDS
- Host-based intrusion prevention system—CSA
- Identity—Access control solutions
- Security management—Cisco IP Solution Center and VMS
- Cisco AVVID
- Summary

Objectives

This topic lists the lesson's objectives.

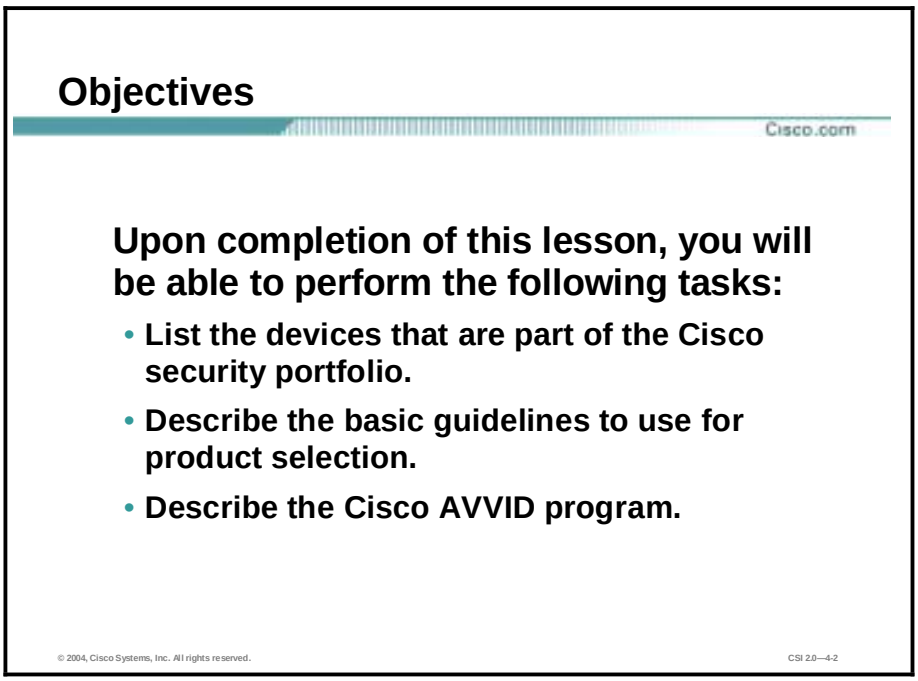

Cisco Security Portfolio Overview

To successfully use network technologies, you increasingly need to protect valuable data and network resources from corruption and intrusion. The Cisco security solutions provide the services necessary to achieve this. This topic covers the security solutions that Cisco offers.

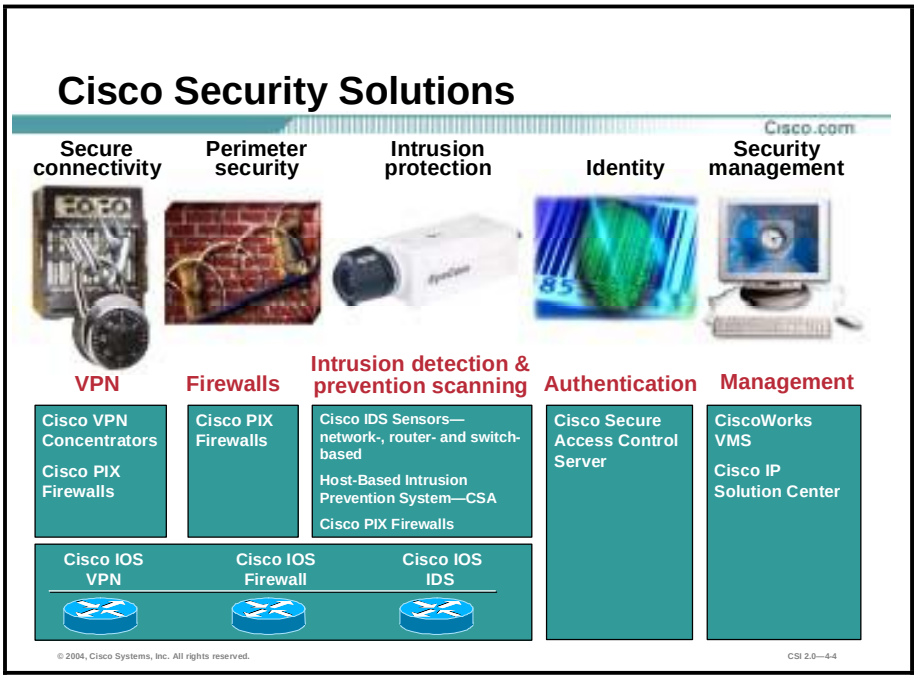

Cisco offers a wide variety of security solutions:

- Secure connectivity—Virtual private network (VPN)
  - Cisco VPN Concentrators
  - Cisco PIX Firewalls
  - Cisco IOS VPN
- Perimeter security—Firewalls
  - Cisco PIX Firewalls
  - Cisco IOS Firewalls
- Intrusion protection—Intrusion detection and scanning
  - Cisco network-based intrusion detection system (NIDS) Sensor
  - Cisco IOS-based intrusion detection
  - Cisco Intrusion Detection System Module (IDSM2)
  - Network Module-Cisco Intrusion Detection System (NM-CIDS) for access routers
  - PIX Firewall-based intrusion detection
  - Host-based intrusion prevention system—Cisco Security Agent (CSA)
- Identity—Authentication (Cisco Secure Access Control Server [ACS])
- Security management—Policy (CiscoWorks VPN/Security Management Solution [VMS])

Cisco offers a wide variety of value-added benefits:

- Breadth of solutions—Cisco offers the widest range of security and VPN products available in the market today. These products span multiple technology categories—including firewalls, intrusion detection systems, Concentrators, and routers—and are scaled to meet your business needs, from enterprise gateways to remote-office connections that permit you to deploy customized network security solutions from a single partner.
- Industry leadership and recognition—The PIX Firewall family has gained worldwide market leadership, according to the IDC analyst group. The Cisco IDS is the market leader, according to Frost & Sullivan. Cisco access control lists (ACLs) represent the most widely deployed security technology in the world. In addition, the Cisco VPN 3060 Concentrator was named "Hardware Product of the Year" by *Network Computing* magazine.
- Non-stop technical support—Cisco security and VPN products receive the same legendary technical support as other Cisco networking gear, permitting users to gain assistance day or night. Few other security or VPN companies are large enough to provide such critical, full-time assistance. Cisco support and service also includes the tools, expertise, and resources needed to quickly install, maintain, and enhance Cisco security and VPN products to protect the business network as effectively as possible.
- Unsurpassed interoperability—Instead of deploying a patchwork of different security technologies from different vendors promising interoperability, Cisco provides you with the confidence that all its Cisco security and VPN products have been thoroughly tested for compatibility. In addition, the Security Associate program provides a formal, third-party testing ground for other vendors to prove their interoperability with Cisco products, as opposed to making you rely on ambiguous marketing claims.
- Unsurpassed integration—Cisco owns and controls the industry and has, therefore, built its security into the router infrastructure, which is the very foundation of the Cisco network. No other vendor has such a unique perspective and ability to provide you with security at the core of the network.
- Management prowess—Cisco provides a reliable, available, and proven foundation for managing your networks efficiently and cost effectively.

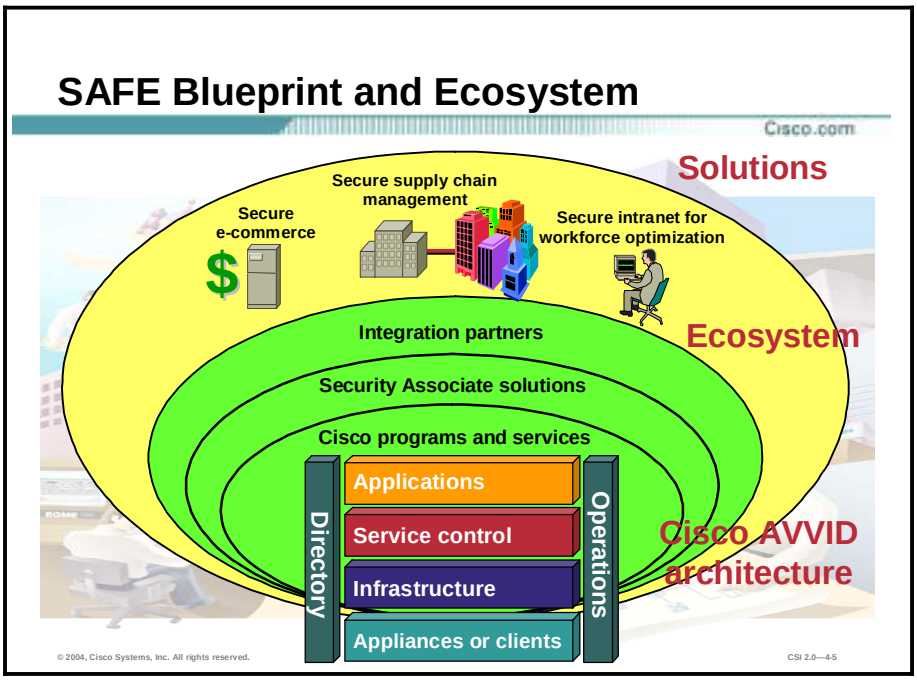

Cisco has opened its Cisco Architecture for Voice, Video, and Integrated Data (AVVID) and SAFE Blueprint to key third-party vendors to create a security solutions ecosystem to spur development of best-in-class multiservice applications and products. The Cisco AVVID and SAFE Blueprint provide interoperability for third-party hardware and software using standards-based media interfaces, application-programming interfaces (APIs), and protocols. This ecosystem is offered through the Security and VPN Associate Program, an interoperability solutions program that provides Cisco customers with tested and certified, complementary products for securing their businesses. The ecosystem enables businesses to design and roll out secure networks that best fit their business model and enable maximum agility.

Secure Connectivity—VPN Solutions

Cisco has developed and acquired products and solutions that are optimized for secure connectivity. This topic describes these products and solutions, and the security they provide.

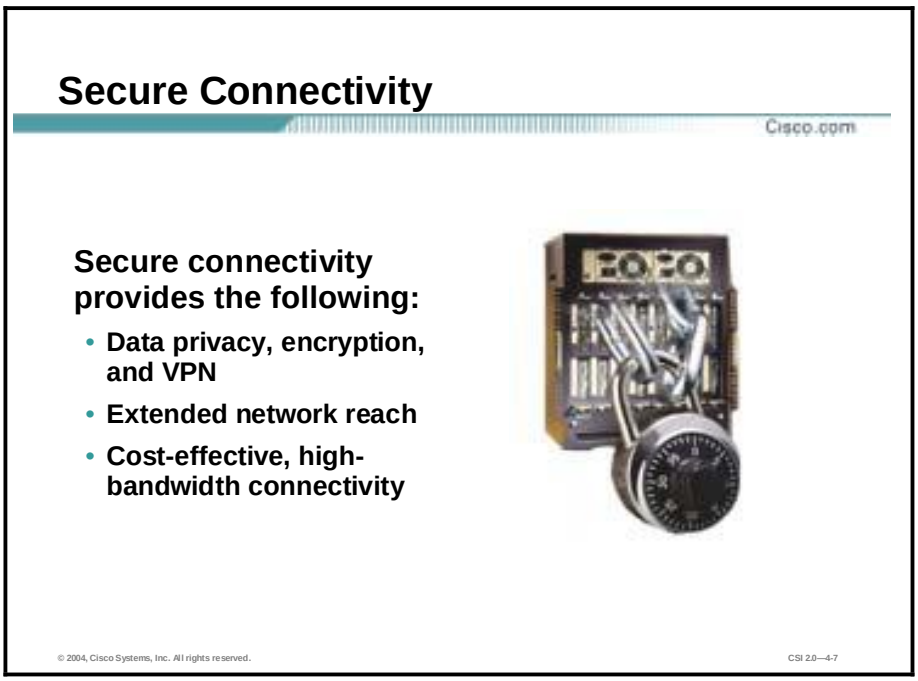

Secure connectivity provides the following:

- Data privacy, encryption, and VPN
  - Provides security over untrusted public networks
  - Provides enhanced transport security for private networks
- Extended network reach
  - Teleworkers
  - New or small sites
  - Partner connectivity
- Cost-effective, high-bandwidth connectivity
  - Reduces transport costs
  - Enables fast broadband telecommuters and remote site connectivity

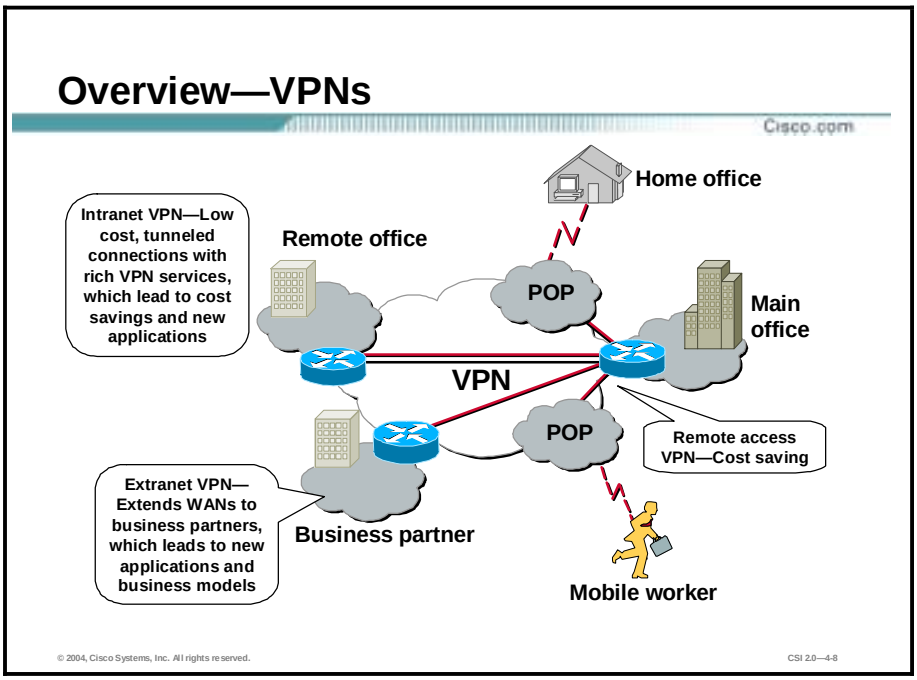

VPN solutions are provided for the following types of implementation:

- Intranet VPN—Links corporate headquarters to remote offices over a shared, prioritized network, and offers an extremely cost-effective alternative to dedicated WANs. Intranet VPNs need to scale easily as the organization grows.
- Extranet VPN—Links network resources with third-party vendors and business partners, extending elements of the corporate intranet beyond the organization. To keep pace with rapidly changing business climates, extranet VPN access needs to be able to be turned on and off on the fly.
- Remote access VPN—Connects telecommuters and mobile users securely and cost-effectively to corporate network resources from anywhere in the world over any access technology. Because this traffic may run on untrusted segments outside the service provider's network, it must be encrypted to ensure privacy and security.

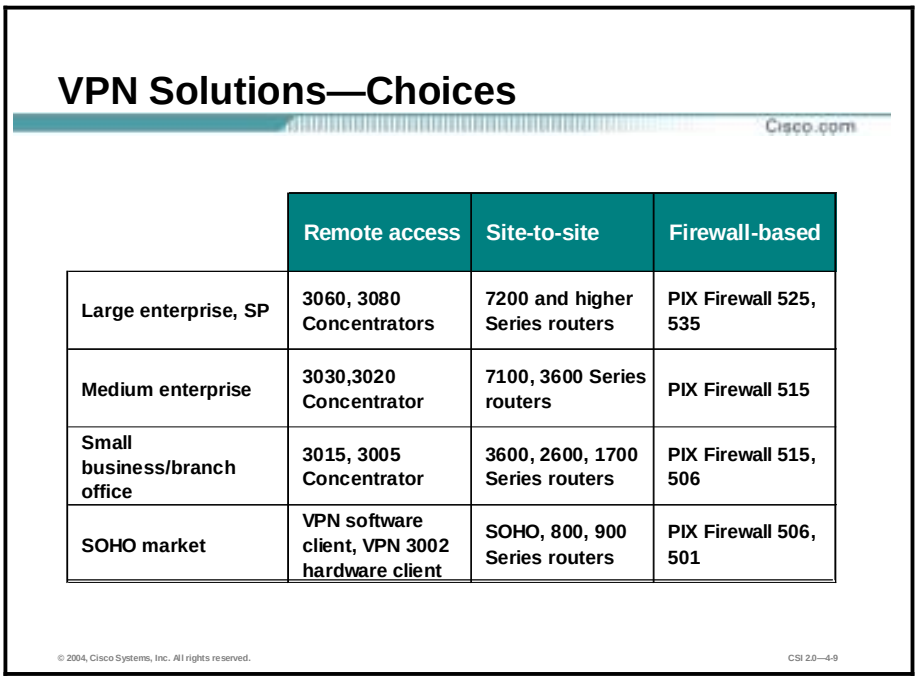

Cisco provides VPN solutions for all network sizes. The information in the figure indicates the platforms that can support each size network most effectively. You can use this information as a starting point to choose which device best fits your environment.

Secure Connectivity—The VPN 3000 Concentrator Series

The Cisco VPN 3000 Series Concentrator is a family of purpose-built, remote-access VPN platforms and client software that incorporates high availability, high performance, and scalability with the most advanced encryption and authentication techniques available today. This topic describes the Cisco VPN 3000 Series Concentrator.

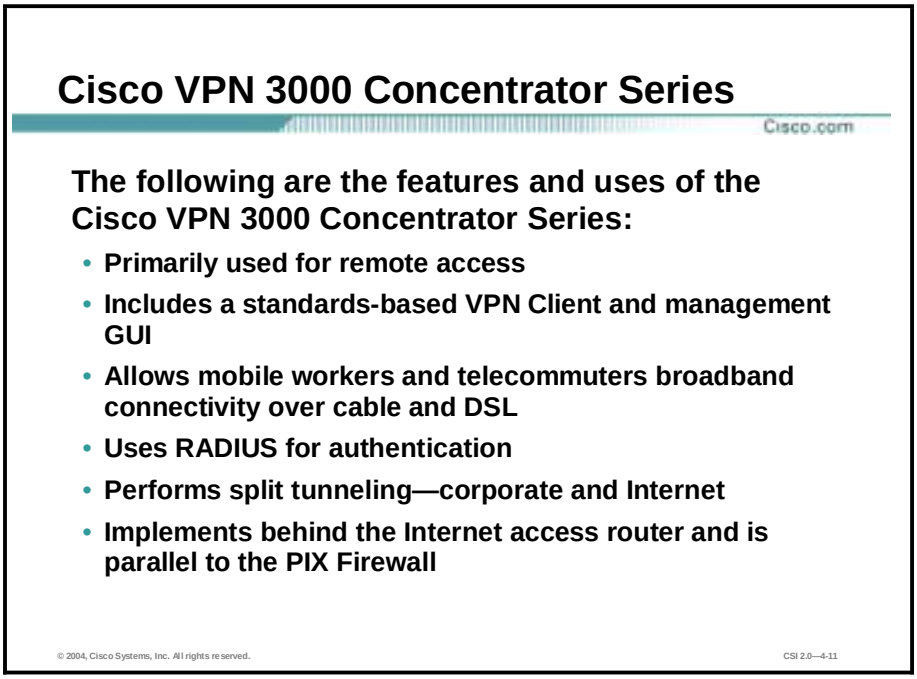

The following are the features and uses for the Cisco VPN 3000 Concentrator Series:

- Primarily used for remote access
- Includes a standards-based VPN Client and management GUI
- Allows mobile workers and telecommuters broadband connectivity over cable and DSL
- Uses Remote Access Dial-In User Service (RADIUS) for authentication
- Split tunneling—corporate and Internet
- Implements behind the Internet access router and is parallel to the PIX Firewall

With the Cisco VPN 3000 Series Concentrator, customers can take advantage of the latest VPN technology to vastly reduce their communications expenditures. It is the only scalable platform to offer field-swappable and customer-upgradeable components. These components, called Scalable Encryption Processing (SEP) modules, enable users to easily add capacity and throughput.

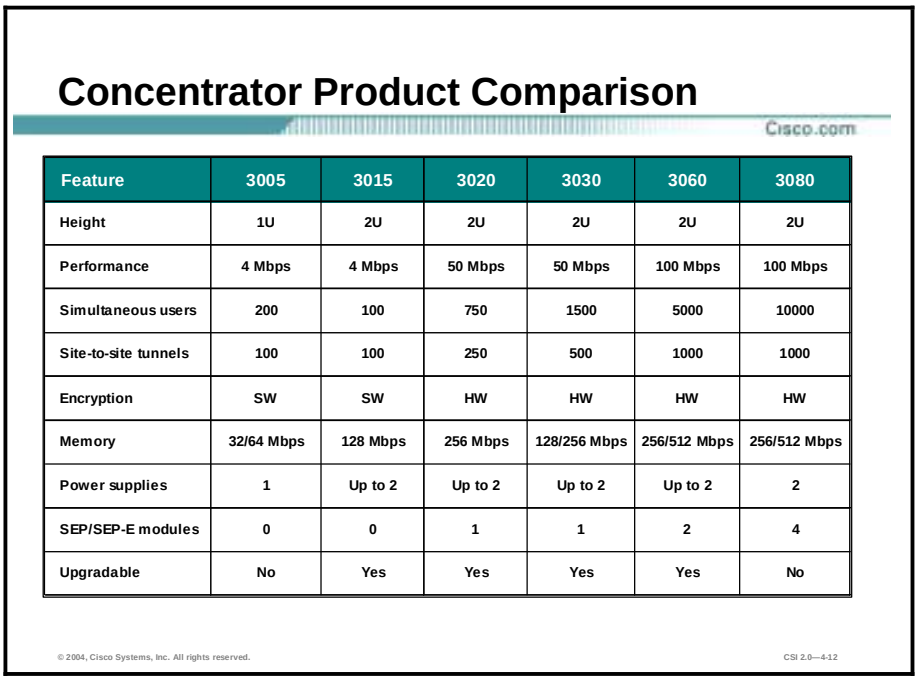

The Cisco VPN 3000 Series Concentrator includes models to support a range of enterprise customers, from small businesses with 100 or fewer remote-access users, to large organizations with up to 10,000 simultaneous remote users. The Cisco VPN Client is provided with all versions of the Cisco VPN 3000 Series Concentrator, and includes unlimited distribution licensing. The Cisco VPN 3000 Series Concentrator is available in both nonredundant and redundant configurations. The table in the figure can assist engineers in choosing the most scaleable, cost-effective, and redundant solution for their networks.

The Cisco Secure VPN Client Framework

Cisco.com

- Connectivity between all clients and all Cisco central-site VPN gear
- Centralized push policy technology
  - Simplifies user experience
  - Provides more control for companies
  - Reduces complexity of VPN deployments
- Implemented across all Cisco VPN Concentrators, Cisco IOS routers, and PIX Firewalls
  - Includes non-Windows operating systems (Linux, Mac, and Solaris)
  - Substantial savings
  - Reduced support expense
  - Consolidated hardware
  - Reduced administration in the central site at the central site

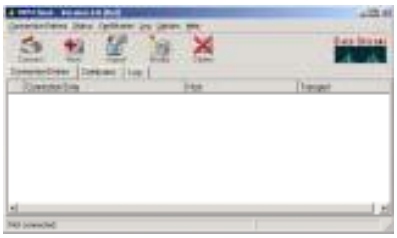

CSI 2.0—4-13

The Cisco VPN Client 1.x strategy is a new framework with a specification to enable VPN connectivity between all desktop, laptop, and personal digital assistant clients and the full range of Cisco VPN-enabled Concentrators, routers, and firewalls. Using “push policy” capabilities, the VPN Client framework allows customers to centrally manage security policies, while easily delivering large-scale VPN connectivity to remote users. All of the Cisco IPSec-based VPN products for the enterprise and service providers support the VPN Client framework.

The following are the features and uses of the Cisco Secure VPN Client framework:

- Connectivity between all clients and all Cisco central-site VPN gear
- Centralized push policy technology
  - Simplifies user experience
  - Provides more control for companies
  - Reduces complexity of VPN deployments
- Implemented across all Cisco VPN Concentrators, IOS Routers, and PIX Firewalls
  - Includes non-Windows operating systems (Linux, Mac, and Solaris)
  - Substantial savings
  - Reduced support expense
  - Consolidated hardware
  - Reduced administration in the central site at the central site

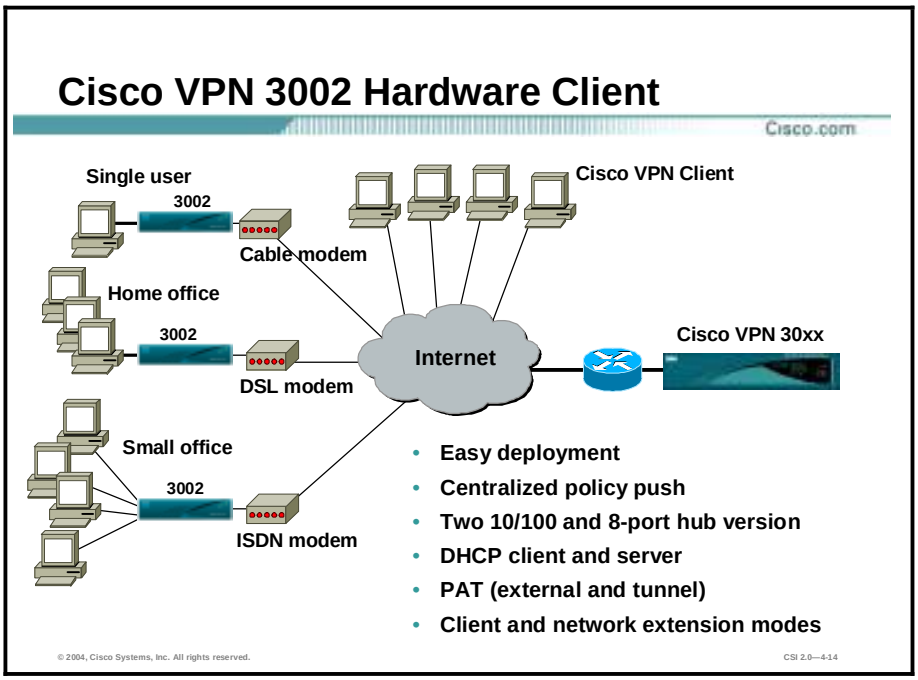

Based on the unified VPN Client framework, the Cisco VPN 3002 Hardware Client combines the best features of a software client, including scalability and ease-of-deployment, with the stability and independence of a hardware platform. The Cisco VPN 3002 Hardware Client works with all operating systems and does not interfere with the operation of the PC because it is a separate hardware appliance.

The Cisco VPN 3002 Hardware Client is a small, highly cost-effective appliance. It is ideal for organizations where thousands of remote end-users might be tunneling into corporate networks from large numbers of geographically dispersed branch or home office sites.

Other features and uses of the Cisco 3002 Hardware VPN Client are as follows:

- Easy to deploy and scalable
- Centralized push for easy policy deployment
- Two 10/100 and 8-port hub version
- Supports DHCP client and server functionality
- Allows Port Address Translation (PAT) (external and tunnel)
- Supports client and network extension modes

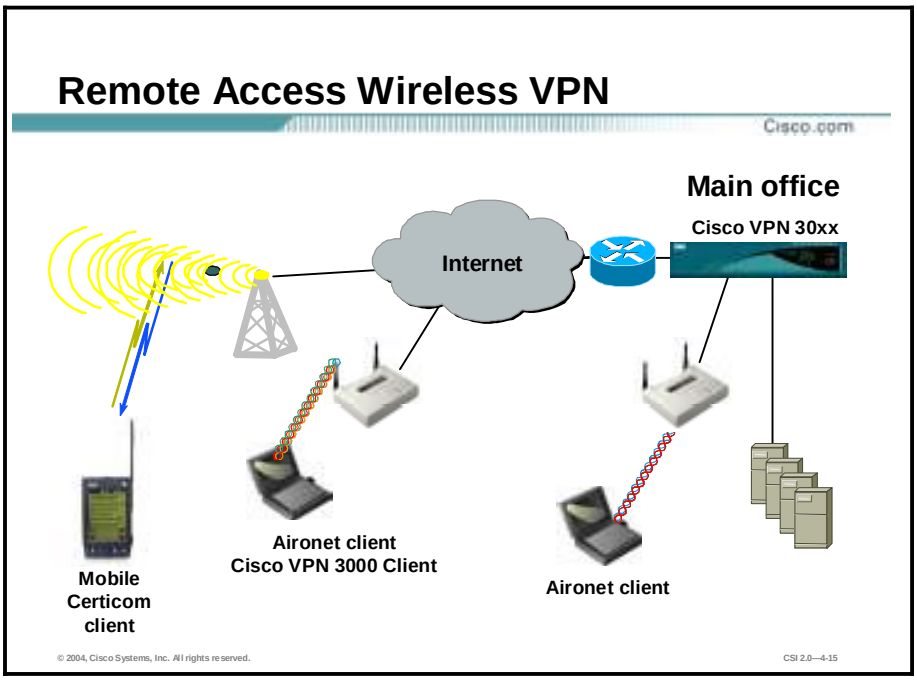

Remote access wireless VPN solutions are available for the VPN Concentrator via the Cisco AVVID partner program. With release 3.0, all Cisco VPN 3000 Concentrators support Elliptic Curve Cryptography (ECC). This is a new Diffie-Hellman (DH) group that allows for much faster processing of keying information by devices with limited processing power such as Personal Digital Assistants (PDAs) and Smart Phones. Cisco VPN 3000 Concentrators can now securely terminate tunnels from IP-enabled wireless devices, allowing a whole new class of users to securely access enterprise information while preserving the investment in VPN termination equipment in the enterprise data center.

Secure Connectivity—Cisco VPN-Optimized Routers

This topic describes the solutions provided by Cisco routers for secure connectivity.

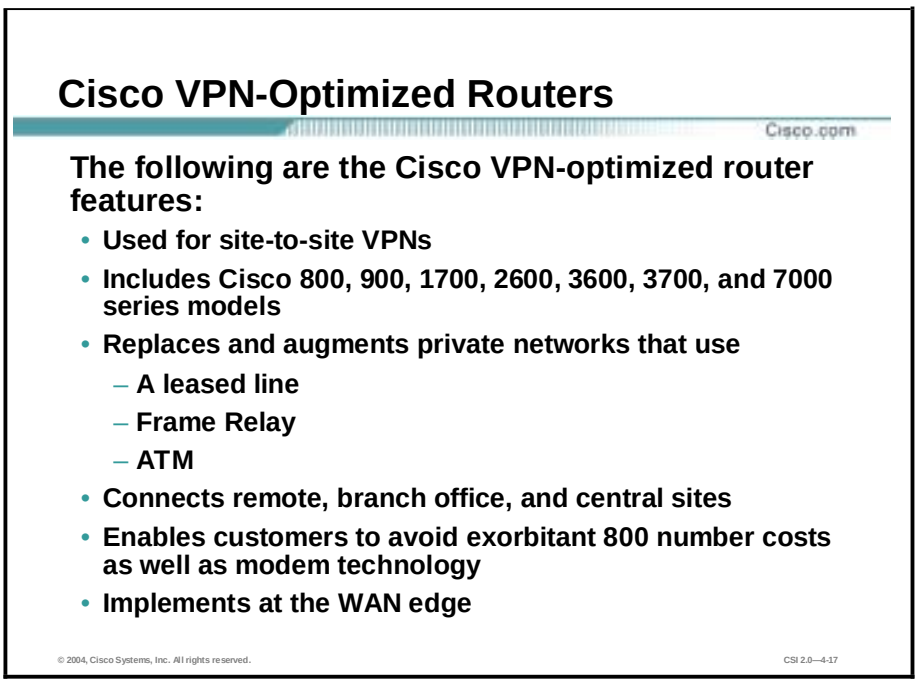

Site-to-site VPNs are an alternative WAN infrastructure that are used to connect branch offices, home offices, or business partners' sites to all or portions of a company's network. VPNs do not inherently change private WAN requirements, such as support for multiple protocols, high reliability, and extensive scalability, but instead meet these requirements more cost-effectively and with greater flexibility. Site-to-site VPNs use the most pervasive transport technologies available today, such as the Internet or service providers' IP networks, by employing tunneling and encryption for data privacy and Quality of Service (QoS) for transport reliability.

The following are the Cisco VPN-optimized router features:

- Used for site-to-site VPNs
- Includes Cisco 800, 900, 1700, 2600, 3600, 3700, and 7000 series models
- Replaces and augments private networks that use
  - A leased line
  - Frame Relay
  - ATM
- Connects remote, branch office, and central sites
- Enables customers to avoid exorbitant 800 number costs as well as modem technology
- Implements at the WAN edge

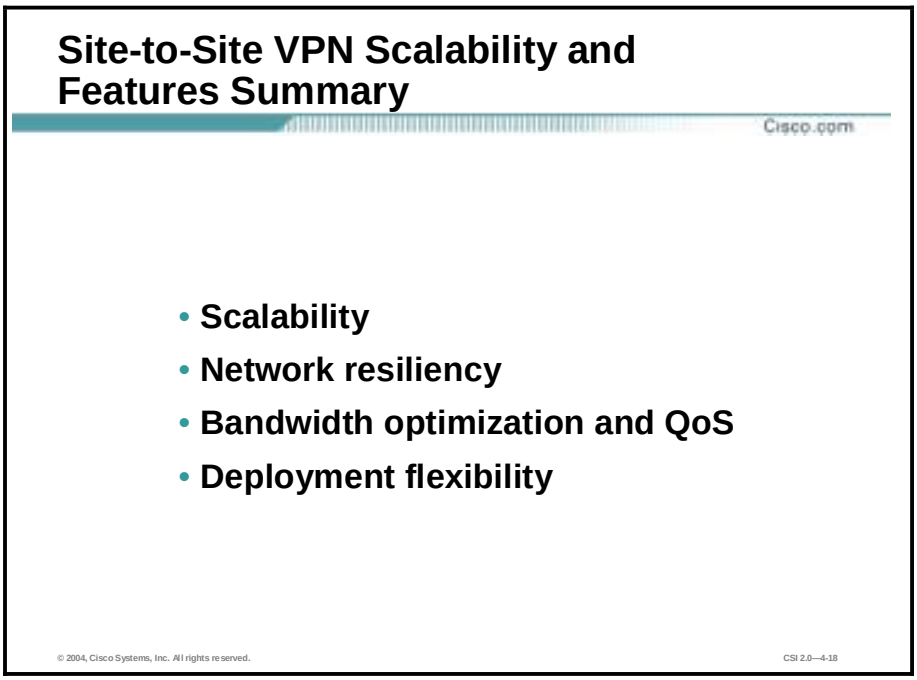

Cisco VPN-optimized routers include high-performance, hardware-based IPSec encryption, multiple WAN interfaces, and the entire Cisco IOS Software feature set. Using Cisco IOS Software, Cisco VPN Routers also provide a comprehensive feature set to meet the most diverse networking requirements, including support for routing, multiprotocol, and multicast across the VPN, as well as enhanced features like firewall capabilities and QoS.

The following is the site-to-site VPN scalability and features summary for Cisco VPN optimized routers:

- Scalability—Up to 140 Mbps of Triple-Data Encryption Standard (3DES) throughput and 3000 tunnels
- Network resiliency
  - Dynamic Router Recovery—Using routing protocols through IPSec-secured Generic Routing Encapsulation (GRE) tunnels
  - Dynamic Tunnel Recovery—Using IPSec (Internet Key Exchange [IKE]) keepalives
- Bandwidth optimization and QoS
  - Application-aware bandwidth allocation, queuing, policing, and traffic shaping
  - Ensured quality of latency-sensitive traffic
- Deployment flexibility
  - Interface flexibility for combined WAN and VPN or behind-edge VPN
  - Use as standalone VPN device or integrated multi-function device

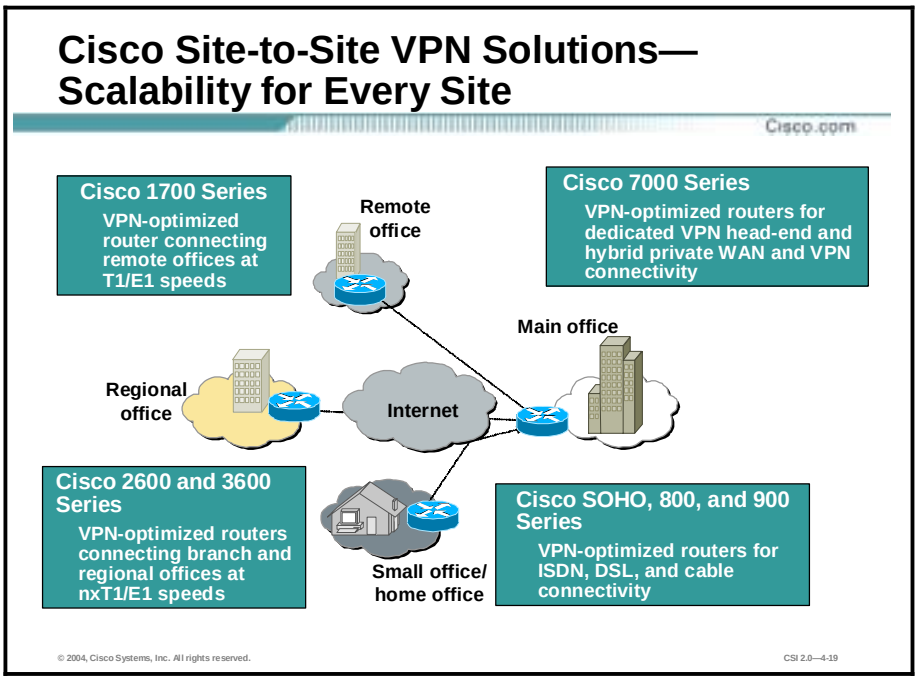

Site-to-site VPNs are best constructed using a wide variety of Cisco VPN routers. VPN routers provide scalability through optional encryption acceleration. The Cisco VPN router portfolio provides solutions for small office/home office (SOHO) access through central-site VPN aggregation, including platforms for fast-emerging cable and DSL access technologies.

The following are recommendations for scalability for site-to-site VPN solutions:

- Remote office—Cisco 1700 Series router, which is a VPN-optimized router connecting remote offices at T1/E1 speeds
- Regional office—Cisco 2600 and 3600 Series routers, which are VPN-optimized routers connecting branch and regional offices at nxT1/E1 speeds
- Small Office/Home Office (SOHO)—Cisco 800 and 900 Series routers, which are VPN-optimized routers for ISDN, DSL, and cable connectivity
- Main Office—Cisco 7000 Series routers, dedicated VPN head-end and hybrid private WAN and VPN connectivity

VAM2—For Cisco 7100, 7200, and 7400 Series Routers

Cisco.com

Hardware acceleration for

- **IPSec encryption—Up to 145 Mbps of VPN performance and 5000 tunnels**
- **RSA—Faster tunnel-recovery key generation and authentication**
- **IPPCP LZS compression**

CSI 2.0—4-20

The VPN Acceleration Module 2 (VAM2) is a single-width acceleration module. It provides high-performance, hardware-assisted tunneling and encryption services suitable for VPN remote-access, site-to-site intranet, and extranet applications. It also provides platform scalability and security while working with all services necessary for successful VPN deployments—security, QoS, firewall and intrusion detection, service-level validation, and management. The VAM off-loads IPSec processing from the main processor, thus freeing resources on the processor engines for other tasks.

Perimeter Security Firewalls—Cisco PIX Firewall and Cisco IOS Firewall

This topic describes the Cisco perimeter security solutions.

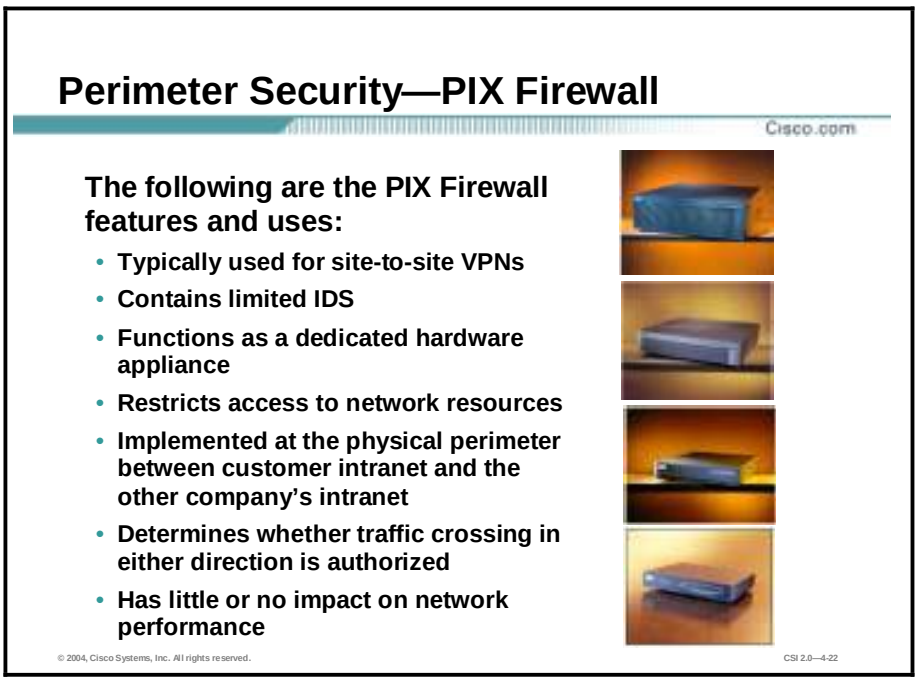

Globally networked businesses rely on their networks to communicate with employees, customers, partners, and suppliers. While immediate access to information and communication is an advantage, it raises concerns about security—protecting access to critical network resources. Network administrators need to know who is accessing what resources and establish clear perimeters to control that access. An effective security policy balances accessibility with protection. Security policies are enforced at network perimeters. Often people think of a perimeter as the boundary between an internal network and the Internet, but a perimeter can be established anywhere within a private network, or between your network and a partner's network. A solid perimeter security solution enables communications across it as defined by the security policy, yet protects network resources from breaches or attacks. It controls multiple network entry and exit points. It also increases user assurance by implementing multiple layers of security.

The following are the PIX Firewall features and uses:

- Typically used for site-to-site VPNs
- Contains limited IDS
- Functions as a dedicated hardware appliance
- Restricts access to network resources
- Implemented at the physical perimeter between customer intranet and the other company's intranet
- Determines whether traffic crossing in either direction is authorized

- Has little or no impact on network performance

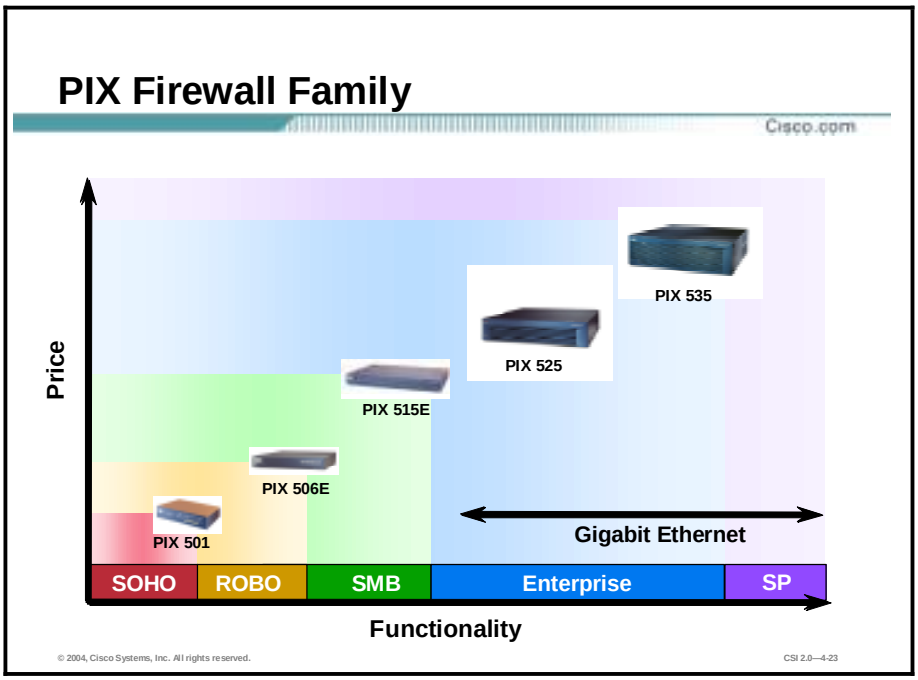

The Cisco PIX Firewall 500 series scales to meet a range of requirements and network sizes, and currently consists of five models: the PIX Firewall 501, 506E, 515E, 525, and 535. The PIX Firewall 501 has an integrated 10/100BASE-T port (100BASE-T option available in release 6.3) and an integrated four-port 10/100 switch. The PIX Firewall 506E has dual integrated 10/100BASE-T ports (100BASE-T option available in release 6.3 for 506E only). The PIX Firewall 515E supports single-port or four-port 10/100 Ethernet cards. The PIX Firewall 525 supports single-port or four-port 10/100 Fast Ethernet and Gigabit Ethernet. The PIX Firewall 535 supports Fast Ethernet and Gigabit Ethernet. The PIX Firewall 515E, 525, and 535 models come with an integrated VPN Accelerator Card (VAC).

The PIX Firewall is secure right out of the box. The PIX Firewall default settings allow all connections from the inside interface access to the outside interface, and block all connections from the outside interface to the inside interface. After a few installation procedures and an initial configuration with six general commands, your PIX Firewall is operational and protecting your network.

Note	Prior to PIX Firewall Software Release 6.3, the PIX Firewall 501 outside interface and 506E outside and inside interfaces operated at 10BASE-T. With the upgrade to software release 6.3, the PIX Firewall 501 outside interface and 506E outside and inside interfaces can operate at 10/100BASE-T. To enable the speed change on the interface requires a software upgrade only.

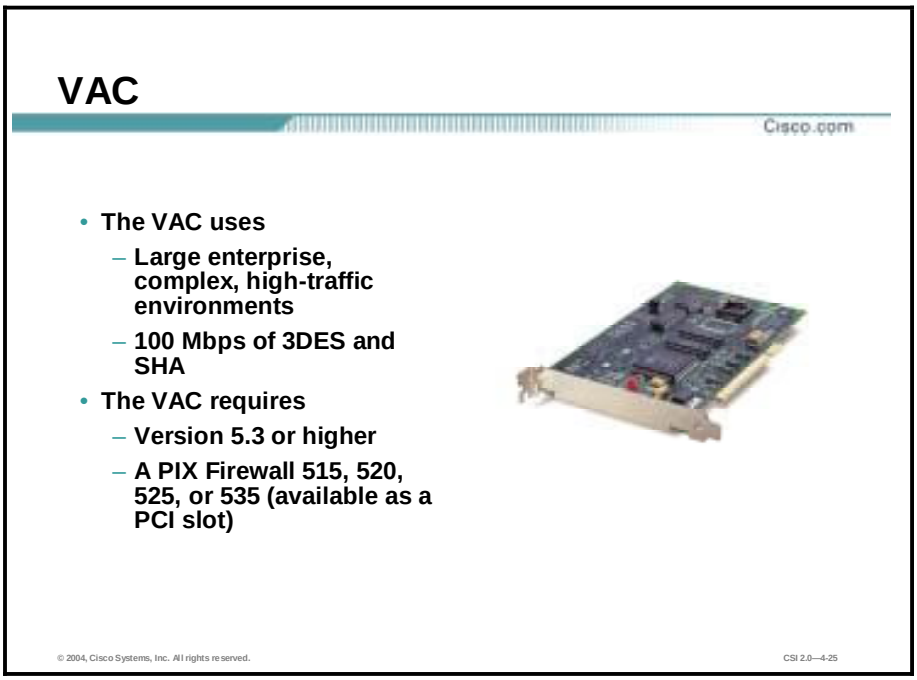

The VPN Accelerator Card (VAC) for the PIX Firewall series provides high-performance tunneling and encryption services suitable for site-to-site and remote-access applications. This hardware-based VPN accelerator is optimized to handle the repetitive but voluminous mathematical functions required for IPSec. Offloading encryption functions to the card not only improves IPSec encryption processing, but also maintains high-end firewall performance. As an integral component of the Cisco VPN solution, the VAC provides platform scalability and security while working seamlessly with services necessary for successful VPN deployments—encryption, tunneling, and firewall.

The VAC, which fits in a PCI slot inside the PIX Firewall chassis, encrypts data using the 56-bit Data Encryption Standard (DES) or 168-bit 3DES algorithms at speeds up to 100 Mbps. A PIX Firewall equipped with a VAC supports as many as 2000 encrypted tunnels for concurrent sessions with mobile users or other sites. In addition to encryption, the card handles a variety of other IPSec-related tasks—hashing, key exchange, and storage of security associations—that free the PIX Firewall main processor and memory to perform other perimeter security functions. The following are features of the VAC:

- Encryption—DES and 3DES encryption are very CPU-intensive, potentially impacting firewall performance in high-throughput configurations. The VAC makes it possible to send DES- or 3DES-encrypted data at high speeds while still providing the full range of perimeter security services available from the PIX Firewall.
- Authentication—Rivest, Shamir, and Adleman (RSA) and Diffie-Hellman (DH) are CPU-intensive protocols that are used when a new IPSec tunnel is established. RSA authenticates the remote device while DH exchanges keys that will be used for DES or 3DES encryption. The VAC implements these protocols in specialized hardware ensuring fast tunnel setup and high overall encryption throughput.
- Tunneling—The PIX Firewall and VAC support IPSec tunneling protocols, which enable high-performance and flexible network designs for both remote-access and site-to-site VPNs. Site-to-site solutions can be designed with the PIX Firewall, or combinations of PIX Firewalls with Cisco VPN appliances or VPN-enabled multiservice routers. Remote-access

solutions can use the Cisco VPN Client or other third-party clients supporting the IPSec tunneling protocol.

Firewall Services Module

Cisco.com

- **Designed for high-end enterprise and service providers**
- **Runs in Catalyst 6500 Series switches and 7600 Series routers**
- **Based on PIX Firewall technology**
- **Includes PIX Firewall 6.0 feature set (some 6.2)**
- **Supports multiple performance and redundancy features**

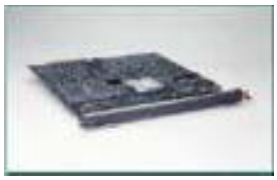

CSI 2.0—4-24

The Firewall Services Module (FWSM) is a multigigabit integrated firewall module for the Cisco Catalyst 6500 Series switch and the Cisco 7600 Series Internet router. It is fabric-enabled and capable of interacting with the bus and the switch fabric. Based on PIX Firewall technology, FWSM provides stateful firewall functionality in these switches and routers.

The following are the key features of FWSM:

- Includes entire PIX Firewall 6.0 software feature set and the following PIX Firewall 6.2 software features:
  - Command authorization
  - Object grouping
  - Internet Locator Service (ILS)/NetMeeting setup
  - URL filtering enhancement
- Support for 100 VLANs
- High-performance–5 Gbps / three million pps throughput, full-duplex firewall functionality
- One million concurrent connections
- LAN failover—Active or standby, and interchassis or intrachassis
- Dynamic routing with Open Shortest Path First (OSPF) and passive RIP
- Supports multiple modules per chassis

VAC

Cisco.com

- The VAC uses
  - Large enterprise, complex, high-traffic environments
  - 100 Mbps of 3DES and SHA
- The VAC requires
  - Version 5.3 or higher
  - A PIX Firewall 515, 520, 525, or 535 (available as a PCI slot)

Perimeter Security—Cisco IOS Firewall

Cisco.com

The following are the Cisco IOS Firewall features and uses:

- **Integrated software solution**
- **Limited IDS**
- **Add-on module to Cisco IOS software**
- **Cost effective**
- **Highly scalable**
- **Home office to enterprise**
- **Intranet protection**
- **Familiar Cisco IOS configuration**
- **CBAC**
- **Authentication Proxy**

As network security becomes increasingly critical to securing business transactions, businesses must integrate security into the network design and infrastructure. Security policy enforcement is most effective when it is an inherent component of the network.

Cisco IOS Software runs on more than 80 percent of Internet backbone routers, making it the most fundamental component of today's network infrastructure. Cisco IOS software-based security offers the best solution for end-to-end Internet, intranet, and remote-access network security.

The Cisco IOS Firewall is a security-specific option for Cisco IOS software. It integrates robust firewall functionality and intrusion detection for every network perimeter and enriches existing Cisco IOS security capabilities. It adds greater depth and flexibility to existing Cisco IOS security solutions—such as authentication, encryption, and failover—by delivering state-of-the-art security features, such as stateful, application-based filtering; dynamic per-user authentication and authorization; defense against network attacks; Java blocking; and real-time alerts. When combined with Cisco IOS IPSec software and other Cisco IOS software-based technologies, such as Layer 2 Tunneling Protocol (L2TP) and QoS, the Cisco IOS Firewall provides a complete, integrated VPN solution.

The following are the Cisco IOS Firewall features and uses:

- Integrated software solution
- Limited IDS support
- Add-on module to Cisco IOS software
- Cost effective
- Highly scalable solution
- Home office to enterprise solution
- Provides Intranet protection features

- Familiar Cisco IOS configuration
- Supports CBAC
- Supports authentication Proxy

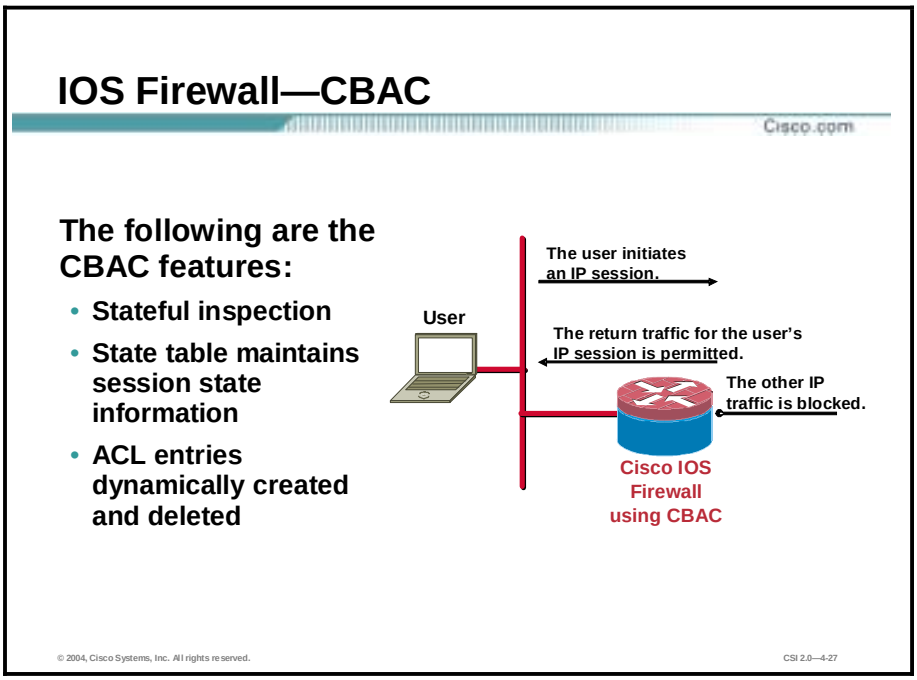

Context-Based Access Control (CBAC) intelligently filters TCP and UDP packets based on application-layer protocol session information, and can be used for intranets, extranets, and the Internet. You can configure CBAC to permit specified TCP and UDP traffic through a firewall only when the connection is initiated from within the network you want to protect (that is, CBAC can inspect traffic for sessions that originate from the external network). However, while this example discusses inspecting traffic for sessions that originate from the external network, CBAC can inspect traffic for sessions that originate from either side of the firewall.

Without CBAC, traffic filtering is limited to ACL implementations that examine packets at the network layer, or at most, the transport layer. However, CBAC examines not only network-layer and transport-layer information, but also examines the application-layer protocol information (such as FTP connection information) to learn about the state of the TCP or UDP session. This allows support of protocols that involve multiple channels created as a result of negotiations in the control channel. Most of the multimedia protocols as well as some other protocols (such as FTP, Remote Procedure Call [RPC], and SQL*Net) involve multiple channels.

CBAC inspects traffic that travels through the firewall to discover and manage state information for TCP and UDP sessions. This state information is used to create temporary openings in the firewall's ACLs to allow return traffic and additional data connections for permissible sessions (sessions that originated from within the protected internal network).

CBAC also provides the following benefits:

- Java blocking
- Denial of service (DoS) prevention and detection
- Real-time alerts and audit trails

Intrusion Protection—IDS

This topic provides an overview and product information for the Cisco intrusion detection system (IDS).

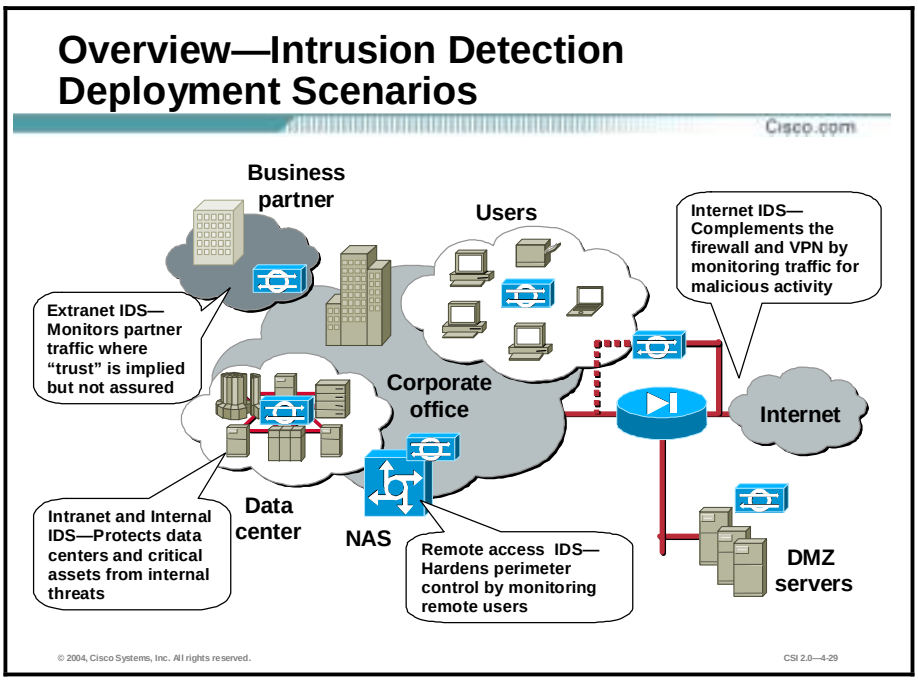

The Cisco IDS is an enterprise-class, network-based intrusion detection system. Designed to address the increased requirements for security visibility, DoS protection, hacking detection, and e-commerce business defenses, the Cisco IDS family leads the market in innovative security monitoring solutions. Sensor devices detect unauthorized activity traversing the network, such as attacks by hackers, by analyzing traffic in real time, enabling users to quickly respond to security breaches. When unauthorized activity is detected, Cisco IDS Sensors can send alarms to a management console with details of the activity, and can control other systems, such as routers, to terminate the unauthorized sessions.

There are four recommended deployment scenarios:

- Extranet IDS—IDS deployment to an extended network
- Internet IDS—IDS deployment to a public network
- Intranet and internal IDS—IDS deployment to an internal network
- Remote access IDS—IDS deployment to a remote-access network

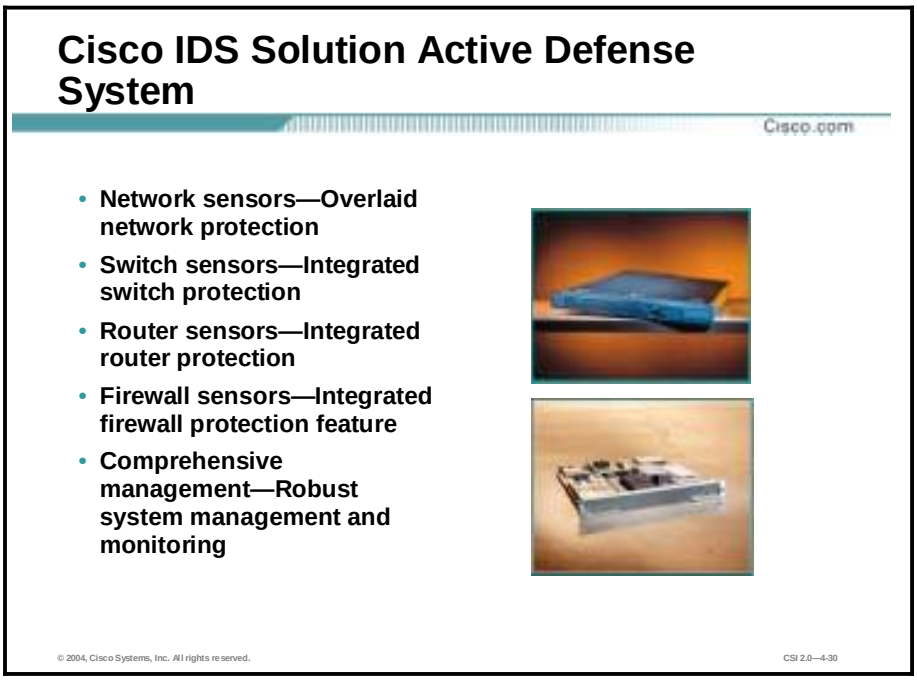

Cisco provides a complete product portfolio that enables customers to implement and manage active defense systems. The Cisco IDS products include the following:

- Network Sensors—Network Sensors provide a dedicated intrusion detection appliance with the ability to monitor and protect network segments.
- Switch Sensors—Switch Sensors are integrated into the switch fabric to provide seamless intrusion detection.
- Router Sensors—Router Sensors provide intrusion detection for deployments that require basic intrusion detection features.
- Firewall Sensors—Firewall Sensors provide intrusion detection for deployments that require basic intrusion detection features.
- Comprehensive management—A comprehensive management solution provides robust system management and monitoring.

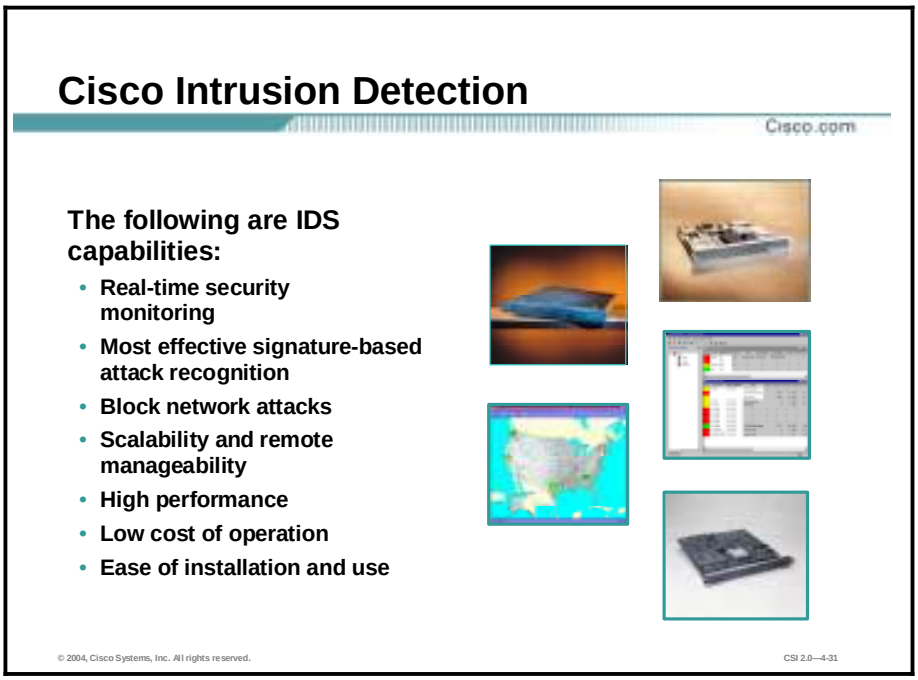

Cisco IDSs provide the following:

- Real-time security monitoring—IDS monitors security by capturing and analyzing packets.
- Effective signature-based attack recognition—When the Cisco IDS analyzes network data, it looks for patterns of misuse. Patterns can be as simple as an attempt to access a specific port on a specific host (an “atomic pattern”), or as complex as sequences of operations distributed across multiple hosts over an arbitrary period of time (a “composite pattern”).
- Blocking—The Cisco IDS uses another network device to deny entry to a specific network host or an entire network. To implement blocking, the Sensor dynamically reconfigures and reloads a network device’s ACLs. This type of automated response by the Sensor should only be configured for attack signatures with a low probability of false-positive detection, such as an unambiguous SATAN attack.
- Scalability and remote manageability—The Cisco IDS solutions provide a robust solution that is scalable to even the largest enterprise networks.
- High performance—Depending on the needs of a network, the Cisco IDS portfolio is designed to provide above-industry-standard performance for each platform, whether host-based or network-based.
- Low cost of operation—Delivering the lowest cost-of-ownership with network-integrated and hardware-based solutions, Cisco reduces the cost of advanced intrusion protection.
- Ease of installation and use—Because of the management platforms and menu-based configuration tools, the Cisco IDS is easily configured and maintained.

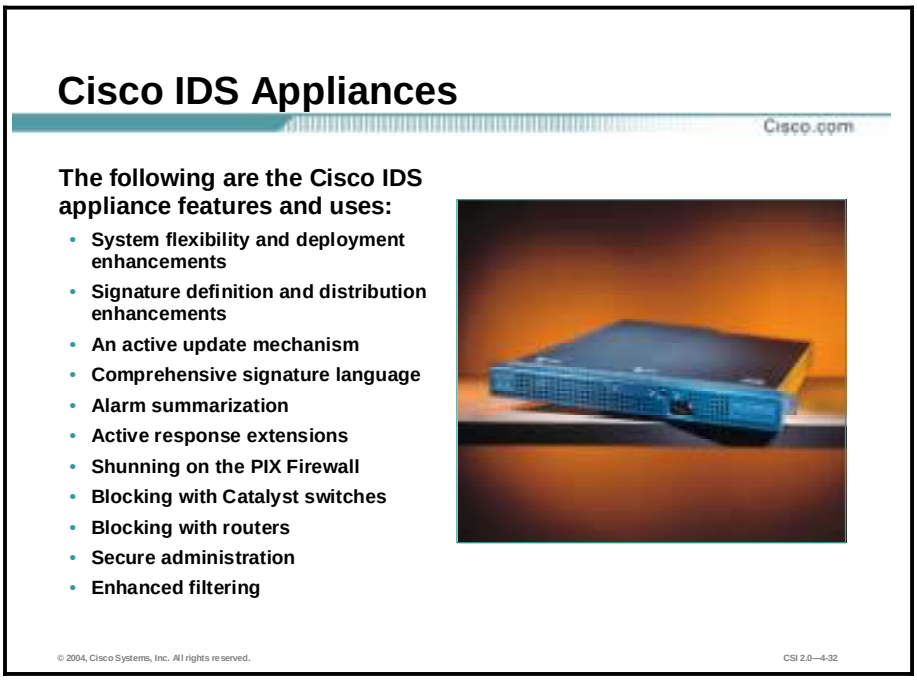

The following are the Cisco IDS appliance features and uses:

- System flexibility and deployment enhancements
- Signature definition and distribution enhancements
- An active update mechanism
- Comprehensive signature language
- Alarm summarization
- Active response extensions
- Shunning with the PIX Firewall
- Blocking with Catalyst switches
- Blocking with routers
- Secure administration
- Enhanced filtering

The complete line of market-leading Cisco IDS 4200 Series appliances delivers performance-optimized intrusion protection within an integrated, turnkey solution.

Note	The Sensor does not modify ACLs on the PIX Firewall. The **shun** command is used on the PIX Firewall to enforce blocking.

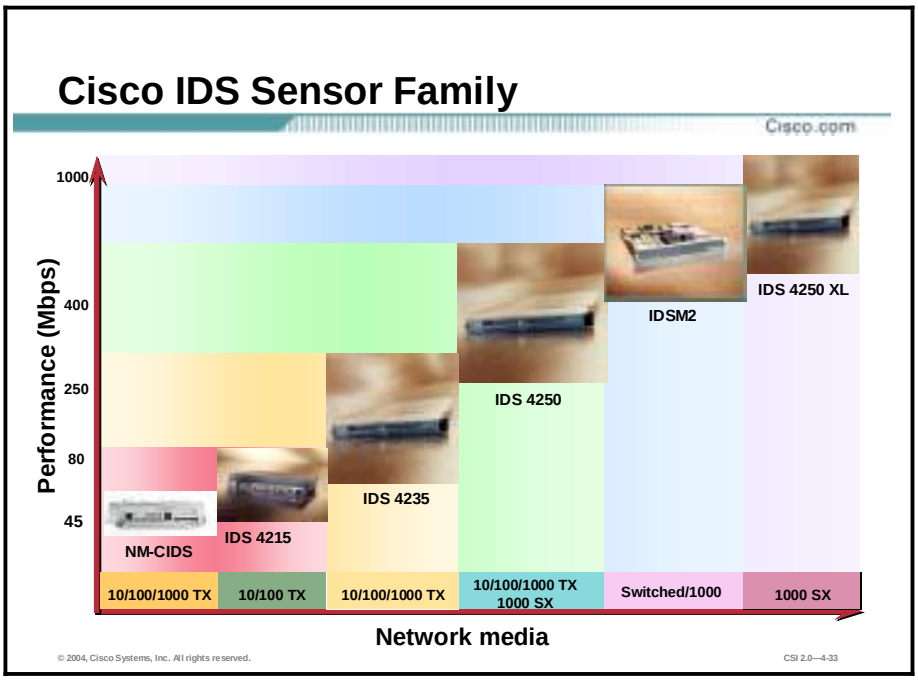

The figure refers to products that run IDS code 4.0 or higher, as follows:

The Cisco IDS 4200 Series Sensors are purpose-built, high-performance network security appliances that protect against unauthorized, malicious activity traversing the network (for example, attacks by hackers). Cisco IDS Sensors analyze traffic in real time, enabling users to quickly respond to security breaches.

The IDSM2 is a switching module that is easy to install and maintain in the Catalyst 6000 family switch. The IDSM2 performs network sensing—real-time monitoring of network packets through packet capture and analysis. The IDSM2 captures network packets and then reassembles them and compares this data against a rule set indicating typical intrusion activity. Network traffic is either copied to the IDSM2 based on security VLAN access control lists (VACLs) in the switch or is routed to the IDSM2 via the switch's Switch Port Analyzer (SPAN). Both methods allow user-specified kinds of traffic based on switch ports, VLANs, or traffic type to be inspected.

The NM-CIDS can be installed in a Cisco 2600XM, 2691, 3660, or 3700 Series router to provide 45 Mbps of full-featured intrusion protection services within the router. The NM-CIDS provides the ability to inspect all traffic traversing the router and then identify and terminate unauthorized or malicious activity. The NM-CIDS leverages the current Cisco IDS Sensor technology to expand IDS support into the branch office router. It requires an encryption feature set of Cisco IOS Software Release 12.2(15)ZJ or later for the routers.

Host-Based Intrusion Prevention System—CSA

This topic describes the features of host-based intrusion prevention system (HIPS) and introduces the HIPS, which takes host-based intrusion detection a step further.

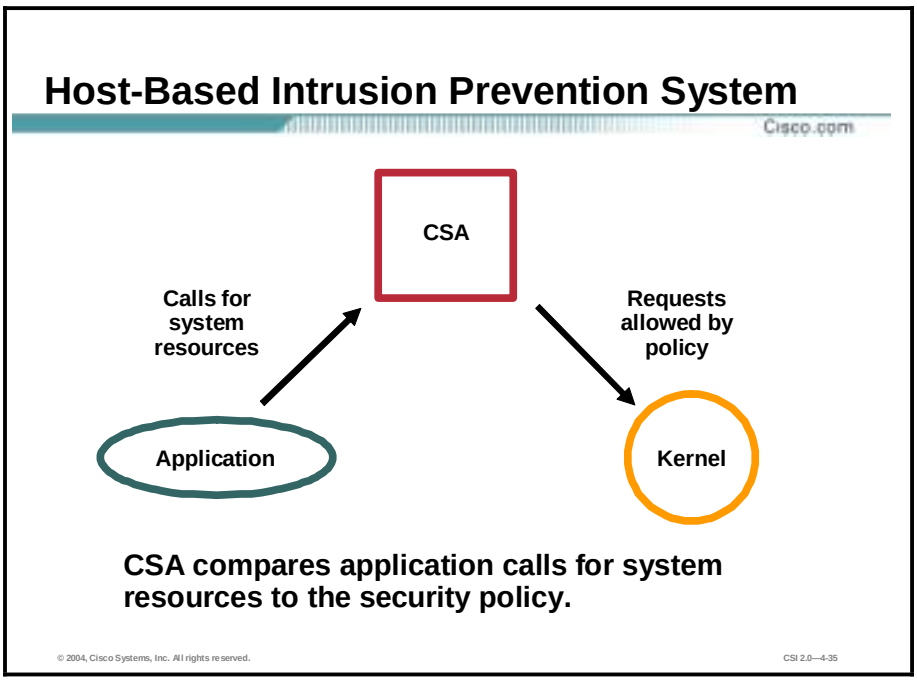

CSA is a HIPS that provides a third layer of depth to the network defense by applying security policy to system behavior at the host level. With this different approach to security, CSA can stop attacks missed at other levels of network security for the following reasons:

- CSA proactively blocks intrusive attacks by comparing all requests for system resources to the behaviors allowed by the security policy.
- CSA is not dependent upon signatures or updates to recognize attacks; in other words, it provides “day zero” protection from previously unknown attacks.
- CSA creates significantly fewer false positive alerts than any IDS; less administrative time is needed.

Note	CSA is a behavior-based intrusion prevention system and the Cisco IDS Sensor is a signature-based intrusion detection system.

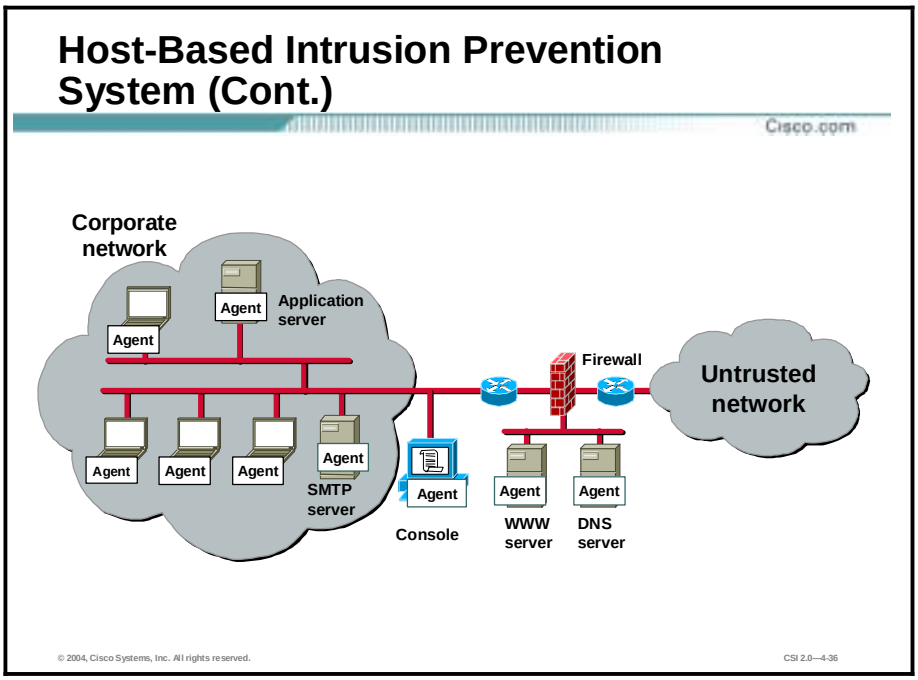

The figure illustrates a typical HIPS deployment. Agents are installed not only on publicly accessible servers, corporate mail servers, and application servers, but also on user desktops. The Agents report events to a central console server located inside the corporate firewall.

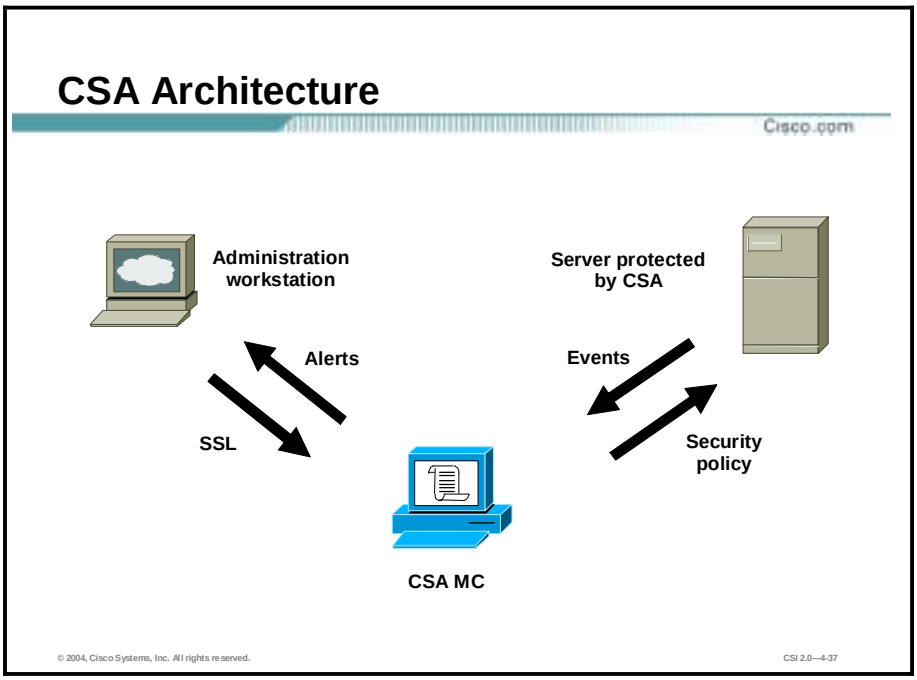

The CSA architecture model consists of:

- Management Center for Cisco Security Agent (CSA MC)—Allows the administrator to divide network hosts into groups by function and security requirements, and then configure security policies for those groups. The CSA MC can maintain a log of security violations and send alerts through e-mail or pager.
- CSA—Software installed in the host systems that continually monitors local system activity and analyzes the operations of that system. CSA takes proactive action to block attempted malicious activity. CSA also polls the CSA MC at configurable intervals for policy updates.
- An administration workstation—Can be any workstation connecting securely to the CSA MC using a Secure Sockets Layer (SSL)-enabled web interface.

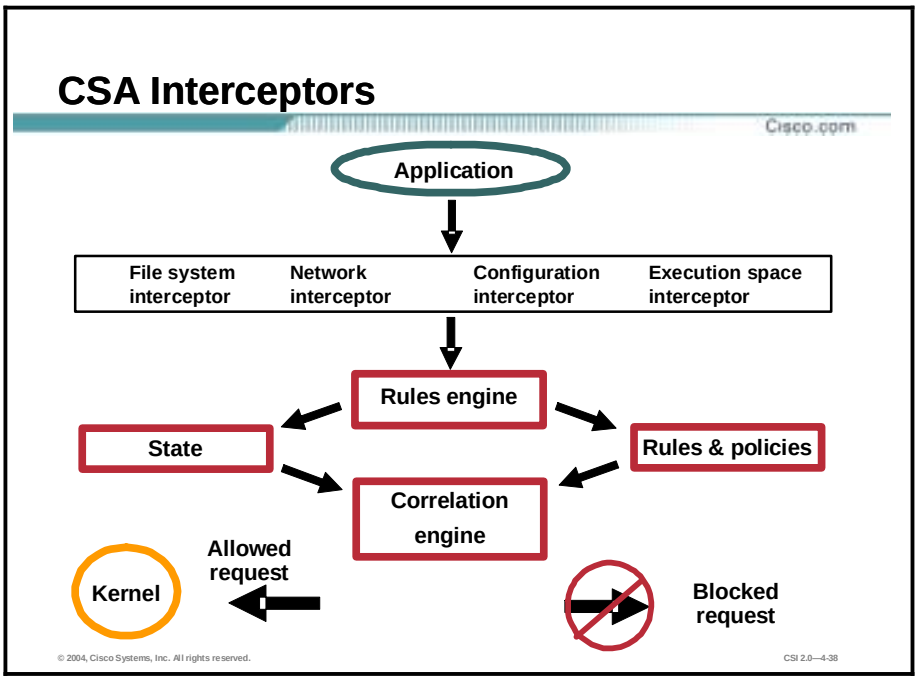

When an application needs access to system resources, it makes an operating system call to the kernel. CSA intercepts these operating system calls and compares them to the cached security policy. If the request does not violate policy, it is passed to the kernel for execution.

If the request does violate policy, then CSA takes the following actions:

- The request is blocked—not passed to the kernel.
- An appropriate error message is passed back to the application.
- An alert is generated and sent to the CSA MC.

CSA correlates this particular operating system call with others made by that application or process and correlates these events to detect malicious activity.

CSA provides protection through deployment of these four interceptors:

- File system interceptor—All file read or write requests are intercepted and allowed or denied based on the security policy.
- Network interceptor—Network driver interface specification (NDIS) changes are controlled and network connections are cleared through the security policy by port/IP address pairs. The number of network connections allowed within a specified time can also be limited to prevent DoS attacks.
- Configuration interceptor—Read/write requests to the registry in Windows or to rc files in UNIX are intercepted. Because modification of operating system configuration is highly unusual, it is tightly controlled by CSA.
- Execution space interceptor—This interceptor deals with maintaining the integrity of each application's dynamic runtime environment. Requests to write to memory not owned by the requesting application are detected and blocked by this interceptor. Attempts by one application to inject code, such as a shared library or dynamic link library (DLL), into another are also detected and blocked. Buffer overflow attacks are detected by this interceptor as well. The result is that not only is the integrity of dynamic resources, such as

the file system and configuration, preserved, but the integrity of highly dynamic resources such as memory and network I/O is also preserved.

CSA Interceptors (Cont.)

Cisco.com

Security application	Network interceptor	File system interceptor	Configuration interceptor	Execution space interceptor
Distributed firewall	X			
Host intrusion detection	X			X
Application sandbox		X	X	X
Network worm prevention	X			X
File integrity monitor		X	X	

By intercepting communication between applications and the underlying system, CSA combines the functionality of the following traditional security approaches:

- Distributed firewall—The network interceptor does the duties of a host firewall.
- HIDS or HIPS—The network interceptor teams with the execution space interceptor to provide the alerting capability of an HIDS or HIPS with the proactive enforcement of security policy.
- Application sandbox—An application sandbox is an execution space in which suspect programs can be run with less than normal access to system resources. This security service is provided by a combination of the file system, configuration, and the execution space interceptors.
- Network worm prevention—The network and execution space interceptors provide Day Zero worm prevention without a need for updates.
- File integrity monitor—The file system and configuration interceptors act as a file integrity monitor.

The default policies preconfigured on CSA implement all of these security features. Customers can easily create or change policies, but the default policies provide all of these protections at once.

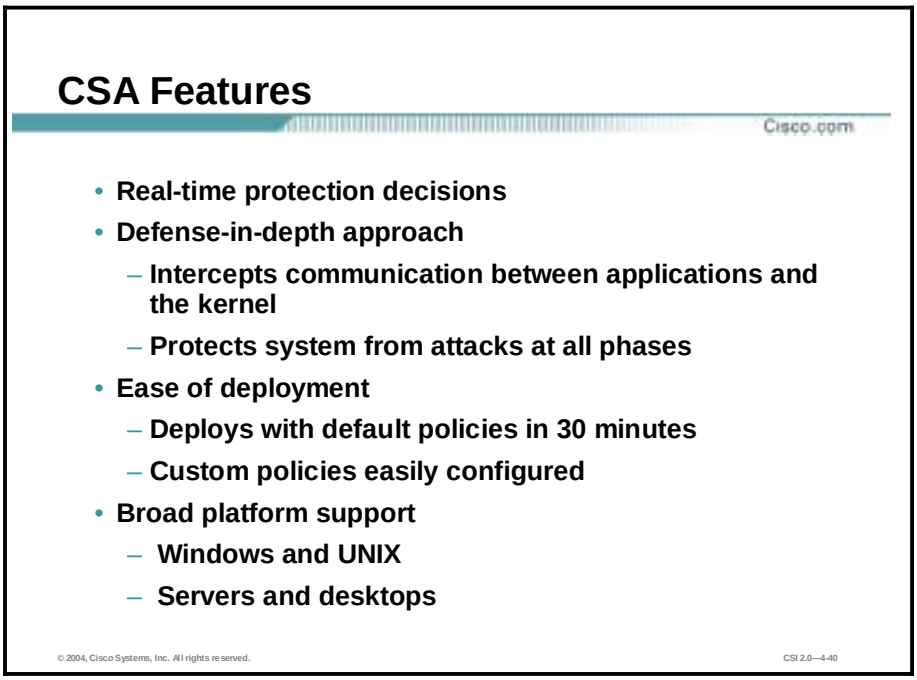

The following are the key features of CSA:

- Real-time protection decisions—Provides protection in real time rather than spotting attacks after they have happened.
- Defense-in-depth approach—More than a network perimeter defense or an attempt to detect attacks inside the network, CSA controls access to host system resources for complete protection.
  - Intercepts communication between applications and the kernel
  - Protects system from attacks at all phases
    - Network
    - File system
    - Configuration
    - Execution space
    - Ease of deployment
  - Deploys with default policies in 30 minutes
  - Custom policies easily configured
- Broad platform support
  - Windows and UNIX
  - Servers and desktops

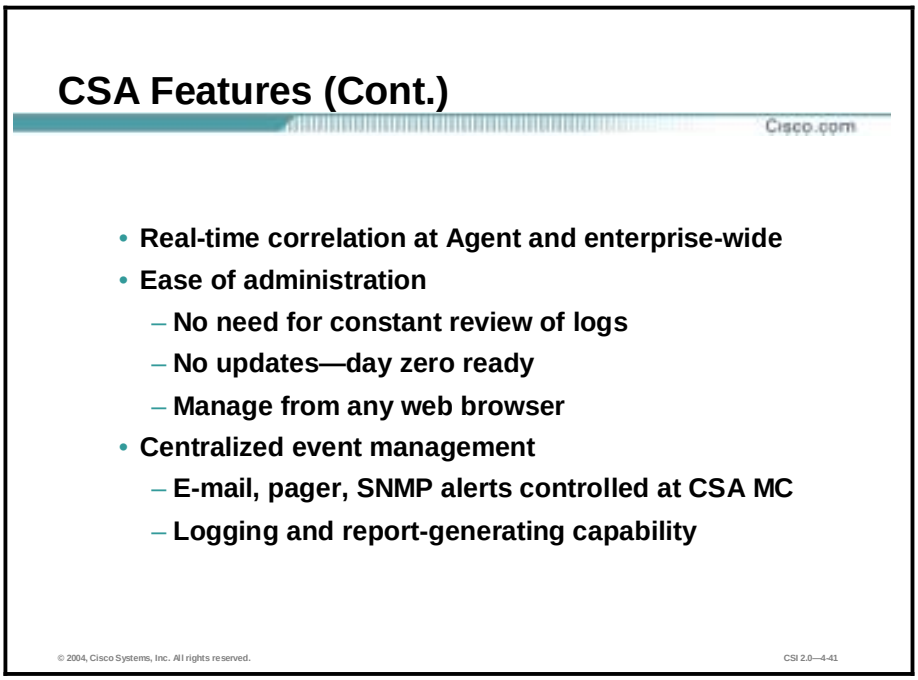

The following are additional CSA features:

- Real-time correlation at Agent and enterprise-wide—Reduces false positives and allows adaptability to new threats enterprise-wide
  - A network scan over multiple systems within a configured time period logs network events.
  - Worm events on multiple systems cause all systems to quarantine the contaminated files.
  - NT Event Logs and virus scanner logs can be correlated across the enterprise.
- Ease of administration
  - Less need for constant review of logs—Proactive defense approach minimizes requirement for administrator involvement.
  - No updates—“Day zero” ready.
  - Manage securely from any web browser.
- Centralized event management
  - E-mail, pager, and Simple Network Management Protocol (SNMP) alerts controlled at the CSA MC.
- Logging and report-generating capability.

Identity—Access Control Solutions

This topic describes the available Cisco access control solutions.

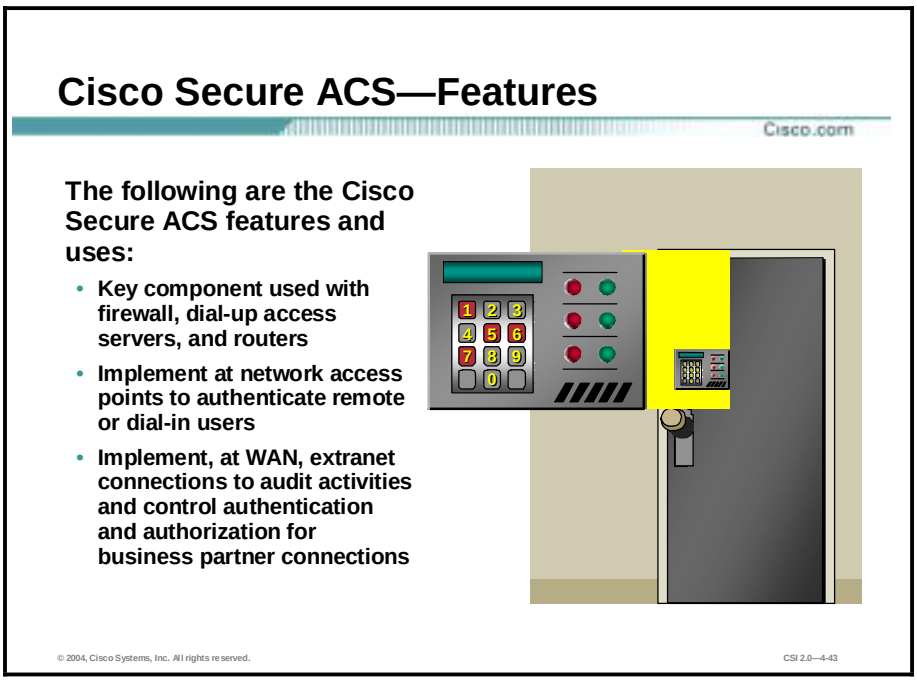

The features and uses of the Cisco Secure Access Control Server (ACS) are as follows:

- Key component used with firewall, dial-up access servers, and routers
- Implement at network access points to authenticate remote or dial-in users
- Implement extranet connections at WAN to audit activities and control authentication and authorization for business partner connections

You can leverage the same Cisco Secure ACS access framework to control administrator access and configuration for all network devices in your network that are enabled by RADIUS and Terminal Access Controller Access Control System Plus (TACACS+). Advanced features of the Cisco Secure ACS include the following:

- Automatic service monitoring
- Database synchronization and importation of tools for large-scale deployments
- Lightweight Directory Access Protocol (LDAP) user authentication support
- User and administrative access reporting
- Restrictions such as time of day and day of week
- User and device group profiles

Cisco Secure ACS—Product Summary

Cisco.com

The following is the Cisco Secure ACS product summary:

- **Easy-to-use web GUI**
- **Full RADIUS and TACACS+ user and administrator access control**
- **High performance (500+ authorizations per second)**
- **Supports LDAP, NDS, and ODBC datastores**
- **Scalable data replication and redundancy services**
- **Full accounting and user reporting features**

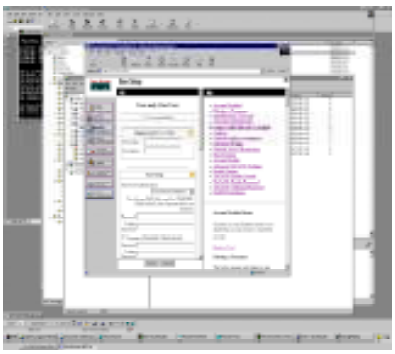

Cisco Secure ACS features include the following:

- Easy-to-use GUI
- Full RADIUS and TACACS+ user and administrator access control
- High performance (500+ authorizations per second)
- Supports LDAP, Novell Directory Service (NDS), and Open Database Connectivity (ODBC) datastores
- Scalable data replication and redundancy services
- Full accounting and user reporting features

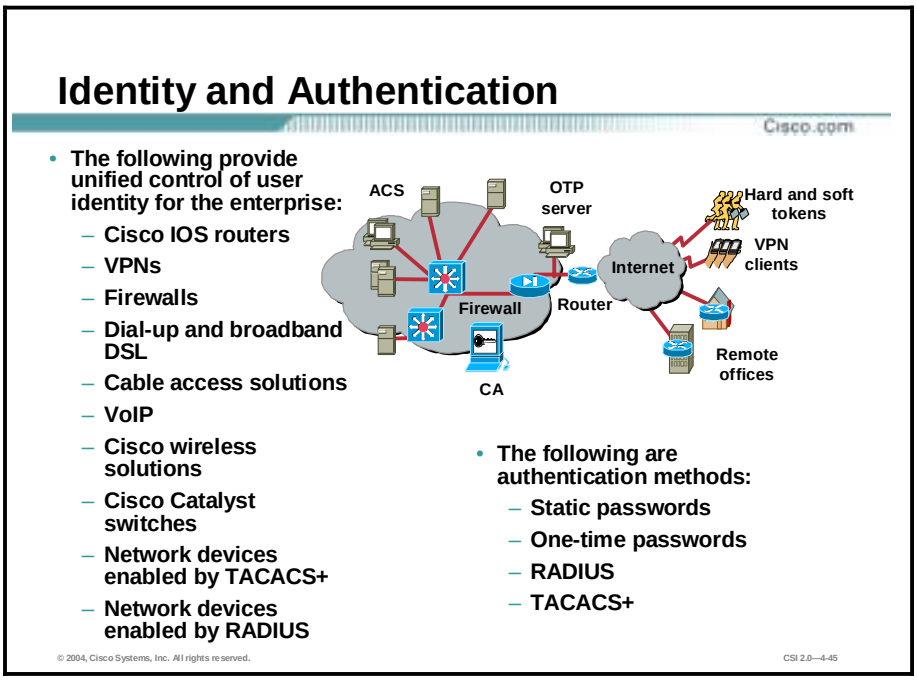

The Cisco Secure ACS is a high-performance, highly scalable, centralized user access control framework. Cisco Secure ACS offers centralized command and control for all user authentication, authorization, and accounting activities. Cisco Secure ACS also distributes those controls to hundreds or thousands of access gateways in your network. Authentication verifies user identity. Authorization configures integrity, such as user access rights. Accounting assists with auditing by logging user activities.

With Cisco Secure ACS you can manage and administer user access for the following:

- Cisco IOS routers
- VPNs
- Firewalls
- Dial-up and broadband DSL
- Cable access solutions
- Voice over IP (VoIP)
- Cisco wireless solutions
- Cisco Catalyst switches via IEEE 802.1x access control
- Network devices enabled by TACACS+
- Network devices enabled by RADIUS

The following authentication methods are employed:

- Static passwords
- One-time passwords
- RADIUS
- TACACS+

Security Management—Cisco IP Solution Center and VMS

This topic describes the Cisco security management solutions.

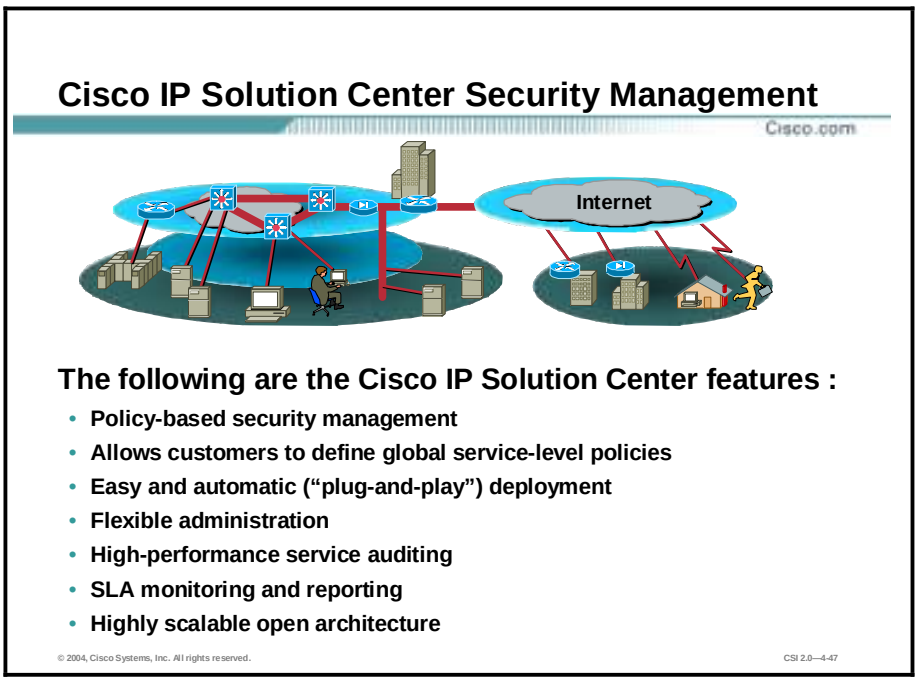

The Cisco IP Solution Center implements a business-centric, policy-level management model that allows our customers to define high-level security policies, while the application of those policies to specific network devices is off-loaded to the Cisco IP Solution Center software. The Cisco IP Solution Center Security Management application provides full support for the provisioning and management of LAN-to-LAN VPN, remote-access VPN, zero-touch deployment, Cisco Easy VPN, and Dynamic Multipoint VPN (DMVPN), firewall, Network Address Translation (NAT) and Quality of Service (QoS) technologies for the following Cisco security devices: Cisco IOS, PIX Firewall, VPN 3000 Concentrator, and so on. Following are Cisco IP Solution Center features:

- Policy-based security management—The Cisco IP Solution Center centrally manages the configuration of firewall and site-to-site VPN, network-based VPN, remote-access VPN, DMVPN, and Easy VPN devices. This allows customers to effectively deploy hundreds of thousands of security policies to their networks.

- Allows customers to define global service-level policies—The software will then automatically generate the device-level commands and provision the network accordingly. Once defined, global policies can be reused across multiple networks. This powerful management platform enables customers to:

  — Easily manage full-mesh, hub-and-spoke, or partial-mesh VPN topologies

  — Efficiently deploy site-to-site, network-based VPN, remote-access VPN, DMVPN, and Easy VPN technologies

  — Manage integrated GRE

- Design and deploy complex firewall rules
- Automate failover and load-balancing configuration
- Enable large-scale NAT configuration
- Manage integrated QoS service

- Easy and automatic, or "plug-and-play" deployment—As business increases, companies typically add new security devices to their networks. Cisco IP Solution Center, working in collaboration with embedded Cisco Networking Services (CNS) Intelligent Agents, can detect and manage newly added security devices dynamically and automatically. Once a new device is added to the network, the intelligent, embedded CNS Agent informs the Cisco CNS 2100 Series server, which operates Cisco IP Solution Center software, in real-time of all the latest information about that particular device. By subscribing to the Cisco CNS Message Bus, the Cisco IP Solution Center is able to dynamically manage the security policy, which applies to each new device, accordingly. Because of the dynamic nature of networks, a device configuration or status can be changed at any time. The intelligent, embedded CNS Agent can notify the Cisco CNS 2100 Series server of all the changes in network security devices—such as the change of the Dynamic Host Configuration Protocol (DHCP)-assigned IP address, loop-back interface, and so on—creating a network security management environment that does not require human intervention.
- Flexible administration—Cisco IP Solution Center provides role-based access control (RBAC) administration to enable granular management privileges control over network devices, services, provision actions, user groups, and all other possible components. Users can define administrative roles once and easily assign these roles to multiple users and user groups.
- High-performance service auditing—Cisco IP Solution Center Service Auditor validates IP service configurations and identifies faults to ensure high network integrity and service quality. The Cisco IP Solution Center also generates reports about the status of service deployment: requested, pending, deployed, or operating. Service assurance features ensure that IP service target devices remain provisioned correctly and that the service itself is operational.
- Service level agreement (SLA) monitoring and reporting—Cisco IP Solution Center SLA Manager monitors IP-aware SLAs for round-trip times, availability, and usage. Thresholds can be configured that allow violations to be reported and recorded for billing purposes.
- Highly scalable open architecture—Cisco IP Solution Center is a highly scalable, open security management platform. The system's four-tier architecture, consisting of client, interface, control, and distribution tiers, means it can manage tens of thousands of security systems and devices.

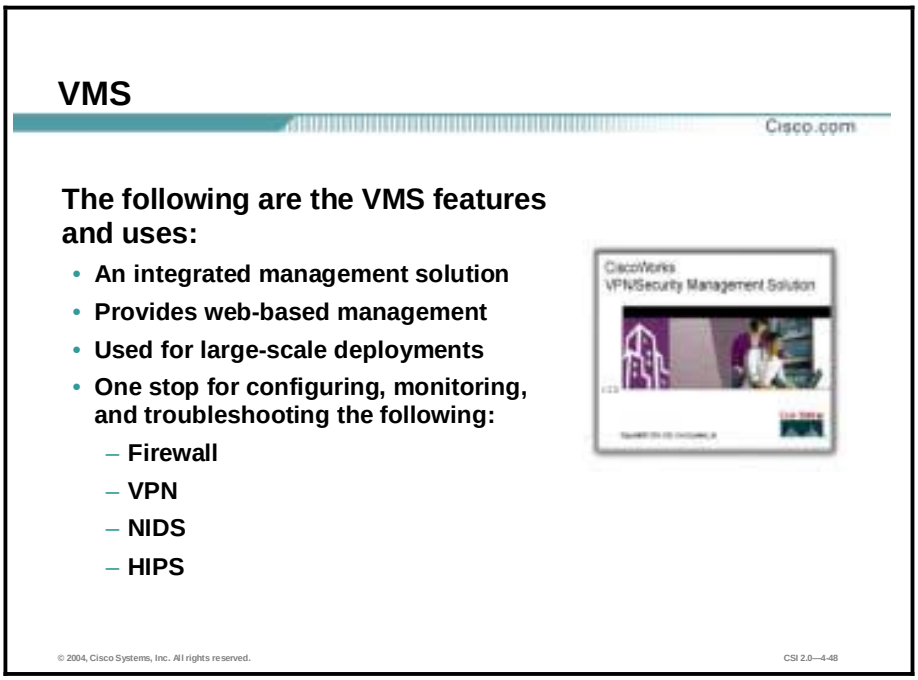

CiscoWorks VPN/Security Management Solution (VMS) is an integral element of the SAFE enterprise and contributes to organizational productivity by combining web-based tools for configuring, monitoring, and troubleshooting VPNs, firewalls, NIDS, and host intrusion prevention systems. Integrated with other CiscoWorks products, CiscoWorks VMS also includes network device inventory, change audit, and software distribution features.

CiscoWorks VMS 2.2 provides security management for your overall security needs. It includes the following applications, organized by functional area:

- Firewall management—Enables the large-scale deployment of Cisco firewalls. Smart Rules is an innovative feature that allows a security policy to be consistently applied to all firewalls. Smart Rules allows a user to define common rules once, reducing configuration time and resulting in fewer administrative errors.
- NIDS management—Offers efficient deployment of hundreds of Sensors using group profiles. Additionally, powerful signature management helps to increase the accuracy and specificity of detection.
- Host intrusion prevention system management—Scalable to thousands of endpoints per manager to support large enterprise deployments. The open and extensible architecture offers the capability to define and enforce security according to corporate policy. Offers "zero update" prevention for known and unknown attacks.
- VPN router management—Provides functions for the setup and maintenance of large deployments of VPN connections and Cisco IOS Firewalls on Cisco routers and Cisco Catalyst 6000 IPSec VPN Service Modules.
- Security monitoring—Provides integrated monitoring to help administrators have a comprehensive view of security across the network, with event correlation to detect threats not apparent with individual events.
- Performance monitoring—Provides functions for monitoring and troubleshooting services that contribute to enterprise network security.

- VPN monitoring—Allows network administrators to collect, store, and view information on VPN connections for remote-access or site-to-site VPN terminations.
- Operational management—Allows network managers to build a complete network inventory, report on hardware and software changes, and manage software updates to multiple devices.

Cisco AVVID

This topic discusses Cisco Architecture for Voice, Video, and Integrated Data (AVVID).

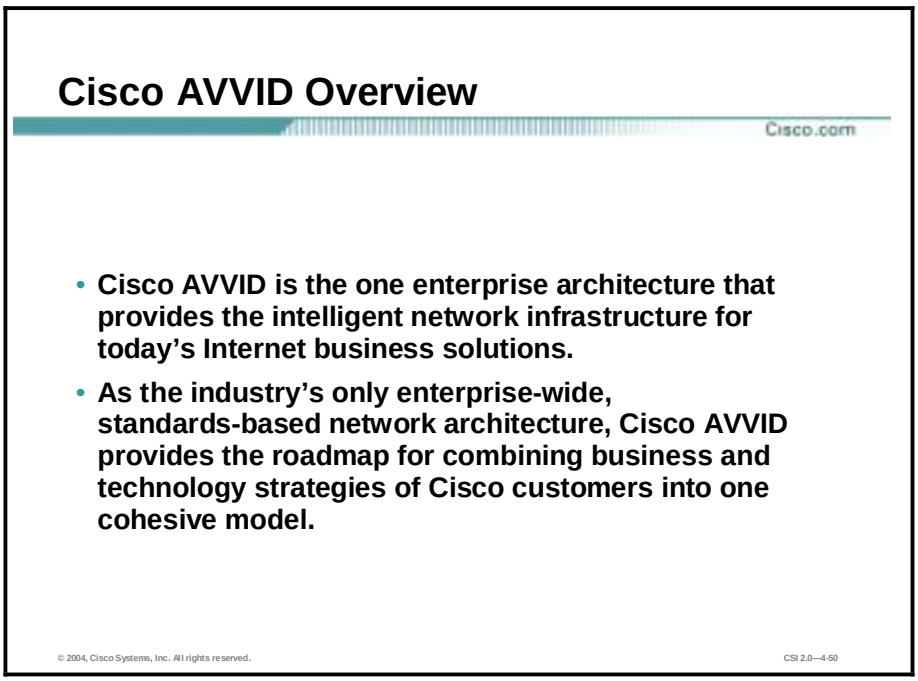

The Internet is creating tremendous business opportunities for Cisco and Cisco customers. Internet business solutions such as e-commerce, supply chain management, e-learning, and customer care are dramatically increasing productivity and efficiency.

Cisco AVVID is the one enterprise architecture that provides the intelligent network infrastructure for today's Internet business solutions. As the industry's only enterprise-wide, standards-based network architecture, Cisco AVVID provides the roadmap for combining customers' business and technology strategies into one cohesive model.

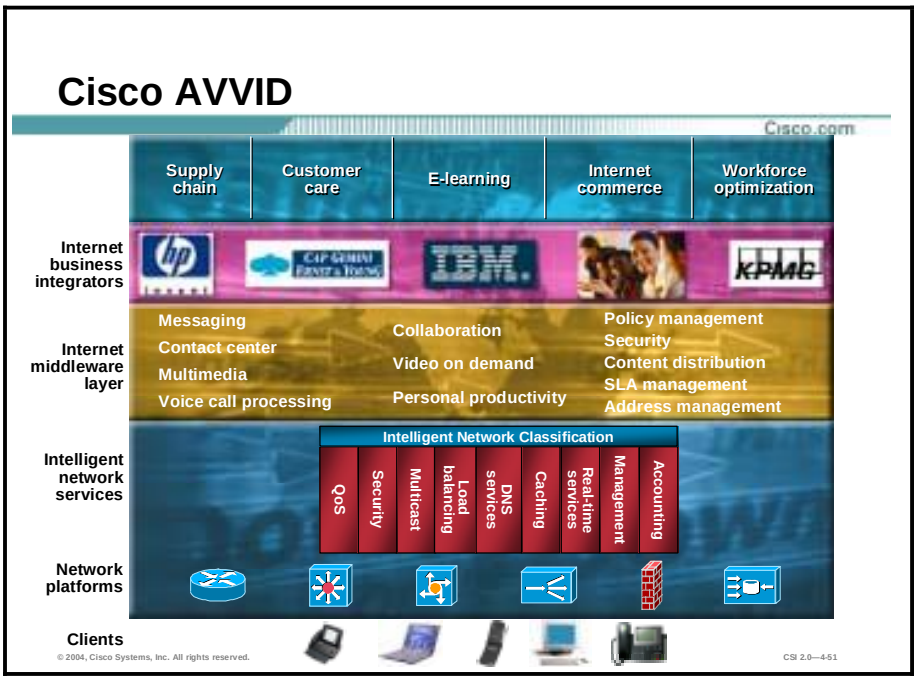

Cisco AVVID can be viewed as a framework to describe a network optimized for the support of Internet business solutions, and as a best practice or roadmap for network implementation. The following are the different parts of the Cisco AVVID:

- Clients
- Network platforms
- Intelligent network services
- Internet middleware layer
- Internet business integrators
- Internet business solution

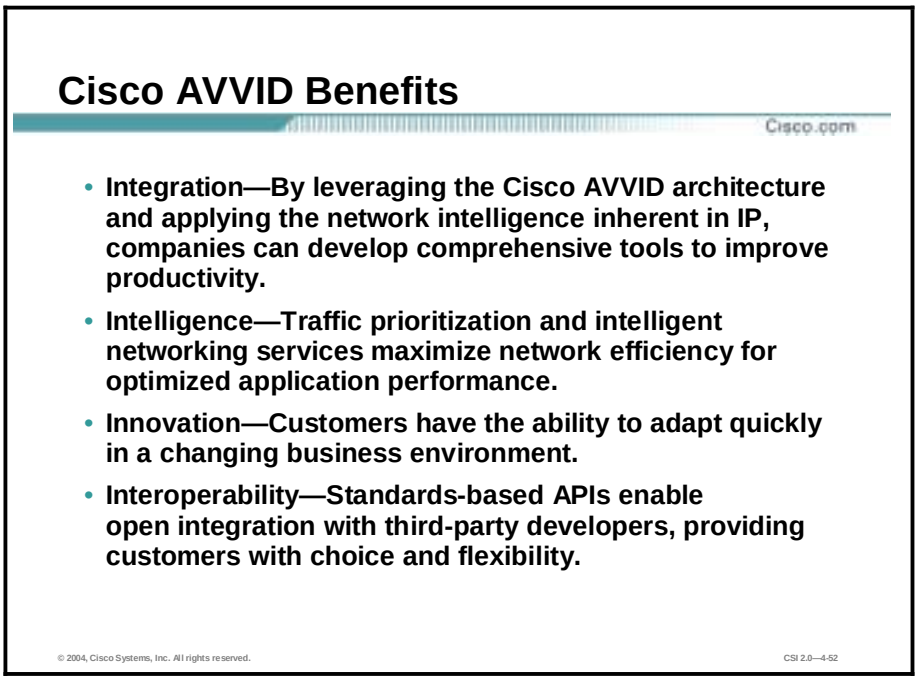

With Cisco AVVID, customers have a comprehensive roadmap for enabling Internet business solutions and creating a competitive advantage. There are four Cisco AVVID benefits:

- Integration—By leveraging the Cisco AVVID and applying the network intelligence inherent in IP, companies can develop comprehensive tools to improve productivity.
- Intelligence—Traffic prioritization and intelligent networking services maximize network efficiency for optimized application performance.
- Innovation—Customers have the ability to adapt quickly in a changing business environment.
- Interoperability—Standards-based APIs enable open integration with third-party developers, providing customers with choice and flexibility.

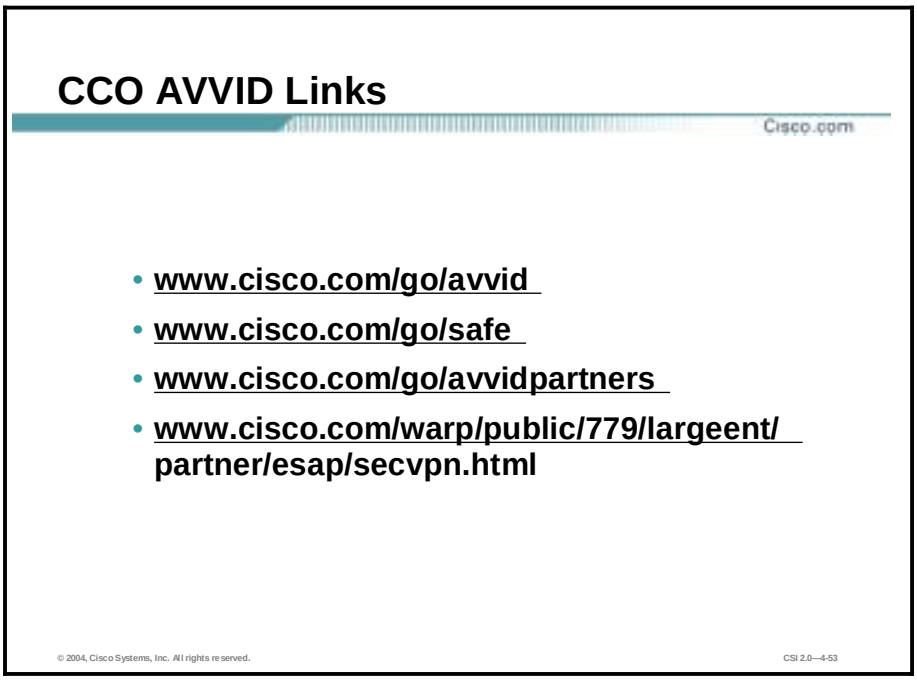

The Cisco AVVID program changes frequently with new partners and products being introduced on an ongoing basis. The links in the figure can be used to get the latest information about the AVVID program and products offered.

Summary

This topic summarizes the information you learned in this lesson.

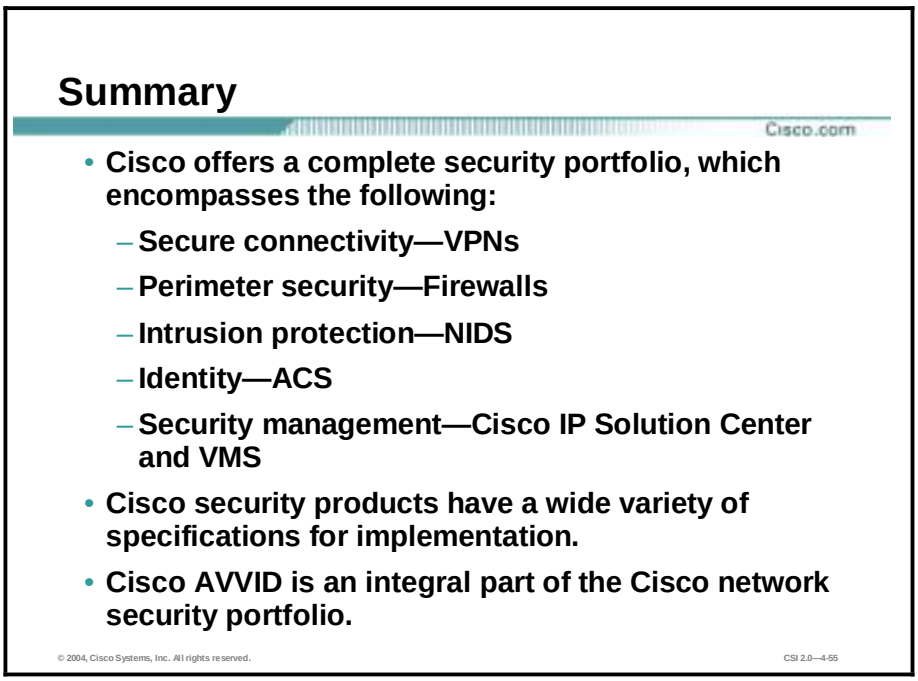

5

SAFE Small Network Design

Overview

This lesson describes the SAFE small network design. It includes the following topics:

- Objectives
- Small network design overview
- Small network corporate Internet module
- Small network campus module
- Implementation—ISP router
- Implementation—Cisco IOS Firewall features and configuration
- Implementation—PIX Firewall
- Implementation—CSA
- Summary
- Lab exercise

Objectives

This topic lists the lesson's objectives.

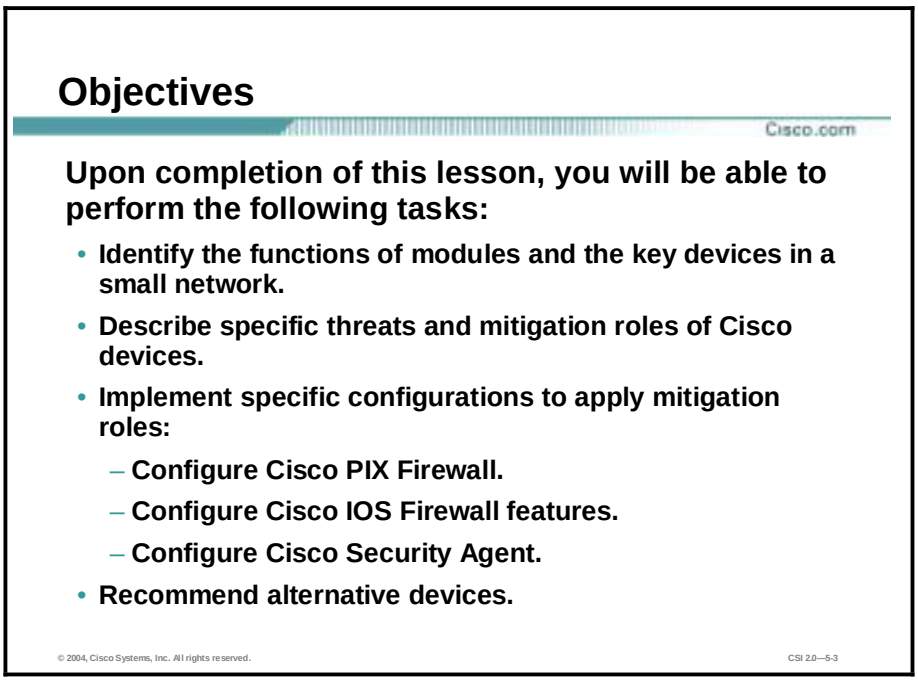

Small Network Design Overview

This topic provides an overview of the SAFE Extending the Security Blueprint to Small, Midsize, and Remote-User Networks (SAFE SMR) small network design.

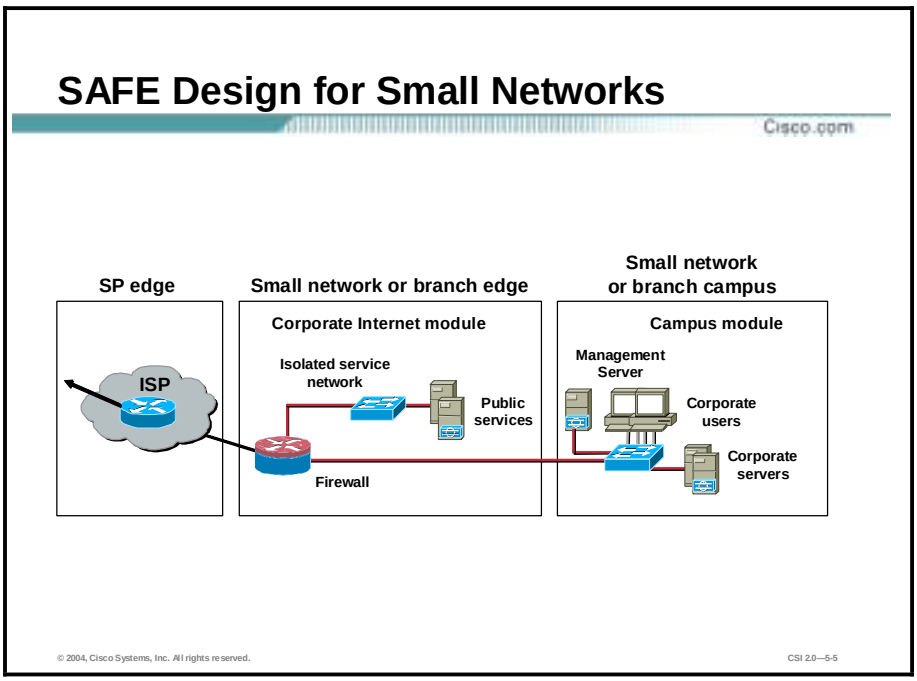

The small network design has two modules:

- Corporate Internet module—This module has connections to the Internet and also terminates virtual private network (VPN) and public services (Domain Name System [DNS], HTTP, FTP, Simple Mail Transfer Protocol [SMTP]) traffic.
- Campus module—This module contains the Layer 2 switching and all the users, as well as the management and intranet servers. (Most of the discussion in this lesson for this design is based on the small network operating as the head-end for a corporation. Specific design changes when used as a branch are also included.)

Small Network Corporate Internet Module

This topic discusses the corporate Internet module.

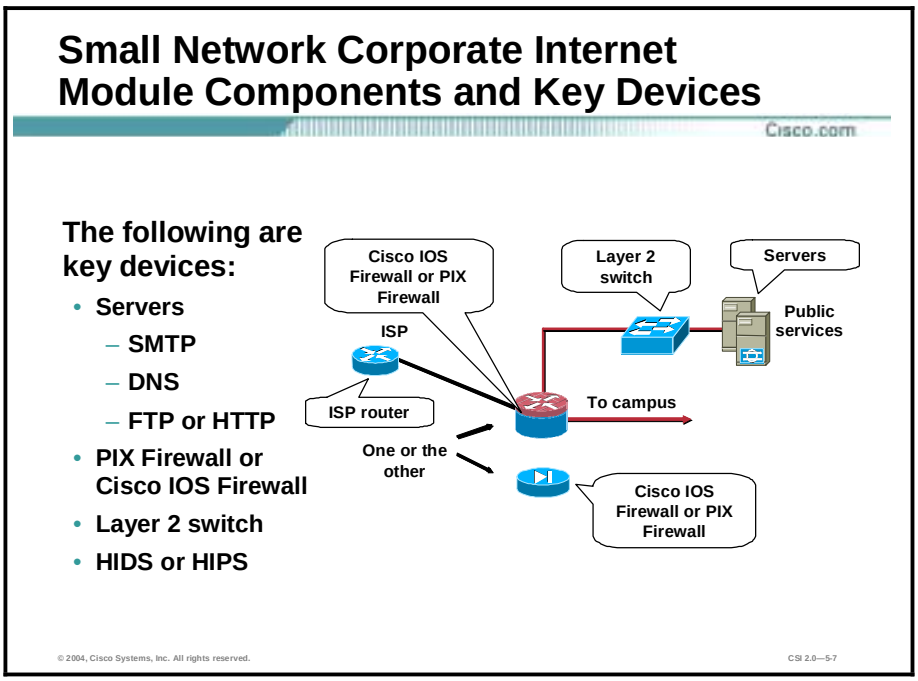

The corporate Internet module provides internal users with connectivity to Internet services and Internet users access to information on public servers. VPN access is also provided to remote locations and telecommuters. This module is not designed to serve e-commerce type applications.

The following are key devices in the corporate Internet module:

- SMTP server—Acts as a relay between the Internet and the intranet mail servers.
- DNS server—Serves as authoritative external DNS server for the small network. It relays internal requests to the Internet.
- FTP or HTTP server—Provides public information about the organization.
- Firewall or Cisco IOS Firewall router—Provides network-level protection of resources and stateful filtering of traffic and provides differentiated security for remote access users. It authenticates trusted remote sites and provides connectivity using IPSec tunnels.
- Layer 2 switch (with private VLAN support)—Provides Layer 2 connectivity for devices.
- Host-based intrusion detection sysem (HIDS) or host-based intrusion prevention system (HIPS)—Cisco Security Agent (CSA) provides host-level intrusion prevention.

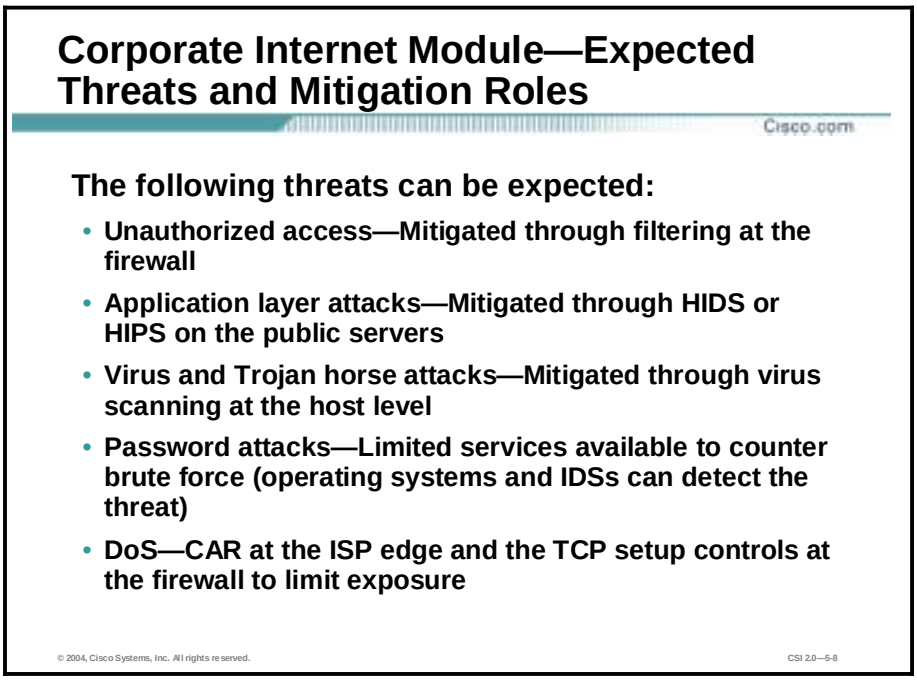

Publicly addressable servers are the most likely points of attacks. The following are expected threats to those publicly addressable servers in a corporate Internet module:

- Unauthorized access—Mitigated through filtering at the firewall
- Application layer attacks—Mitigated through HIDS or HIPS on the public servers
- Virus and Trojan horse attacks—Mitigated through virus scanning at the host level
- Password attacks—Limited services available to counter brute force (operating systems and intrusion detection systems [IDSs] can detect the threat)
- Denial of service (DoS)—Committed access rate (CAR) at the ISP edge and TCP setup controls at the firewall to limit exposure

From a threat perspective, a small or midsize network is like most networks connected to the Internet: there are internal users who need access out and external users who need access in. Several common threats can generate the initial compromise that a hacker needs to further penetrate the network with secondary exploits.

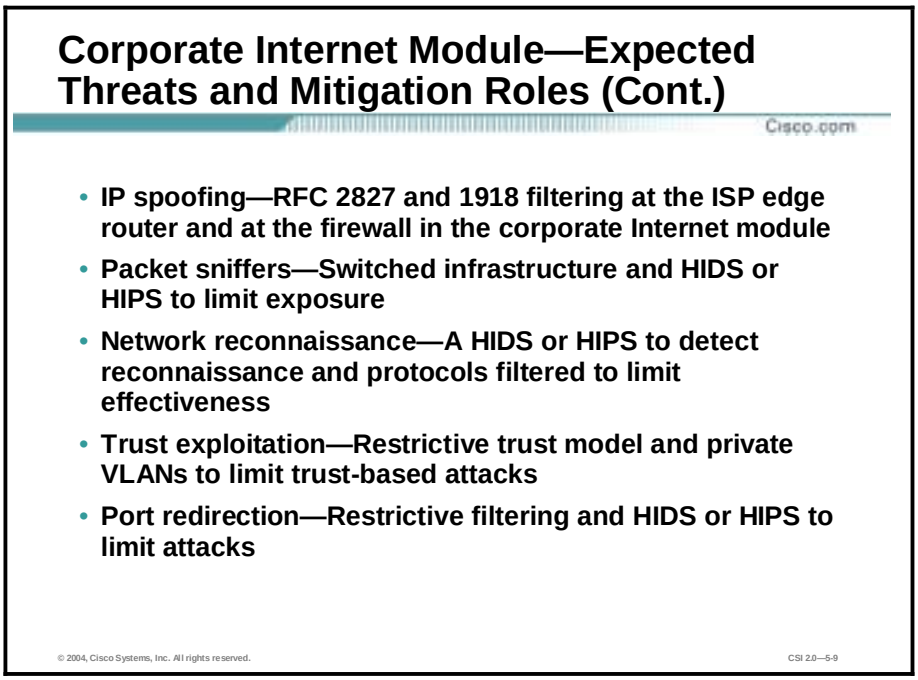

Mitigation roles to threats to devices in corporate internet module include the following:

- IP spoofing—RFC 2827 and RFC 1918 filtering at the ISP edge and local firewall
- Packet sniffers—Switched infrastructure and a HIDS or HIPS to limit exposure
- Network reconnaissance—A HIDS or HIPS to detect reconnaissance, and protocols filtered to limit effectiveness
- Trust exploitation—Restrictive trust model and private VLANs to limit trust-based attacks
- Port redirection—Restrictive filtering and a HIDS or HIPS to limit attacks

Though statistics vary on the percentage, it is an established fact that most attacks come from the internal network. Disgruntled employees, corporate spies, visiting guests, and bumbling users are all potential sources of such attacks. When designing security, you must be aware of the potential for internal threats.

There is also the risk of threat to the publicly addressable hosts that are connected to the Internet. These systems will likely be attacked with application-layer vulnerabilities and DoS attacks.

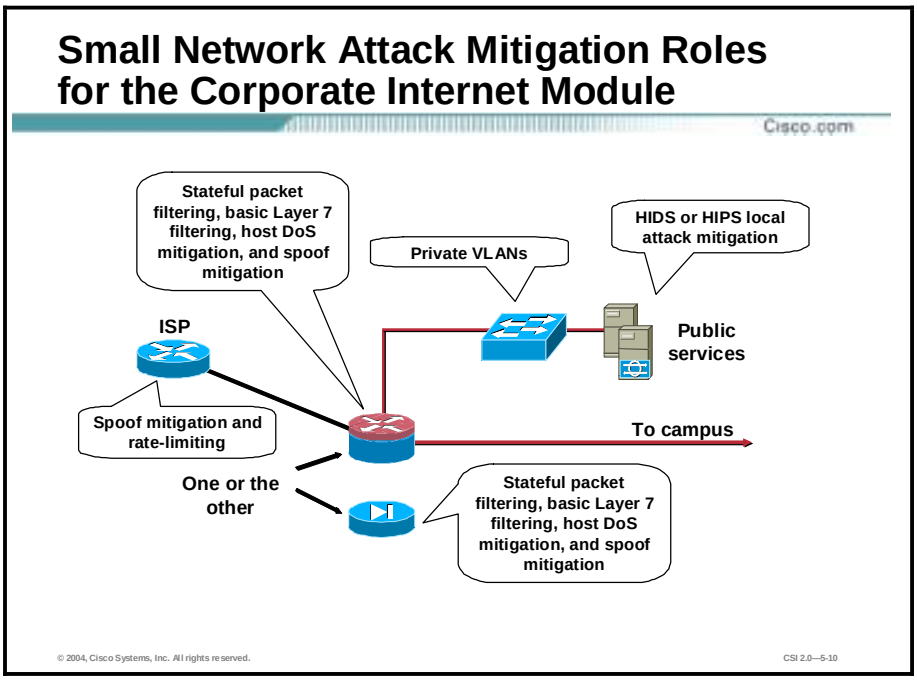

The attack mitigation roles for each device in the SAFE SMR small network corporate Internet module are as follows:

- ISP router
  - Spoof mitigation
  - Rate-limiting
- Firewall
  - Stateful packet filtering
  - Basic Layer 7 filtering
  - Host DoS mitigation
  - Spoof mitigation
- Layer 2 switches (with private VLAN support)
  - Private VLANs
- HIDS or HIPS—Local attack mitigation

At the ingress of the firewall, RFC 1918 and RFC 2827 filtering are provided as verification of the ISP's filtering. Because of the enormous security threat fragmented packets create, the firewall is configured to drop most fragmented packets. Any legitimate traffic lost due to filtering is considered acceptable when compared to the risk of allowing such traffic to traverse the network. Traffic destined to the firewall from outside the network is limited to IPSec traffic and necessary routing protocols.

The firewall provides connection-state enforcement and detailed filtering for sessions initiated through the firewall. From a filtering standpoint, in addition to limiting traffic on the public services segment to relevant addresses and ports, the firewall also provides filtering in the opposite direction. If an attack compromises one of the public servers (by circumventing the firewall and HIDS or HIPS), that server should not be able to further attack the network. To

mitigate this type of attack, specific filtering prevents any unauthorized requests from being generated by the public servers to any other location. For example, the web server should be filtered so that it cannot originate requests of its own, but merely respond to requests from clients. This setup helps prevent a hacker from downloading additional utilities to the compromised device after the initial attack. It also helps stop unwanted sessions from being triggered by the hacker during the primary attack. An attack that generates an Xwindows terminal session from the web server through the firewall to the hacker's machine is an example of such an attack.

In addition, private VLANs on the Demilitarized Zone (DMZ) switch prevent a compromised public server from attacking other servers on the same segment. This traffic is not even detected by the firewall, a fact that explains why private VLANs are critical. Finally, publicly addressable servers have some protection against TCP SYN floods through mechanisms such as the use of half-open connection limits on the firewall.

From a host perspective, each of the servers on the public services segment has host intrusion detection software to monitor against any rogue activity at the operating system level, as well as activity in common server applications (HTTP, FTP, SMTP, and so forth). The DNS host should be locked down to respond only to desired commands and eliminate any unnecessary responses that might assist hackers in network reconnaissance. This includes preventing zone transfers from anywhere except legitimate secondary DNS servers. For mail services, the firewall itself filters SMTP messages at Layer 7 to allow only necessary commands to the mail server.

Firewalls and firewall routers generally have some limited network-based intrusion detection system (NIDS) capability within their security functions. This capability affects the performance of the device, but does provide some additional attack visibility in the event you are under attack. Remember that you are trading performance for attack visibility. Many of these attacks can be dropped without the use of an IDS, but the monitoring station will not be aware of the specific attack being launched.

The VPN connectivity is provided through the firewall or firewall and router. Remote sites authenticate each other with pre-shared keys, and remote users are authenticated through the access control server in the campus module.

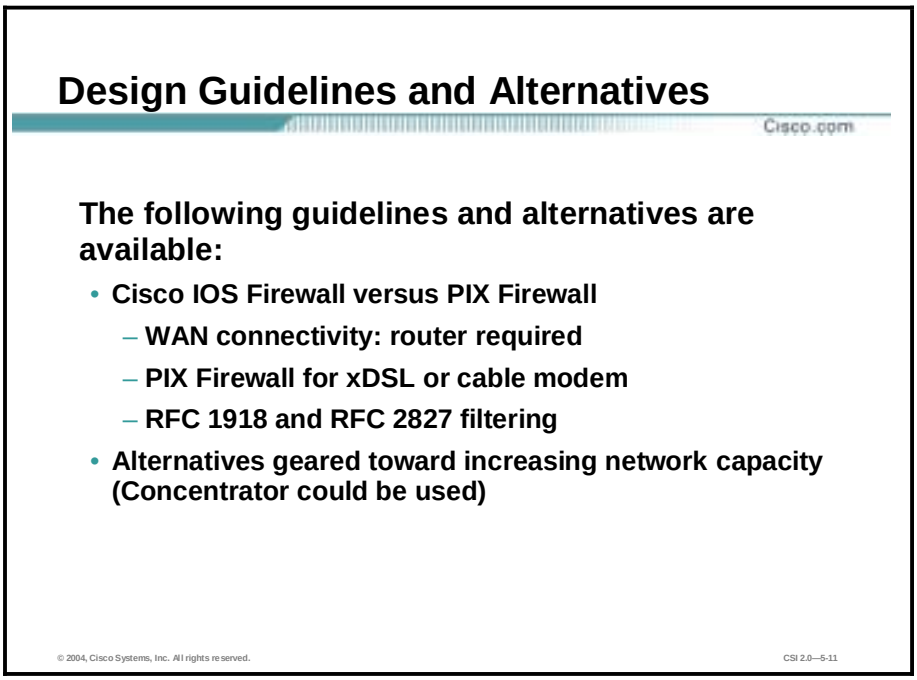

The SAFE SMR small network design alternatives represent the ultimate in scaled-down, security-conscious network design, where all the security and VPN services are compressed into a single device. There are two choices when deciding how to implement this network design:

- Use a router with firewall and VPN functionality—This setup yields the greatest flexibility for the small network because the router will support all the advanced services (Quality of Service [QoS], routing, multiprotocol support, and so on) that may be necessary in today's networks.
- Use a dedicated firewall instead of the router—This setup places some restrictions on the deployment. First, firewalls are generally Ethernet-only, requiring some conversion to the appropriate WAN protocol. In today's environments, most cable and DSL routers and modems are provided by the service provider and can be used to connect to the firewall over Ethernet. If WAN connectivity on the device is required (such as with a circuit from a telco provider), then a router must be used. Using a dedicated firewall does have the advantage of easier configuration of security services, and a dedicated firewall can provide improved performance when doing firewall functions. Whatever the selection of device, stateful inspection is used to examine traffic in all directions, ensuring that only legitimate traffic crosses the firewall. Before the traffic even reaches the firewall, ideally, some security filtering has already occurred at the ISP. Remember that routers tend to start out permitting traffic, whereas firewalls tend to deny traffic by default. At the ingress of the firewall, RFC 1918 and RFC 2827 filtering is provided as supplement to ISP's filtering.
- Any deviation from the SAFE SMR small network design would be geared toward increasing the capacity of the network, or separating the various security functions onto distinct devices. In doing this, the design starts to look more and more like the medium network design discussed later in this lesson. Instead of adopting the complete medium design, you might consider the addition of a dedicated remote access Cisco VPN Concentrator to increase the manageability of the remote-user community.

Small Network Campus Module

This topic describes the small network campus module and its key devices.

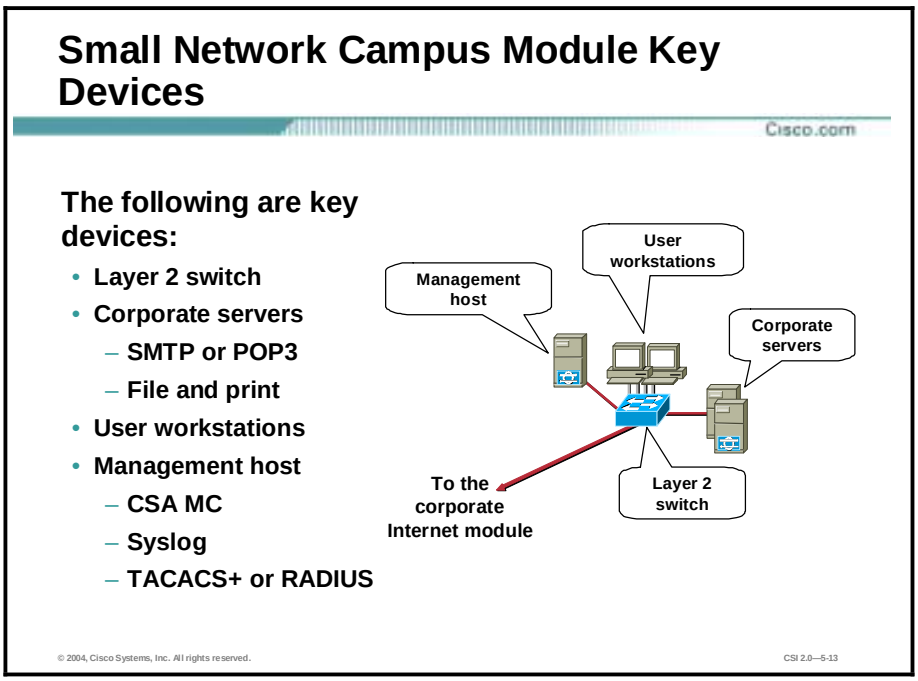

The campus module contains end-user workstations, corporate intranet servers, management servers, and the associated Layer 2 infrastructure required to support the devices. Within the small network design, this Layer 2 functionality has been combined into a single switch.

The following are key devices for the SAFE SMR small network campus module:

- Layer 2 switch (with private VLAN support)—Provides Layer 2 services to user workstations
- Corporate servers—Provides e-mail (SMTP and Post Office Protocol Version 3 [POP3]) services to internal users, as well as delivering file, print, and DNS services to workstations
- User workstations—Provides data services to authorized users on the network
- Management host—Provides Management Center for Cisco Security Agent (CSA MC) to deploy and manage CSA on end-points, Syslog, Terminal Access Controller Access Control System Plus (TACACS+) or Remote Access Dial-In User Service (RADIUS), and general configuration management

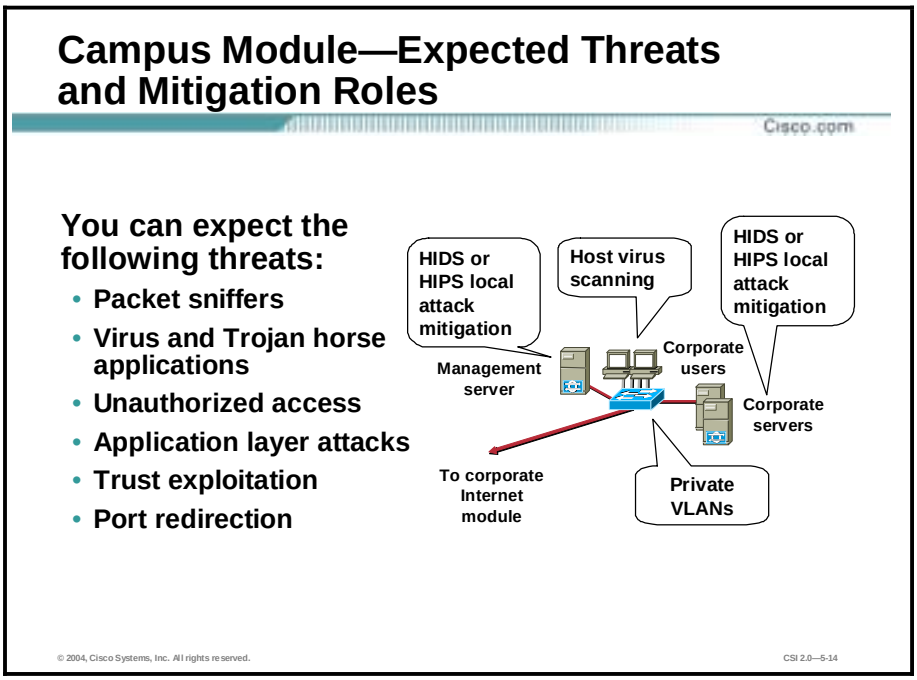

The following are expected threats and mitigation of those threats to the SAFE SMR small network campus module:

- Packet sniffers—A switched infrastructure limits the effectiveness of sniffing.
- Virus and Trojan horse applications—Host-based virus scanning prevents most viruses and many Trojan horses.
- Unauthorized access—This type of access is mitigated through the use of host-based intrusion detection and application access control.
- Application layer attacks—Operating systems, devices, and applications are kept up-to-date with the latest security fixes, and they are protected by a HIDS or HIPS.
- Trust exploitation—Private VLANs prevent hosts on the same subnet from communicating unless necessary.
- Port redirection—A HIDS or HIPS prevents port redirection agents from being installed.

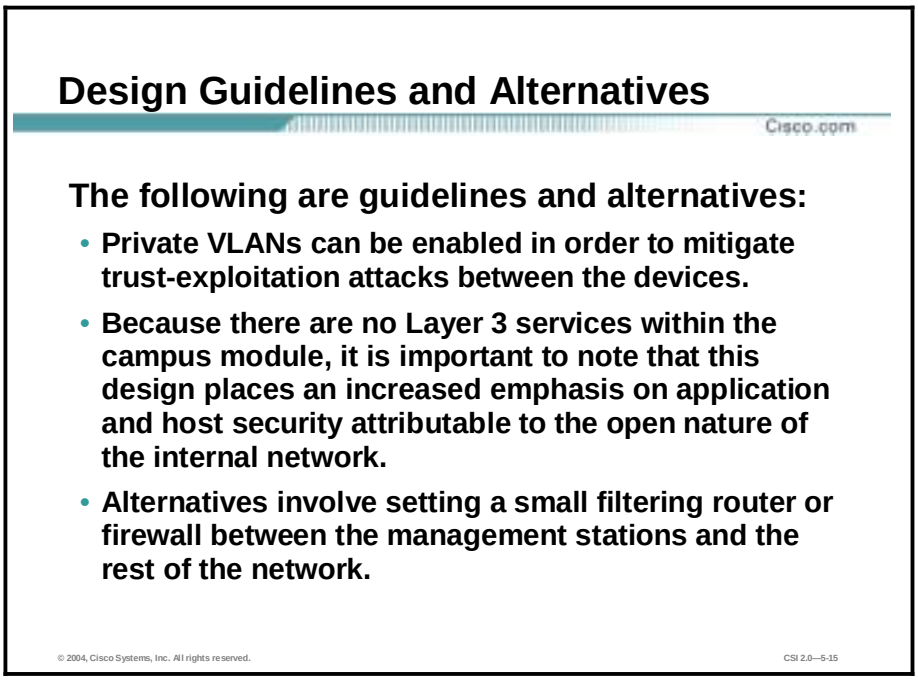

The primary functions of the campus switch are to switch production and management traffic and to provide connectivity for the corporate and management servers and users. Within the switch, private VLANs can be enabled to mitigate trust-exploitation attacks between the devices. For instance, the corporate users might need to be able to talk to the corporate servers but may not have any requirement to communicate with each other.

Because there are no Layer 3 services within the campus module, it is important to note that the SAFE SMR small network campus module design places an increased emphasis on application and host security because of the open nature of the internal network. Therefore, a HIDS or HIPS was also installed on key systems within the campus, including the corporate servers and management systems.

Setting a small filtering router or firewall between the management stations and the rest of the network can improve overall security. This setup allows management traffic to flow only in the specific direction deemed necessary by the administrators. If the level of trust within the organization is high, a HIDS or HIPS can potentially be eliminated, though this is not recommended.

Implementation—ISP Router

This topic details the implementation of the ISP router.

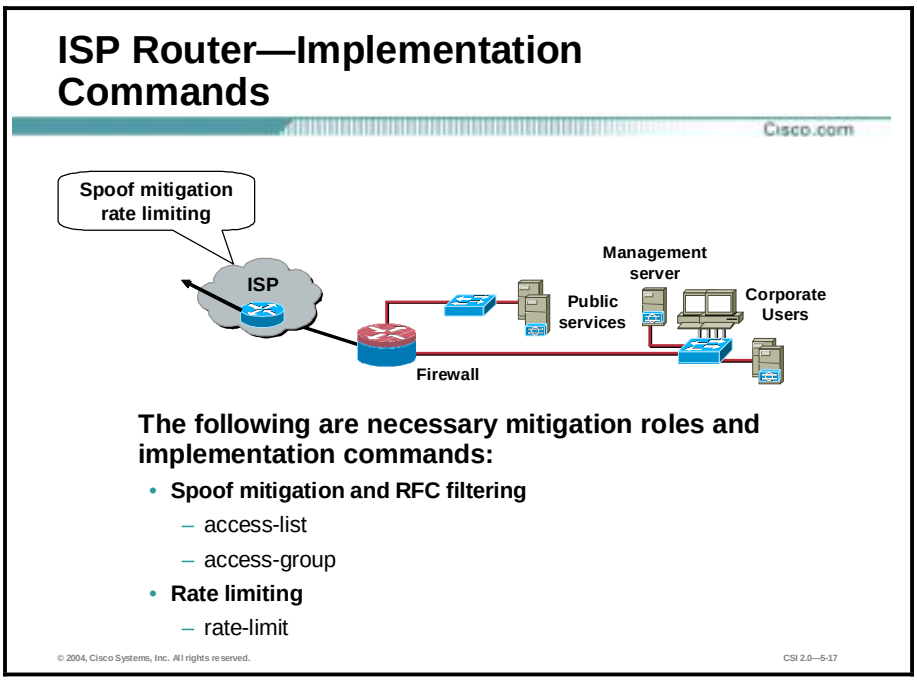

The primary function of the customer-edge router in the ISP is to provide connectivity to the Internet or ISP network. The egress of the ISP router rate-limits nonessential traffic that exceeds pre-specified thresholds in order to mitigate distributed denial of service (DDoS) attacks. At the egress of the ISP router, RFC 1918 and RFC 2827 filtering is configured to mitigate source-address spoofing of local networks and private address ranges.

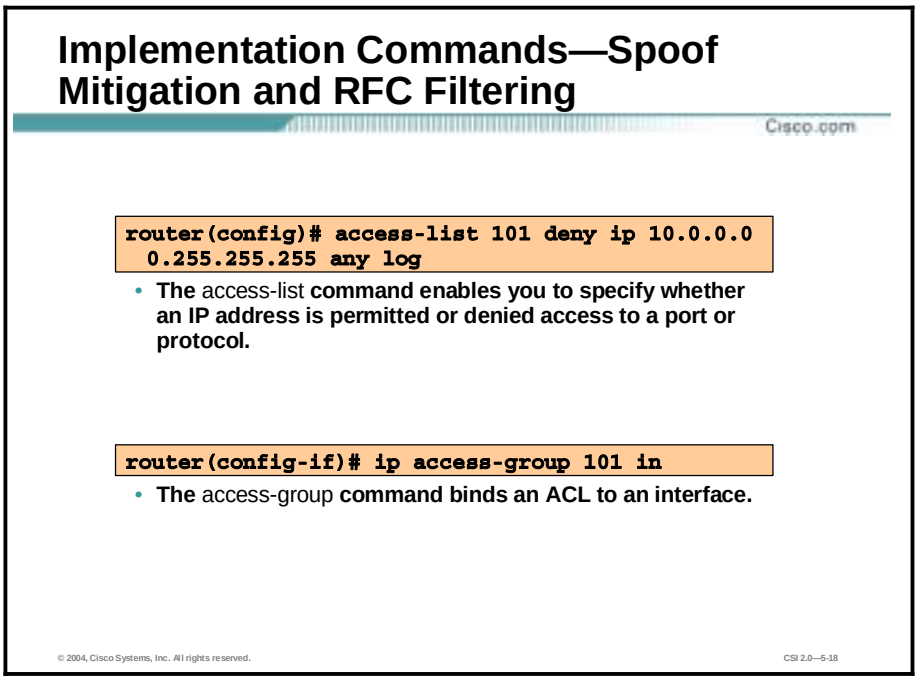

Cisco provides basic traffic filtering capabilities with access control lists (ACLs). ACLs can be configured for all routed network protocols (IP, AppleTalk, and so on) to filter the packets of those protocols as the packets pass through a router.

You can configure ACLs at your router to control access to a network. ACLs can prevent certain traffic from entering or exiting a network.

ACLs should be used in Cisco IOS Firewall routers, which are often positioned between your internal network and an external network (for example, the Internet). You can also use ACLs on a router positioned between two parts of your network, to control traffic entering or exiting a specific part of your internal network.

To provide the security benefits of ACLs, you should at a minimum configure ACLs on border routers—routers situated at the edges of your networks. This provides a basic buffer from the outside network, or from a less controlled area of your own network into a more sensitive area of your network.

On these routers, you should configure ACLs for each network protocol configured on the router interfaces. You can configure ACLs so that inbound traffic, outbound traffic, or both are filtered on an interface.

ACLs must be defined on a per-protocol basis. In other words, you should define ACLs for every protocol enabled on an interface if you want to control traffic flow for that protocol.

For inbound ACLs, after receiving a packet, the Cisco IOS software checks the source address of the packet against the ACL. If the ACL permits the address, the software continues to process the packet. If the ACL rejects the address, the software discards the packet and returns an Internet Control Message Protocol (ICMP) host unreachable message.

For outbound ACLs, after receiving and routing a packet to a controlled interface, the software checks the source address of the packet against the ACL. If the ACL permits the address, the

software sends the packet. If the ACL rejects the address, the software discards the packet and returns an ICMP host unreachable message.

When you apply an ACL that has not yet been defined to an interface, the software acts as if the ACL has not been applied to the interface and accepts all packets. Remember this behavior if you use undefined ACLs as a means of security in your network.

Implementation—Cisco IOS Firewall Features and Configuration

This topic details the implementation of the Cisco IOS Firewall features and configuration.

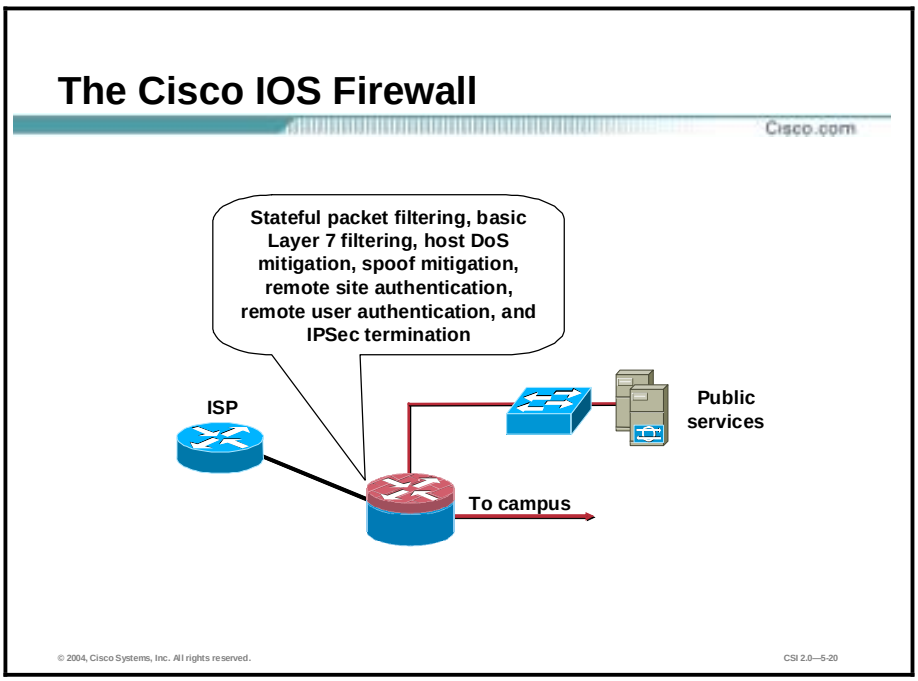

The primary function of the Cisco IOS Firewall is to provide connection-state enforcement and detailed filtering for sessions initiated through the Cisco IOS Firewall. The Cisco IOS Firewall also acts as a termination point for site-to-site IPSec VPN tunnels for both remote site production and remote site management traffic.

There are two segments of the Cisco IOS Firewall. The first segment is the public services segment, which contains all the publicly addressable hosts. The second segment is for remote access VPN and dial-in, which are discussed in a later lesson. Publicly addressable servers have some protection against TCP SYN floods through mechanisms such as the use of half-open connection limits on the Cisco IOS Firewall.

From a filtering standpoint, in addition to limiting traffic on the public services segment to relevant addresses and ports, filtering in the opposite direction also occurs. If an attack compromises one of the public servers (by circumventing the Cisco IOS Firewall and the HIDS or HIPS), that server should not be able to further attack the network. To mitigate this type of attack, specific filtering prevents any unauthorized requests from being generated by the public servers to any other location.

As an example, the web server should be filtered so that it cannot originate requests of its own, but merely respond to requests from clients. This setup helps prevent a hacker from downloading additional utilities to the compromised box after the initial attack. It also helps stop unwanted sessions from being triggered by the hacker during the primary attack. An attack that generates an xterm from the web server through the firewall to the hacker’s machine is an example of such an attack. In addition, private VLANs prevent a compromised public server from attacking other

servers on the same segment. This traffic is not even detected by the Cisco IOS Firewall, a fact that explains why private VLANs are critical.

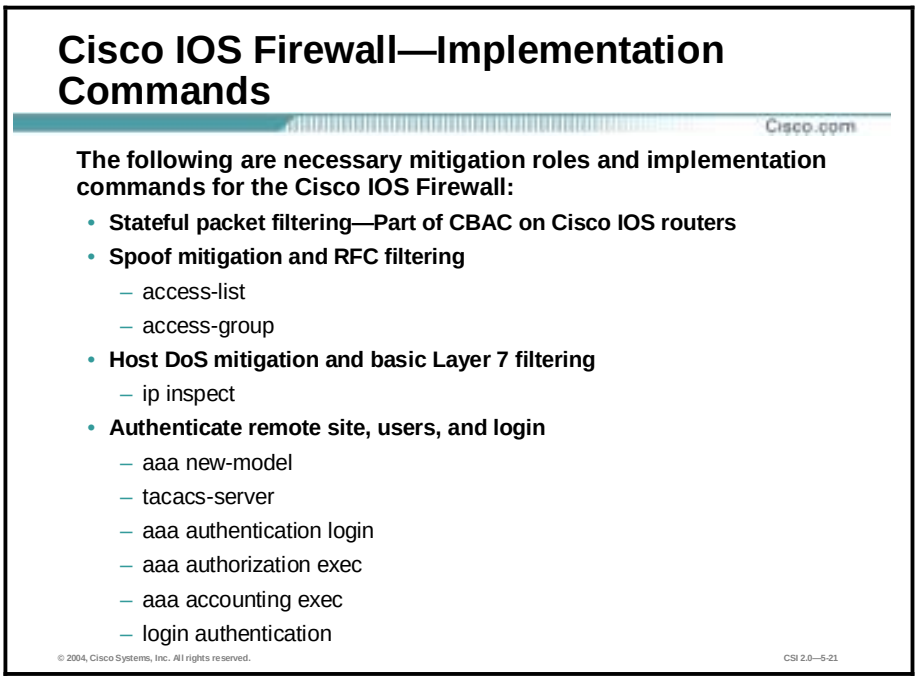

The following are necessary mitigation roles and implementation commands for the Cisco IOS Firewall:

- Stateful packet filtering—Part of Context-Based Access Control (CBAC) on Cisco IOS routers
- Spoof mitigation and RFC filtering
  - **access-list**—Enables you to specify whether an IP address is permitted or denied access to a port or protocol
  - **access-group**—Binds an ACL to an interface
- Host DoS mitigation and basic Layer 7 filtering
  - **ip inspect**—Defines the application protocols to inspect
- Authenticate remote site and login
  - **aaa new-model**—Defines a set of inspection rules for each protocol, using the same inspection-name
  - **tacacs-server**—Defines the TACACS server
  - **aaa authentication login**—Enables authentication, authorization, and accounting (AAA) authentication at login
  - **aaa authorization exec**—Restricts network access to a user
  - **aaa accounting exec**—Runs accounting for EXEC shell session
  - **login authentication**—Specifies the name of a list of AAA authentication methods to try at login

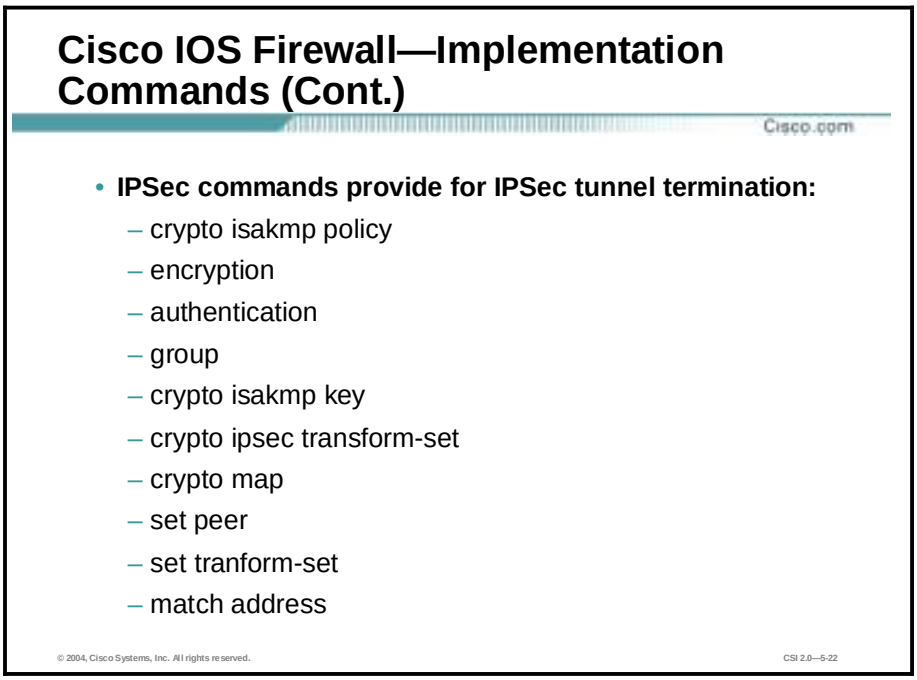

The following are implementation commands for the Cisco IOS Firewall IPSec:

- **crypto isakmp policy**—Specifies the parameters to be used during an Internet Key Exchange (IKE) negotiation
- **encryption**—Sets the algorithm to be negotiated
- **authentication**—Specifies the authentication method within an IKE policy
- **group**—Specifies the Diffie-Hellman (DH) group identifier within an IKE policy
- **crypto isakmp key**—Configures pre-shared authentication keys
- **crypto ipsec transform-set**—An acceptable combination of security protocols, algorithms, and other settings to apply to IP Security-protected traffic
- **crypto map**—Configures filtering, classifies traffic to be protected, and defines the policy to be applied to that traffic
- **set peer**—Specifies an IPSec peer for a crypto map
- **set transform-set**—Specifies which transform sets to include in a crypto map entry
- **match address**—Specifies an extended access list for a crypto map entry

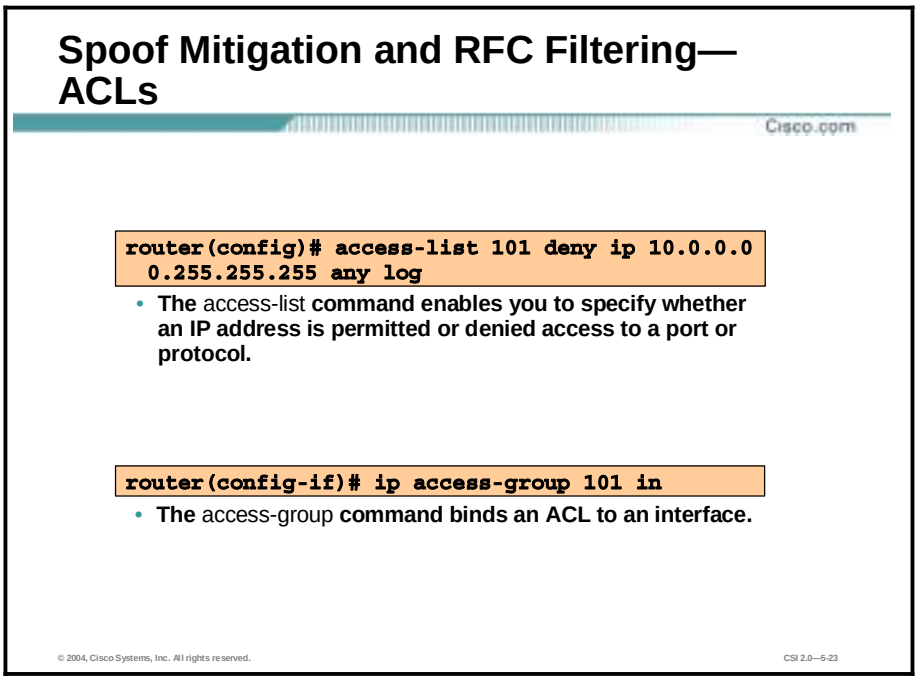

You can use ACLs to control the transmission of packets on an interface, control virtual teletype (VTY) access, and restrict the contents of routing updates. The Cisco IOS software stops checking the extended ACL after a match occurs.

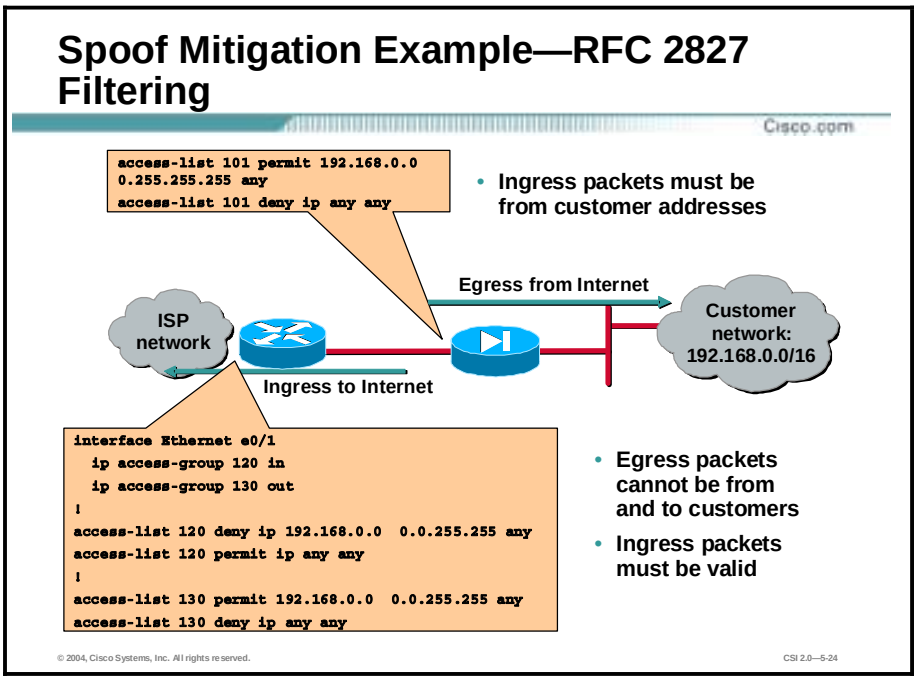

A resurgence of DoS attacks aimed at various targets in the Internet have produced new challenges within the ISP and network security communities to find new and innovative methods to mitigate these types of attacks. While the filtering method in RFC 2827 does absolutely nothing to protect against flooding attacks which originate from valid prefixes (IP addresses), it will prohibit an attacker within the originating network from launching an attack of this nature using forged source addresses that do not conform to ingress filtering rules.

All providers of Internet connectivity are urged to implement filtering described in RFC 2827 to prohibit attackers from using forged source addresses which do not reside within a range of legitimately advertised prefixes. In other words, if an ISP is aggregating routing announcements for multiple downstream networks, strict traffic filtering should be used to prohibit traffic that claims to have originated from outside of these aggregated announcements.

An additional benefit of implementing this type of filtering is that it enables the originator to be easily traced to its true source, because the attacker would have to use a valid, and legitimately reachable, source address.

Unauthorized Access—Cisco IOS Firewall Intrusion Detection

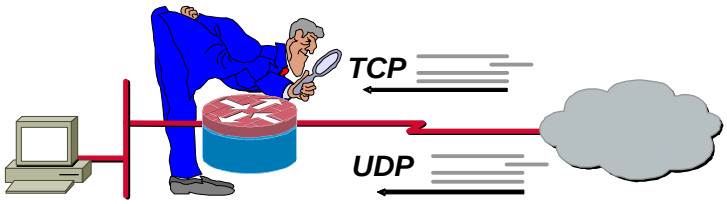

The following are the Cisco IOS Firewall intrusion detection features:

- **Acts as an in-line intrusion detection sensor**
- **When a packet or packets match a signature, it can perform any of the following configurable actions:**
  - **Alarm—Sends an alarm to a Cisco IDS Director or Syslog server**
  - **Drop—Drops the packet**
  - **Reset—Sends TCP resets to terminate the session**
- **Detects, reports, and acts upon many common attacks**

Cisco IOS Firewall and intrusion detection solutions are designed to meet security requirements within a network, whether the requirement is for multilevel security with a dedicated appliance (Cisco PIX Firewall or Cisco IDS Sensor) or for integrated security wherever a router is deployed (Cisco IOS Firewall with intrusion detection). The standalone appliance and integrated Cisco solutions meet the various network security needs in a network, based on an organization's security policy, network security risk, vulnerability, level of performance requirements at the site, network segmentation, network media or interface, and routing requirements.

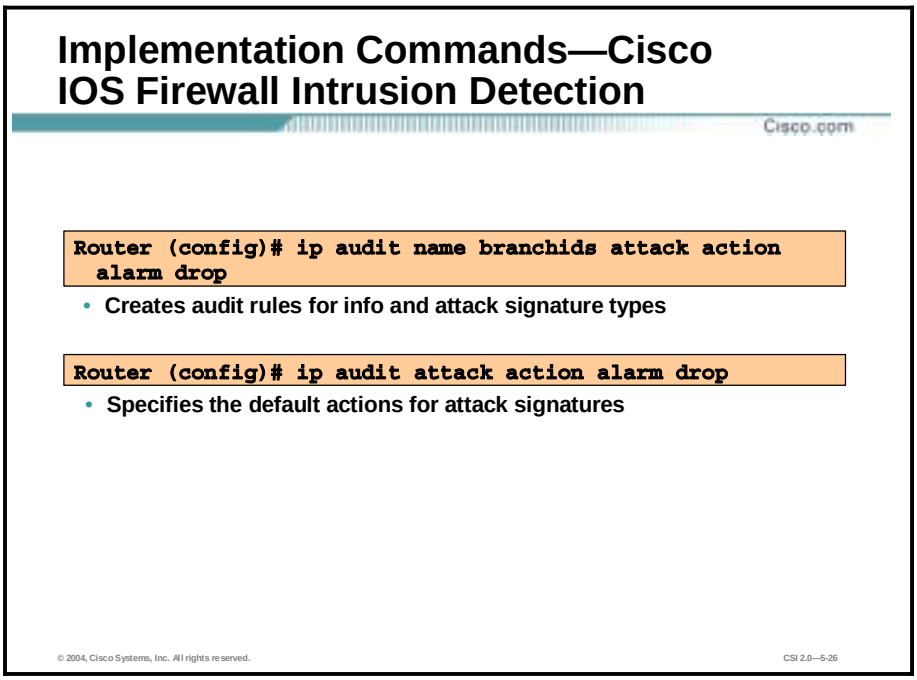

Use the **ip audit name** global configuration command to create audit rules for info and attack signature types. Use the **no** form of this command to delete an audit rule.

Use the **ip audit attack** global configuration command to specify the default actions for attack signatures. Use the **no** form of this command to set the default action for attack signatures.

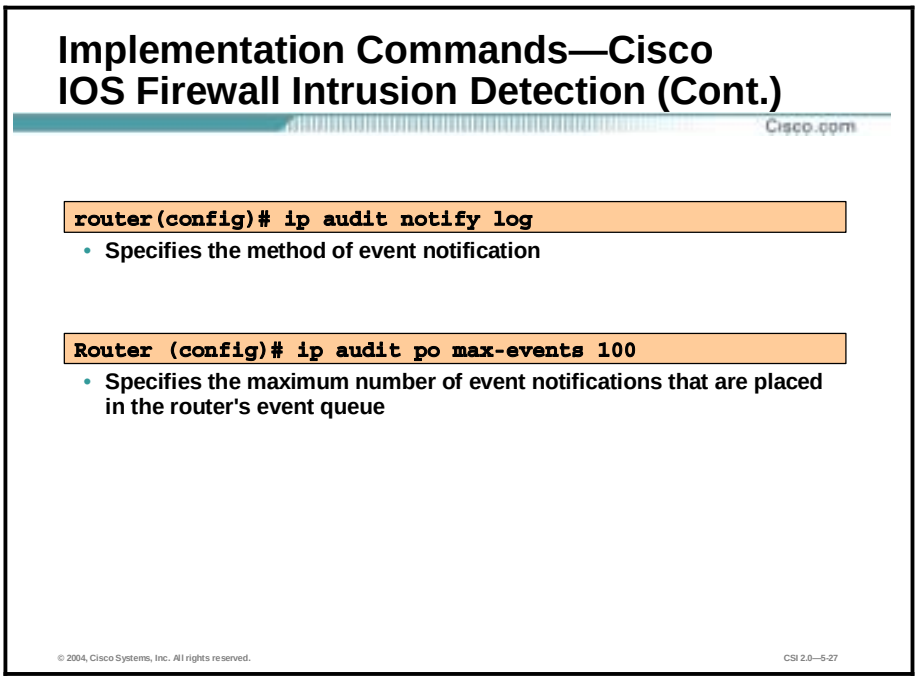

Use the **ip audit notify** global configuration command to specify the method of event notification. Use the **no** form of this command to disable event notifications.

Use the **ip audit po max-events** global configuration command to specify the maximum number of event notifications that are placed in the router's event queue. Use the **no** version of this command to set the number of recipients to the default setting.

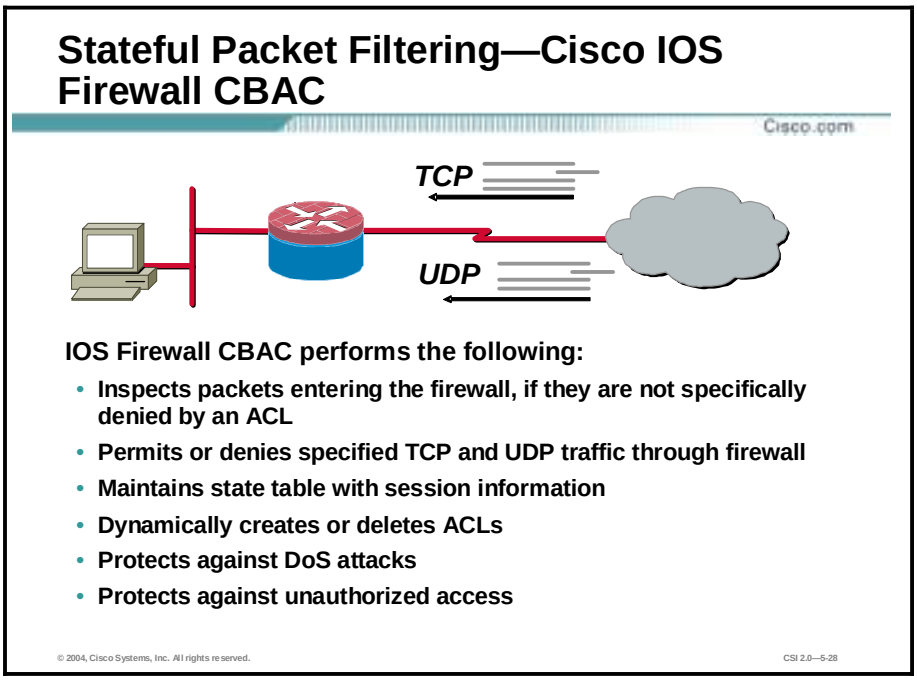

CBAC adds inspection intelligence to ACL capabilities by performing stateful packet inspection for application status information for applications using TCP and UDP packets for transport. Using this information, CBAC creates a temporary or dynamic, session-specific ACL entry, permitting return traffic into the trusted network. This dynamic ACL effectively opens a door in the Cisco IOS Firewall. When a session times out or ends, the ACL entry is deleted and the door closes to additional traffic. To perform this function, a state table is maintained for session information.

Standard and extended ACLs cannot create temporary ACL entries, so, until now, administrators have been forced to weigh security risks against information access requirements. Advanced applications that select from multiple channels for return traffic have been difficult to secure using standard or extended ACLs.

CBAC is more secure than current ACL-only solutions because it accounts for application type in deciding whether to allow a session through the IOS Firewall, and determines whether it selects from multiple channels for return traffic. Before CBAC, administrators could permit advanced application traffic only by writing permanent ACLs that essentially left Cisco IOS Firewall doors open; so most administrators opted to deny all such application traffic. With CBAC, they can now securely permit multimedia and other application traffic by opening the Cisco IOS Firewall as needed, and closing it at all other times. For example, if CBAC is configured to allow Microsoft NetMeeting, when an internal user initiates a connection, the Cisco IOS Firewall permits return traffic. However, if an external NetMeeting source initiates a connection with an internal user, CBAC denies entry and drops the packets.

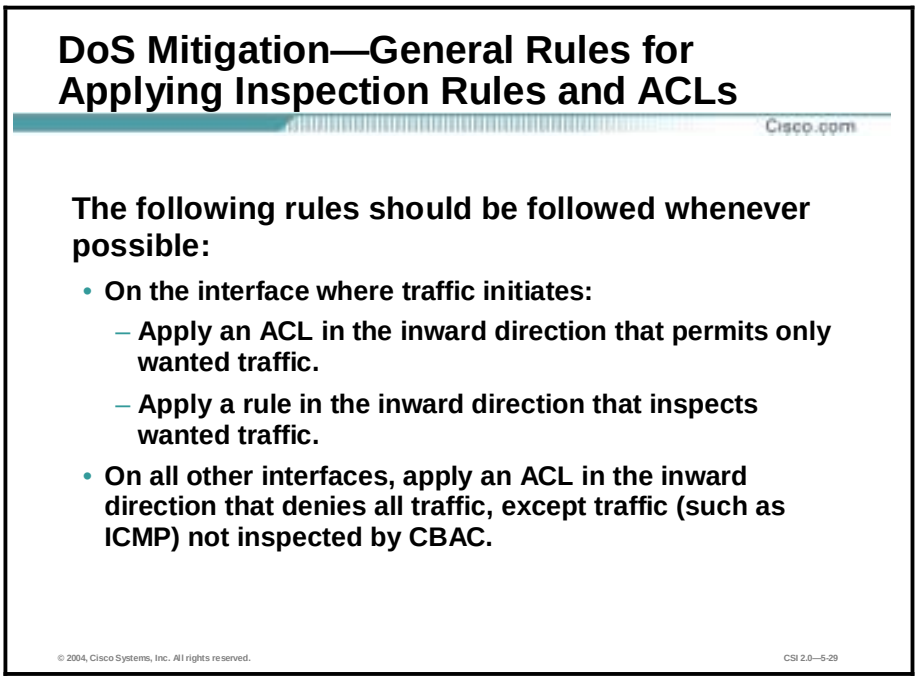

On the Cisco IOS Firewall interface where traffic initiates, ACLs should be applied on the inward direction that only permit wanted traffic, and on the inward direction that inspects wanted traffic. On all other interfaces, the ACL should be applied on the inward direction that denies all traffic, except traffic (such as ICMP) not inspected by CBAC.

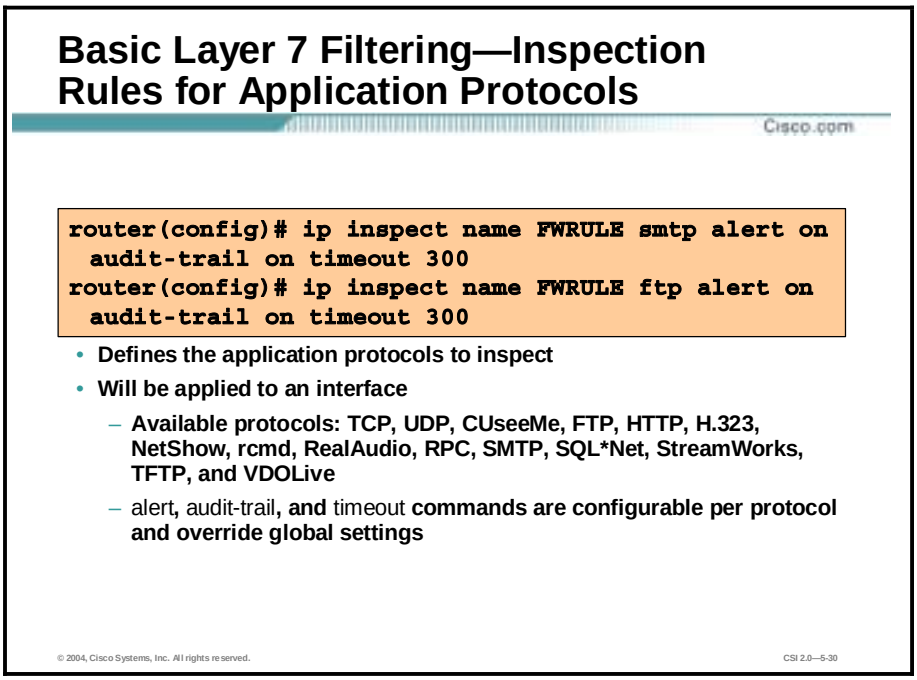

You can configure TCP and UDP inspections to permit TCP and UDP packets to enter the internal network through the firewall, even if the application-layer protocol is not configured to be inspected. However, TCP and UDP inspections do not recognize application-specific commands, and therefore might not permit all return packets for an application, particularly if the return packets have a different port number than the previous exiting packet.

Any application-layer protocol that is inspected will take precedence over the TCP or UDP packet inspection. For example, if inspection is configured for FTP, all control channel information will be recorded in the state table, and all FTP traffic will be permitted back through the firewall if the control channel information is valid for the state of the FTP session. The fact that TCP inspection is configured is irrelevant to the FTP state information.

With TCP and UDP inspections, packets entering the network must exactly match the corresponding packet that previously exited the network. The entering packets must have the same source or destination addresses, and source or destination port numbers as the exiting packet (but reversed); otherwise, the entering packets will be blocked at the interface. Also, all TCP packets with a sequence number outside of the window are dropped.

With UDP inspection configured, replies will only be permitted back in through the firewall if they are received within a configurable time after the last request was sent out.

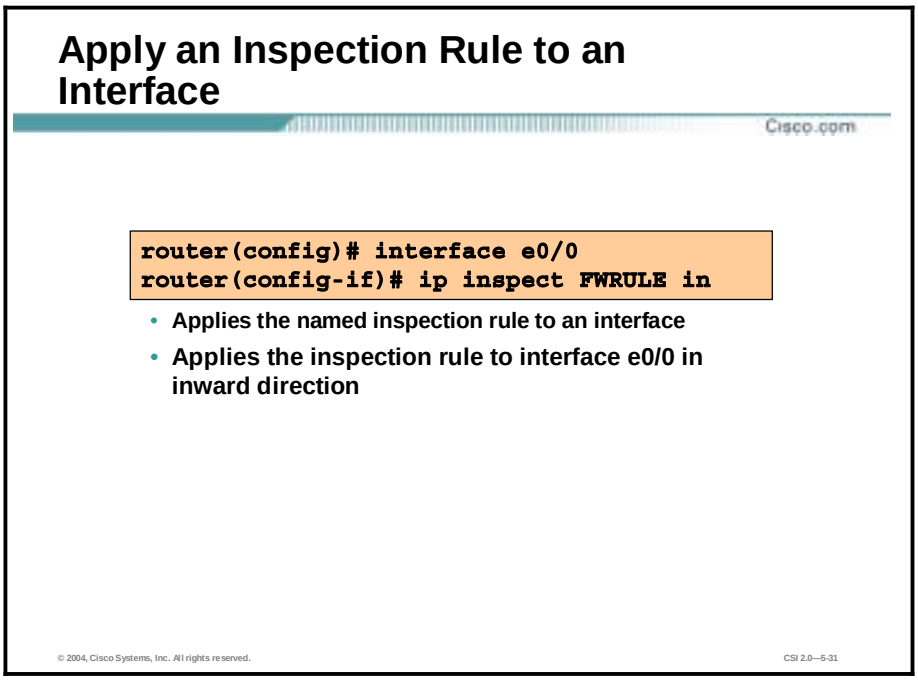

After you define an inspection rule, you apply it to an interface.

Usually, you apply only one inspection rule to one interface. The only exception might occur if you want to enable CBAC in two directions. For CBAC configured in both directions at a single firewall interface, you should apply two rules, one for each direction, as follows:

- If you are configuring CBAC on an external interface, apply the rule to outbound traffic.
- If you are configuring CBAC on an internal interface, apply the rule to inbound traffic.

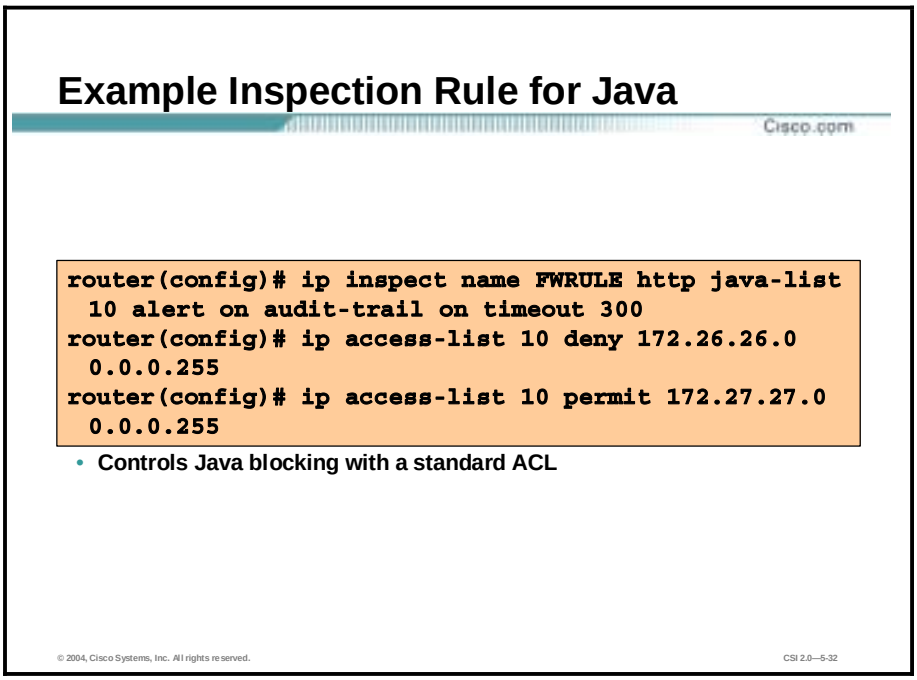

Java applet filtering distinguishes between trusted and untrusted applets by relying on a list of external sites that you designate as “friendly.” If an applet is from a friendly site, the firewall allows the applet through. If the applet is not from a friendly site, the applet is blocked. (Alternately, you could permit applets from all external sites except those you specifically designate as “hostile.”)

Note	CBAC does not detect or block encapsulated Java applets. Therefore, Java applets that are wrapped or encapsulated, such as applets in .zip or .jar format, are not blocked at the firewall. CBAC also does not detect or block applets loaded from FTP, gopher, HTTP on a nonstandard port, and so forth.

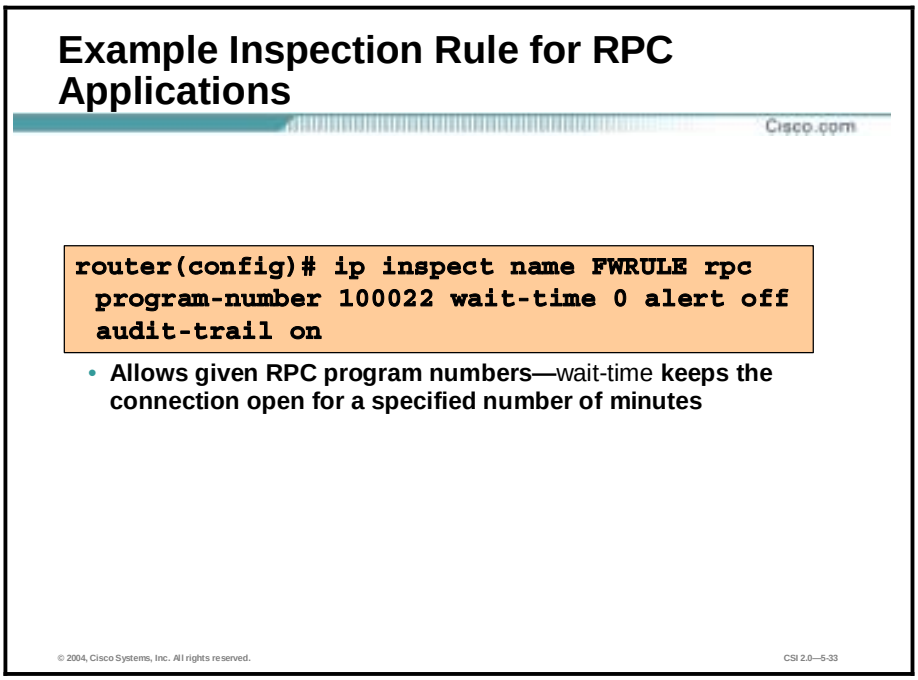

Remote Procedure Call (RPC) inspection enables the specification of various program numbers. You can define multiple program numbers by creating multiple entries for RPC inspection, each with a different program number. If a program number is specified, all traffic for that program number is permitted. If a program number is not specified, all traffic for that program number is blocked. For example, if you created an RPC entry with the Network File System (NFS) program number, all NFS traffic is allowed through the firewall.

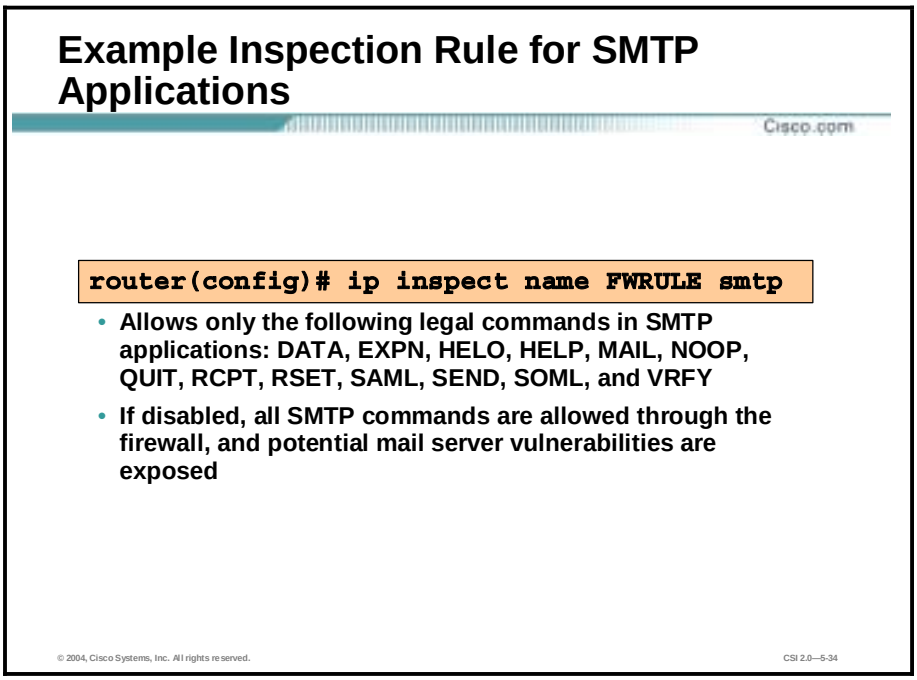

SMTP inspection causes SMTP commands to be inspected for illegal commands. Any packets with illegal commands are dropped, and the SMTP session hangs and eventually times out.

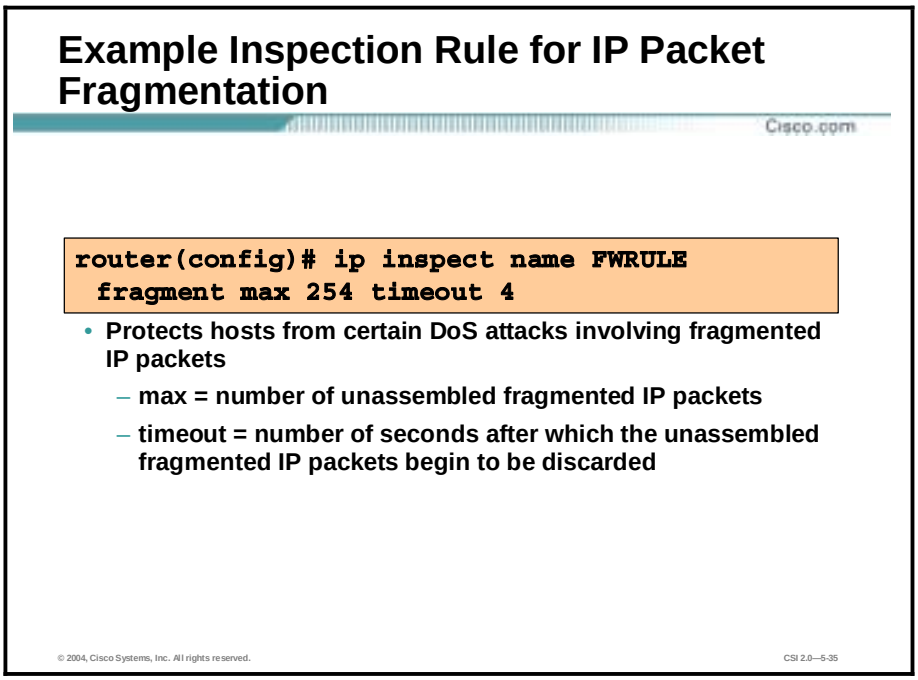

CBAC inspection rules can help protect hosts against certain DoS attacks involving fragmented IP packets. Even though the firewall keeps an attacker from making actual connections to a given host, the attacker may still be able to disrupt services provided by that host. This is done by sending many non-initial IP fragments or by sending complete fragmented packets through a router with an ACL that filters the first fragment of a fragmented packet. These fragments can tie up resources on the target host as it tries to reassemble the incomplete packets.

Using fragmentation inspection, the firewall maintains an interfragment state (structure) for IP traffic. Non-initial fragments are discarded unless the corresponding initial fragment was permitted to pass through the firewall. Non-initial fragments received before the corresponding initial fragments are discarded.

Note	Fragmentation inspection can have undesirable effects in certain cases, because it can result in the firewall discarding any packet whose fragments arrive out of order. There are many circumstances that can cause out-of-order delivery of legitimate fragments. Apply fragmentation inspection in situations where legitimate fragments, which are likely to arrive out of order, might have a severe performance impact.

Because routers running Cisco IOS software are used in a very large variety of networks, and because the CBAC feature is often used to isolate parts of internal networks from one another, the fragmentation inspection feature is not enabled by default. Fragmentation detection must be explicitly enabled for an inspection rule using the **ip inspect name** command. Unfragmented traffic is never discarded because it lacks a fragment state. Even when the system is under heavy attack with fragmented packets, legitimate fragmented traffic, if any, still gets some fraction of the firewall's fragment state resources, and legitimate, unfragmented traffic can flow through the firewall unimpeded.

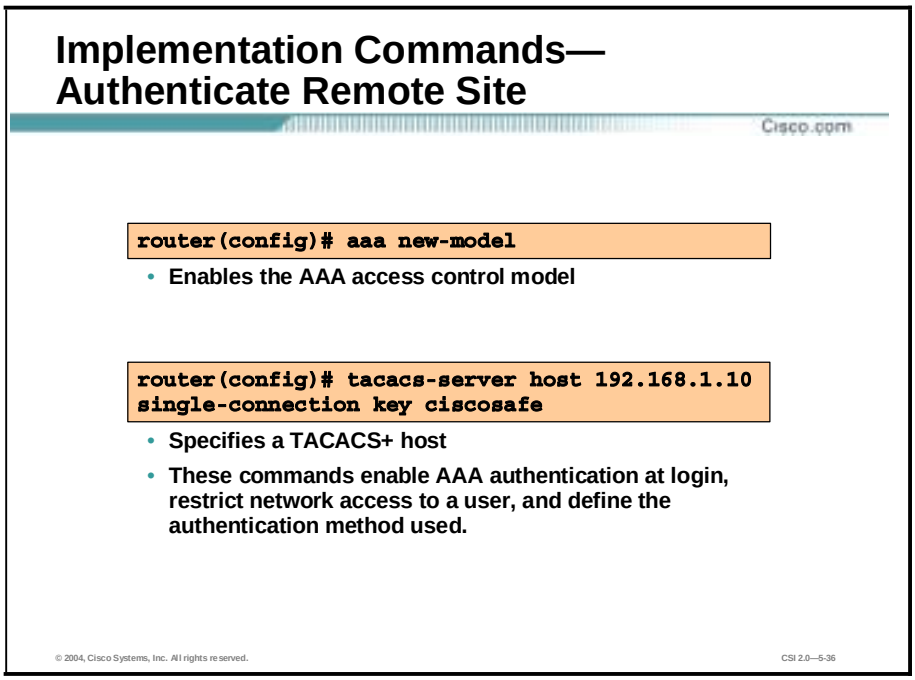

You must configure AAA network security services using the **aaa new-model**, **aaa authentication**, **aaa authorization**, and **aaa accounting** global configuration commands.

You also need to configure your network access server to communicate with the applicable security server, either an extended TACACS or RADIUS daemon.

Complete these to configure AAA authentication:

Step 1 Enable AAA by using the **aaa new-model** global configuration command.

Step 2 Configure security protocol parameters, such as RADIUS, TACACS+, or Kerberos if you are using a security server.

Define the method lists for authentication by using an **aaa authentication** command. A method list is a named list describing the authorization methods to be used (such as RADIUS or TACACS+), in sequence. Method lists enable you to designate one or more security protocols to be used for authorization, thus ensuring a backup system in case the initial method fails.

Step 3 Apply the method lists to a particular interface or line, if required.

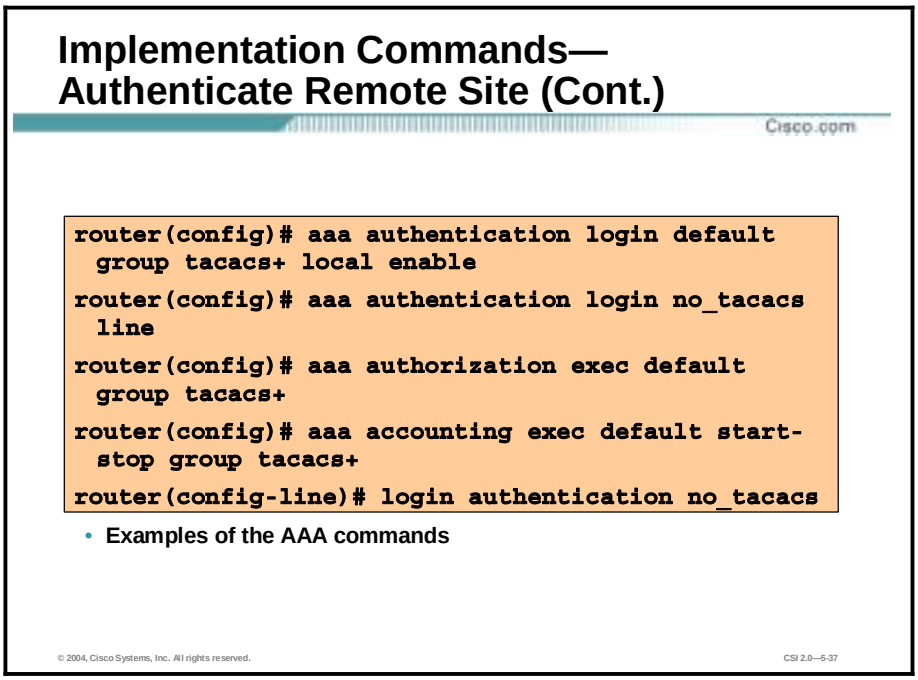

Create a method list by entering the **aaa authentication login *list-name method*** command for a particular protocol, where list-name is any character string used to name this method list (for example, MIS-access). The method argument identifies the list of methods that the authentication algorithm tries in the given sequence.

Use the **login authentication** command with the default argument followed by the methods you want to use in default situations to create a default method list that is used if no method list is assigned to a line.

The additional methods of authentication are used only if the previous method returns an error, not if it fails. To ensure that the authentication succeeds even if all methods return an error, specify **none** as the final method in the command line.

If authentication is not specifically set for a line, the default is to deny access and no authentication is performed. Use the **more system:running-config** command to display currently configured lists of authentication methods.

Use the **aaa authorization** command to enable authorization and to create named method lists, defining authorization methods that can be used when a user accesses the specified function. Method lists for authorization define the ways authorization is performed and the sequence in which these methods are performed.

Cisco IOS software uses the first method listed in your method list to authorize users for specific network services; if that method fails to respond, the Cisco IOS software selects the next method listed in the method list. This process continues until there is successful communication with a listed authorization method, or all methods defined are exhausted.

The command **login authentication** is a per-line command used with AAA that specifies the name of a list of AAA authentication methods to try at login. If no list is specified, the default list is used (whether or not it is specified in the command line).

Authentication—Cisco IOS Firewall Authentication Proxy

Cisco.com

The Cisco IOS Firewall authentication proxy provides the following:

- **HTTP-based authentication**
- **Dynamic, per-user authentication and authorization via TACACS+ and RADIUS protocols**

CSI 2.0—5-38

The HTTP-based authentication proxy capability in the Cisco IOS Firewall feature set provides dynamic, per-user authentication and authorization for network users. Previously, user identity and related authorized access were determined by a user IP address, or the security policy had to be applied to a user group or subnet. Now, a per-user policy can be downloaded dynamically to the router from the TACACS+ or RADIUS authentication server using Cisco IOS software authentication, authorization, and accounting (AAA) services.

Implementation—PIX Firewall

This topic describes the detailed configuration and implementation of the Cisco PIX Firewall.

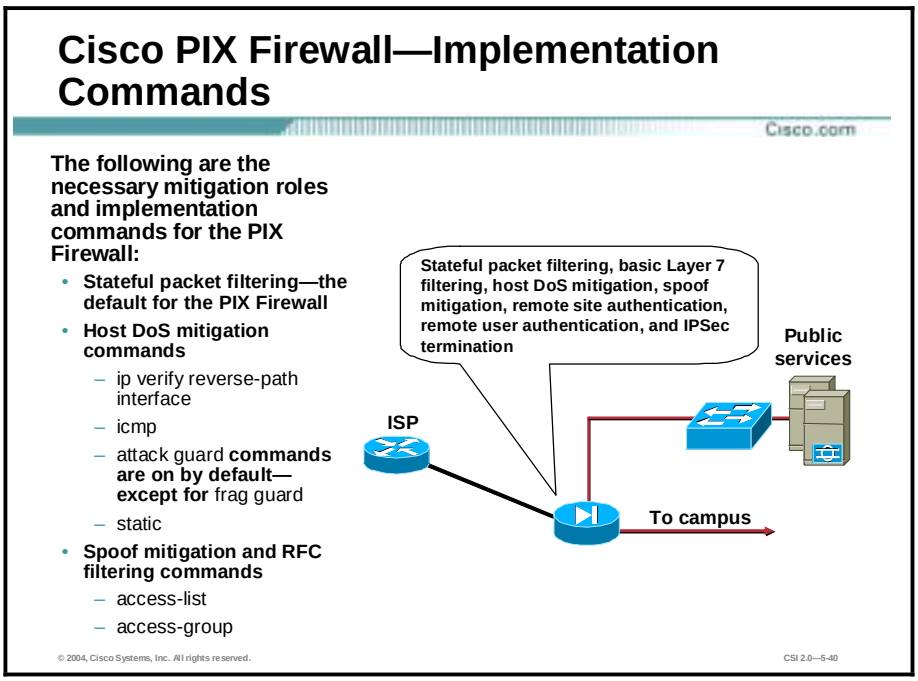

The following mitigation roles and commands are used to implement the policy on the PIX Firewall:

- Stateful packet filtering—The default for the PIX firewall
- Host DoS mitigation commands
  - **ip verify reverse-path interface**—Implements Unicast Reverse Path Forwarding (RPF) IP spoofing protection
  - **icmp**—Enables or disables pinging to an interface
  - **attack guard**—Enabled by default
  - **static**—Used to set maximum connections and maximum embryonic connections
- Spoof mitigation and RFC filtering commands
  - **access-list**—Creates an ACL
  - **access-group**—Binds the ACL to an interface

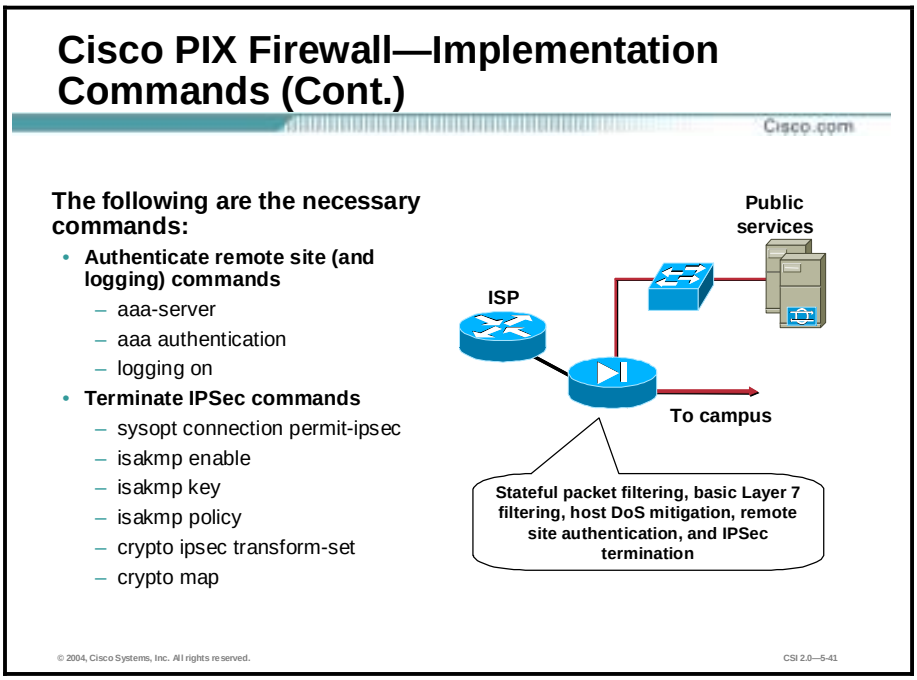

The following are necessary implementation commands for the Cisco PIX Firewall:

- Authenticate remote site and login commands
  - **aaa-server**—Specifies an AAA server
  - **aaa authentication**—Enables, disables, or views LOCAL, TACACS+, or RADIUS user authentication
  - **logging on**—Enables or disables Syslog and SNMP logging
- Terminate IPSec commands:
  - **sysopt connection permit-ipsec**—Implicitly permits any packet that came from an IPSec tunnel
  - **isakmp enable**—Enables Internet Security Association Key Management Protocol (ISAKMP) negotiation on the interface on which the IPSec peer communicates with the PIX Firewall
  - **isakmp key**—Specifies the authentication pre-shared key
  - **isakmp policy**—Uniquely identifies the IKE policy and assigns a priority to the policy
  - **crypto ipsec transform-set**—Creates, views, or deletes IPSec Security Associations (SAs), SA global lifetime values, and global transform sets
  - **crypto map**—Creates, modifies, views, or deletes a crypto map entry

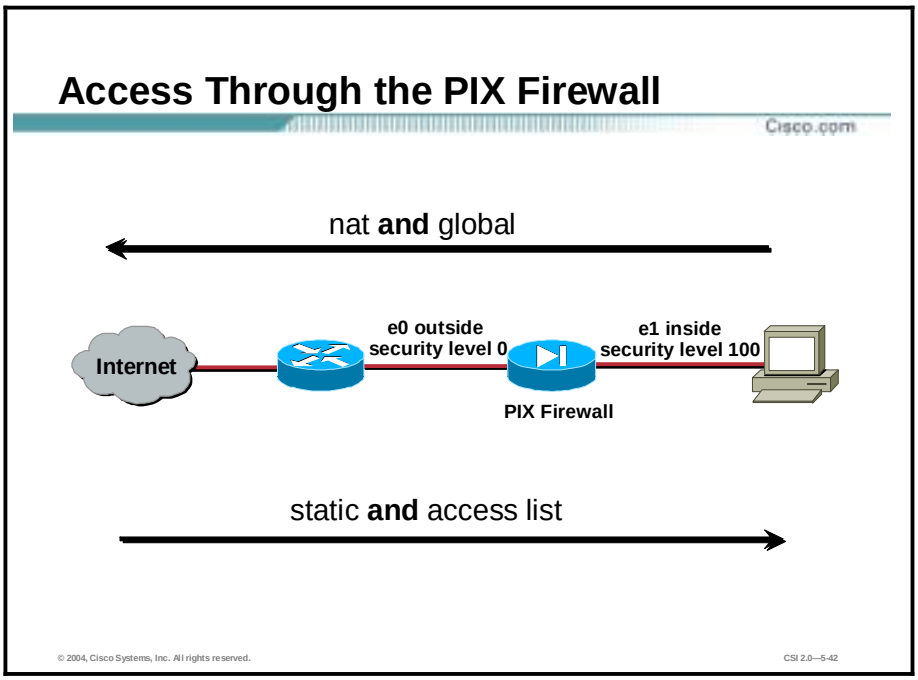

By default, the PIX Firewall denies access to an internal or perimeter (more secure) network from an external (less secure) network. You specifically allow inbound connections by using ACLs. ACLs permit access on a first-match basis. For inbound access, you must deny access first and then permit access.

Static Network Address Translation (NAT) creates a permanent, one-to-one mapping between an address on an internal network (a higher security level interface) and a perimeter or external network (lower security level interface). For example, to share a web server on a perimeter interface with users on the public Internet, use static address translation to map the server's actual address to a registered IP address. Static address translation hides the actual address of the server from users on the less secure interface, making casual access by unauthorized users less likely. Unlike NAT or Port Address Translation (PAT), it requires a dedicated address on the outside network for each host, so it does not save registered IP addresses.

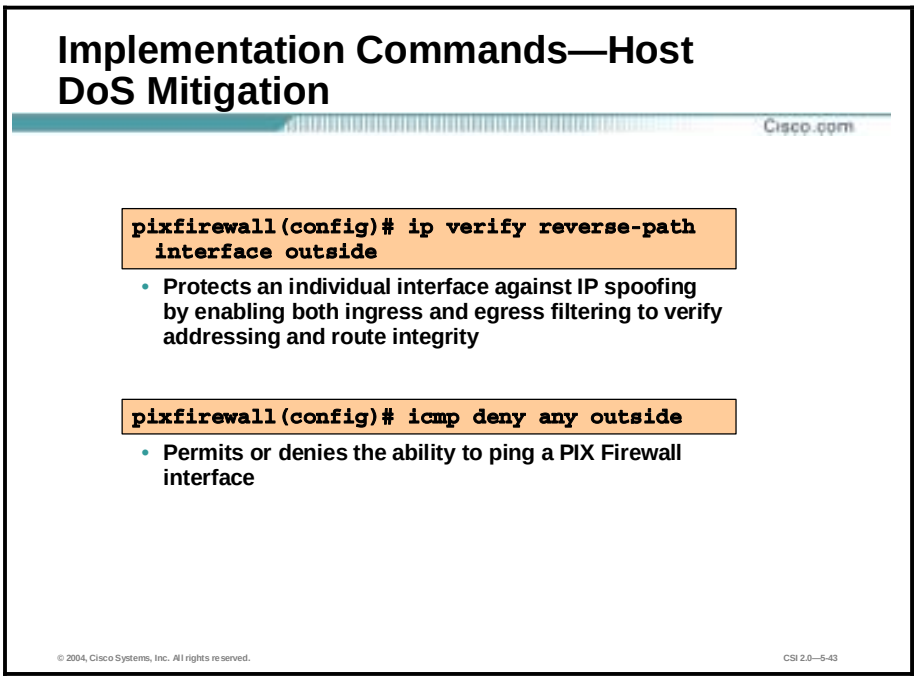

The **ip verify reverse-path** command implements the following:

- Performs a route lookup based on the source address. Usually, the route lookup is based on the destination address. This is why it is called “reverse path forwarding.” With this command enabled, packets are dropped if there is no route found for the packet or the route found does not match the interface on which the packet arrived.
- Specifies which interfaces to protect from an IP spoofing attack using network ingress and egress filtering, which is described in RFC 2267. This command is disabled by default and provides Unicast Reverse Path Forwarding (Unicast RPF) functionality for the PIX Firewall.
- Provides ingress filtering, which checks inbound packets for IP source address integrity, and is limited to addresses for networks in the enforcing entity’s local routing table. If the incoming packet does not have a source address represented by a route, then it is impossible to know whether the packet has arrived on the best possible path back to its origin. This is often the case when routing entities cannot maintain routes for every network.
- Provides egress filtering, which verifies that packets destined for hosts outside the managed domain have IP source addresses verifiable by routes in the enforcing entity’s local routing table. If an exiting packet does not arrive on the best return path back to the originator, then the packet is dropped and the activity is logged. Egress filtering prevents internal users from launching attacks using IP source addresses outside of the local domain, because most attacks use IP spoofing to hide the identity of the attacking host. Egress filtering makes the task of tracing the origin of an attack much easier. When employed, egress filtering enforces what IP source addresses are obtained from a valid pool of network addresses. Addresses are kept local to the enforcing entity and are therefore easily traceable.

The **clear ip verify** command removes **ip verify** commands from the configuration. Unicast RPF is a unidirectional input function that screens inbound packets arriving on an interface. Outbound packets are not screened.

Because of the danger of IP spoofing in the IP protocol, measures need to be taken to reduce this risk when possible. Unicast RPF, or reverse route lookup, prevents such manipulation under certain circumstances.

Unicast RPF is implemented as follows:

- ICMP packets have no session, so each packet is checked.
- UDP and TCP have sessions, so the initial packet requires a reverse route lookup. Subsequent packets arriving during the session are checked using an existing state maintained as part of the session. Non-initial packets are checked to ensure they arrived on the same interface used by the initial packet.

Note	Before using the **ip verify** command, add **static route** command statements for every network that can be accessed on the interfaces you wish to protect. Only enable **ip verify** if routing is fully specified. Otherwise, the PIX Firewall will stop traffic on the interface you specify if routing is not in place.

Use the **show interface** command to view the number of dropped packets, which appears in the "unicast rpf drops" counter.

The **icmp** command implements the following:

- The **icmp** command controls ICMP traffic that terminates on the PIX Firewall. If no ICMP control list is configured, then the PIX Firewall accepts all ICMP traffic that terminates at the interface.
- The **icmp deny** command disables pinging to an interface, and the **icmp permit** command enables pinging to an interface. With pinging disabled, the PIX Firewall cannot be detected on the network. This is also referred to as configurable proxy pinging.

For traffic that is routed through the PIX Firewall only, you can use the **access-list** or **access-group** commands to control the ICMP traffic routed through the PIX Firewall.

It is recommended that you grant permission for ICMP unreachable message type (type 3). Denying ICMP unreachable messages disables ICMP Path maximum transmission unit (MTU) discovery, which can halt IPSec and Point-to-Point Tunneling Protocol (PPTP) traffic. See RFC 1195 and RFC 1435 for details about Path MTU Discovery.

If an ICMP control list is configured, then the PIX Firewall uses a first match to the ICMP traffic followed by an implicit deny all. That is, if the first matched entry is a permit entry, the ICMP packet continues to be processed. If the first matched entry is a deny entry or an entry is not matched, the PIX Firewall discards the ICMP packet and generates the %PIX-3-313001 Syslog message. An exception is when an ICMP control list is not configured; in that case, a permit is assumed.

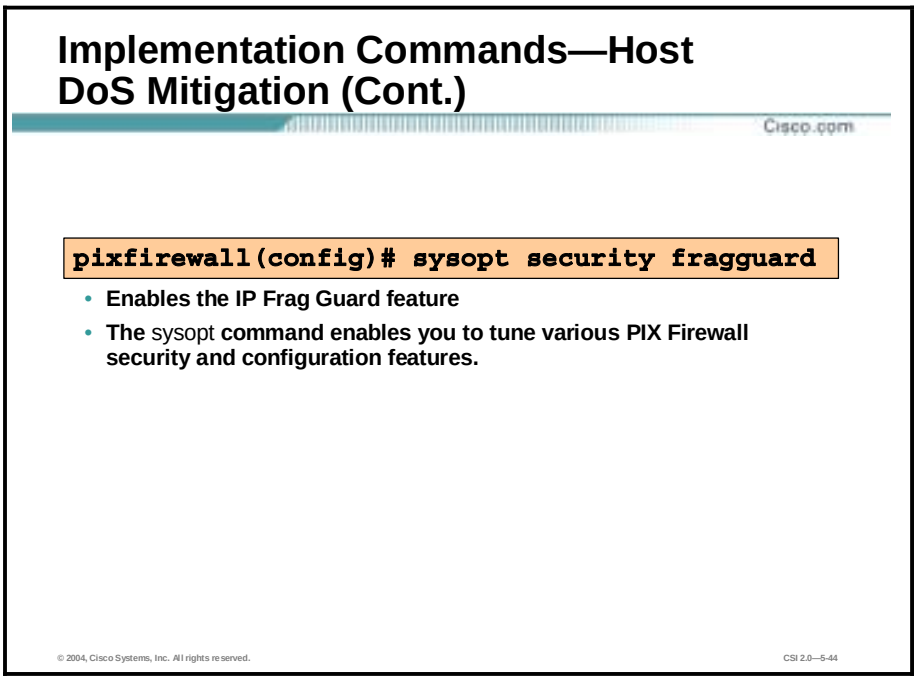

The **sysopt security fragguard** command enables the IP Frag Guard feature. This feature is disabled by default. The IP Frag Guard feature enforces two security checks in addition to the security checks recommended by RFC 1858, against the many IP fragment style attacks: teardrop, land, and so on. Each non-initial IP fragment is required to be associated with an already-seen valid initial IP fragment. IP fragments are rated to 100 full IP fragmented packets per second to each internal host.

The IP Frag Guard feature operates on all interfaces in the PIX Firewall and cannot be selectively enabled or disabled by interface.

The PIX Firewall uses the **security fragguard** command to enforce the security policy determined by an **access-list permit** or **access-list deny** command to permit or deny packets through the PIX Firewall.

Note	Use of the **sysopt security fragguard** command breaks normal IP fragmentation conventions. However, not using this command exposes the PIX Firewall to the possibility of IP fragmentation attacks. It is recommended that packet fragmentation not be permitted on the network if at all possible.

Disable the **security fragguard** command feature if the PIX Firewall is used as a tunnel for Fiber Distributed Data Interface (FDDI) packets between routers.

Note	Because Linux sends IP fragments in reverse order, fragmented Linux packets will not pass through the PIX Firewall with the **sysopt security fragguard** command enabled.

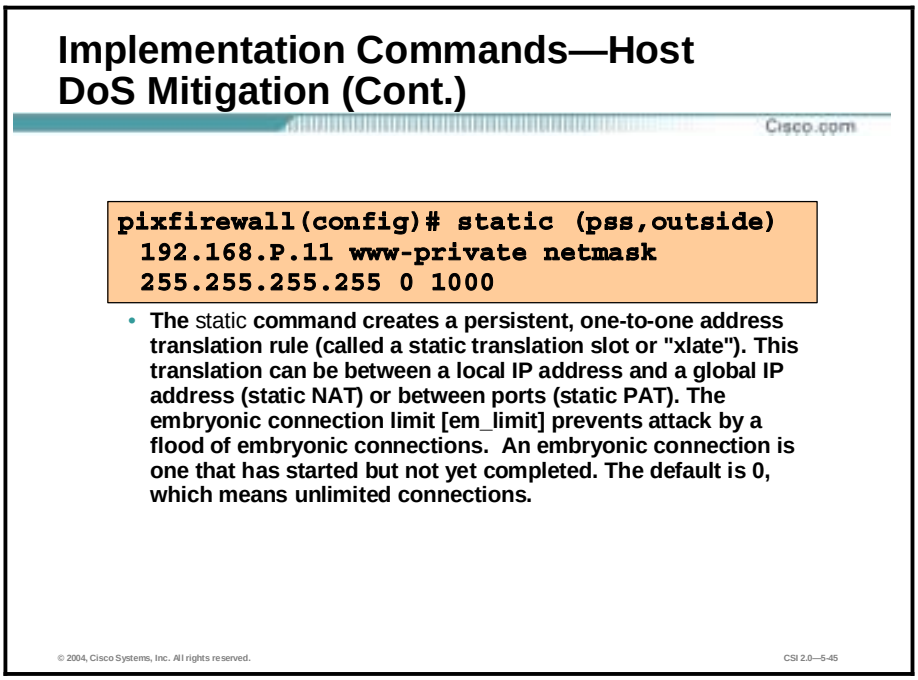

The **static** command creates a persistent, one-to-one address translation rule (called a static translation slot or “xlate”). This translation can be between a local IP address and a global IP address (static NAT) or between ports (static PAT).

For an external host to initiate traffic to an inside host, a static translation rule needs to exist for the inside host. This can also be done using a **nat 0 access-list** address translation rule. Without the persistent translation rule, the translation cannot occur.

You can use the static and access-list commands when you are accessing the interface of a higher security level from an interface of a lower security level (for example, when accessing the inside from a perimeter or the outside interface).

Prior to version 5.3, the PIX Firewall offered no mechanism to protect systems reachable via a static and TCP conduit from TCP SYN attacks. Previously, if an embryonic connection limit was configured in a static command statement, the PIX Firewall dropped new connection attempts after the embryonic threshold was reached. Given this, a modest attack could stop an institution’s web traffic. For static command statements without an embryonic connection limit, the PIX Firewall passes all traffic. If the affected system does not have TCP SYN attack protection (and most operating systems do not offer sufficient protection), then the affected system’s embryonic connection table overloads and all traffic stops.

With the new TCP intercept feature, after the optional embryonic connection limit is reached, and until the embryonic connection count falls below this threshold, every SYN bound for the affected server is intercepted. For each SYN, the PIX Firewall responds on behalf of the server with an empty SYN/ACK segment. The PIX Firewall retains pertinent state information, drops the packet, and waits for the client’s acknowledgement. If the ACK is received, then a copy of the client’s SYN segment is sent to the server and the TCP three-way handshake is performed between the PIX Firewall and the server. If this three-way handshake completes, the connection resumes as normal. If the client does not respond during any part of the connection phase, the PIX Firewall retransmits the necessary segment using exponential back-offs.

This feature requires no change to the PIX Firewall command set, only that the embryonic connection limit on the static command now has a new behavior.

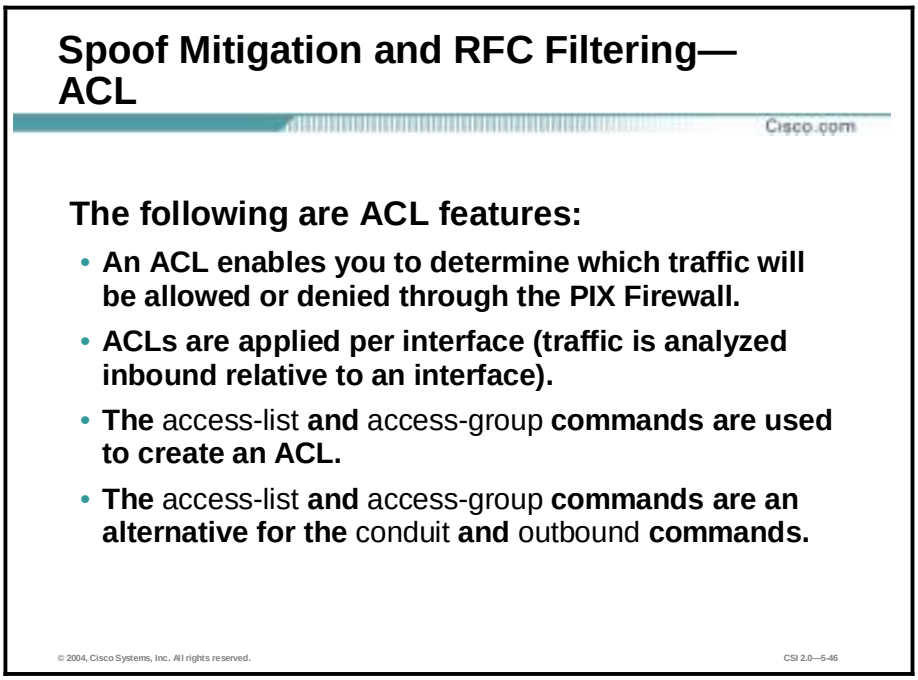

The **access-list** command allows any outside host access via a specified port. You use the **access-list** and **access-group** commands to permit access based on source or destination IP address, or by the protocol port number. Use the **access-list** command to create a single ACL entry, and use the **access-group** command to bind one or more ACL entries to a specific interface. Only specify one **access-group** command for each interface.

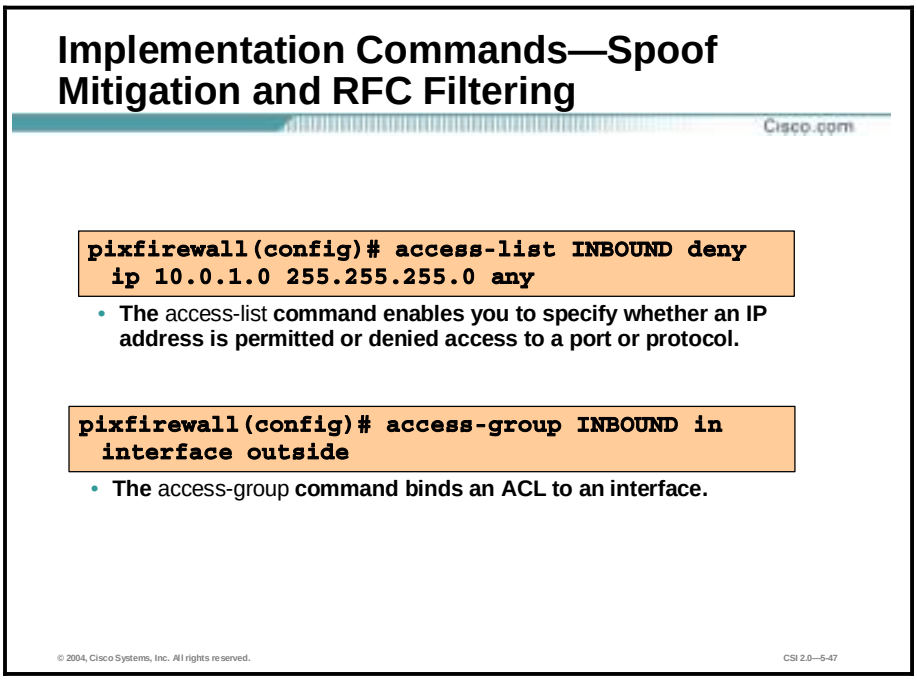

The **access-list** command enables you to specify if an IP address is permitted or denied access to a port or protocol. One or more **access-list** command statements with the same name are referred to as an "ACL." ACLs associated with IPSec are known as "crypto ACLs."

By default, all **access-list** commands have an implicit deny unless you explicitly specify permit. In other words, by default, all access in an ACL is denied unless you explicitly grant access using a **permit** statement.

The **access-group** command binds an ACL to an interface. The ACL is applied to traffic inbound to an interface. If you enter the **permit** option in an **access-list** command statement, the PIX Firewall continues to process the packet. If you enter the **deny** option in an **access-list** command statement, the PIX Firewall discards the packet.

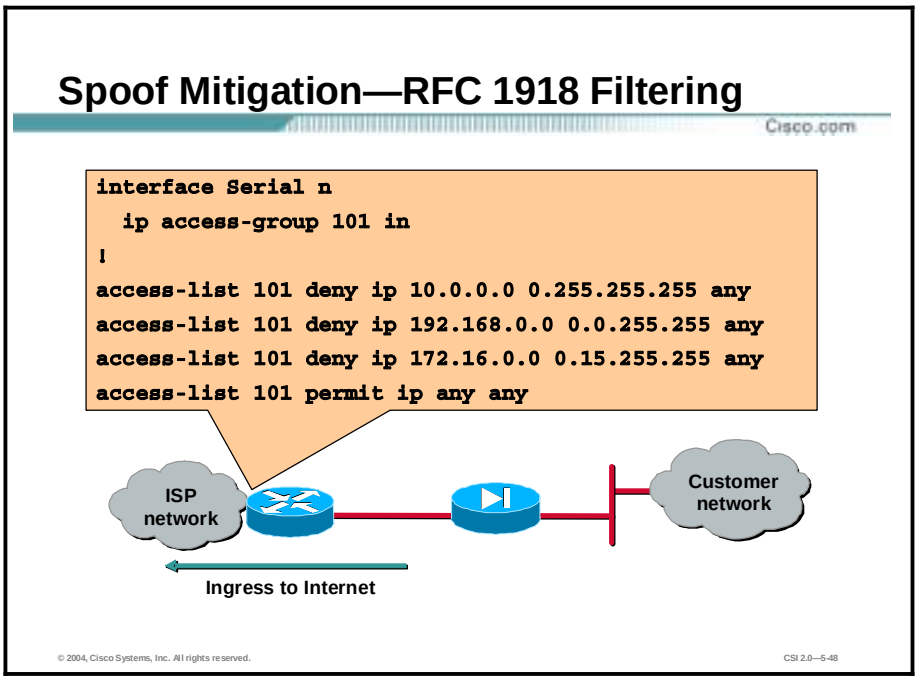

The Internet Assigned Numbers Authority (IANA) has reserved the following three blocks of the IP address space for private Internets:

- 10.0.0—10.255.255.255 (10/8 prefix)
- 172.16.0.0—172.31.255.255 (172.16/12 prefix)
- 192.168.0.0—192.168.255.255 (192.168/16 prefix)

SAFE SMR small network design recommends that these networks be filtered as RFC 1918 designates them private and they are not to be routed on the Internet. It is recommended that this filtering be completed on the ISP router as well as the PIX Firewall.

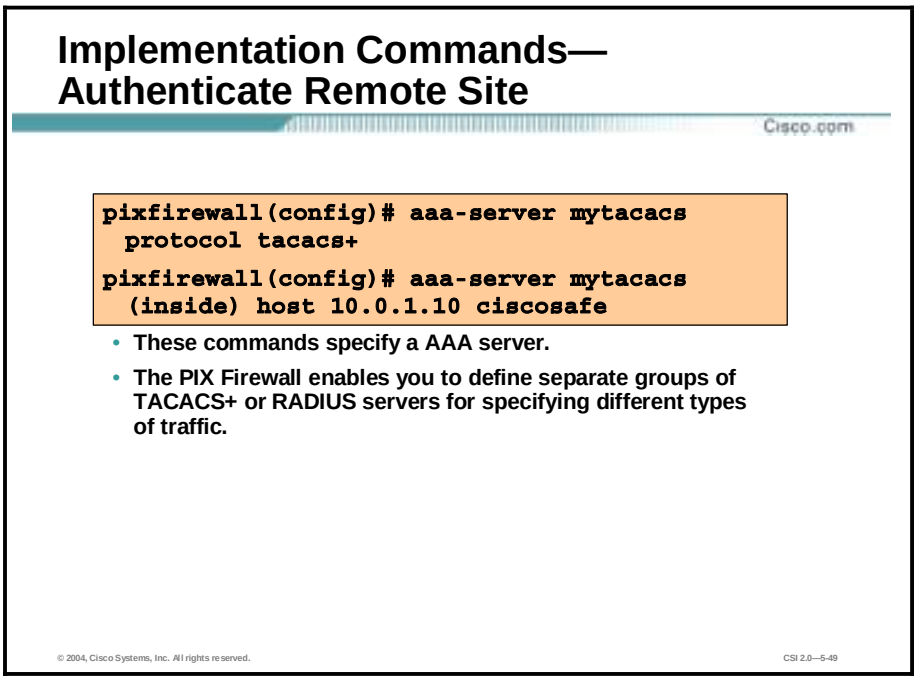

The **aaa-server** command enables you to specify AAA server groups. The PIX Firewall enables you to define separate groups of TACACS+ or RADIUS servers for specifying different types of traffic (for example, a TACACS+ server for inbound traffic and another for outbound traffic). Another use is where all outbound HTTP traffic is authenticated by a TACACS+ server, and all inbound traffic uses RADIUS.

AAA server groups are defined by a tag name that directs different types of traffic to each authentication server. If the first authentication server in the list fails, the AAA subsystem fails over to the next server in the tag group. You can have up to 14 tag groups and each group can have up to 14 AAA servers for a total of up to 196 AAA servers.

If your RADIUS server uses ports 1812 for authentication and 1813 for accounting, you are required to reconfigure the PIX Firewall to use ports 1812 and 1813.

If accounting is in effect, the accounting information goes only to the active server.

If you are upgrading from a previous version of the PIX Firewall and have AAA command statements in your configuration, using the default server groups enables you to maintain backward compatibility with the AAA command statements in your configuration.

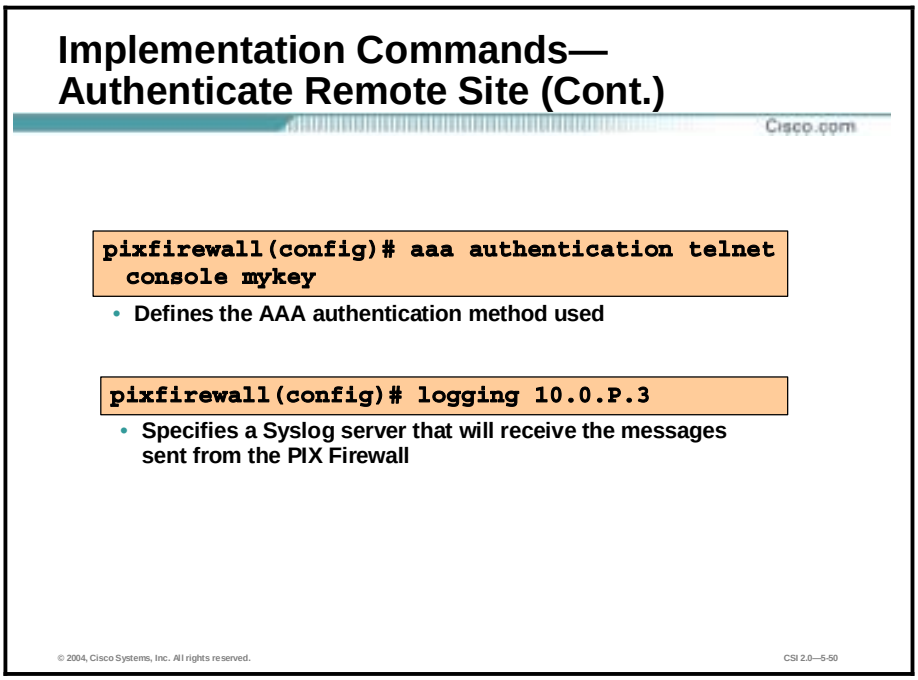

To use the **aaa authentication** command, you must first designate an authentication server with the **aaa-server** command. Also, for each IP address, one **aaa authentication** command is permitted for inbound connections and one for outbound connections.

The **aaa authentication** command is not intended to mandate your security policy. The authentication servers determine whether a user can or cannot access the system, what services can be accessed, and what IP addresses the user can access. The PIX Firewall interacts with FTP, HTTP (web access), and Telnet to display the credentials prompts for logging in to the network or logging in to exit the network. You can specify that only a single service be authenticated, but this must agree with the authentication server to ensure that both the firewall and server agree.

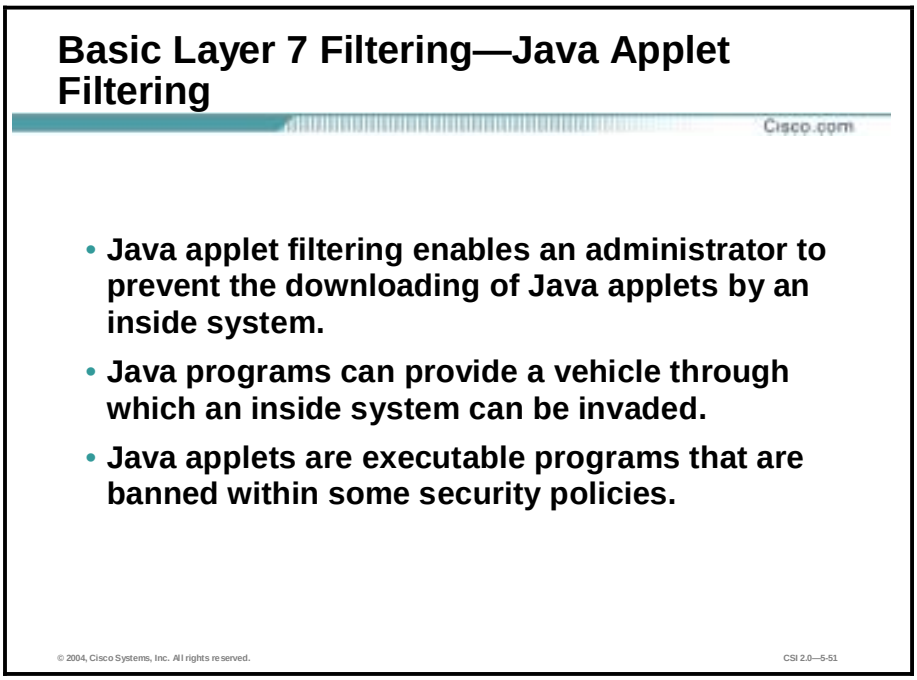

Java applet filtering distinguishes between trusted and untrusted applets by relying on a list of external sites that you designate as “friendly.” If an applet is from a friendly site, the firewall allows the applet through. If the applet is not from a friendly site, the applet is blocked. (Alternately, you could permit applets from all external sites except for those you specifically designate as “hostile.”)

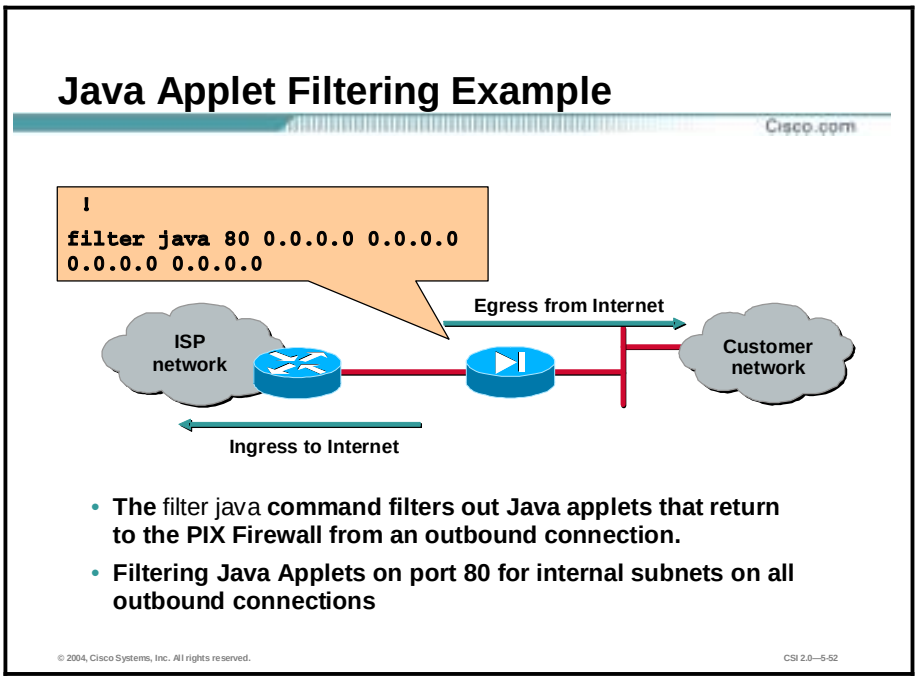

The **filter java** command filters out Java applets that return to the PIX Firewall from an outbound connection. The user still receives the HTML page, but the web page source for the applet is commented out so that the applet cannot execute. Use **0** for the local IP or foreign IP addresses to mean all hosts.

Note	If Java applets are known to be in <object> tags, use the **filter activex** command to remove them.

RFC 2827 describes a straightforward method for using ingress traffic filtering to prohibit DoS attacks that use forged IP addresses to be propagated from behind an ISP's aggregation point.

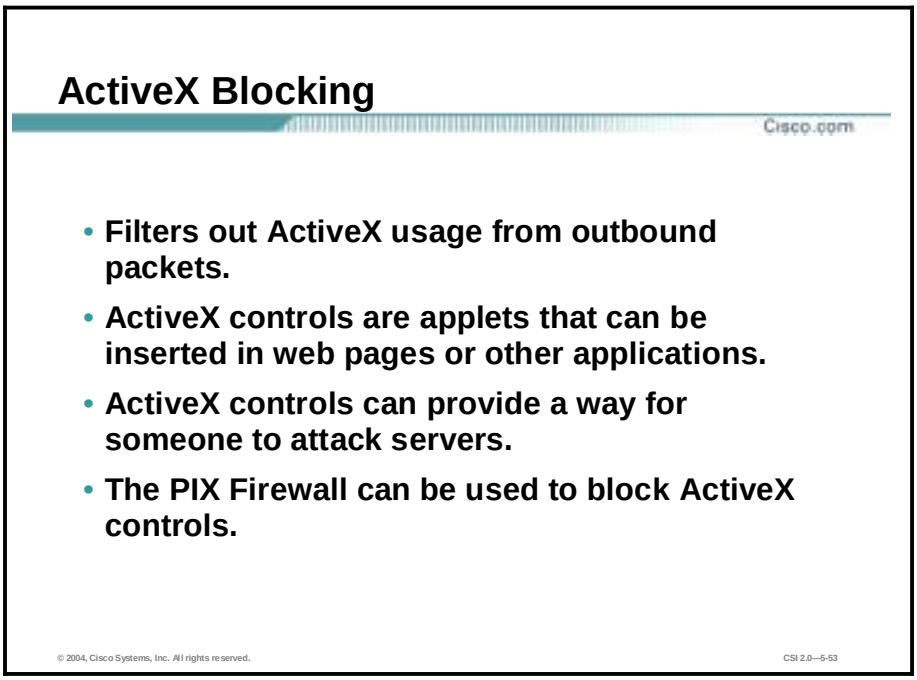

The **filter activex** command filters out ActiveX, Java applets, and other HTML <object> usages from outbound packets. ActiveX controls, formerly known as Object Linking and Embedding (OLE) or Object Linking and Embedding Control (OCX) controls, are components you can insert in a web page or other application. These controls include custom forms, calendars, or any of the extensive third-party forms for gathering or displaying information.

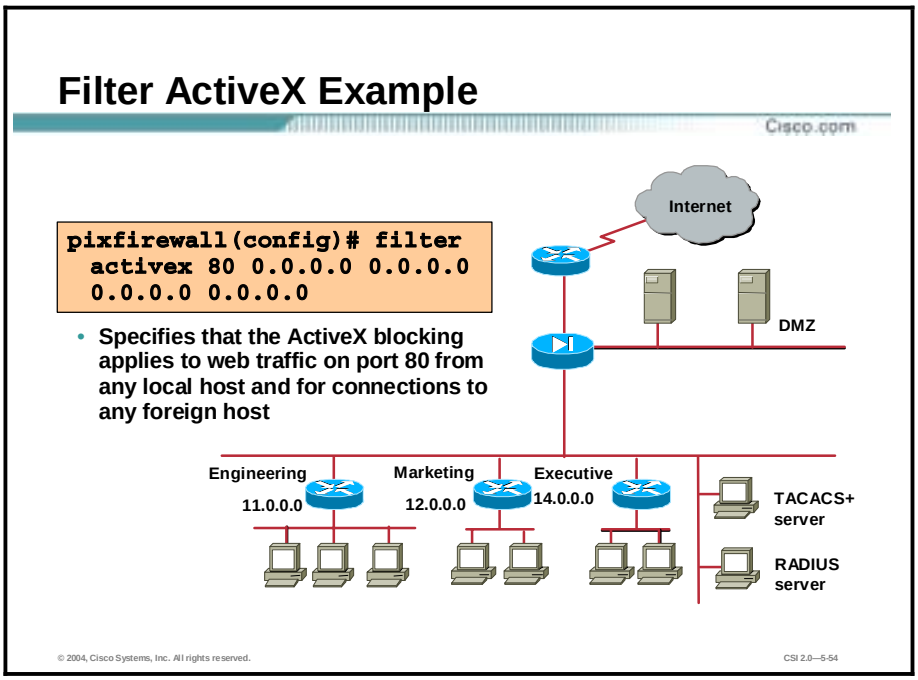

As a technology, ActiveX creates many potential problems for the network clients, including causing workstations to fail, introducing network security problems, or being used to attack servers. This feature blocks the HTML <object> tag and comments it out within the HTML web page.

Note	The <object> tag is also used for Java applets, image files, and multimedia objects, which is also blocked by the **filter activex** command. If the <object> or </object> HTML tags split across network packets or if the code in the tags is longer than the number of bytes in the MTU, the PIX Firewall cannot block the tag.

Implementation—CSA

This topic describes the detailed configuration and implementation of the Cisco Security Agent.

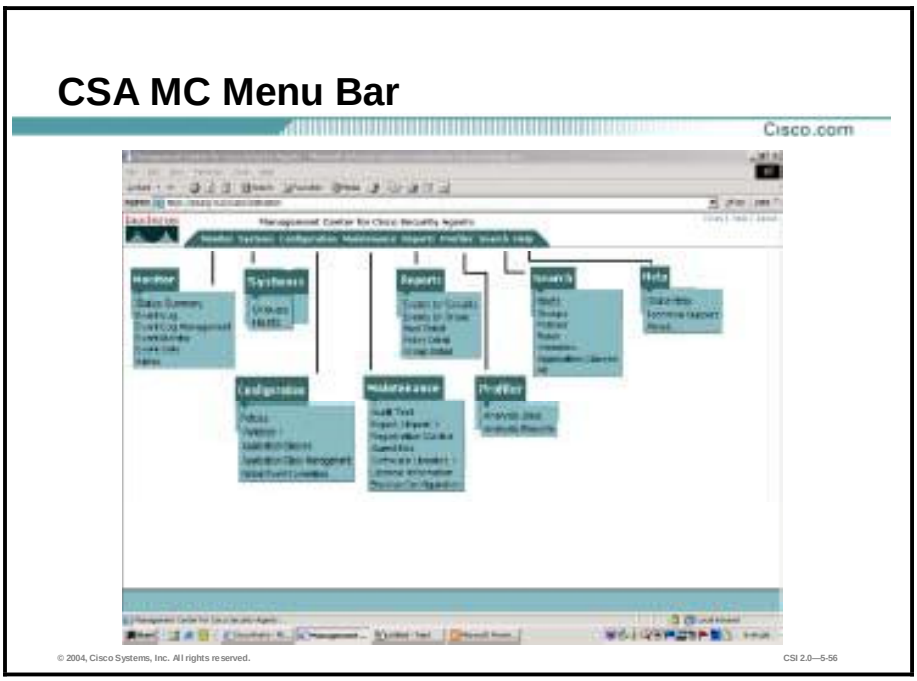

The menu bar at the top of the Management Center for Cisco Security Agent (CSA MC) window provides links to all configuration windows and list views. Arrows indicate that there are subcategories that you can choose from the top-level items. The subcategories appear when you move the mouse over the main item. The configuration options available from each menu bar item are as follows:

- Monitor—The Monitor drop-down list provides tools for viewing system status and log files. You can also set alerts and alert parameters from here.
- Systems—The Systems drop-down list lets you configure the groups into which agent host systems are placed when they register with the CSA MC.
- Configuration—The Configuration drop-down list allows you to access most of the windows you need to configure your policies for Agents. This list provides links to the rule windows you use to develop your policies, as well as links to application classes and variables. Variables such as file sets and network addresses are the building blocks for policies. They are accessible from the cascading menu that appears when you move your mouse over the Variables option in the Configuration drop-down list.
- Maintenance—The Maintenance drop-down list lets you build Agent kits, import and export configuration files, distribute software updates, and back up your database configuration. When you move your mouse over the Export/Import and Software Updates options, you can select further options from the cascading menus that appear.
- Reports—The Reports drop-down list lets you generate reports by various categories such as event severity level, the group or groups that generated the event, or by individual host systems.

- Profiler—The Profiler drop-down list lets you configure analysis jobs for the purpose of analyzing applications and creating policies.
- Search—Use the Search drop-down list options to search for a specific configuration item in the CSA MC database. You can specify a search of Hosts, Groups, Policies, Rules, Variables, Application Classes, or All by selecting one of those options from the Search drop-down list. Each option has its own criteria by which you can search.

All CSA MC action items appear in a frame at the bottom of the CSA MC window. The buttons in this frame change in accordance with the actions available for the window that you are viewing.

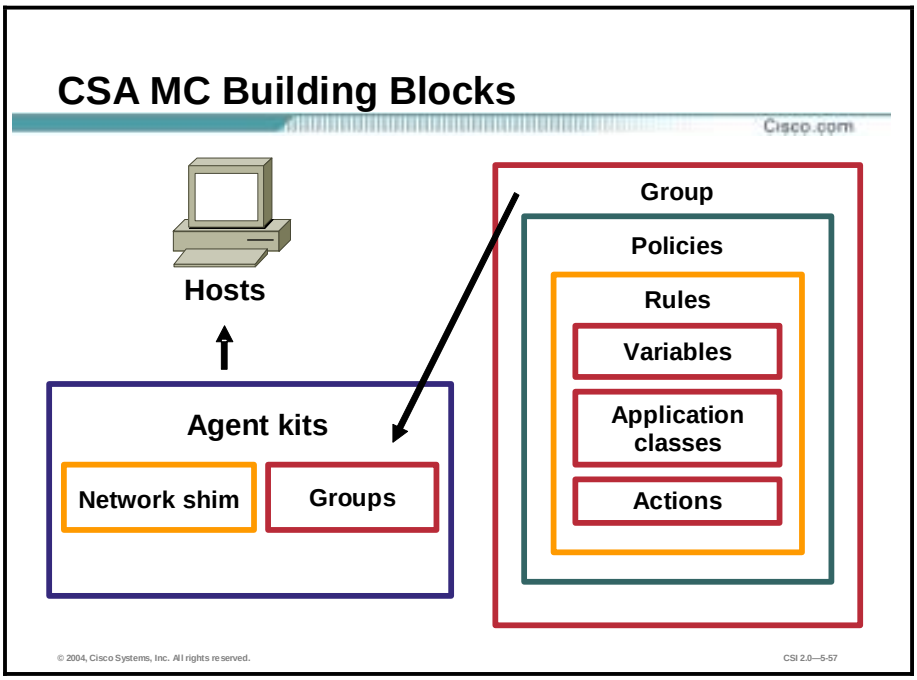

The figure illustrates the following components that build together to create Agent kits:

- Variables, application classes, and actions—Combine to create rules
- Rules—Contain variables, application classes, and actions; combine to form policies
- Policies—Contain rules; are applied to a group or multiple groups
- Groups—Contain associations with policies; accept hosts as members
- Agent kits—Contain groups and (optionally) the network shim. Agent kits are deployed to hosts, installing the CSA software and all of the policies and rules that have been built into them.

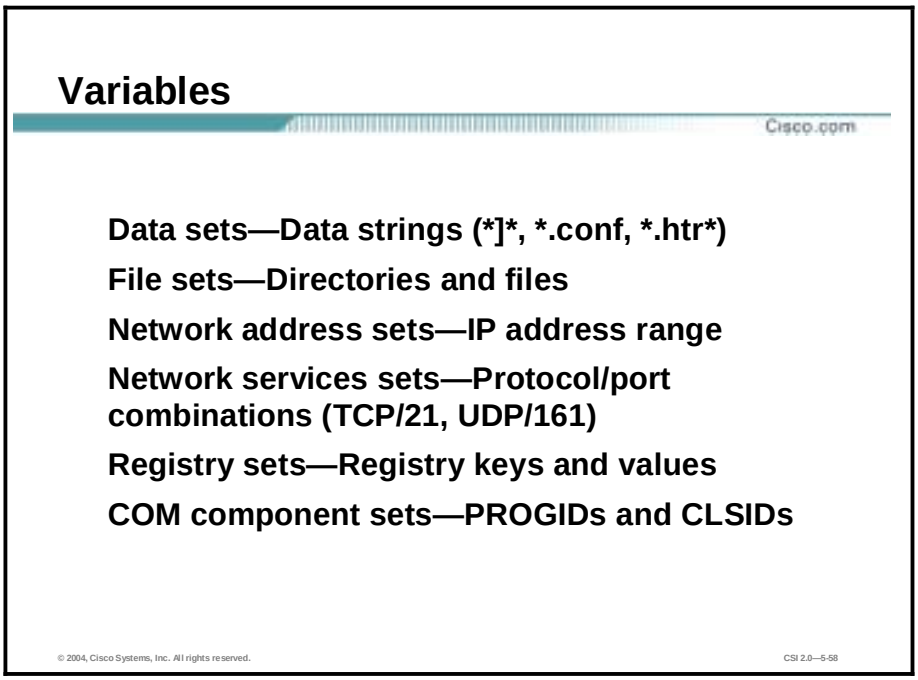

Configuration variables are configuration data items, such as files, network addresses, and network services, that you create for repeated use in configuring other items, such as file access control rules, network access control rules, and alerts. Once you have configured them, you enter these global variables in the corresponding fields for other CSA MC items.

You use configuration variables to help build the rules that form your policies. Using variables makes it easy for you to maintain policies by letting you make any modifications in one place and have those changes take effect across all rules and policies.

To configure variables, select **Configuration>Variables>(variable set type)**.

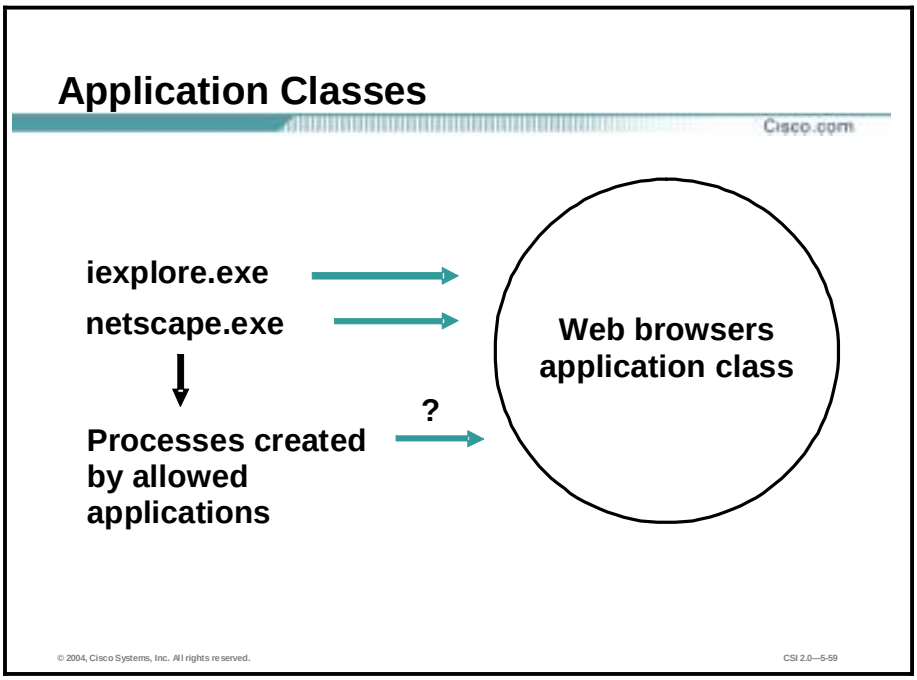

Access control rules are application-centric. When you write your rules, you should understand that the applications you select are really the heart of each rule. In your file, network, registry, and COM rules, you are controlling what applications can do to the files, addresses, registry keys, and COM components that you specify. So, when you begin creating rules, think in terms of the applications that your enterprise as a whole uses and the manner in which you want to limit the ability of an application to perform undesired actions.

Application classes are groupings of application executable files that you combine under one name, generally as part of a file set variable. For example, you can enter netscape.exe and iexplore.exe in an application class that you name Web Browsers. Then you can select Web Browsers in the application field for a rule and apply restrictions to the actions that both Netscape and Internet Explorer can perform on specified resources.

The CSA MC ships with several preconfigured application classes. These nonconfigurable application classes include the following:

- Network applications
- Processes created by network applications
- Processes executing downloaded content
- Processes created by servers (TCP and UDP)
- Processes with elevated privileges
- Remote clients (for file access control and COM component access control rules only)
- System process (network access control rules only)

Application classes can be static and dynamic. In a static application class a process is added to the class based on the name of its executable file (or the process name). Alternatively, you can build an application class based on the behavior of an application rather than building the class based on a specific application executable name. This class would be a dynamic application class

defined by process behavior on a system. To configure application classes, select **Configuration>Application Classes**.

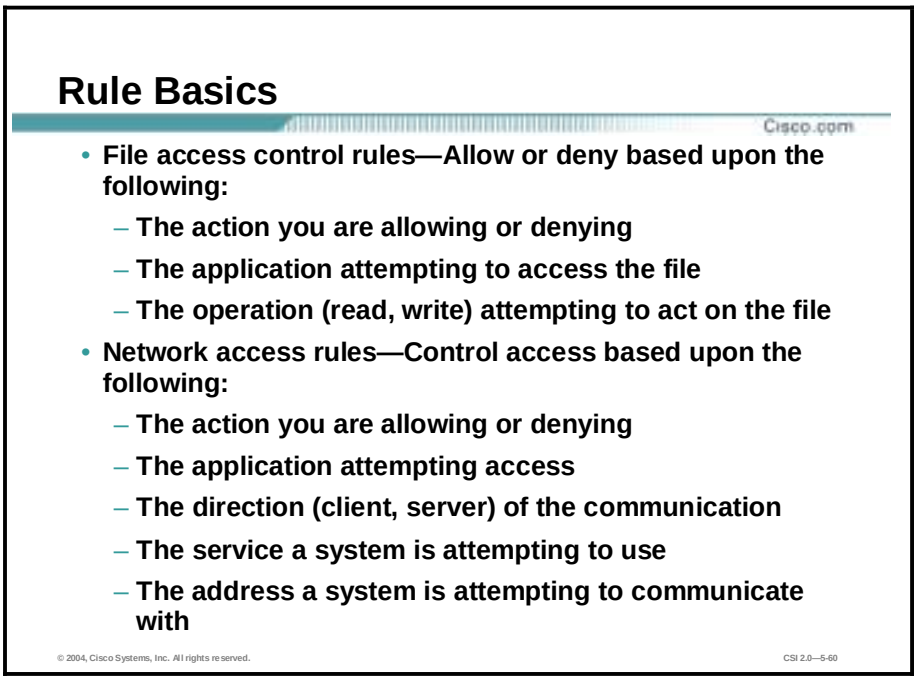

Rules are the foundation of your security policies. Creation of each rule type requires you to enter information specifying the desired behavior, as shown:

- Use file access control rules to allow or deny the operations (read, write) that the selected applications can perform on files, depending upon the following:
  - The action you are allowing or denying
  - The application attempting to access the file
  - The operation (read, write) attempting to act on the file
- Use network access rules to control access to specified network services according to the following:
  - The action you are allowing or denying
  - The application attempting to access the service or address
  - The direction (client, server) of the communication
  - The service a system is attempting to use
  - The address a system is attempting to communicate with

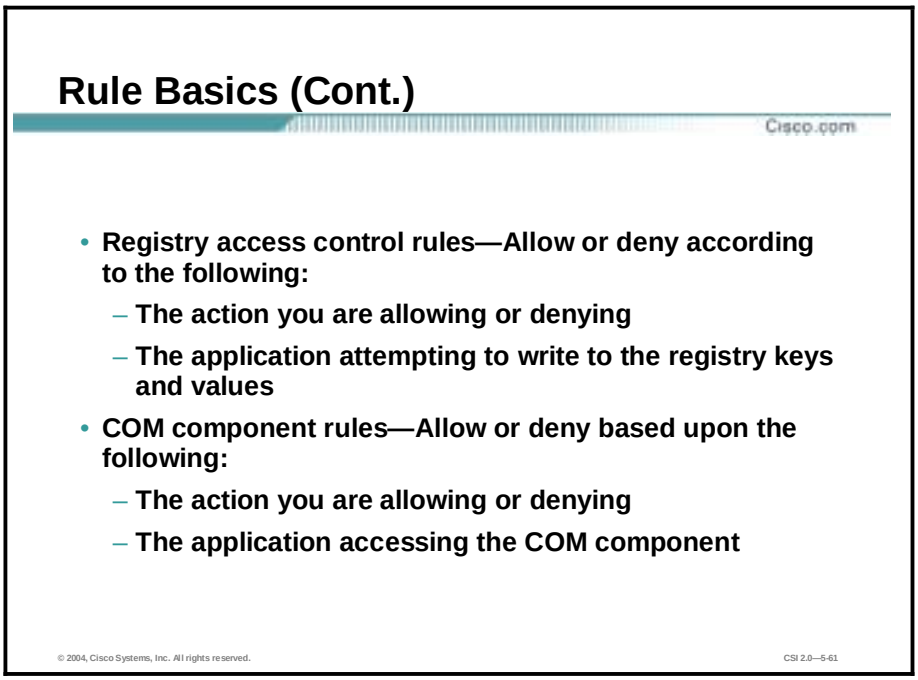

- Use registry access control rules (Windows only) to allow or deny writing to specified registry keys by selected applications, according to the following:
  - The action you are allowing or denying
  - The application attempting to write to the registry keys and values
- Use Component Object Model (COM) component access control rules (Windows only) to allow or deny access to specified COM components by selected applications according to the following:
  - The action you are allowing or denying
  - The application accessing the COM component

Other types of policies shipped with the CSA MC provide event correlation and heuristic features that can be enabled on a per-group basis, such as port scan detection, SYN flood protection, the prevention of predictable TCP sequence numbers, and the blocking of malformed IP packets. These features are especially useful for network servers.

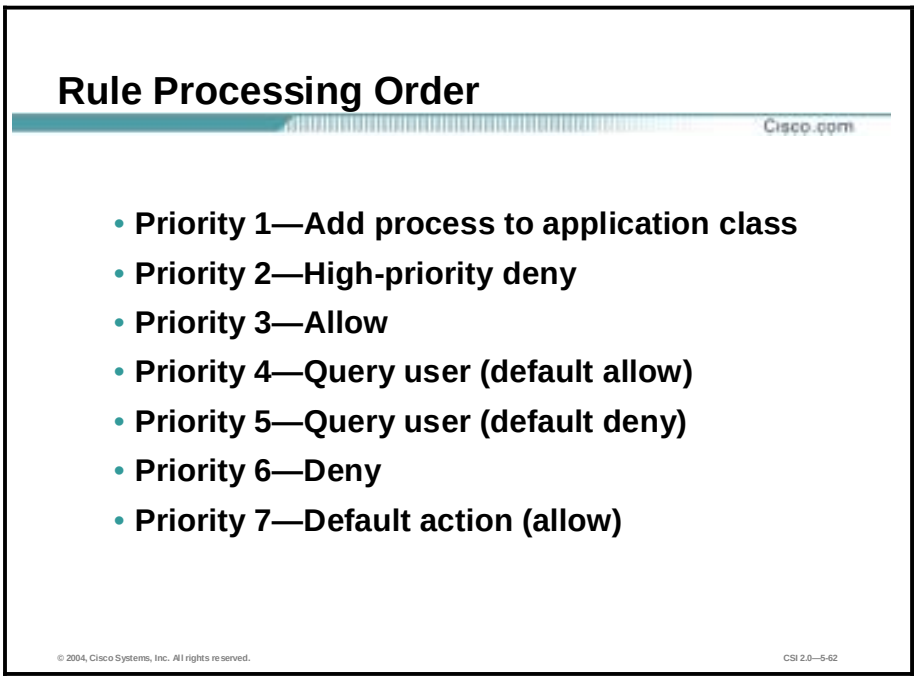

When you write rules, you create them as allow or deny actions. When you add your rules to policies, the CSA MC orders them in the following manner within each policy:

- Priority 1—Add process to application class
- Priority 2—High-priority deny
- Priority 3—Allow
- Priority 4—Query user (default allow)
- Priority 5—Query user (default deny)
- Priority 6—Deny
- Priority 7—Default action (allow)

The priority numbers indicate the order in which the CSA MC processes rules. All priority 1 rules (add to application class) are checked first, and priority 6 rules (deny) are checked last. The priority 6 rules are checked only if no other, higher-priority rules have already been triggered by a system action.

Note	An “add process to application class” rule type takes precedence over all other types. But the only action of this rule is to build a dynamic application class for any rules that make use of it. The application-builder rule does *not* override other rules, as allow, deny, and query rules do when triggered.

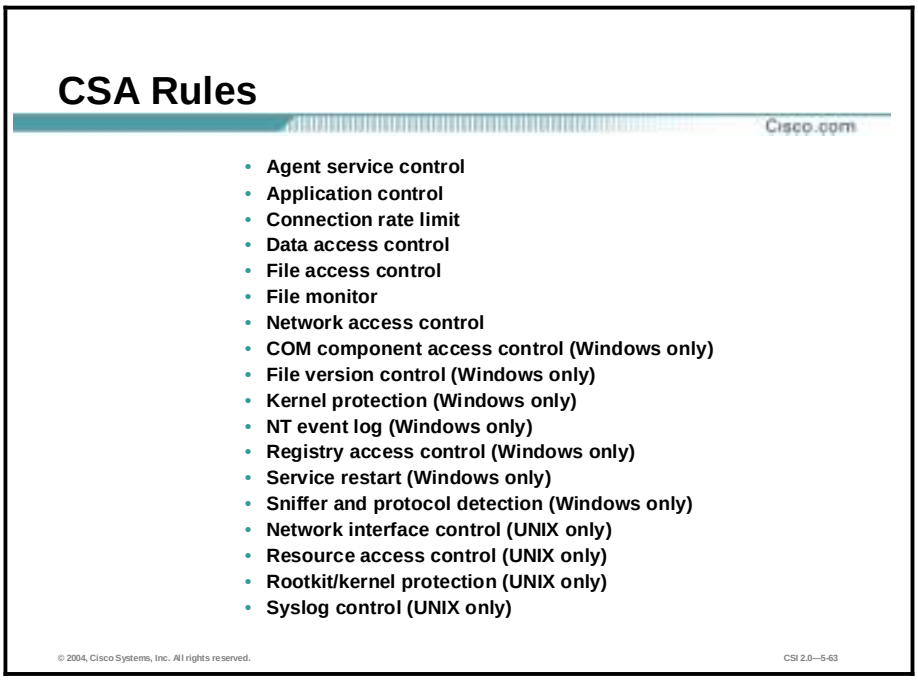

Many types of CSA rules can be built from variables, application classes, and the directed actions. The following rules can be built and then added to policies:

- Agent service control—Use the Agent service control rule to control whether users (including users who are not administrators) are allowed to suspend Agent security through the CSA user interface or by stopping the CSA service.
- Application control—Use application control rules to control which applications can run on designated Agent systems. If you deny an application class in this rule, users cannot use any application in that class. With this rule, you can also prevent an application from running only if that application was invoked by another application you specify.
- Connection rate limit—Use the connection rate limit rule to control the number of network connections that can be sent or received by systems within a specified time frame. This rule is useful in preventing DoS attacks aimed at bringing down system services. Connection rate limiting is also useful in preventing the propagation of DoS attacks.
- Data access control—Use data access control rules on web servers to detect clients making malformed web server requests where such requests could crash or hang the server. This rule detects and stops such web server attacks by examining the uniform resource identifier (URI) portion of the HTTP request.
- File access control—Use file access control rules to allow or deny the operations (read and write) that selected applications can perform on files.
- File monitor—Use the file monitor rule to track read and write access to a specified file. Unlike a file access control rule, a file monitor rule will not deny access to a file, but will log an event when a user or an application attempts read or write access of a file.
- Network access control—Use network access control rules to control access to specified network services and network addresses.
- COM component access control—Use COM component access control rules to allow or deny applications access to specified COM components.

Note	The CSA MC provides a COM component import utility that installs with each CSA. Running this utility extracts all COM component program identifiers (PROGIDs) and class identifiers (CLSIDs) for software running on the system.

- File version control—Use the file version control rule to prevent users from running specified versions of applications on their systems.
- Kernel protection—Use the kernel protection rule to prevent unauthorized access to the operating system. In effect, this rule prevents drivers from dynamically loading after system startup.
- NT event log—Use the NT Event Log rule to have specified Windows NT event log items appear in the CSA MC Event Log for selected groups.
- Registry access control—Use registry access control rules to allow or deny writing to specified registry keys by applications.
- Service restart—Use the service restart rule to have the CSA restart Windows NT services that have gone down on a system or that are not responding to service requests.
- Sniffer and protocol detection—Use the sniffer and protocol detection rule to cause an event to be logged when non-IP protocols and packet sniffer programs are detected running on systems.
- Network interface control—Use the network interface control rule to specify whether an application can open a device and act as a sniffer (promiscuous mode).
- Resource access control—Use the resource access control rule to protect systems from symbolic link attacks by preventing suspicious symbolic links from being followed.
- Rootkit/kernel protection—Use the rootkit/kernel protection rule to prevent unauthorized access to the operating system. In effect, this rule prevents drivers from dynamically loading after boot time.
- Syslog control—Use the Syslog control rule to have specified Solaris Syslog items appear in the CSA MC Event Log for selected groups.

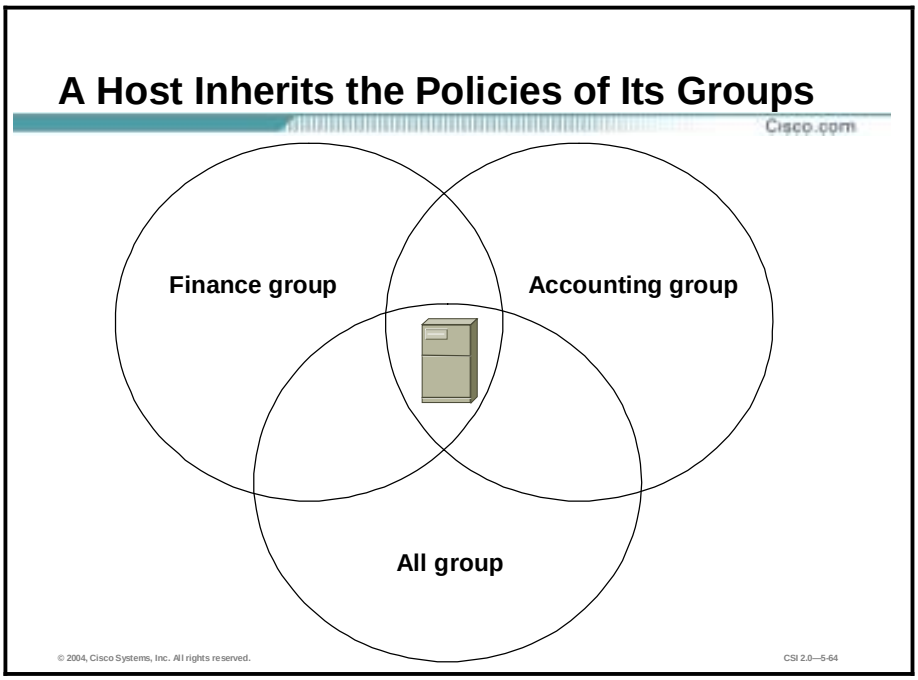

CSA provides overall system protection by tying together the control of various system components while operating under the direction of assigned security policies.

You can attach multiple policies to a group. A host can belong to multiple groups, and the host then inherits policies from all of them. For example, a desktop can belong to the Finance group and inherit the Accounting group policy. It can also belong to the All group, through which it receives the corporate mail policy.

When more than one policy is associated with a host, the rules in the individual policies are merged as though they were all defined within a single policy. In particular, the rules are ordered in the same sequence as they would be within a single policy.

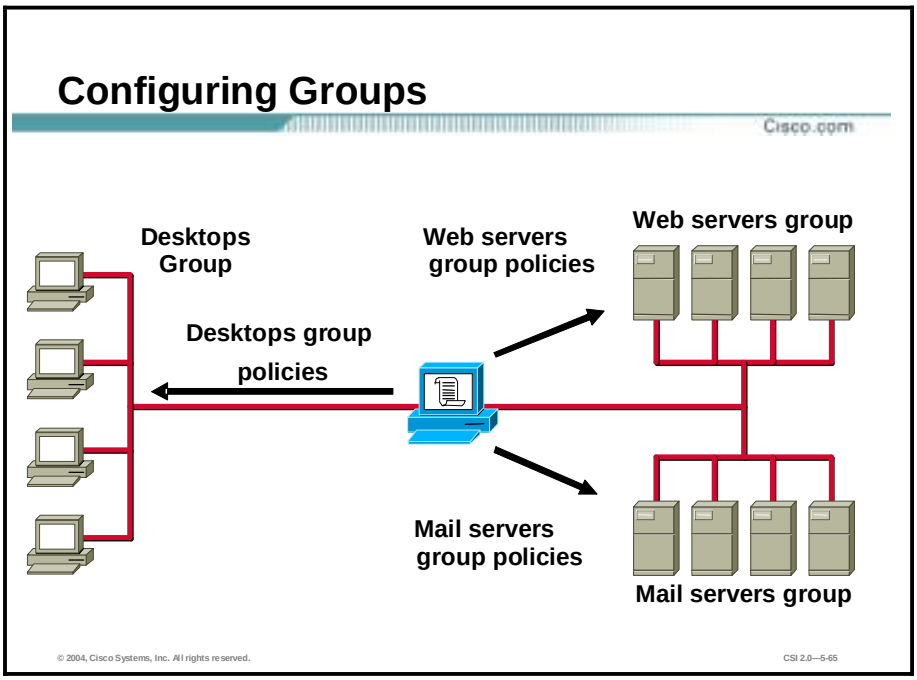

The system hosts across your network, including mobile systems in the field, must download CSA software and register with the CSA MC to receive the security policies configured for them. When you are ready to apply policies to the hosts running Agents, having those hosts placed into common groups streamlines the process of assigning policies to several hosts at once. Using groups can reduce the administrative burden of managing a large number of Agents.

In order to place hosts into groups, you must first analyze the security needs of each host system and map out a security plan. Hosts with similar requirements can then be grouped together.

The CSA MC ships with several preconfigured groups you can use. If the included groups do not suit your needs, use the instructions in this lesson to configure new groups or to edit existing ones.

To view or configure groups select **Systems>Groups**.

Event Log

As security events occur on CSA-protected hosts, CSA will take the action configured in the security policy and log the event to the CSA MC. To view details of host events, choose **Monitor>Event Log** from the main menu bar. The Event Log lets you view details of host events, including the following:

- Date—Date and time of the event
- Host—Name of the host reporting the event
- Severity—Severity level of the event
- Event—Explanation of the event and the following links
- NT events—Events from the host's NT Event Log (if an NT event log rule is configured)

The Event Log provides these tools to research and analyze a security event:

- Host link—Access information on the security policy configuration on the host
- Details link—A granular description of the event
- Rule (number)—A view of the construction of the rule involved in the logged event
- Find Similar link—Can help assess the impact of an event on the network by filtering for similar events with the same host, policy rule, severity level, event type, time frame, or any combination of these factors
- Wizard link—Can modify CSA's response to the event by allowing the action, or stop logging similar events, or begin a Profiler analysis job of the process that triggered the event

Configuring Alerts

You can configure the CSA MC to send various types of alerts to specified recipients when a policy triggers an event. Available alert types include e-mail, pager, SNMP, log to file, and a custom program that you provide.

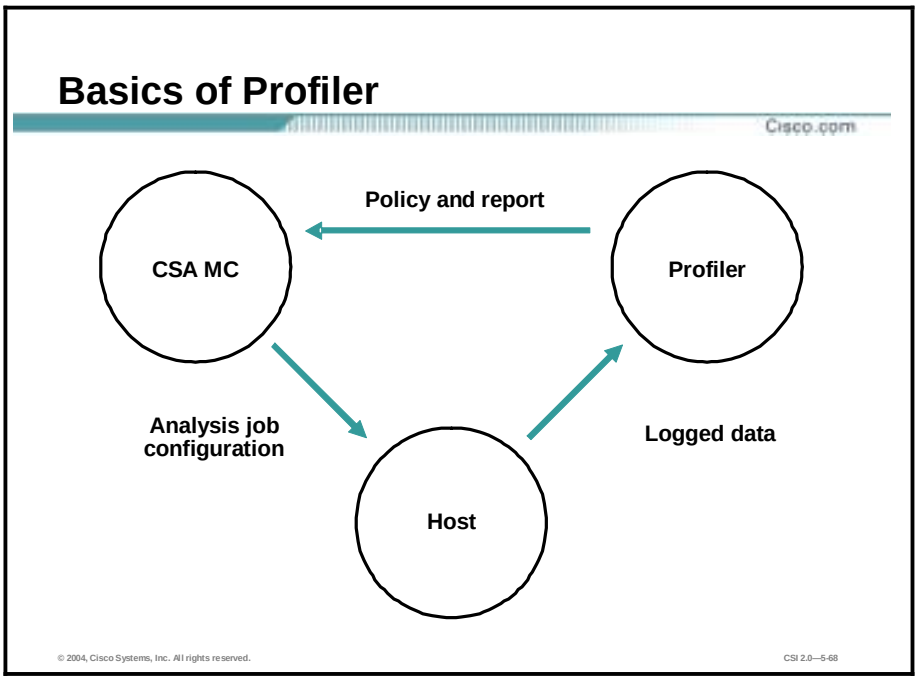

Writing effective CSA policies requires understanding the resources that applications require for normal operations. The Profiler can provide that information by analyzing applications as they operate in a normal environment and generating useful policies based on that analysis.

When deployed on a system running a CSA, the Profiler monitors the actions of designated applications on that system, logging all resource access attempts made by the application. The Profiler then analyzes the logging data that it collects and develops a policy for the application. This policy enforces what is determined to be normal application behavior while restricting all other behavior. This other behavior could be construed as abnormal or suspicious based on the analysis.

The Analysis Process

The application analysis and policy creation process is performed by three contributing components: the CSA MC, the CSA (logging agent), and the Profiler.

- CSA MC—The CSA MC is used to designate which application you want to analyze. You must also select an Agent host on which the analysis is to take place and a time frame within which the analysis will be completed. This analysis configuration is then sent to the Agent on the selected host in the same way that policies are sent to Agents.
- CSA—The Agent receives the analysis configuration information when it next polls in to CSA MC. This Agent now becomes the logging agent in this process. It logs all operations performed by the designated application, based on the assumption that the application is being thoroughly exercised in a normal operating environment. When the analysis is complete, the logged data is sent to the Profiler analysis software. The CSA MC imports the policy created by the Profiler.
- Profiler—When the Profiler license is copied to the CSA MC, a new menu item, Profiler, is added to the CSA MC menu bar. You can configure parameters for analyzing a particular application by choosing **Profiler>Analysis Jobs** in the main menu bar. The Profiler examines all the logged data it receives from the logging agent. When the analysis is

complete, the Profiler creates a policy for the application and generates reports containing information on all resources accessed by the application. The policy enforces the normal operations seen in the log file and will deny any operation attempts by the application that do not align with this normal behavior.

When you are ready to configure an analysis job for an application, you must have the following information:

- The application you want to analyze—You must have an appropriate application class configured for the analysis.
- The host running CSA that will become the logging agent—You should have an appropriate host running the application to be analyzed.

Summary

This topic summarizes the key points discussed in this lesson.

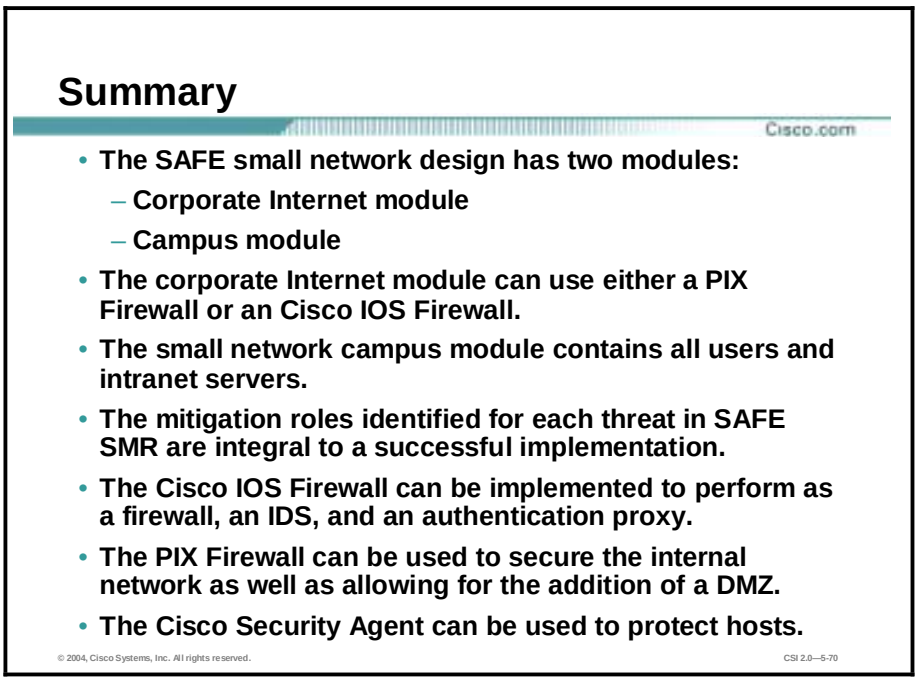

6

SAFE Midsize Network Design

Overview

This lesson describes the SAFE Midsize network design. It includes the following topics:

- Objectives
- Midsize network corporate Internet module
- Midsize network corporate Internet module design guidelines
- Midsize network campus module
- Midsize network campus module design guidelines
- Midsize network WAN module
- Implementation—ISP router and edge router
- Implementation—NIDS
- Implementation—VPN 3000 Concentrator
- Implementation—Layer 3 switch
- Summary
- Lab exercise

Objectives

This topic lists the lesson's objectives.

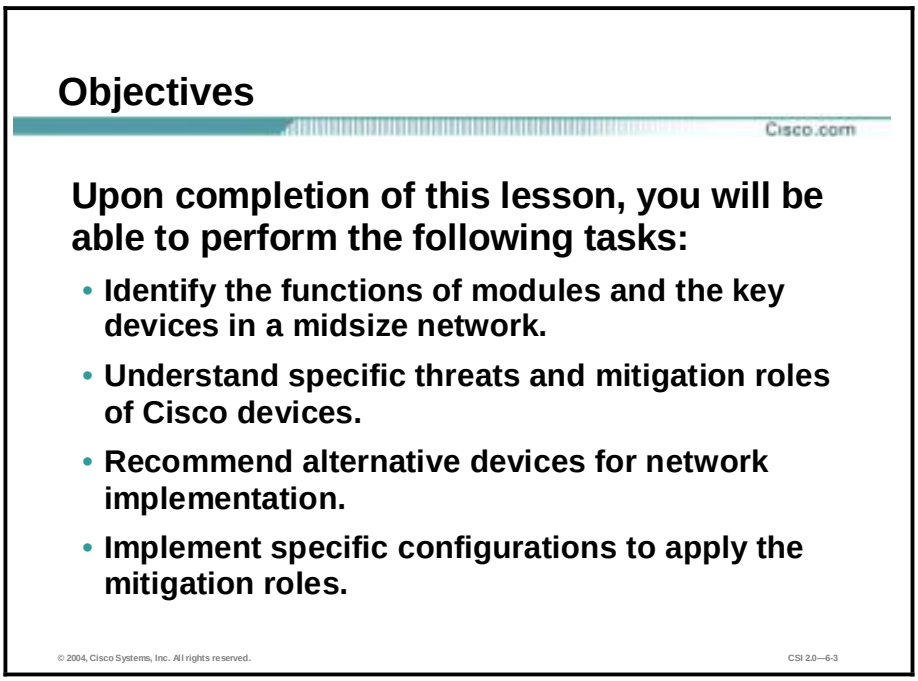

Midsize Network Corporate Internet Module

The SAFE midsize network design consists of three modules: the corporate Internet module, the campus module, and the WAN module.

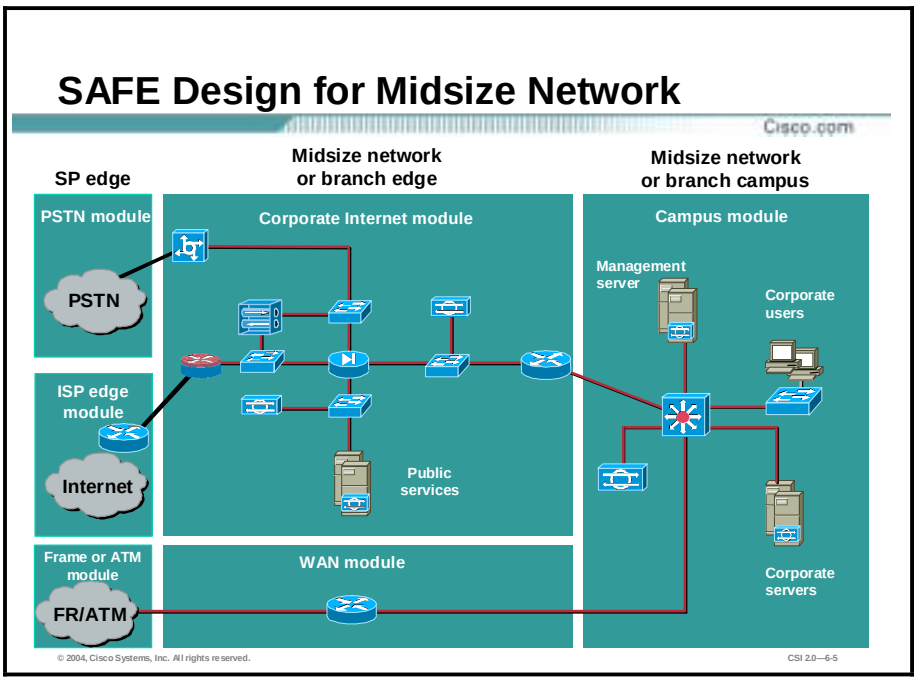

As in the small network design, the corporate Internet module has the connection to the Internet, and terminates VPN traffic and public-services traffic (Domain Name System [DNS], HTTP, FTP, and Simple Mail Transfer Protocol [SMTP]). Dial-in traffic also terminates at the corporate Internet module. The campus module contains the Layer 2 and Layer 3 switching infrastructure along with all the corporate users, management servers, and intranet servers.

From a WAN perspective, there are two options for the remote sites connecting into the midsize design. The first is a private WAN connection using the WAN module; the second is an IPSec virtual private network (VPN) into the corporate Internet module. Most of the discussion about this design is based on the midsize network operating as the head-end for a corporation. Specific design changes when used as a branch are also included.

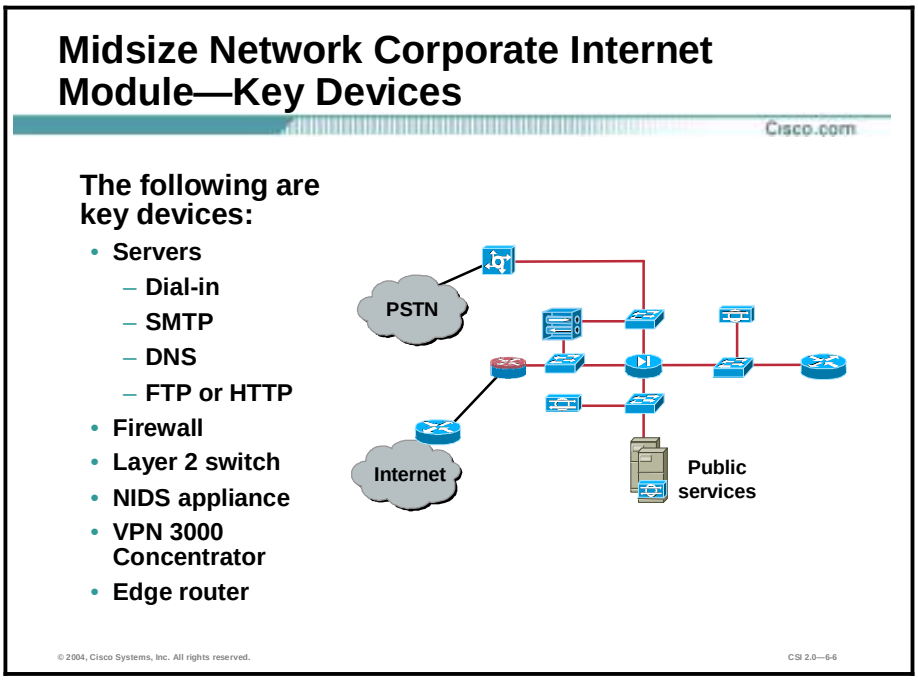

The goal of the corporate Internet module for a midsize network is to provide internal users with connectivity to Internet services and Internet users access to information about the public servers (DNS, FTP, HTTP, and SMTP). Additionally, this module terminates VPN traffic from remote users and remote sites as well as traffic from traditional dial-in users.

The following are key devices used in the corporate Internet module for a midsize network:

- Dial-in server—Authenticates individual remote users and terminates their analog connections
- SMTP server—Acts as a relay between the Internet and the intranet mail servers and inspects content
- DNS server—Serves as authoritative external DNS server for the midsize network and relays internal requests to the Internet
- FTP or HTTP server—Provides public information about the organization
- Firewall—Provides network-level protection of resources and stateful filtering of traffic, provides differentiated security for remote-access users, and authenticates trusted remote sites and provides connectivity using IPSec tunnels
- Layer 2 switches (with private VLAN support)—Provide Layer 2 connectivity for devices
- Network-based intrusion detection system (NIDS) appliance—Provides Layer 4 to Layer 7 monitoring of key network segments in the module
- VPN 3000 Concentrator—Authenticates individual remote users and terminates their IPSec tunnels
- Edge router—Provides basic filtering and Layer 3 connectivity to the Internet

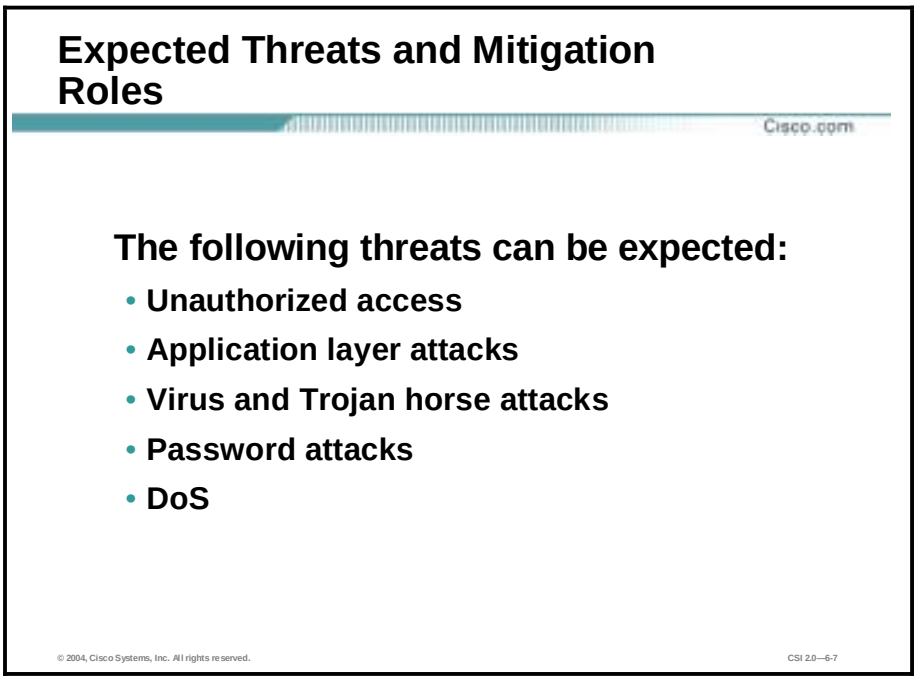

From a threat perspective, a small or midsized network is, like most networks, connected to the Internet—there are internal users who need access out and external users who need access in. Several common threats can generate the initial compromise that a hacker needs to further penetrate the network with secondary exploits.

The following threats can be expected in the SAFE SMR midsize network, and the mitigation roles:

- Unauthorized access—Mitigated through filtering at the ISP, edge router, and corporate firewall
- Application layer attacks—Mitigated through intrusion detection systems (IDSs) at the host and network levels
- Virus and Trojan horse attacks—Mitigated through e-mail content filtering, host-based intrusion detection system (HIDS) or host-based intrusion prevention system (HIPS), and host-based virus scanning
- Password attacks—Limited services available to brute force (operating systems and IDSs can detect the threat)
- Denial of service (DoS)—CAR at the ISP edge and TCP setup controls at the firewall

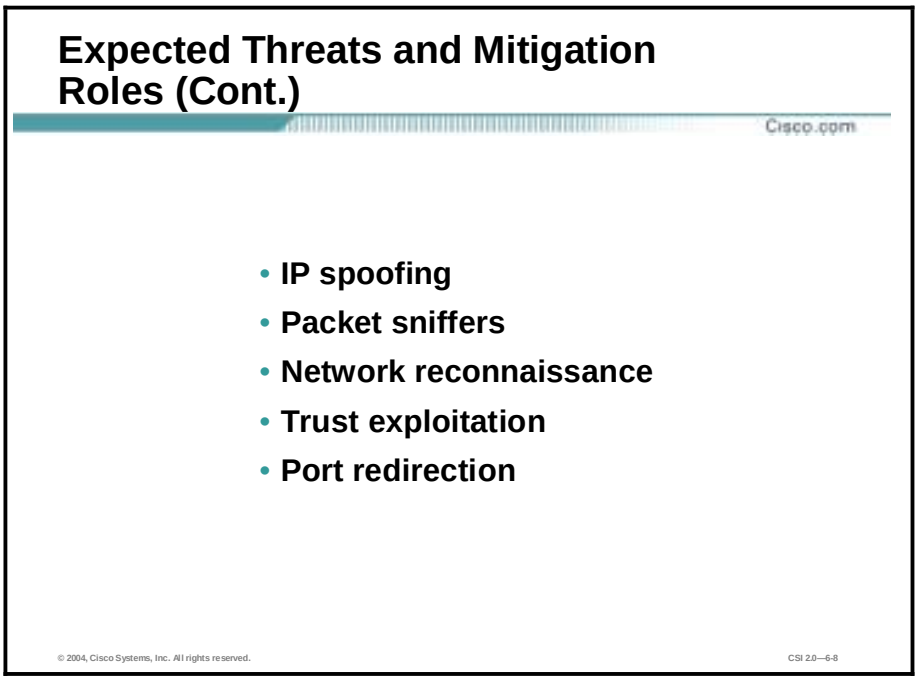

- IP spoofing—RFC 2827 and 1918 filtering at the ISP edge and midsize network edge router
- Packet sniffers—A switched infrastructure and a HIDS or HIPS to limit exposure
- Network reconnaissance—An IDS detects reconnaissance and filters protocols to limit effectiveness
- Trust exploitation—Restrictive trust model and private VLANs to limit trust-based attacks
- Port redirection—Restrictive filtering and a HIDS or HIPS to limit attacks

The publicly addressable servers are likely points of attack within the SAFE SMR Midsize corporate Internet module. These systems will likely be attacked with application layer vulnerabilities and DoS attacks.

The remote access and site-to-site VPN services are also points of attack within this module, with these expected threats and mitigation roles:

- Network topology discovery—Access control lists (ACLs) on the ingress router limit access to the VPN 3000 Concentrator and firewall (when used to terminate IPSec tunnels from remote sites), to Internet Key Exchange (IKE), and to Encapsulating Security Payload (ESP) from the Internet.
- Password attack—One-time passwords (OTPs) mitigate brute force password attacks.
- Unauthorized access—Firewall services after packet decryption prevent traffic on unauthorized ports.
- Man-in-the-middle attacks—These attacks are mitigated through encrypted remote traffic.
- Packet sniffers—A switched infrastructure limits the effectiveness of sniffing.

Though statistics vary on the percentage, it is an established fact that most attacks come from the internal network. Disgruntled employees, corporate spies, visiting guests, and inadvertent bumbling users are all potential sources of such attacks. When designing security, you must be aware of the potential for internal threats.

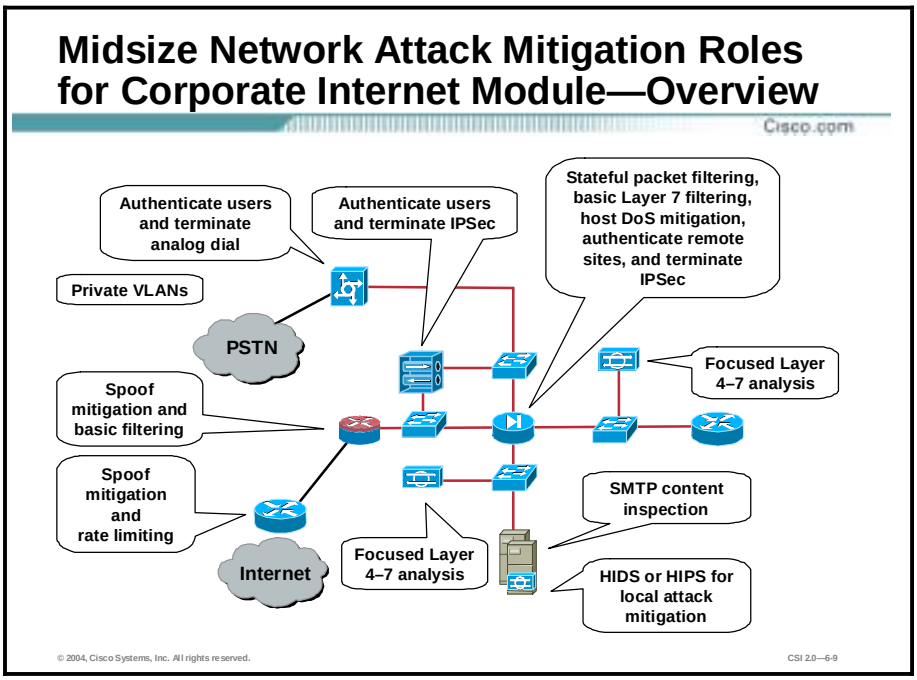

The overall roles and relative location of each device are detailed in the figure. Each device will be discussed in detail in the following topics.

Midsize Network Corporate Internet Module Design Guidelines

This topic details the functionality of each of the devices within the midsize network corporate Internet module.

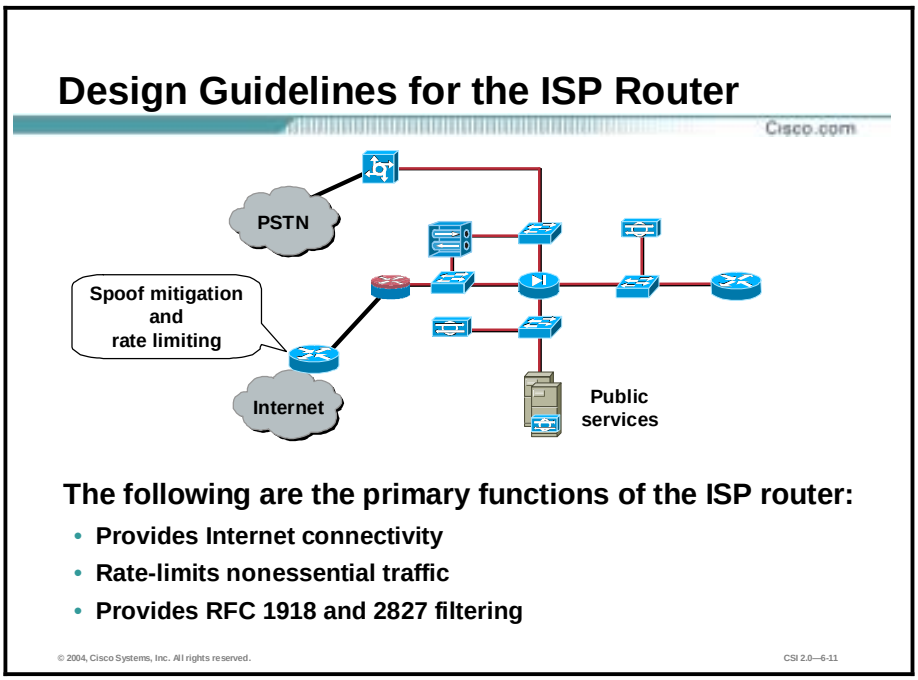

The primary function of the customer-edge router in the ISP (ISP router) is to provide connectivity to the Internet or ISP network. The egress of the ISP router rate-limits nonessential traffic that exceeds pre-specified thresholds in order to mitigate against distributed denial of service (DDoS) attacks. In addition, at the egress of the ISP router, RFC 1918 and RFC 2827 filtering is configured to mitigate against source-address spoofing of local networks and private address ranges.

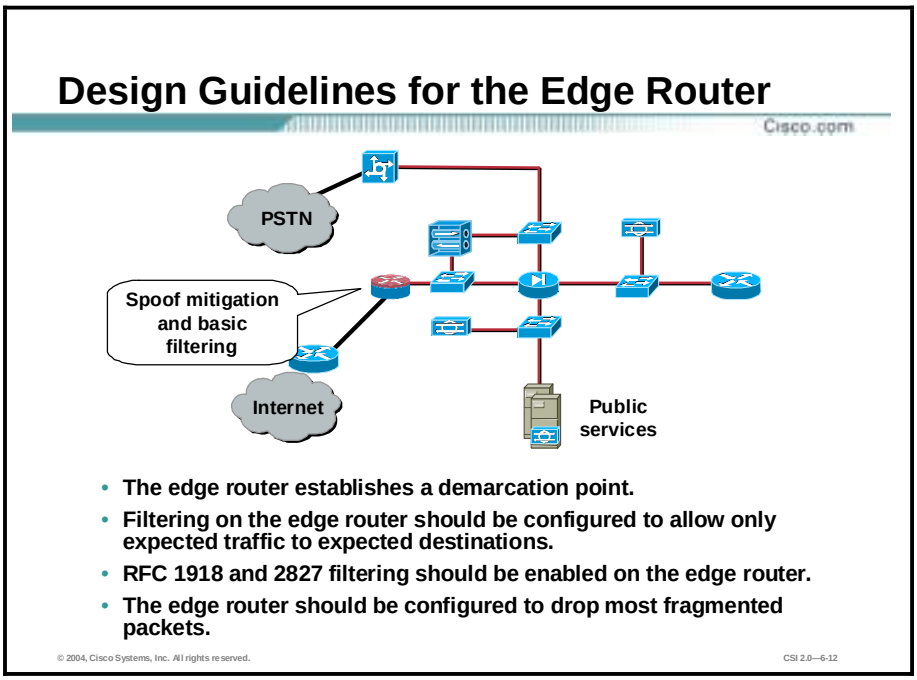

The function of the edge router on the midsize network is to provide the demarcation point between the ISP network and the midsize network. At the ingress of the edge router on the midsize network, basic filtering limits access to allow only expected IP traffic. This provides a coarse filter for the most basic attacks.

The edge router also provides RFC 1918 and RFC 2827 filtering as a verification of the ISP's filtering. In addition, the router is configured to drop most fragmented packets that are not standard traffic types on the Internet, because of the enormous security threat that they create. The loss of any legitimate traffic because of this filtering is considered acceptable when compared to the risk of allowing such traffic.

Any IPSec traffic destined for the VPN 3000 Concentrator or the firewall is allowed through. Filtering on the router is configured to allow only IKE and IPSec traffic to reach the VPN 3000 Concentrator or firewall. With remote-access VPNs, because the IP address of the remote system is not generally known, the filtering can be specified only to the head-end peer (VPN 3000 Concentrator) with which the remote users are communicating. With site-to-site VPNs, the IP address of the remote site is usually known; therefore, filtering may be specified for VPN traffic to and from both peers.

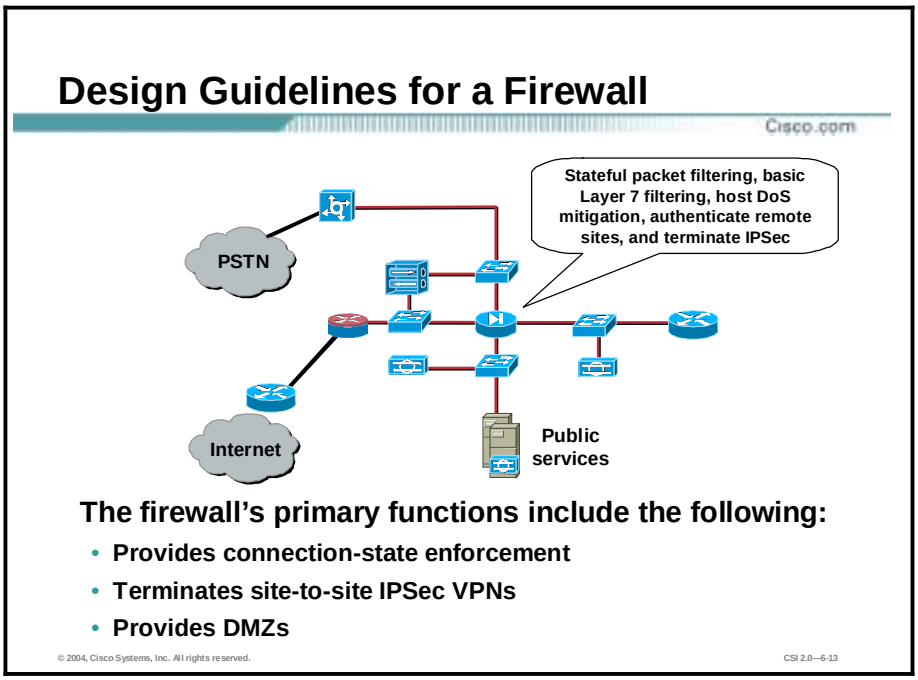

The primary function of the PIX Firewall is to provide connection-state enforcement and detailed filtering for sessions initiated through the firewall. The firewall also acts as a termination point for site-to-site IPSec VPN tunnels for both remote site production and remote site management traffic.

There are multiple segments off the firewall. The first is the public services segment, or Demilitarized Zone (DMZ), which contains all the publicly addressable hosts. The second is the segment for remote-access VPN and dial-in.

DNS should be locked down to respond only to desired commands and eliminate any unnecessary responses that might assist hackers in network reconnaissance. This includes preventing zone transfers from anywhere except legitimate secondary DNS servers.

The SMTP server includes mail-content inspection services that mitigate virus attacks and Trojan horse attacks generated against the internal network, which are usually introduced through the mail system. The firewall itself filters SMTP messages at Layer 7 to allow only necessary commands to the mail server.

Publicly addressable servers have some protection against TCP SYN floods through mechanisms such as the use of half-open connection limits on the firewall. From a filtering standpoint, in addition to limiting traffic on the public services segment to relevant addresses and ports, filtering in the opposite direction also occurs. If an attack compromises one of the public servers (by circumventing the firewall, HIDS or HIPS, and network-based intrusion detection system [NIDS]), that server should not be able to further attack the network.

To mitigate such an attack on one of the public servers, specific filtering prevents any unauthorized requests from being generated by the public servers to any other location. For example, the web server should be filtered so that it cannot originate requests of its own, but merely respond to requests from clients. This setup helps prevent a hacker from downloading additional utilities to the compromised device after the initial attack. It also helps stop unwanted sessions from being triggered by the hacker during the primary attack. An attack that generates

an xterm from the web server through the firewall to the hacker's machine is an example of such an attack.

Private VLANs (PVLANs) prevent a compromised public server from attacking other servers on the same segment. This traffic is not even detected by the firewall, a fact that explains why private VLANs are critical.

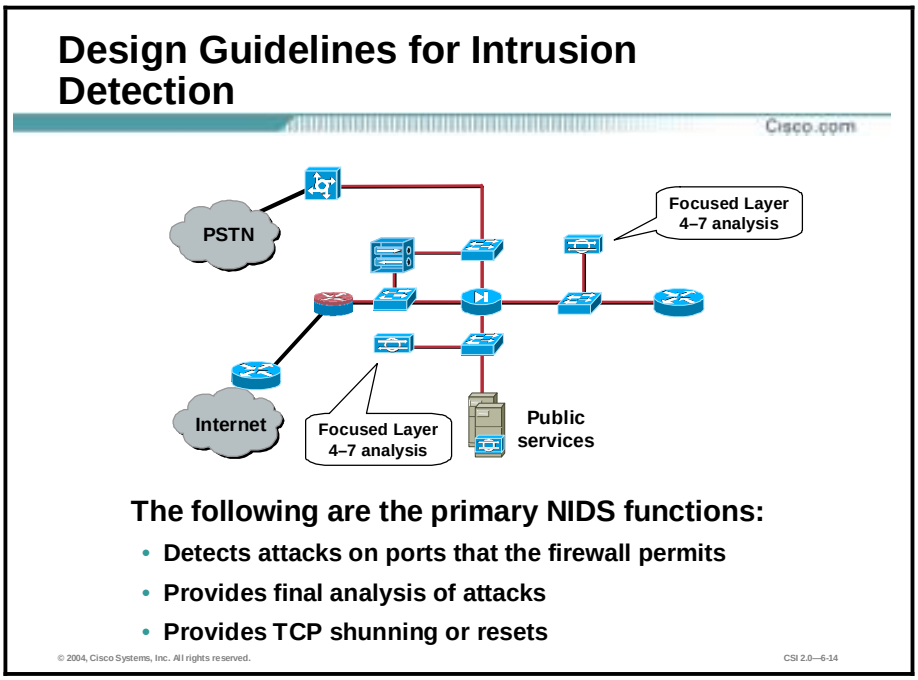

The public services segment includes an NIDS appliance. The primary function of an NIDS is to detect attacks on ports that the firewall is configured to permit. These attacks are most often application layer attacks against specific services. The NIDS on the public services segment should be set in a restrictive stance because signatures matched there have successfully passed through the firewall already.

Each of the servers has a HIDS or HIPS on it as well. The primary function of a HIDS or HIPS is to monitor against any rogue activity that occurs at the operating-system level as well as in common server applications (FTP, HTTP, SMTP, and so on).

The NIDS appliance between the private interface of the firewall and the internal router provides a final analysis of attacks. Very few attacks should be detected on this segment because only responses to initiated requests, a few select ports from the public services segment, and traffic from the remote-access segment are allowed to the inside. Only sophisticated attacks should be seen on this segment because they could mean that a system on the public services segment has been compromised and the hacker is attempting to take advantage of this foothold to attack the internal network.

For example, if the public SMTP server was compromised, a hacker might try to attack the internal mail server over TCP port 25, which is permitted to allow mail transfer between the two hosts. If attacks are seen on this segment, the responses to those attacks should be more severe than those on other segments because they probably indicate that a compromise has already occurred. The use of TCP resets or shunning to thwart, for example, the SMTP attack mentioned previously, should be seriously considered.

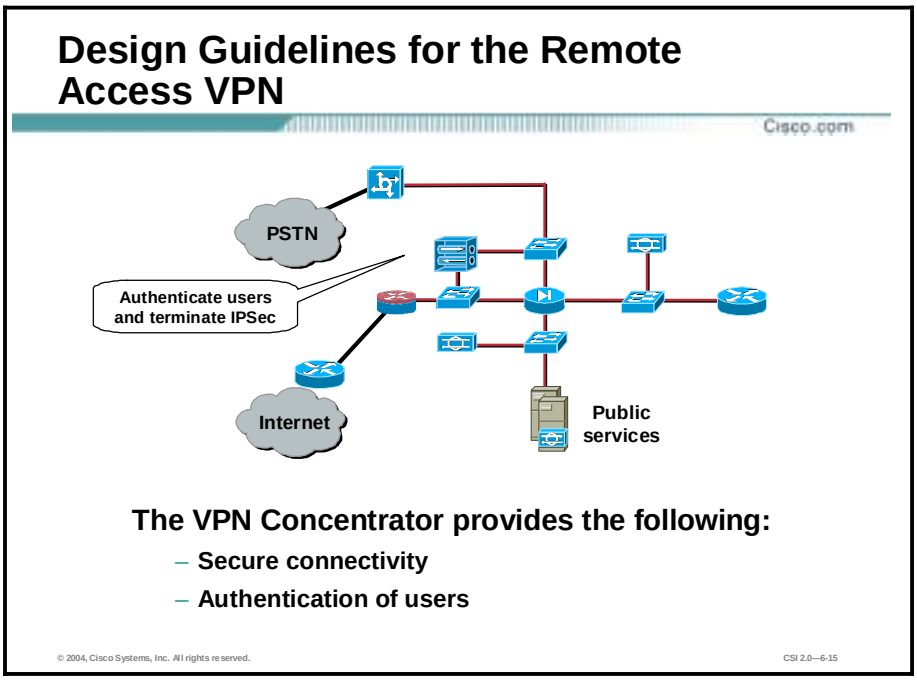

The primary function of the remote-access VPN 3000 Concentrator is to provide secure connectivity to the midsize network for remote users.

Following termination of the VPN tunnel, traffic is sent through a firewall to ensure that VPN users are appropriately filtered. This setup also enables IDS shunning to take place on the firewall. This scenario is in contrast to many deployments today that place the firewall in front of the VPN device. When placed in front, no visibility into the specific types of user traffic is possible because the traffic is still encrypted.

Via IPSec policy sent from the VPN 3000 Concentrator to the client, users are prevented from enabling split tunneling, thereby forcing users to access the Internet via the corporate connection. The IPSec parameters used are Triple-Data Encryption Standard (3DES) for encryption and Secure Hash Algorithm (SHA) and hash-based message authentication code (HMAC) for data integrity.

Design Guidelines for Dial-In Access Users

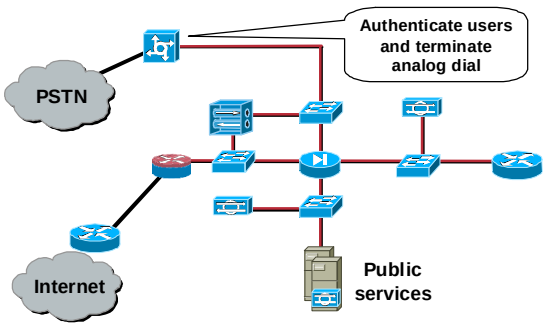

Traditional dial-in users are terminated on an access router with built-in modems. CHAP is used to authenticate the user along with the AAA server.

The traditional dial-in users are terminated on an access router with built-in modems. When the Layer 2 connection is established between the user and the server, three-way Challenge Handshake Authentication Protocol (CHAP) is used to authenticate the user. As in the remote-access VPN service, the authentication, authorization, and accounting (AAA) server is used for authentication. When authenticated, the users are provided with IP addresses from an IP pool.

Design Guidelines for the Inside Router

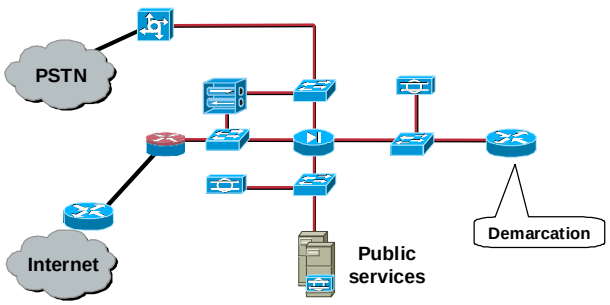

- **The primary function of the inside router is to provide Layer 3 separation and routing between the corporate Internet module and the campus module.**
- **The inside router functions strictly as a router with no ACLs restricting traffic across either interface.**

The primary function of the inside router is to provide Layer 3 separation and routing between the corporate Internet module and the campus module. This device functions strictly as a router with no ACLs restricting traffic across either interface.

Because routing information itself can be used in a DoS attack, authentication of routing updates between devices may be used to mitigate such an attack. The inside router provides a final point of demarcation between the routed intranet and the outside world. Because most firewalls are configured without routing protocols, it is important to provide a point of routing within the corporate Internet module that does not rely on the rest of the network.

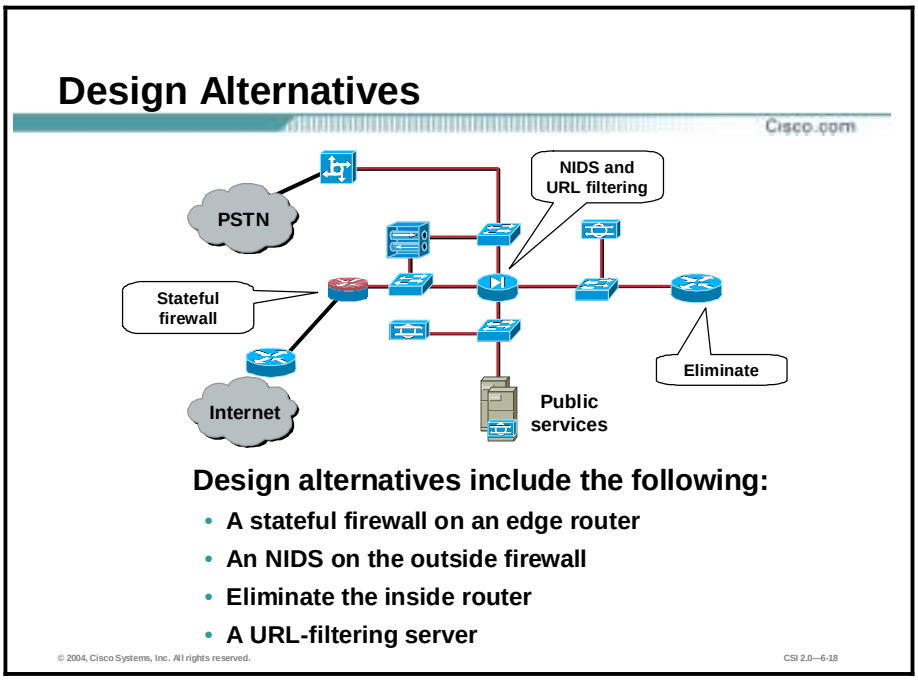

The SAFE SMR corporate Internet module has the following alternative designs:

- Stateful firewall—Rather than implementing basic filtering on the edge router to the midsize network, a network administrator may choose to implement a stateful firewall on this device as well. Having two stateful firewalls provides more of a defense-in-depth approach to security within the module.
- NIDS—Depending on the network administrator's attitude toward attack awareness, an NIDS appliance might be required in front of the firewall. With the appropriate basic filters, the IDS outside the firewall can provide important alarm information that would otherwise be dropped by the firewall.

  Because the number of alarms generated on the segment in front of the firewall is probably large, alarms generated there should have a lower severity than alarms generated behind a firewall. Also, consider logging alarms from the segment in front of the firewall to a separate management station to ensure that legitimate alarms from other segments get the appropriate attention. With the visibility that an NIDS outside the firewall provides, evaluation of the attack types your organization is attracting can be better seen. In addition, evaluation of the effectiveness of ISP and enterprise edge filters can be performed.
- Elimination of the inside router—Another alternative is the elimination of the router between the firewall and the campus module. Although its functions can be integrated into the campus module Layer 3 switch, this setup would eliminate the ability of the corporate Internet module to function without relying on Layer 3 services from another area of the network.
- URL-filtering server—The addition of content inspection beyond the mail-content inspection already specified could also be used. For example, a URL-filtering server could be placed on the public services segment to filter the types of web pages that employees can access.

Midsize Network Campus Module

This topic describes the midsize network campus module.

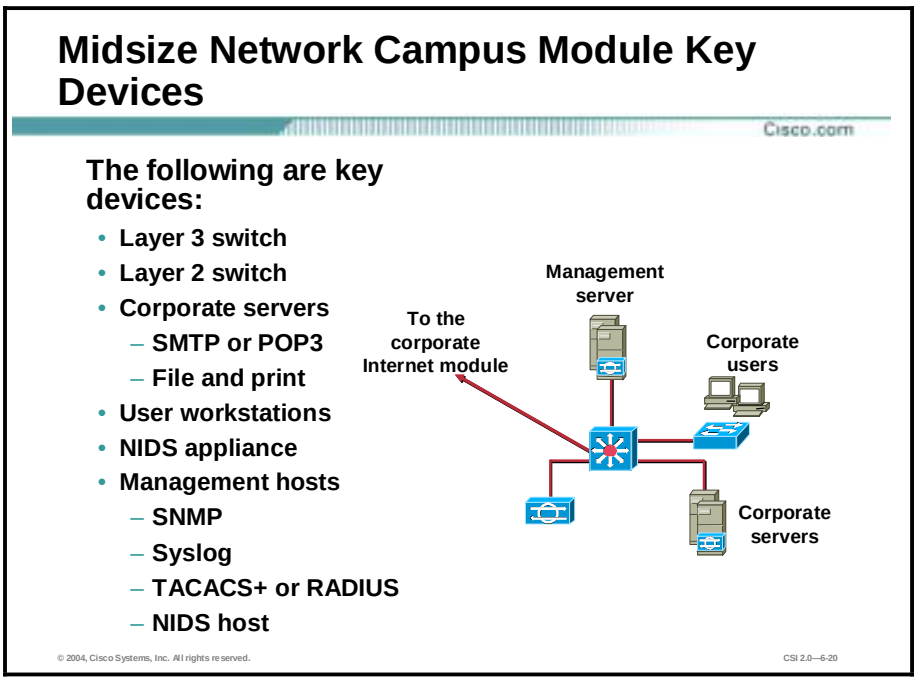

The midsize network campus module contains end-user workstations, corporate intranet servers, management servers, and the associated Layer 2 and Layer 3 infrastructure required to support the devices. All the campus modules from SAFE Enterprise have been combined into a single module. This setup more accurately reflects the smaller size of midsize networks and reduces the overall cost of the design. As in the corporate Internet module, the redundancy that would normally be found in an enterprise design has been removed from the midsize network design.

The following are the midsize network campus module key devices:

- Layer 3 switches—Route and switch production and management traffic within the campus module, provide distribution layer services to the building switches, and support advanced services such as traffic filtering
- Layer 2 switches (with PVLAN support)—Provide Layer 2 services to user workstations
- Corporate servers—Provide e-mail (SMTP and POP3) services to internal users, as well as delivering file, print, and DNS services to workstations
- User workstations—Provide data services to authorized users on the network
- NIDS appliance—Provides Layer 4-to-Layer 7 monitoring of key network segments in the module
- Management hosts—Hosts designated for device management
  - Simple Network Management Protocol (SNMP) host—Provides SNMP management for devices
  - NIDS host—Provides alarm aggregation for all NIDS devices in the network
  - Syslog host—Aggregates log information for firewall and NIDS hosts

- Terminal Access Controller Access Control System Plus (TACACS+)—Provides user authentication for device management
- Remote Access Dial-In User Service (RADIUS)—Provides user authentication for dial-in access

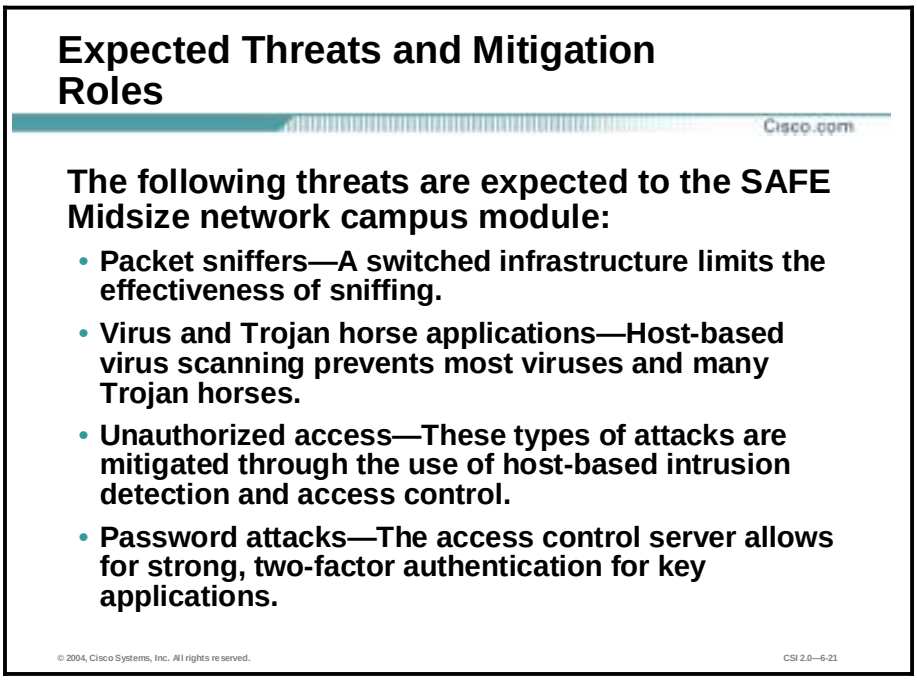

The following threats are expected in the midsize network campus module, with mitigation roles:

- Packet sniffers—A switched infrastructure limits the effectiveness of sniffing.
- Virus and Trojan horse applications—Host-based virus scanning prevents most viruses and many Trojan horses.
- Unauthorized access—These types of attacks are mitigated through the use of host-based intrusion detection and access control.
- Password attacks—The access control server allows for strong two-factor authentication for key applications.

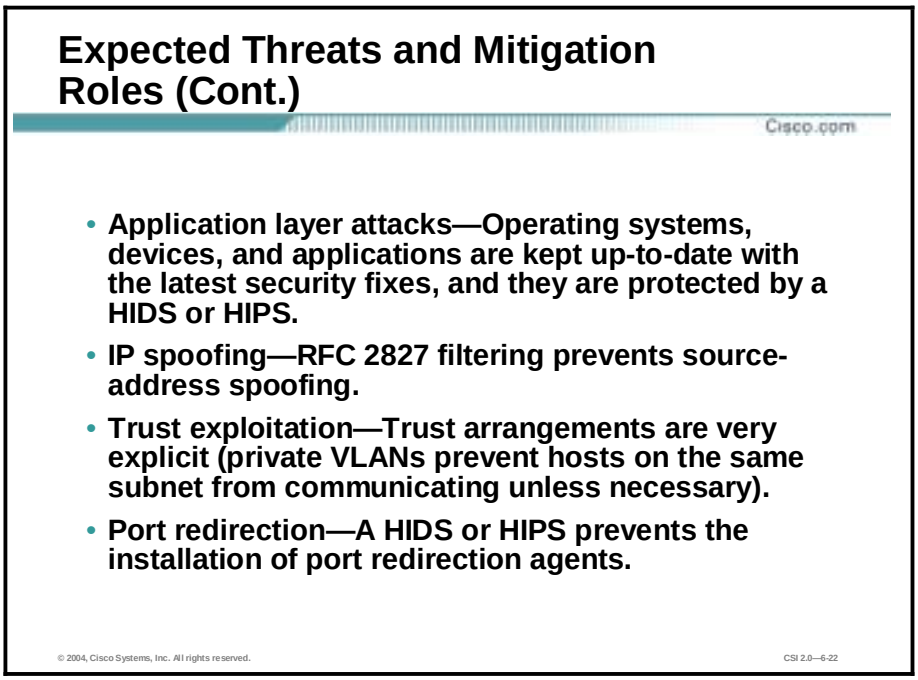

- Application layer attacks—Operating systems, devices, and applications are kept up-to-date with the latest security fixes, and they are protected by a HIDS or HIPS.
- IP spoofing—RFC 2827 filtering prevents source-address spoofing.
- Trust exploitation—Trust arrangements are very explicit (private VLANs prevent hosts on the same subnet from communicating unless necessary).
- Port redirection—A HIDS or HIPS prevents the installation of port redirection agents.

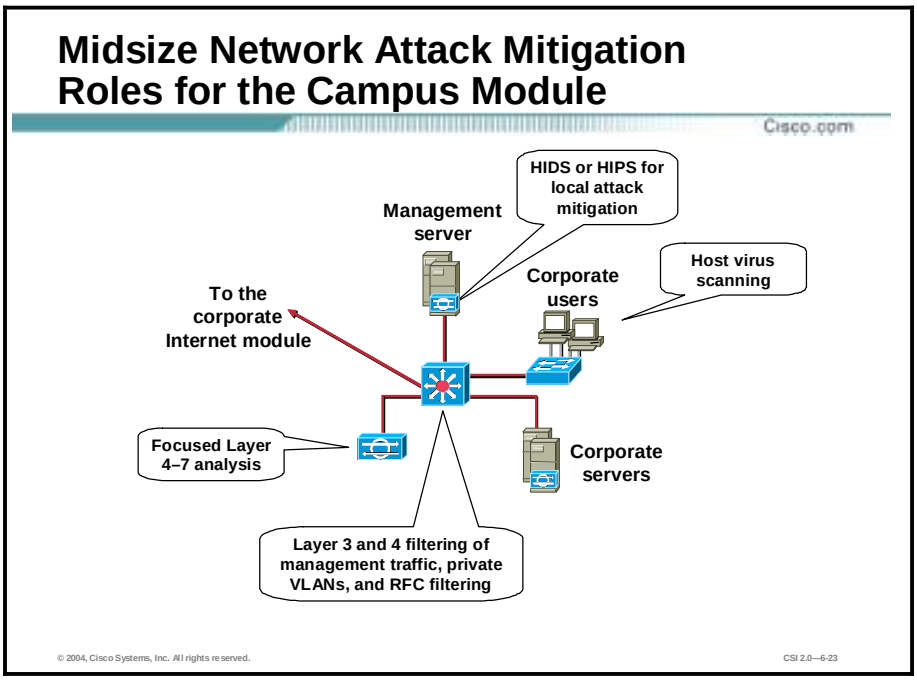

The campus module performs the following actions:

- The NIDS performs focused Layer 4 through 7 analysis.
- The Layer 3 switch provides Layer 3 and 4 filtering of management traffic, private VLANs, and RFC 2827 filtering.
- Hosts provide virus scanning.
- The HIDS or HIPS provides local attack mitigation.

Midsize Network Campus Module Design Guidelines

This topic details the functionality of devices within the campus module.

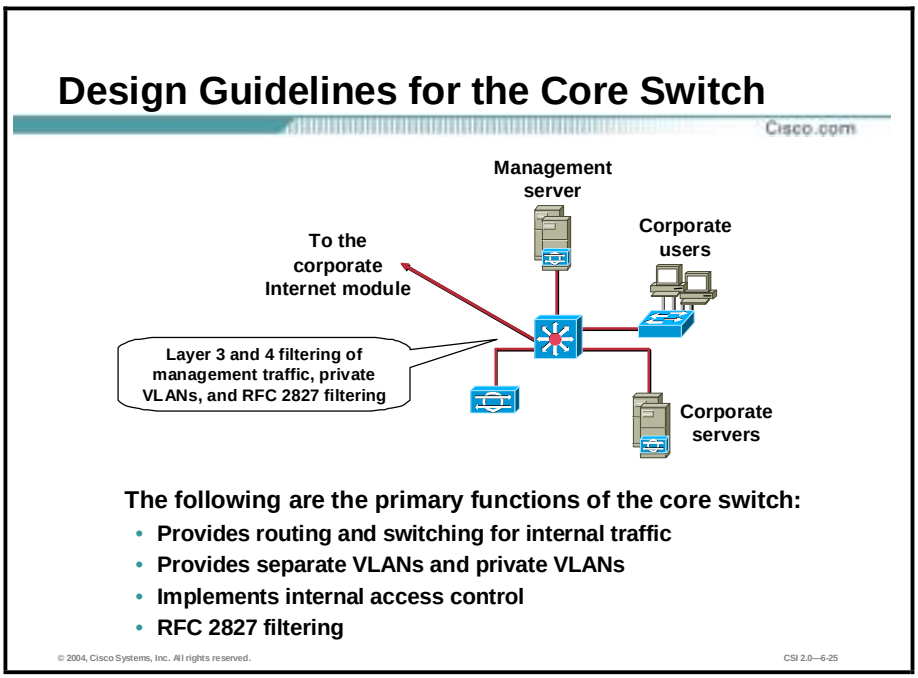

The primary function of the core switch is to provide routing and switching for production and management traffic, distribution layer services (routing, quality of service [QoS], and access control) for the distribution and access switches, connectivity for the corporate and management servers, and advanced services such as traffic filtering between the subnets.

In the figure, a Layer 3 switch is used instead of a Layer 2 switch in order to provide separate VLANs for the following:

- Corporate server segment
- Management server segment
- Corporate user segment
- Connectivity to the WAN module and to the corporate Internet module

The Layer 3 switch provides a line of defense and prevention against internally originated attacks. It can mitigate the chance of a department accessing confidential information on another department's server through the use of access control. For example, a network that contains marketing and research and development (R and D) might segment off the R and D server to a specific VLAN and filter access to it, ensuring that only R and D staffs have access to it. For performance reasons, it is important that this access control be implemented on a hardware platform that can deliver filtered traffic at near wire rates. This setup generally dictates the use of Layer 3 switching, as opposed to more traditional dedicated routing devices. This same access control can also prevent local source-address spoofing through the use of RFC 2827 filtering. RFC 2827 filtering should be implemented on the corporate user and corporate intranet server VLANs.

Within each of the VLANs, private VLANs can be used to mitigate trust-exploitation attacks between the devices. For instance, within the corporate server segment, the individual servers may not have any requirement to communicate with each other. They need to communicate only with devices connected to the corporate user segment.

To provide a further line of defense for the management servers, extensive Layer 3 and Layer 4 filtering is configured outbound on the VLAN interface connecting to the management server segment. The ACL limits connectivity to and from the management servers only to those devices (via IP addresses) under their control, and only for those protocols and services (via port number) that are required. This also includes access control for management traffic destined for the remote site devices. This traffic is encrypted by the firewall and sent to the remote sites. Allowing only established connections back through the ACL further controls access to the managed devices.

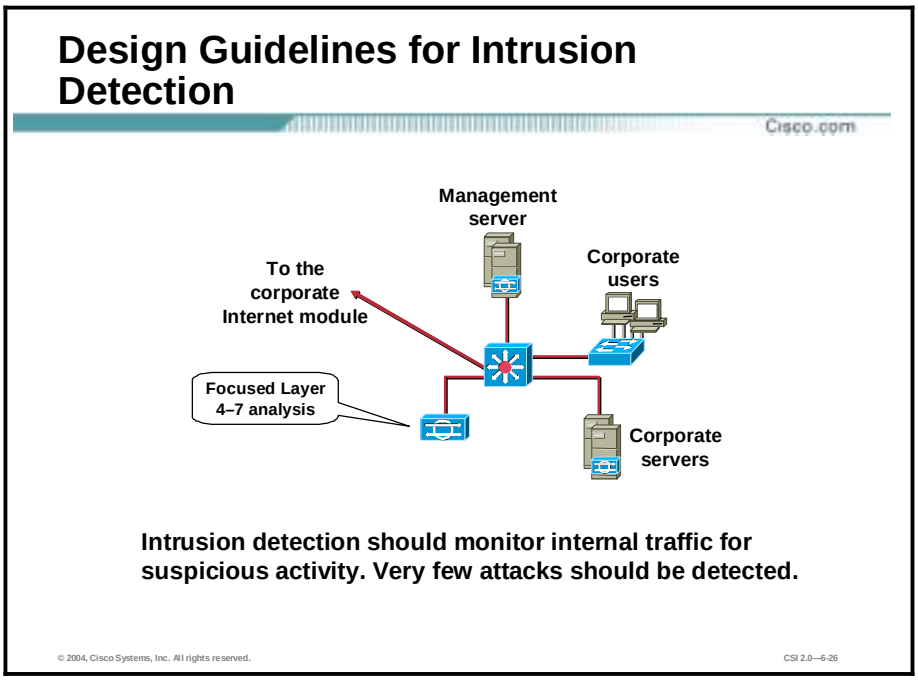

The campus module also includes an NIDS appliance. The switch port that connects to the NIDS appliance is configured such that traffic from all VLANs that require monitoring is mirrored to the monitoring port of the appliance.

Very few attacks should be detected here because this NIDS appliance provides analysis against attacks that may originate from within the campus module itself. For instance, if a user workstation was compromised because of an unknown modem connection to that host, the NIDS could detect suspicious activity originating from within the campus. Other internal attacks could originate from disgruntled employees, workstations left where unauthorized people can gain access to them, or Trojan horse applications inadvertently loaded on portable PCs.

Each of the corporate intranet and management servers also has a HIDS or HIPS installed.

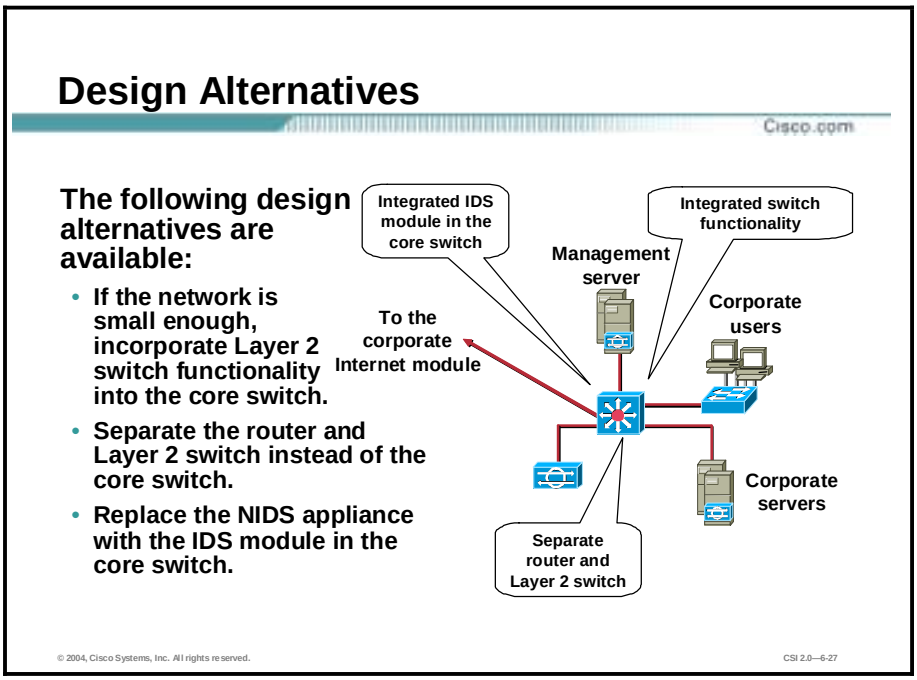

If the midsize network is small enough, the functionality of the distribution and access switches can be rolled into the core switch, and the distribution and access switches can be eliminated. In this case, the end-user workstations would be connected directly to the core switch. Private VLAN functionality would be implemented on the core switch in order to mitigate trust-exploitation attacks.

If the performance requirements of the internal network are not high, a separate router and Layer 2 switch could be used for the core and distribution instead of the higher-performing Layer 3 switch.

If desired, the separate NIDS appliance can be replaced with an integrated IDS that fits into the core switch. This setup provides higher traffic throughput into the IDS because it sits on the backplane of the switch, rather than being connected via a single 10/100-Mbps Ethernet port. ACLs on the switch can be used to control what traffic is sent to the IDS.

Midsize Network WAN Module

This topic explains the midsize network WAN module. The WAN module is included only when connections to remote locations over a private network are required. This requirement may occur when stringent QoS requirements cannot be met by an IPSec VPN, or when legacy WAN connections are in place without a compelling cost justification to migrate to IPSec.

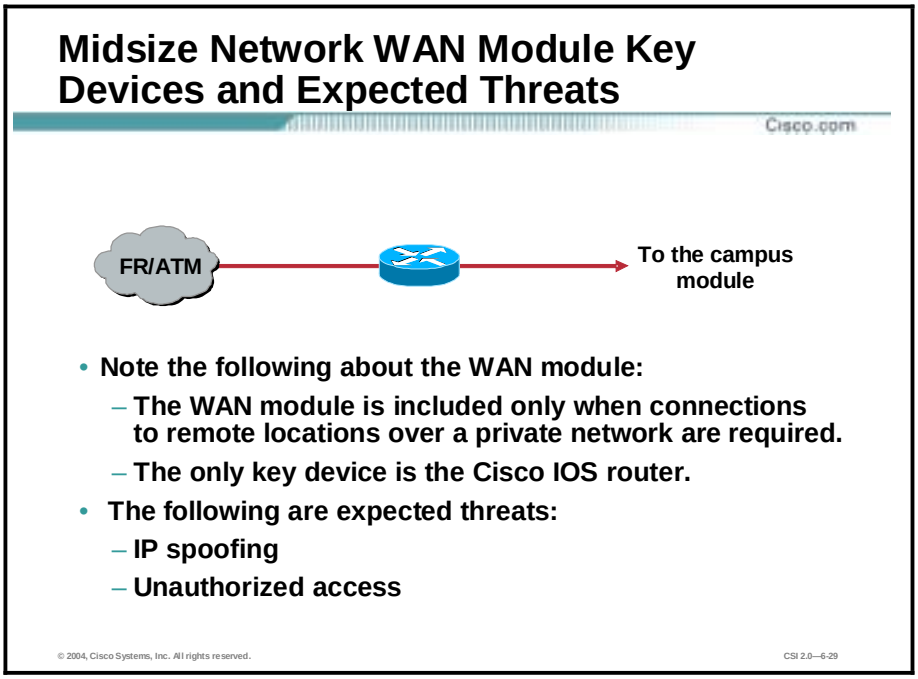

The Cisco IOS router is a key device for the midsize network WAN. It provides routing, access control, and QoS mechanisms to remote locations.

The following threats and mitigation roles can be expected in the midsize network WAN:

- IP spoofing—IP spoofing can be mitigated through Layer 3 filtering.
- Unauthorized access—Simple access control on the router can limit the types of protocols to which branches have access.

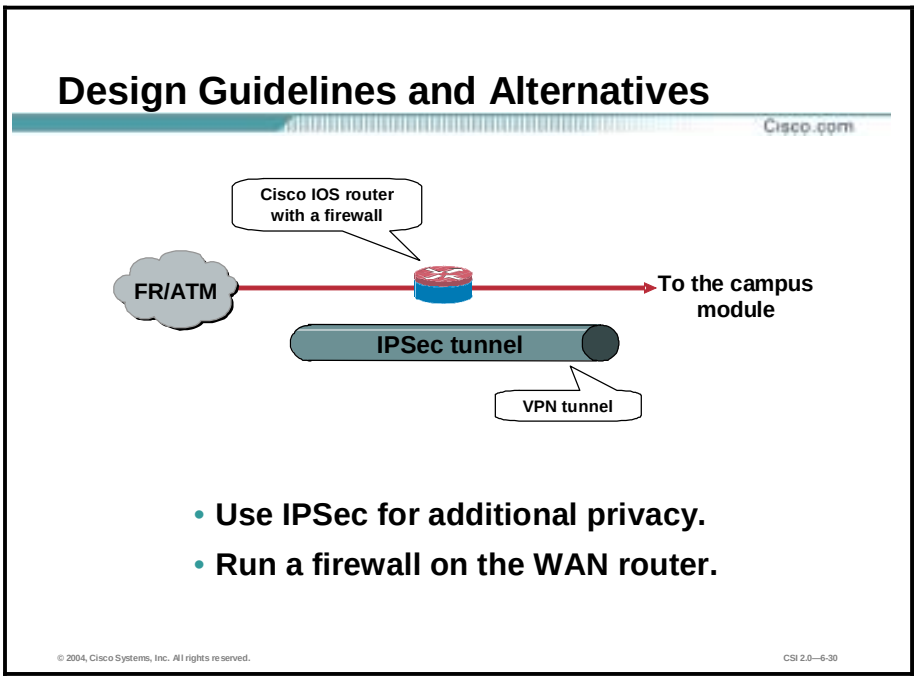

The amount of security placed in the WAN module depends on the level of trust for the remote sites and the ISP to which you are connecting. Security is provided by using Cisco IOS security features. In this design, inbound ACLs applied to the serial interface are used to block all unwanted traffic from accessing the midsize network. Inbound ACLs applied to the Ethernet interface can be used to further limit what traffic passes from the midsize network back to the remote sites.

Alternatives involve organizations that are very concerned about information privacy and encrypt traffic across their classic WAN links. Similar to site-to-site VPNs, IPSec can be used to achieve this level of information privacy. Additionally, running a firewall on the WAN router can provide additional access control options when compared with the basic ACLs used in the SAFE design.

Implementation—ISP Router and Edge Router

This topic covers the necessary steps needed to implement the SAFE guidelines on the ISP router.

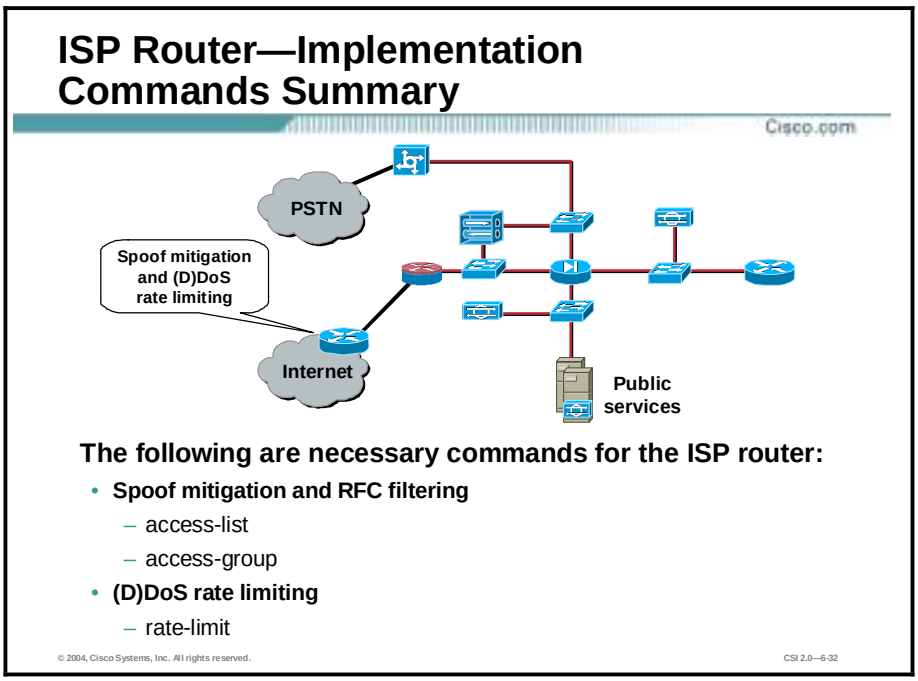

Starting at the customer edge router in the ISP, the egress of the ISP rate limits nonessential traffic that exceeds pre-specified thresholds in order to mitigate DDoS attacks. Also at the egress of the ISP router, RFC 1918 and RFC 2827 filtering mitigates source-address spoofing of local networks and private address ranges.

At the ingress of the firewall, RFC 1918 and RFC 2827 filtering is first provided as a verification of the ISP's filtering. In addition, because of the enormous security threat that fragmented packets create, the firewall is configured to drop most fragmented packets that should not generally be seen for standard traffic types on the Internet. Any legitimate traffic lost because of this filtering is considered acceptable when compared to the risk of allowing such traffic. Traffic destined to the firewall itself from the outside is limited to IPSec traffic and any necessary protocols for routing.

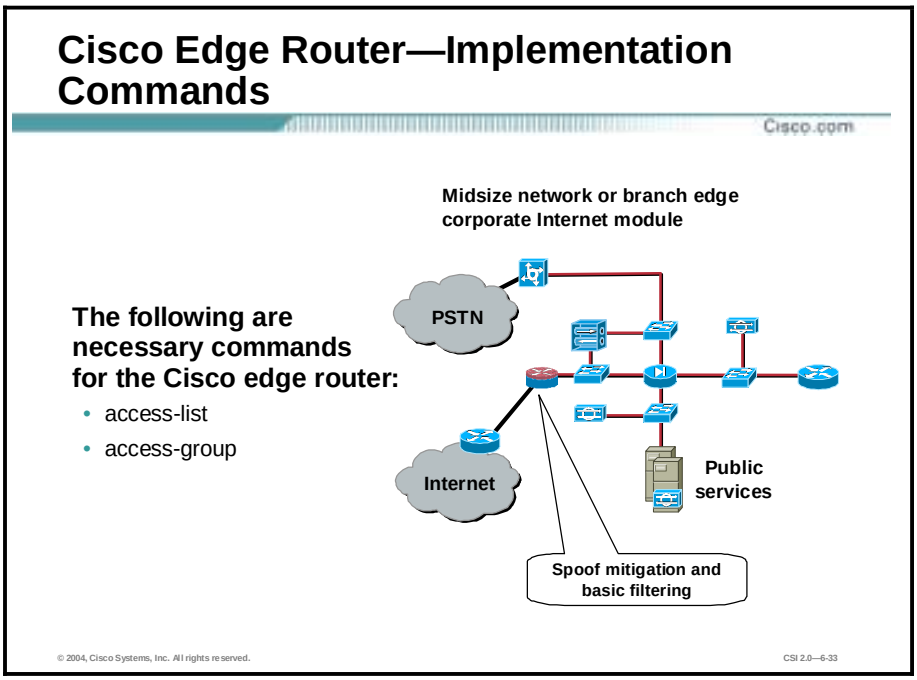

The function of the edge router on the midsize network is to provide the demarcation point between the ISP network and the midsize network. At the ingress of the edge router on the midsize network, basic filtering limits access to allow only expected IP traffic, providing a coarse filter for the most basic attacks. RFC 1918 and RFC 2827 filtering is also provided here as a verification of the ISP's filtering.

In addition, because of the enormous security threat that fragmented packets create, the router is configured to drop most fragmented packets that should not generally be seen for standard traffic types on the Internet. Any legitimate traffic lost because of this filtering is considered acceptable when compared to the risk of allowing such traffic.

Any IPSec traffic destined for the VPN 3000 Concentrator or the firewall is allowed through. Filtering on the router is configured to allow only IKE and IPSec traffic to reach the VPN 3000 Concentrator or firewall. In remote-access VPNs the IP address of the remote system is not generally known, therefore the filtering can be specified only to the head-end peer (VPN 3000 Concentrator) with which the remote users are communicating. With site-to-site VPNs, the IP address of the remote site is usually known; therefore, filtering may be specified for VPN traffic to and from both peers.

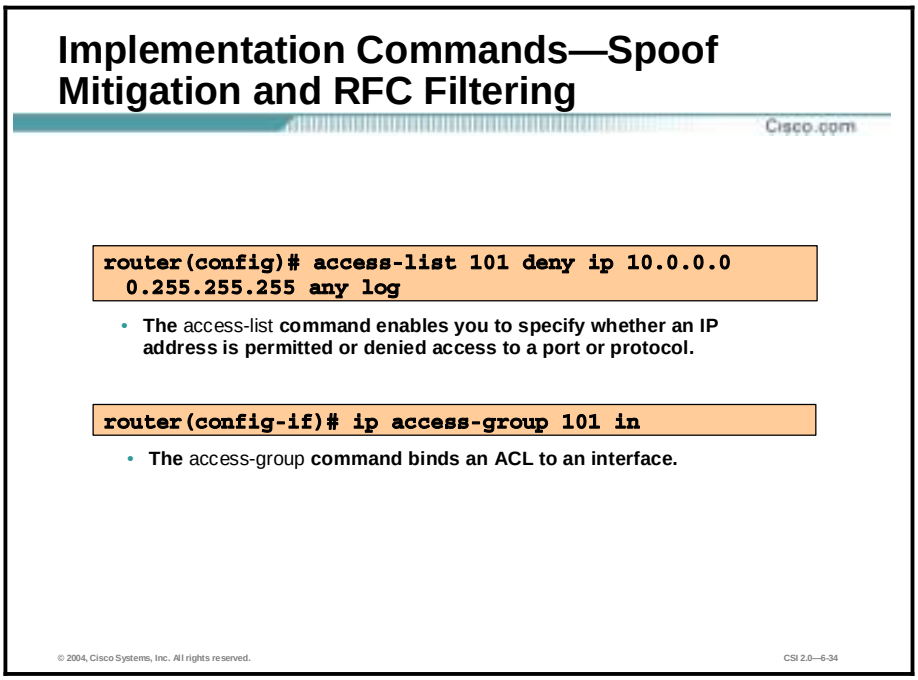

Cisco provides basic traffic filtering capabilities with access control lists (ACLs). ACLs can be configured for all routed network protocols (IP, AppleTalk, and so on) to filter the packets of those protocols as the packets pass through a router.

You can configure ACLs at your router to control access to a network. ACLs can prevent certain traffic from entering or exiting a network.

ACLs should be used in Cisco IOS Firewall routers, which are often positioned between your internal network and an external network (for example, the Internet). You can also use ACLs on a router positioned between two parts of your network, to control traffic entering or exiting a specific part of your internal network.

To provide the security benefits of ACLs, you should at a minimum configure ACLs on border routers—routers situated at the edges of your networks. This provides a basic buffer from the outside network, or from a less controlled area of your own network into a more sensitive area of your network.

On these routers, you should configure ACLs for each network protocol configured on the router interfaces. You can configure ACLs so that inbound traffic, outbound traffic, or both are filtered on an interface.

ACLs must be defined on a per-protocol basis. In other words, you should define ACLs for every protocol enabled on an interface if you want to control traffic flow for that protocol.

For inbound ACLs, after receiving a packet, the Cisco IOS software checks the source address of the packet against the ACL. If the ACL permits the address, the software continues to process the packet. If the ACL rejects the address, the software discards the packet and returns an Internet Control Message Protocol (ICMP) “host unreachable” message.

For outbound ACLs, after receiving and routing a packet to a controlled interface, the software checks the source address of the packet against the ACL. If the ACL permits the address, the

software sends the packet. If the ACL rejects the address, the software discards the packet and returns an ICMP “host unreachable” message.

When you apply an ACL that has not yet been defined to an interface, the software acts as if the ACL has not been applied to the interface and accepts all packets. Remember this behavior if you use undefined ACLs as a means of security in your network.

Implementation—NIDS

This topic explains the Cisco network-based intrusion detection system (NIDS) Sensor interface setup configuration tasks.

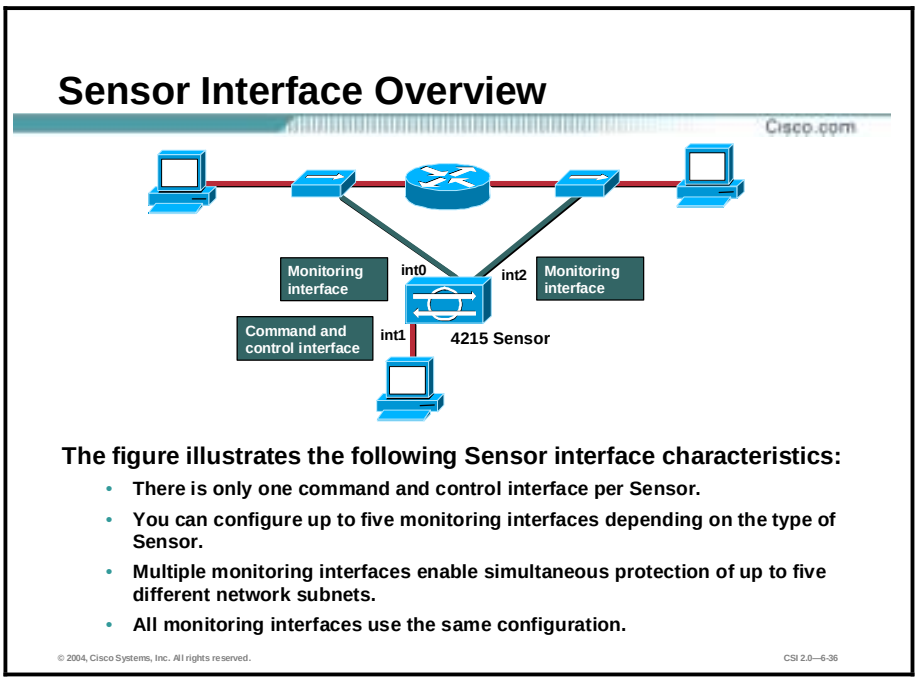

Each Sensor has only one command and control interface, but you can configure up to five monitoring interfaces depending on the type of Sensor you have. Multiple interfaces enable simultaneous protection of up to five different network subnets, which is like having five Sensors in a single appliance.

All monitoring interfaces use the same configuration. There is only one virtual Sensor, so no mapping of virtual Sensor configurations to interfaces is required.

A monitoring interface must be part of Interface Group 0 and must be enabled. Sensors with factory-installed Cisco IDS Version 4.1 are shipped with all monitoring interfaces added to Interface Group 0 and disabled. You must enable the monitoring interfaces in order for the Sensor to monitor your networks. Upgrades from IDS Version 4.0 to 4.1 may leave some interfaces enabled that are not assigned to a group. Either disable these interfaces or add them to Group 0 to prevent inconsistencies in reporting to the Sensor.

You do not need to enable all interfaces. Enable only those interfaces that you want to use.

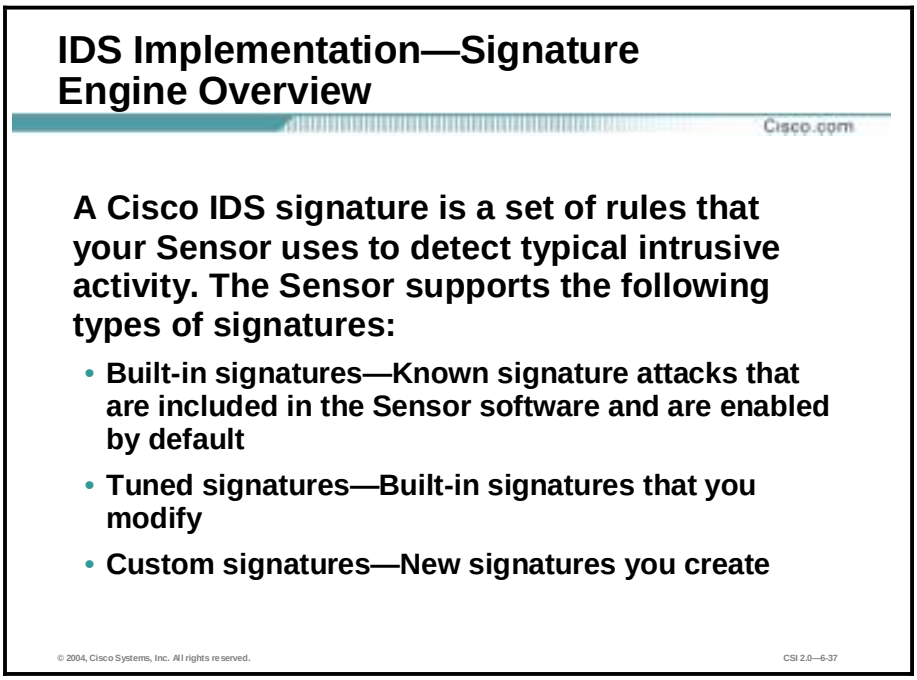

A signature is a set of rules that your Sensor uses to detect typical intrusive activity, such as denial of service (DoS) attacks. As Sensors scan network packets, they use signatures to detect known attacks and respond with actions that you define.

The Sensor compares the list of signatures with network activity. When a match is found, the Sensor logs an event. A Sensor enables you to modify existing signatures and define new ones.

Signature-based intrusion detection can produce false positives because certain normal network activity can be misinterpreted as malicious activity. For example, some network applications or operating systems may send out numerous ICMP messages, which a signature-based detection system might interpret as an attempt by an attacker to map out a network segment. You can minimize false positives by tuning your Sensors.

You must enable the signature to configure a Sensor to monitor network traffic for a particular signature. The most critical signatures are enabled by default. When an attack is detected that matches an enabled signature, the Sensor generates an alert event and stores it in the Event Store. The alert events, as well as other events, may be retrieved from the Event Store by web-based clients. The Sensor logs all informational alarms or higher alarms by default.

Some signatures have subsignatures. This means that the signature is divided into subcategories. When you configure a subsignature, changes made to the parameters of one subsignature apply only to that subsignature.

Built-in signatures are known attack signatures that are included in the Sensor software and are enabled by default. You cannot add to or delete from the list of built-in attack signatures. You also cannot rename them. You can tune built-in signatures by adjusting several signature parameters. Built-in signatures that have been modified are called tuned signatures. You can also create new signatures, which are called custom signatures.

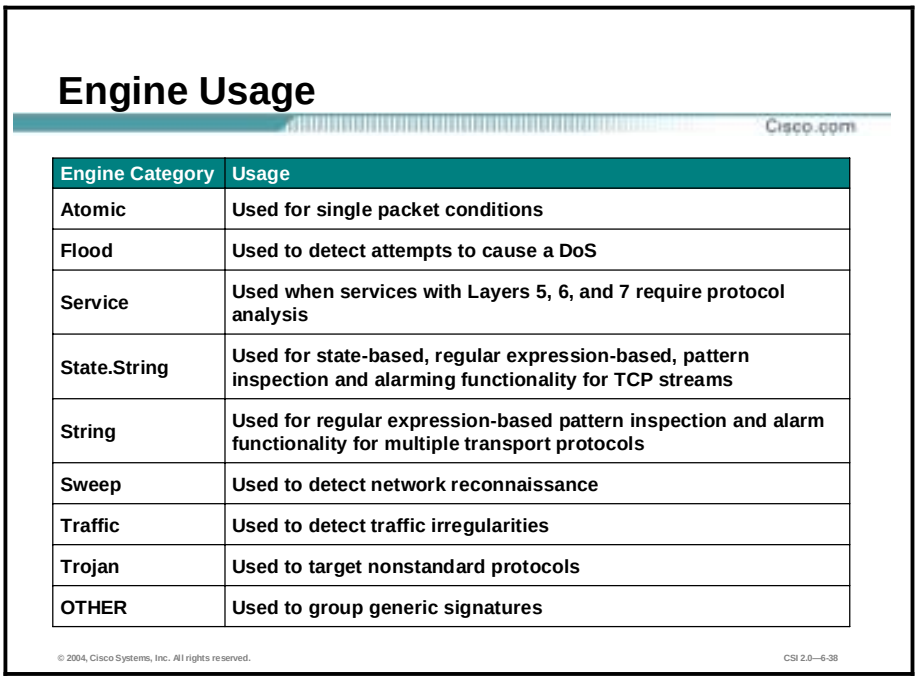

The following are the general categories of Cisco IDS signature engines:

- Atomic—Used to perform per-packet inspection, supporting signatures that trigger based on the analysis of a single packet
- Flood—Used to detect attempts to cause a DoS
- Service—Used when services with Layers 5, 6, and 7 require protocol analysis
- State.String—Used for state-based, regular expression-based, pattern inspection and alarming functionality for TCP streams
- String—Used for regular expression-based pattern inspection and alarm functionality for multiple transport protocols including TCP, UDP, and ICMP
- Sweep—Used to detect network reconnaissance
- Traffic—Identifies traffic irregularities
- Trojan—Used to detect BackOrifice Trojan traffic and Tribal Flood Network 2000 (TFN2K) Trojan or DDoS traffic
- OTHER—Used to group generic signatures so common parameters may be changed

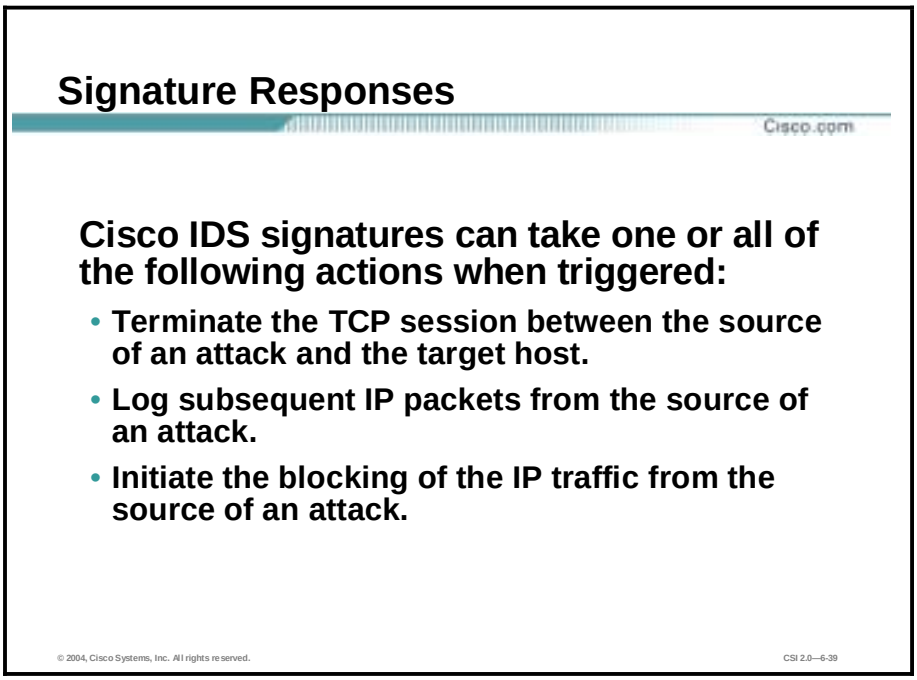

Cisco IDS signatures can take one or all of the following actions when triggered:

- TCP reset—Terminates the TCP session between the source of an attack and the target host
- IP log—Logs subsequent IP packets from the source of an attack
- Block—Initiates the blocking of IP traffic from the source of an attack, either a block on the host or the connection

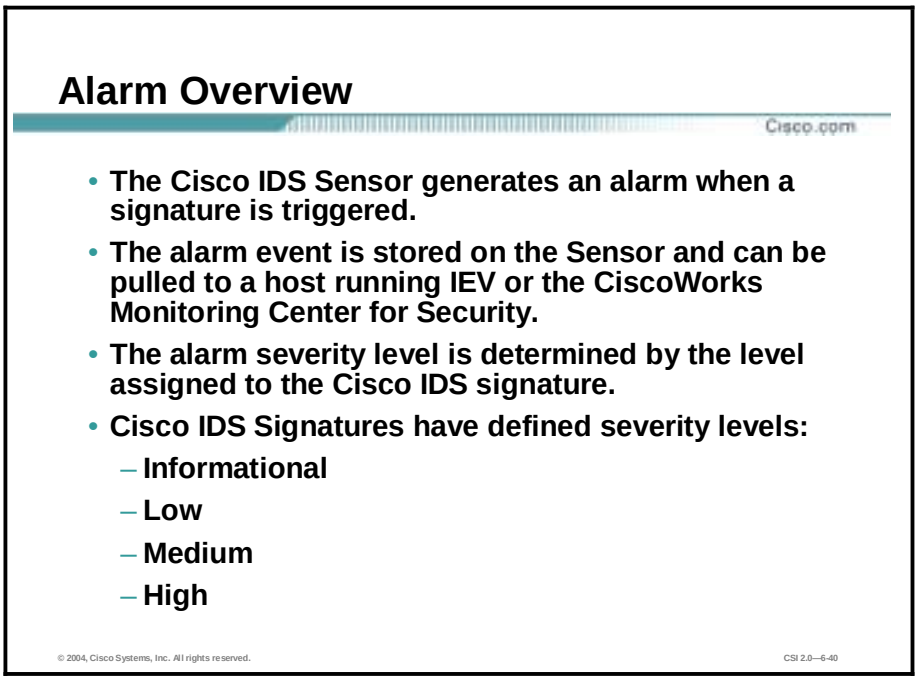

The following information is an overview of alarms for Cisco IDS Sensors:

- The Sensor generates an alarm when an enabled signature is triggered.
- Alarms are stored on the Sensor and a host can pull the alarms off of the Sensor. Pulling alarms from a Sensor allows multiple hosts to subscribe to the event "feed." This allows a host or hosts to subscribe on an as-needed basis.
- The level assigned to the signature determines the alarm severity level. When tuning a signature, you may assign a severity level to the signature which in turn will make the alarm the same severity level as that of the signature.
- A Cisco IDS signature can have one of the following severity levels:
  - Informational—The activity that triggered the signature is not considered an immediate threat, but does provide useful information.
  - Low—Abnormal network activity was detected that could be perceived as malicious, but an immediate threat is not likely.
  - Medium—Abnormal network activity was detected that could be perceived as malicious, and an immediate threat is likely.
  - High—Attacks used to gain access or cause a DoS were detected, and an immediate threat is extremely likely.

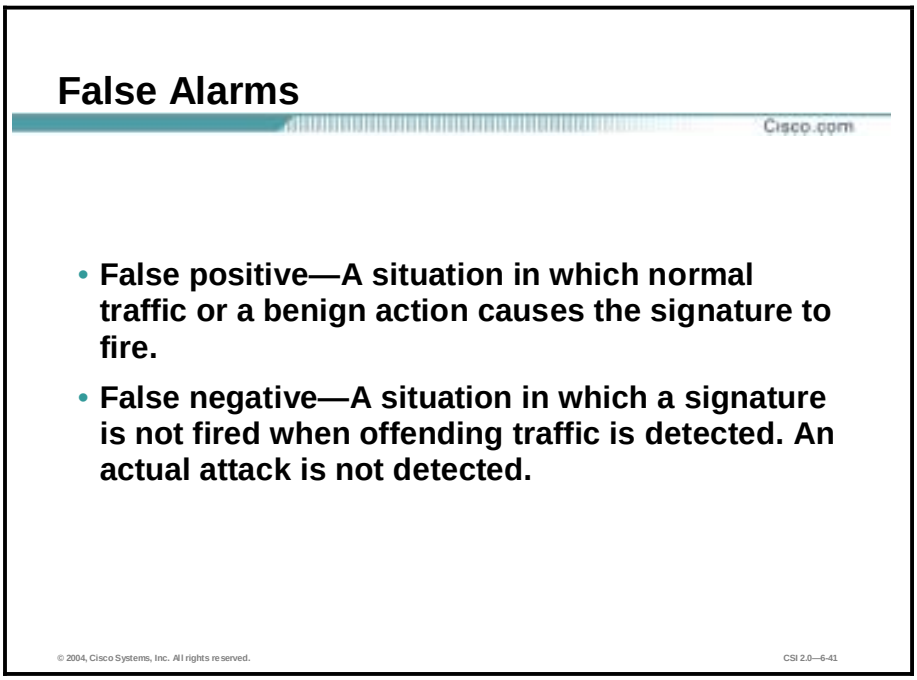

The ability of an intrusion detection product to accurately detect an attack or a policy violation and generate an alarm is critical to its functionality. The two forms of false alarms are false positives and false negatives.

A false positive is a situation in which normal traffic or a benign action causes the signature to fire. Consider the following scenario: a signature exists that generates alarms if the enable password of any network device is entered incorrectly. A network administrator attempts to log in to a Cisco router but enters the wrong password. The IDS cannot distinguish between a rogue user and the network administrator, and it generates an alarm.

A false negative is a situation in which a signature is not fired when offending traffic is detected. Offending traffic can be as simple as someone sending confidential documents outside of the corporate network or as complex as an attack against corporate web servers. False negatives should be considered software bugs and reported in accordance with the software license agreement.

Note	A false-negative should only be considered a software bug if in fact the IDS has a signature that has been designed to detect the offending traffic.

Note	Intelligent Threat Investigation—Cisco Threat Response (CTR) technology works with Cisco Intrusion Detection System (IDS) Sensors to provide an efficient intrusion protection solution. Threat Response technology virtually eliminates false alarms, escalates real attacks, and aids in the remediation of costly intrusions. Unlike other intrusion-management solutions, only Threat Response technology provides an automated, just-in-time analysis of each targeted host to determine whether a compromise has occurred. Only by investigating the host under attack can you efficiently uncover the real intrusions and address them quickly. Threat Response technology's automated, real-time capabilities help protect your network environment around the clock. The result is that false alarms are eliminated and real intrusions are quickly identified and addressed, saving time, resources, and the high costs associated with recovering from a successful attack.

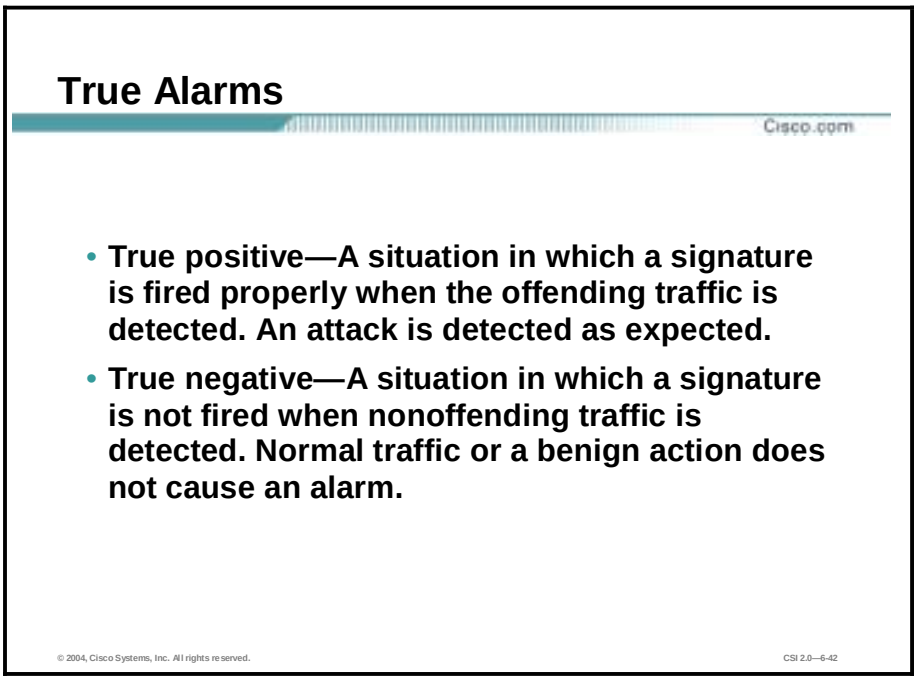

Like false alarms, there are two forms of true alarms. A true positive is a situation in which a signature is fired properly when offending traffic is detected and an alarm is generated. For example, Cisco IDS Sensors have signatures that detect Unicode attacks against Microsoft Internet Information Server (IIS) web servers. If a Unicode attack is launched against Microsoft IIS web servers, the Sensors detect the attack and generate an alarm.

A true negative is a situation in which a signature is not fired when non-offending traffic is captured and analyzed. In other words, the Sensor does not fire an alarm when it captures and analyzes “normal” network traffic.

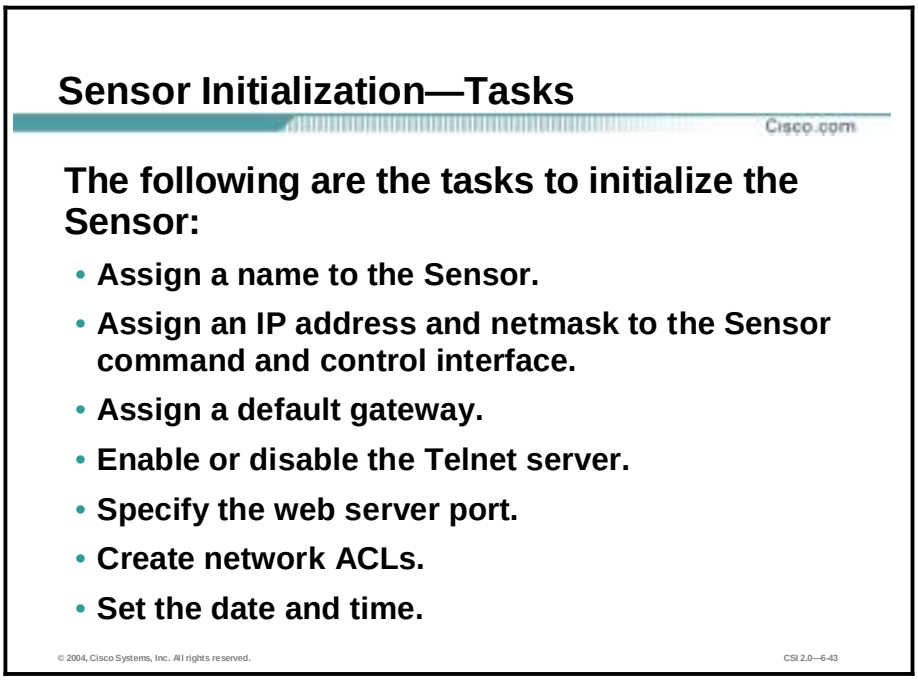

The following Sensor initialization tasks are done via an interactive dialog initiated by the **setup** command:

- Assign the Sensor a hostname.
- Assign an IP address and a subnet mask to the command and control interface.
- Assign a default route.
- Enable or disable the Telnet server.
- Specify the web server port.
- Add and remove ACL entries that specify which hosts are allowed to connect to the Sensor.
- Set the date and time.

Note	If you later change the Sensor's IP address, you will need to generate a self-signed X.509 certificate. This certificate is needed by HTTP secure (HTTPS) communications.

Sensor Initialization—*setup* Command

Cisco.com

Most of the initialization tasks are accomplished by using the Sensor's **setup** command, which helps you to configure the hostname, IP address, netmask, gateway, and communications options. After you enter the **setup** command, the default settings are displayed. Press the **spacebar** to continue. The following question appears: "Continue with configuration dialog?" Answer "yes."

Sensor Management and Monitoring—IDS Device Manager

Cisco.com

- **Web-based device configuration tool**
- **Software installed on the Sensor by default**
- **For small-scale Sensor deployments**

A network Sensor appliance can be managed via the IDS Device Manager (IDM). IDM is a web-based tool that resides on your Sensor and enables you to configure and manage the Sensor. IDM is accessed securely via Secure Sockets Layer (SSL) and Transport Layer Security (TLS) using the Netscape or Internet Explorer web browsers. It is best suited for small-scale Sensor deployments where there are no more than five Sensors.

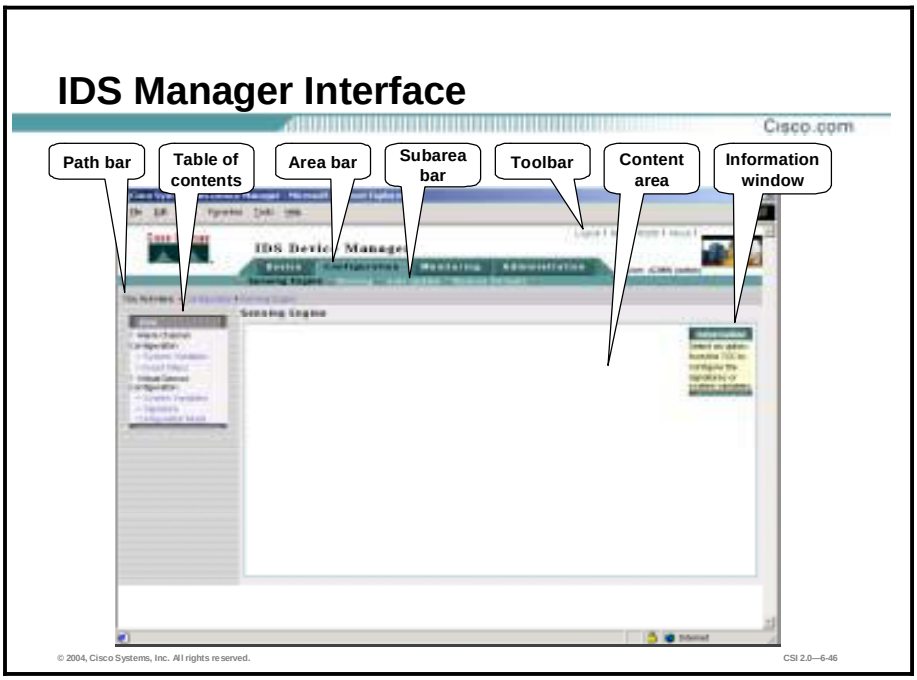

The IDM GUI provides the network security administrator with an intuitive approach to configuring Sensors. The GUI has the following sections:

- Path bar—Displays the current selection. In the figure, the path selected is Configuration>Sensing Engine.
- Table of contents (TOC)—Lists the available options for the item selected from the subarea bar. In the figure, the TOC displays the options for the Sensor engine.
- Area bar—Lists the available Sensor configuration items. The available Sensor configuration items are Device, Configuration, Monitoring, and Administration. Each configuration item has suboptions, which are listed in the subarea bar.
- Subarea bar—Lists the available Sensor configuration suboptions for the item selected from the area bar. In the figure, the available configuration options are Sensing Engine, Blocking, Auto Update, and Restore Defaults.
- Toolbar—Lists the available user functions. The available user functions are Logout, Help, NSDB, and About.
- Content area—Displays the information associated with the option selected or an action associated with a user function.
- Information window—Displays a description associated with the option selected or with the instructions.

IEV

Cisco.com

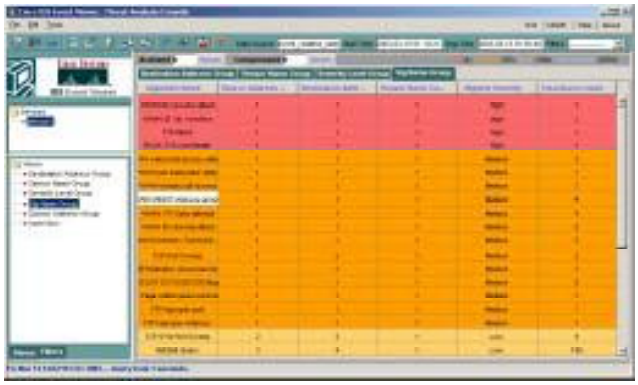

- Windows NT or Windows 2000
- Download from Cisco.com
- Provides event monitoring for up to five Sensors

IEV is a Java-based application that enables you to view and manage alarms for up to five Sensors. You can use the IEV to view alarms in real time or in imported log files. Features and benefits of IEV are as follows:

- The initial view provides an aggregate view of alarm data.
- Views are grouped by signature name, source address, destination address, Sensor identity, and severity levels.
- Each view may have different data sources.
- The level of alarm detail is customizable.
- A graph view displays alarm data in either an area format or a bar graph format.
- The application is downloadable from Cisco.com to an appropriate host.
- IEV provides event monitoring for IDS devices.
- IEV provides a scalable event storage database.

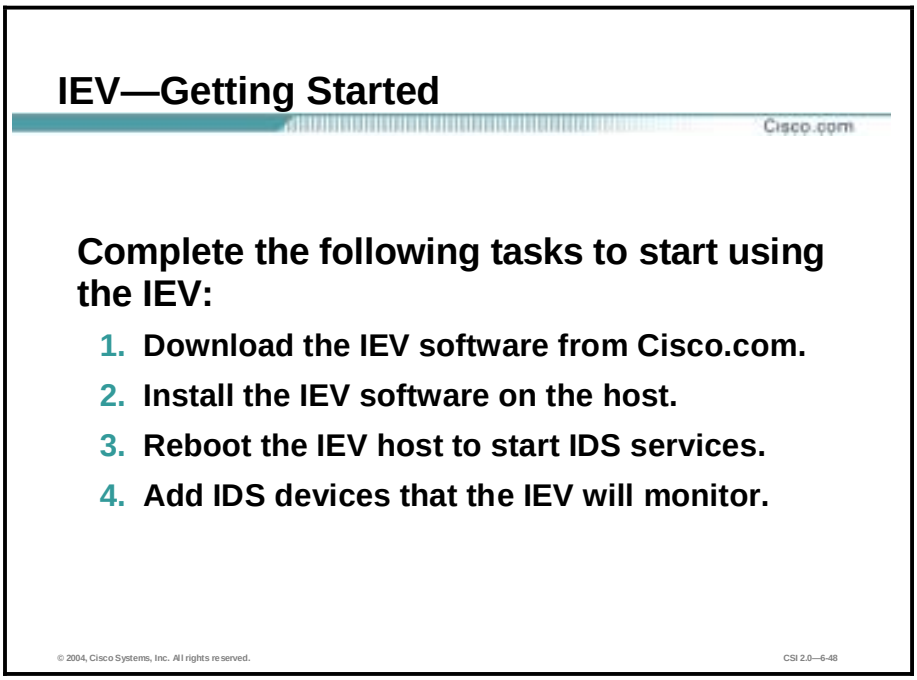

Complete the following tasks to begin using IEV to monitor events from an IDS device:

Step 1 Download the IEV software from Cisco.com.

Step 2 Install the IEV software on the host. This step includes starting the IEV setup program and continuing with the installation wizard.

Step 3 Reboot the IEV host to start the IDS services. This step includes rebooting the IEV host in order to initialize the IDS services needed by IEV.

Step 4 Add IDS devices that the IEV is to monitor. This step includes specifying the IDS devices from which the IEV application accepts events.

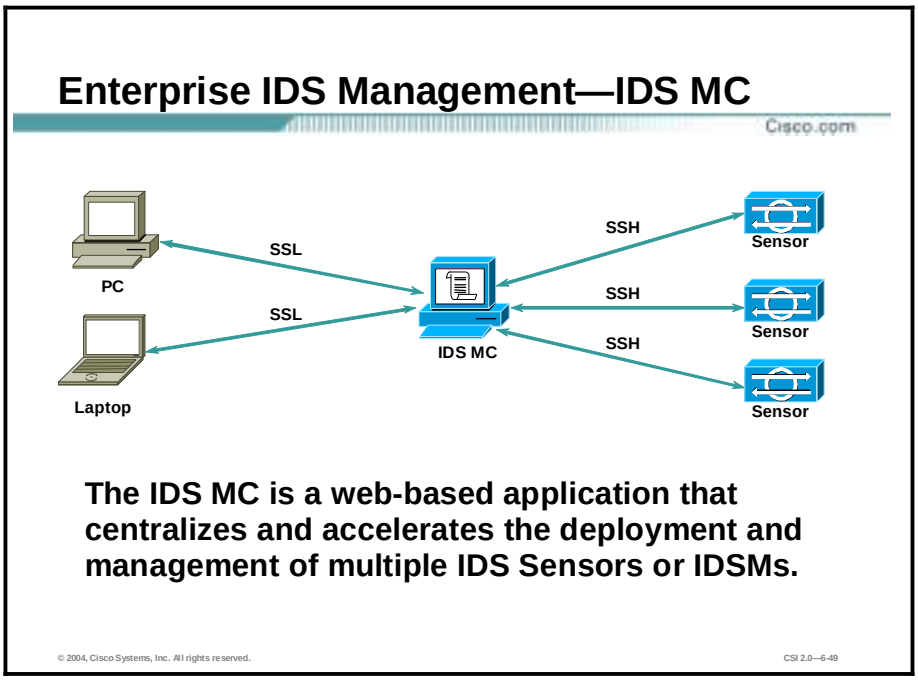

Management Center for IDS Sensors (IDS MC) is a component of the VPN/Security Management Solution (VMS) bundle. The VMS bundle integrates the CiscoWorks Server with a number of individual applications such as VPN Monitor, MCs, and the CiscoWorks Monitoring Center for Security to provide a comprehensive suite of security management tools. Through the CiscoWorks Common Services component, the IDS MC provides a web-based interface for configuring and managing a Sensor or Sensor group. The slide above shows the IDS MC managing multiple Sensors. The Sensor can be accessed and managed either of the following ways:

- You can manage the Sensor directly from the server on which the IDS MC resides. The figure shows that communications between the IDS MC server and the Sensor occur over SSH.
- You can manage the Sensor from a web browser on a client. If a client is used, communications between the client and the IDS MC occur over SSL.

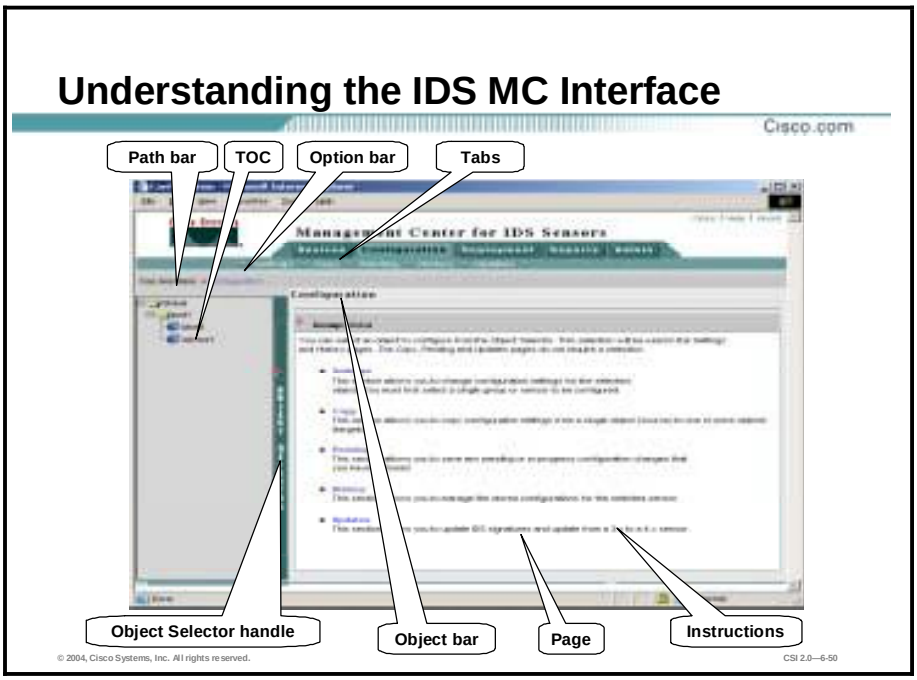

The figure illustrates elements of the IDS MC GUI. The elements are described as follows:

- Path bar—Provides a context for the displayed page, showing tabs, options, and the current page
- TOC (table of contents)—Displays the available suboptions, if required
- Option bar—Displays the options available for the selected tab
- Tabs—Provide access to product functionality
  - Devices tab—Enables you to perform initial setup of devices to be managed by the system
  - Configuration tab—Enables you to perform general configuration tasks
  - Deployment tab—Enables you to generate configuration files, manage Sensor configuration files, and submit or manage new jobs
  - Reports tab—Enables you to generate reports, view scheduled reports, and view reports
  - Admin tab—Enables you to administer system settings
- Instructions—Provides a brief overview of how to use the page
- Page—Displays the area in which you perform application tasks
- Object bar—Displays the object or objects selected in the Object Selector
- Object Selector handle—Opens and closes the Object Selector, which contains devices and device groups from which to select

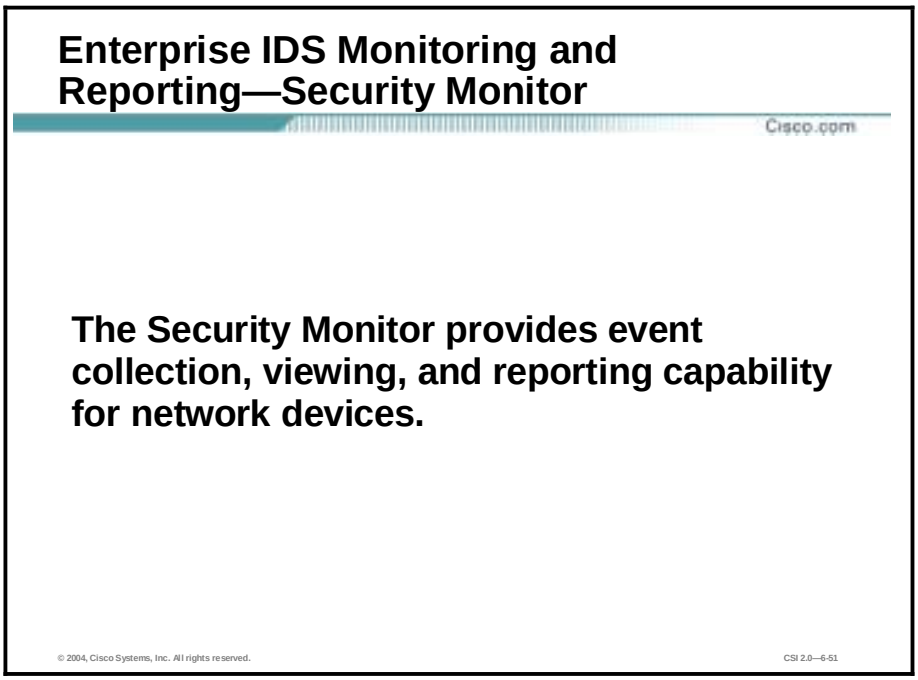

The CiscoWorks Monitoring Center for Security is a component of the Virtual Private Network (VPN)/Security Management Solution (VMS) product. It provides event collection, viewing, and reporting capability for network devices. The VMS product integrates numerous security applications into a single solution, such as the following:

- CiscoWorks Management Center for Firewalls
- CiscoWorks Management Center for IDS Sensors
- CiscoWorks Management Center for Cisco Security Agents
- CiscoWorks Management Center for VPN Routers
- CiscoWorks Monitoring Center for Security
- CiscoWorks Monitoring Center for Performance
- CiscoWorks VPN Monitor
- CiscoWorks Resource Manager Essentials

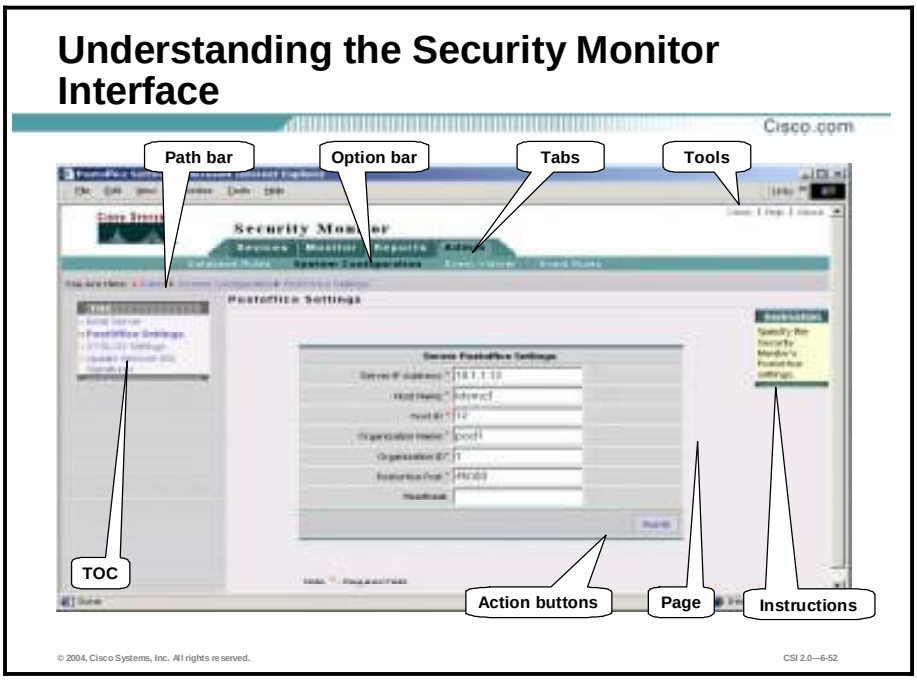

The figure illustrates elements of the Security Monitor GUI. The elements are described as follows:

- Path bar—Provides the location of the current page
- TOC—A menu of choices that is displayed down the left side of the Security Monitor interface, representing the list of suboptions that you can select (based on the option chosen).
- Option bar—Displays the options available for the selected tab
- Configuration tabs—Provides access to product functionality
  - Devices tab—Enables you to perform initial setup of devices to be monitored by Security Monitor
  - Monitor tab—Enables you to monitor information about your devices and launch the Event Viewer
  - Reports tab—Enables you to generate reports, view scheduled reports, and view reports
  - Admin tab—Enables you to administer system and database settings
- Tools—Contains the Close/Logout, Help, and About buttons. The Close/Logout option enables you to close the Security Monitor program. The Help option displays Security Monitor's help information in a separate browser window. Finally, the About option displays the Security Monitor software version.
- Instructions—Provides a brief overview of how to use the page. This information is a quick summary of information provided through the Help option on the Tools bar.

- Page—Displays the area in which you complete application tasks
- Action buttons—Initiates actions or commands for this page. Buttons that do not work on a particular page are grayed out.

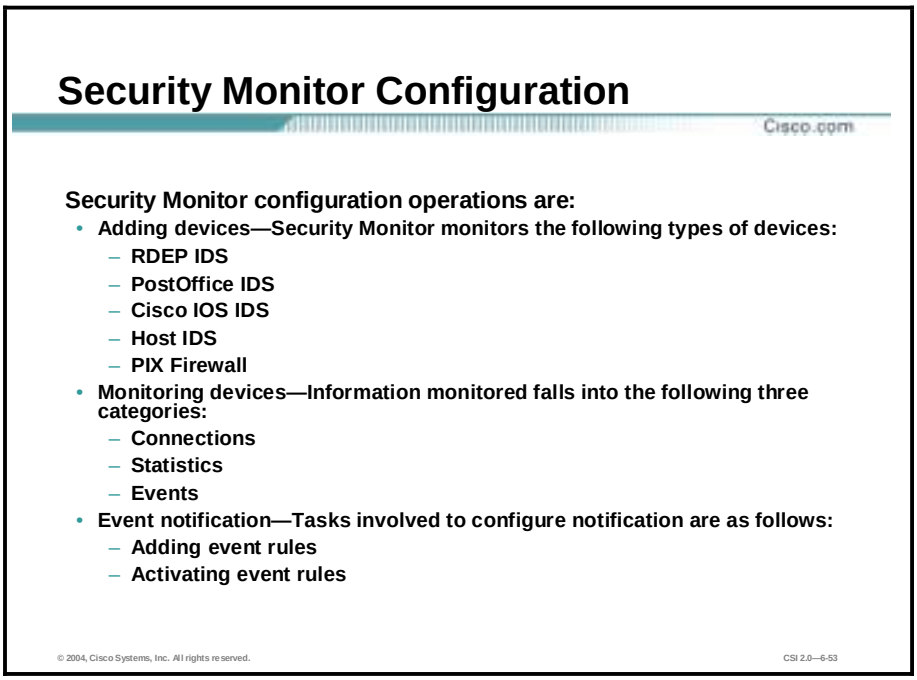

Before you can use Security Monitor to analyze the events from your IDS devices, you must add them to Security Monitor. You can configure the rules that Security Monitor uses to access alerts from the different devices being monitored. For Remote Data Exchange Protocol (RDEP) devices, you can also monitor connection and statistical information. The Security Monitor configuration operations include the following:

- Adding Devices
- Monitoring Devices
- Event Notification

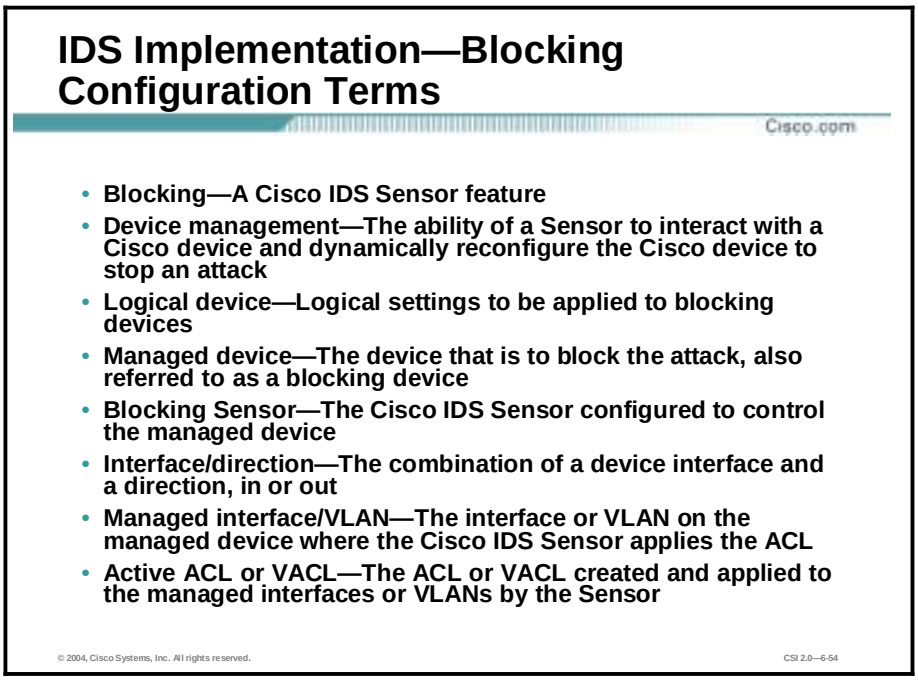

The following terms are used pertaining to the Cisco IDS blocking feature:

- Blocking—A Cisco IDS feature
- Device management—The ability of a Sensor to interact with a Cisco device and dynamically reconfigure the Cisco device to block the source of an attack in real time
- Logical device—Logical settings to be applied to blocking devices
- Managed device—The Cisco device that actually blocks an attack, also referred to as a blocking device
- Blocking Sensor—A Sensor that has been configured to control a managed device
- Interface/direction (ACLs only)—The combination of a device's interface and a direction, in or out, which specifies the blocking of inbound or outbound packets on a particular interface. Blocking is configured separately for each device's interface/direction. The Sensor can be configured to block a total of ten interface/directions across all devices.
- Managed interface/VLAN—The interface or VLAN on the managed device where the Sensor applies the dynamically created ACL or VACL. This interface or VLAN is also referred to as a blocking interface or VLAN.

Note	The Cisco PIX Firewall uses the **shun** command to enforce a block. The PIX Firewall ACLs are not modified.

- Active ACL or VACL—The ACL or VACL that is dynamically created and maintained by the Sensor, which is applied to the managed interface or VLAN.

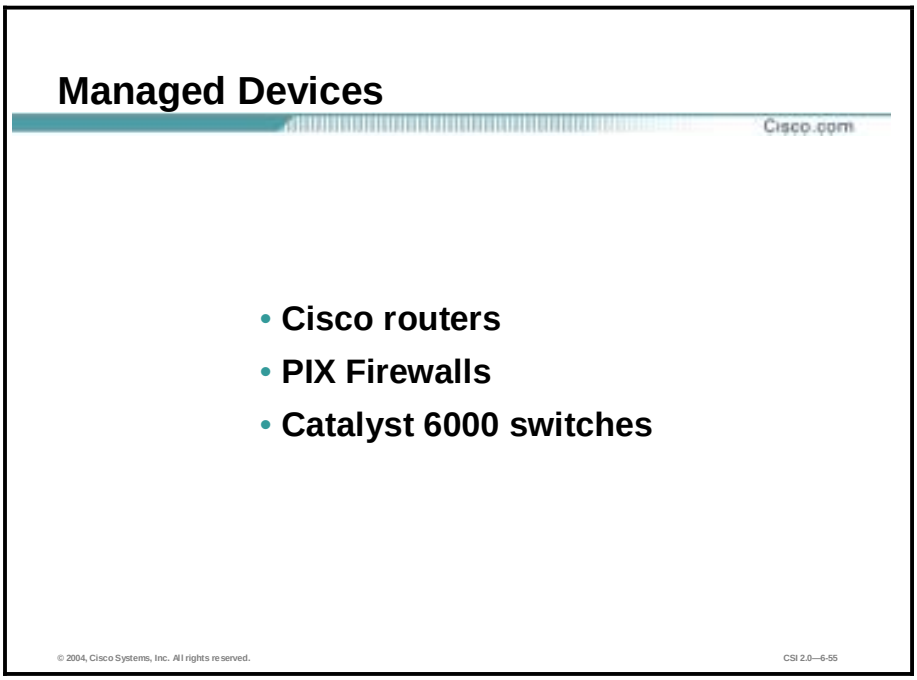

The Sensor network access controller (NAC) service can control up to ten supported devices in any combination. The figure shows a list of blocking devices that have been approved and tested to work with the Sensors and device management, as follows:

- Cisco routers running Cisco IOS Release 11.2 or later using ACLs
- PIX Firewall running version 6.0 or later using the **shun** command—You must use one of the following models:
  - 501
  - 506E
  - 515E
  - 525
  - 535
- Catalyst 6000 switches—Hardware and software requirements for Catalyst 6000 switch blocking devices vary, depending on the following:
  - Switch operating system—Cisco IOS on the Supervisor/Multilayer Switch Feature Card (MSFC [Native]) or CatOS on the Supervisor with Cisco IOS on the MSFC (Hybrid)
  - Sensor type—Appliance or Intrusion Detection System Module 2 (IDSM2)
  - Your choice of blocking method—ACLs or VACLs

Blocking is configured using either ACLs, VACLs, or the **shun** command. All PIX Firewall models that support the **shun** command can be used as blocking devices. The **shun** command was introduced in PIX Firewall operating system 6.0.

The Sensor must be able to communicate with the managed device. The Sensor must have a route to or exist on the same subnet as the managed device. The managed device must be configured to permit Telnet or SSH access from the Sensor.

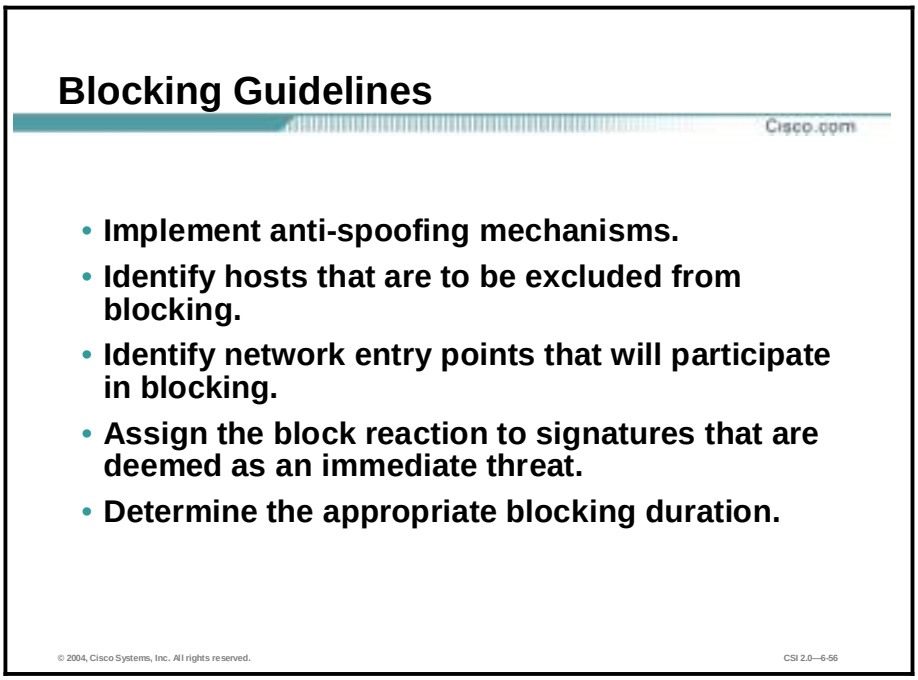

Cisco IDS blocking is a powerful feature that should only be used after thorough planning. The automatic blocking feature generates blocking rules, ACLs, VACLs, and **shun** commands, based solely on the IP addresses of the hosts that generate the alarms. The Sensor cannot determine whether or not the attacking host should be considered a friend or foe. Consequently, it is quite possible that the blocking feature may block legitimate network traffic. The key points to remember when designing and implementing blocking are as follows:

- Anti-spoofing mechanisms—Attackers will forge packets with IP addresses that are either private addresses (RFC 1918) or addresses of your internal network. The attacker's goal may be to elude detection, to gain privileged access through the use of a trusted address, or to cause a DoS if Sensor blocking is configured. By implementing a proper anti-spoofing mechanism and network ingress and egress filtering (RFC 2827), the Sensor will not block possible valid addresses.
- Critical hosts—Each network has critical hosts that should not be blocked. It is important to identify these hosts to prevent possible network disruptions.
- Network topology—Determines which devices should be blocked by which Sensor. Two Sensors cannot control blocking on the same device.
- Entry points—Today's networks have several entry points to provide for reliability, redundancy, and resilience. These entry points are different avenues for the attacker to attack your network. It is important to identify all entry points and decide if the connecting devices should participate in blocking.
- Signature selection—Cisco IDS contains several hundred signatures that can be configured to perform blocking. It is not feasible to perform blocking on all signatures. Identify which signatures are best suited to perform blocking. For example, if you were only allowing web traffic to your server farm, you would identify web-related signatures specific to your web server software. From this list of signatures, you would then identify those signatures whose severity is high and could potentially lead to access. These signatures would be candidates to perform blocking.

- Blocking duration—By default the Sensor will automatically block for 30 minutes. Determine the appropriate time for your network environment.
- Device login information—Before configuring blocking, you must determine any username, password, modal passwords, and connection types needed to log into each blocking device.
- Interface ACL requirements—Each interface/direction can only have one active ACL. Therefore, if an interface needs other ACL entries besides the blocking ACL entries generated by the Sensor, these entries should be configured on the blocking device in the form of a pre-block and post-block ACL. The pre-block and post-block ACLs must be configured on the blocking device independently of the Sensor. These ACLs provide a way to include access rules that a network administrator needs processed before and after the blocking rules are added by the Sensor. When the Sensor NAC service generates an ACL for a device, the NAC first includes all the entries from the pre-block ACL. The NAC then appends its own blocking entries. Finally, the NAC appends the post-block ACL entries to the new ACL of the device. The dynamically created ACL is applied to the specified interface with the specified direction, in or out. When blocking is not in effect, the resulting ACL applied to the interface is simply a combination of the pre- and post-block ACL without any blocking entries inserted.

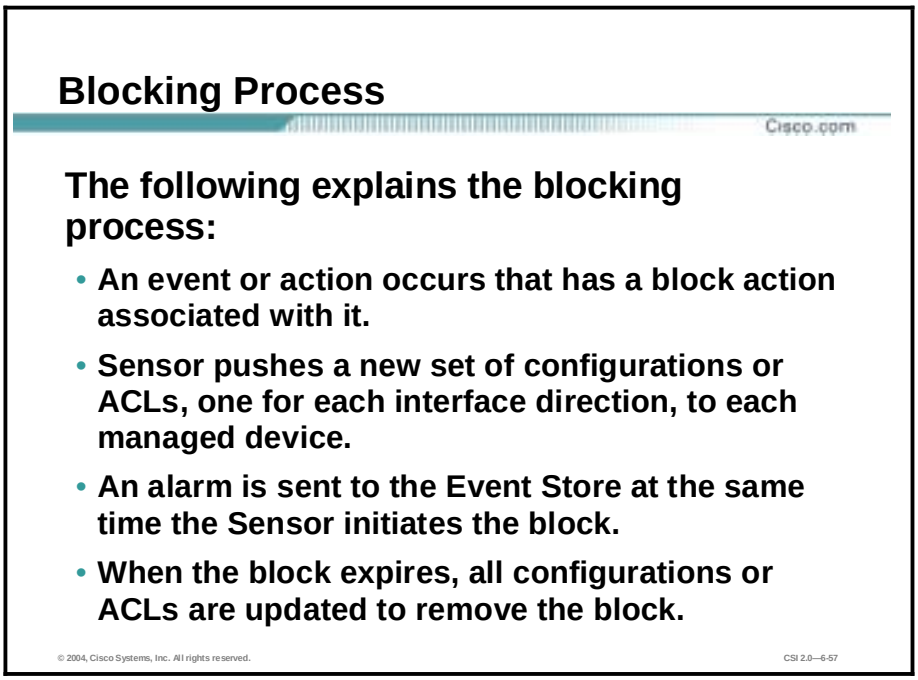

The following explains the blocking process:

- An event or action occurs that has a block action associated with it.

Note	If the NAC is configured to permanently block a specific device, the NAC initiates either a Telnet or an SSH connection with the device and maintains the connection with the device.

- The NAC pushes a new set of configurations or ACLs, one for each interface/direction, to each controlled device. It applies the blocking rules to all configured interface/directions on all devices it is configured to control.
- An alarm is sent to the Event Store at the same time the Sensor initiates the block. The block and alarm occur independently of each other.
- When the blocking event completes, all configurations or ACLs are updated to remove the Sensor's blocking rules.

A time limit can be specified for any manual block except for a permanent block, which is in effect as long as it is configured. The duration of automatic blocks is set globally for all signatures; the default is 30 minutes.

The number of blocking entries that can be active at any given time is configurable with a default limit of 100. The number of blocking entries is not the same as the number of interface/directions. The number of interface/directions corresponds to the number of ACLs that the NAC has to update when a block state changes. The number of blocking entries corresponds to the number of entries in each ACL. The blocking entry in the ACL specifies a host address or network address to be blocked.

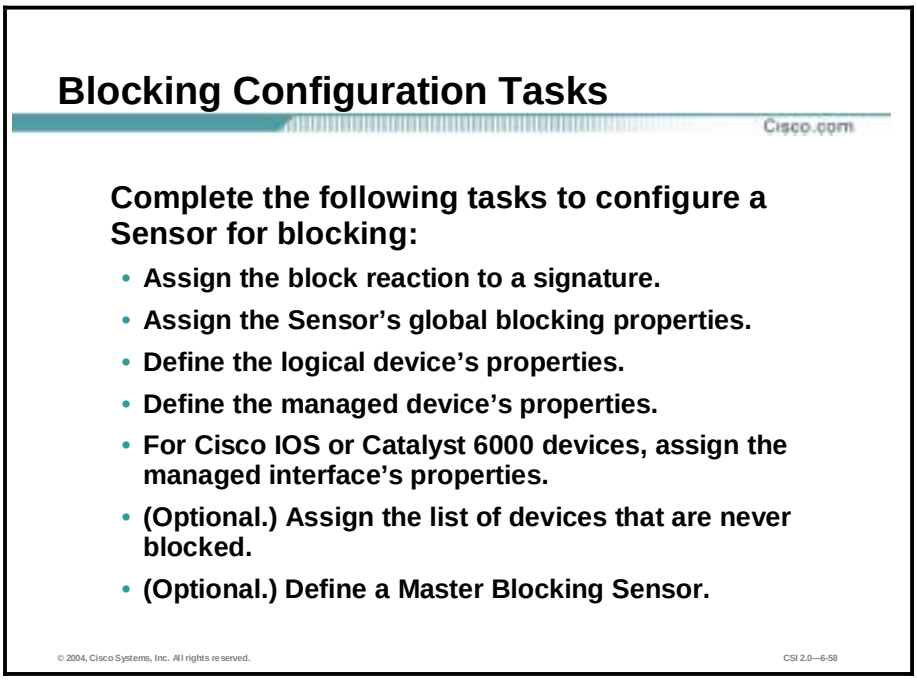

The following are the configuration tasks to configure a Sensor for blocking:

- Assign the block reaction to a selected signature—This task involves using the IDS MC or IDM to configure a signature's action to block.
- Assign the Sensor's global blocking properties—This task involves enabling blocking and defining blocking parameters, such as the block duration, maximum blocking entries, and whether or not to allow the Sensor IP address to be blocked.
- Define the logical device's properties—This task involves defining the username, password, and privileged password for the blocking device.
- Define the managed devices' properties—This task involves defining the blocking device's properties, such as device type, IP address, username, password, and communication method.
- Assign the managed interface's properties for Cisco IOS or Catalyst 6000 devices—This task involves selecting the blocking interface or VLAN and assigning the pre-block and post-block ACLs or VACLs.
- (Optional.) Assign the list of devices that are never blocked—This task involves adding the networks and hosts that the Sensor will never add to the active ACL.
- (Optional.) Define a Master Blocking Sensor—This task involves adding the Sensor that will perform the blocking function for other blocking devices.

Implementation—VPN 3000 Concentrator

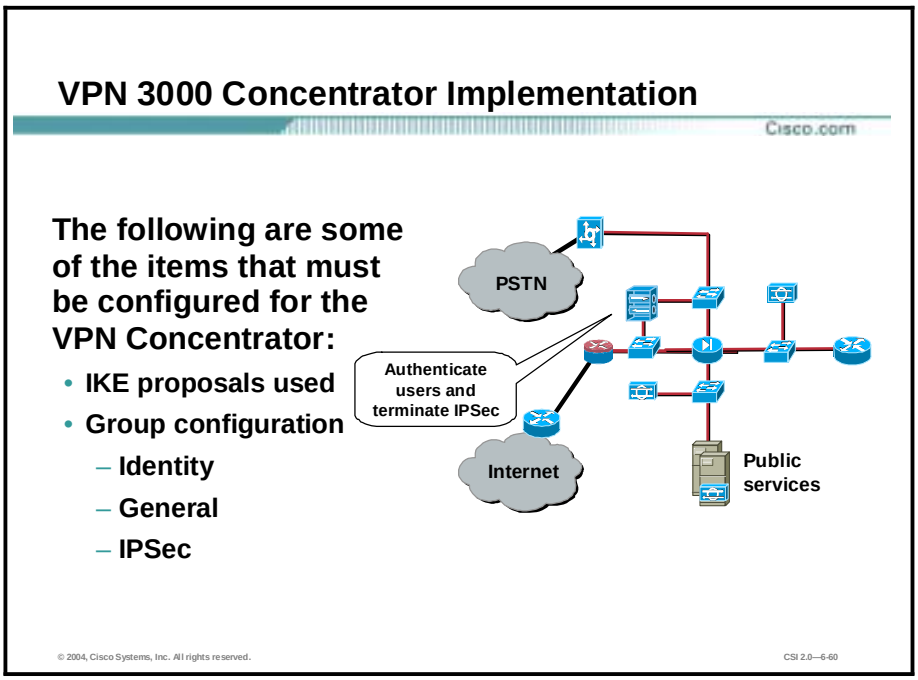

Implementation on the VPN 3000 Concentrator includes but is not limited to setting up the following:

- The IKE proposals that are used
- Group configuration
  - Identity
  - General
  - IPSec

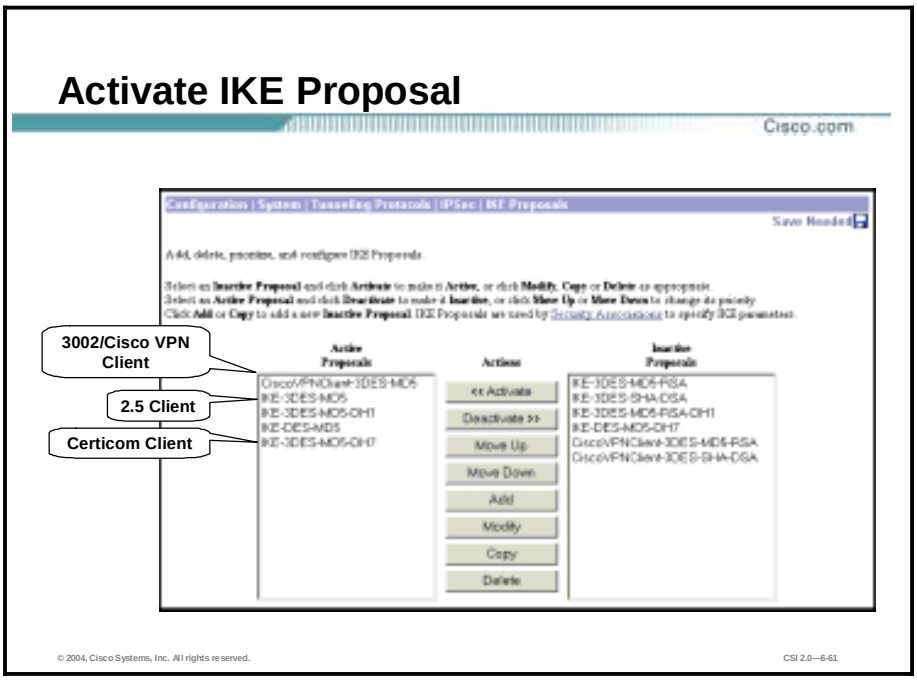

The VPN 3000 Concentrator can handle three types of remote clients: Cisco VPN client, Altiga client, and Certicom client. Before the VPN 3000 Concentrator can interface with these clients, you must make sure that the appropriate IKE proposal is configured, activated, and prioritized.

In remote-access connections, the client sends IKE proposals to the VPN 3000 Concentrator. The VPN 3000 Concentrator functions only as a responder. As the responder, the VPN 3000 Concentrator checks the active IKE proposal list, in priority order, to see if it can find a proposal that matches with parameters in the client's proposed SA. If a match is found, the tunnel establishment continues. If no match is found, the tunnel is torn down.

The IKE proposals are as follows:

- For the Cisco VPN Client, the default is Cisco VPN Client-3DES-Message Digest 5 (MD5). The Cisco VPN Client proposal must be listed first under the Active Proposals list, or your Client will not connect.
- For the Altiga client, use any of the "IKE" proposals, except the IKE proposals that end in DH7.
- For the Certicom client, use a proposal that ends in DH7. The Certicom client requires a proposal that supports Diffie-Hellman group 7 (DH7).

Each IKE proposal in the figure is a template. The parameters assigned to the template are applied to individual remote connections.

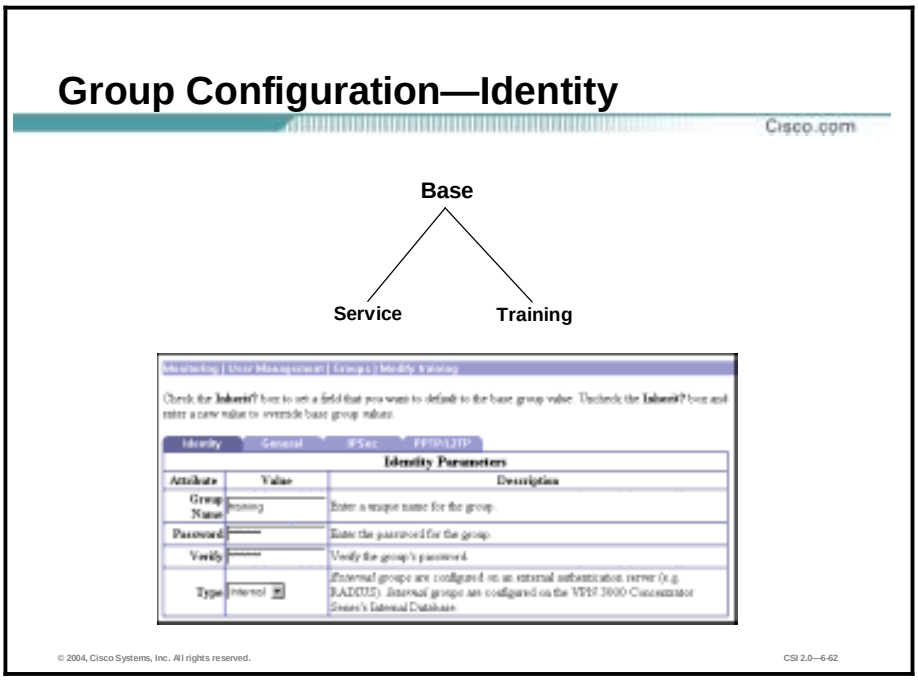

Within the User Management>Groups>Modify training window, you can view or modify individual group parameters. There are four tabs located under User Management>Groups:

- Identity tab—Configure the group name, password, and group authentication server type.
- General tab—Configure access rights and privileges, and access protocols.
- IPSec tab—Configure IPSec tunneling parameters.
- PPTP/L2TP tab—Configure Point-to-Point Tunneling Protocol (PPTP) and Layer 2 Tunneling Protocol (L2TP) tunneling parameters.

In the Identity tab you can set the following Identity parameters:

- Group Name field—Enter a unique name for this specific group. The maximum is 32 characters.
- Password field—Enter a unique password for this specific group. The minimum is 4 and maximum is 32 characters. The field displays only asterisks.
- Verify field—Reenter the group password to verify it. The field displays only asterisks.
- Type drop-down menu—Click the drop-down menu button and choose the type of group:
  - Internal—Use the internal VPN 3000 Concentrator authentication server to authenticate groups for IPSec tunneling. The internal server is the default selection.
  - External—Use an external authentication server to verify this group, such as a RADIUS server.

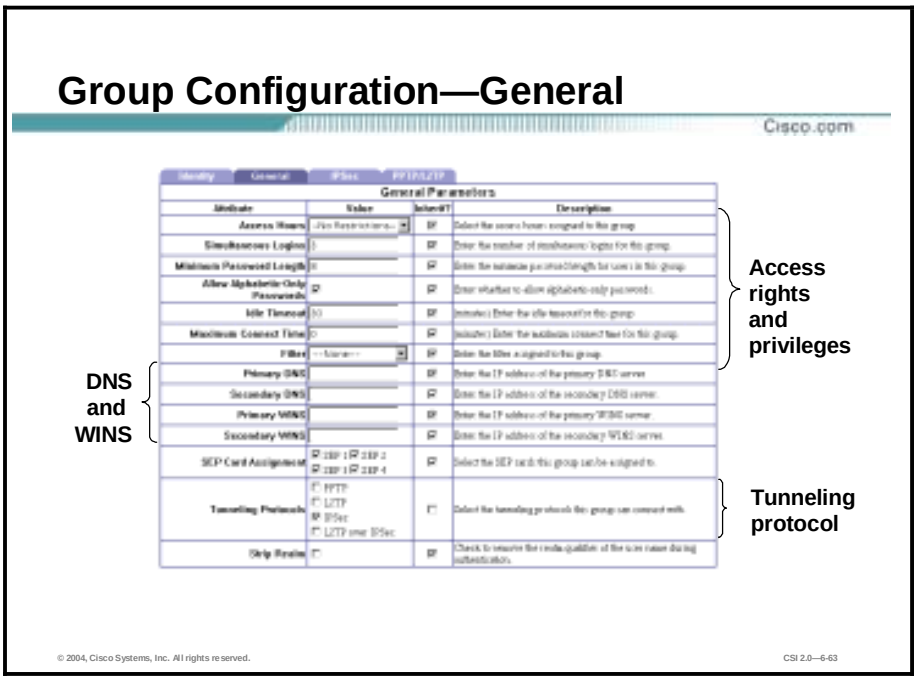

The General tab can be broken down into three sections:

- The top section defines access rights and privileges.
- The center section is used for WINS and DNS information used by the client.
- The bottom section defines which tunneling protocols this group supports.

Within the General tab, you can configure the following parameters:

- Access Hours drop-down menu—Click the drop-down menu button and select the named hours when group users can access the VPN 3000 Concentrator.
  - No Restrictions—No restrictions on access hours.
  - Never—No access at any time.
  - Business hours—Access 9 a.m. to 5 p.m., Monday through Friday.
- Simultaneous Logins field—Enter the number of simultaneous logins that group users are permitted. The minimum is 1 and the default is 3. While there is no maximum limit, allowing several could compromise security and affect performance.
- Minimum Password Length field—Enter the minimum number of characters for group user passwords. The minimum is 1, the default is 8, and the maximum is 32.
- Allow Alphabetic-Only Passwords check box—Select the check box to allow base group user passwords with alphabetic characters only, which is the default. To protect security, it is strongly recommended that you not allow such passwords.
- Idle Timeout field—Enter the base group idle timeout period in minutes. If there is no communication activity on the connection in this period, the system terminates the connection.
- Maximum Connect Time field—Enter the group maximum connection time in minutes. At the end of this time, the system terminates the connection.
- Filter drop-down menu—Select the type of filter you wish to apply to this group.

- Primary DNS field—Enter the primary IP address of the DNS server for this group's users.
- Secondary DNS field—Enter the IP address of the secondary DNS server for this group's users.
- Primary WINS field—Enter the primary IP address of the WINS server for this group's users.
- Secondary WINS field—Enter the secondary IP address of the WINS server for this group's users.
- SEP Card Assignment check boxes—It is recommended that you leave all four check boxes selected.
- Tunneling Protocols check boxes—Select the tunneling protocols the user VPN Clients can use. (Although the VPN 3000 Concentrator can support all four protocols simultaneously, in this lesson's lab exercise, de-select PPTP and L2TP. Select IPSec only.)
- Strip Realm check box—Select this check box only for PPTP, L2TP, or both.

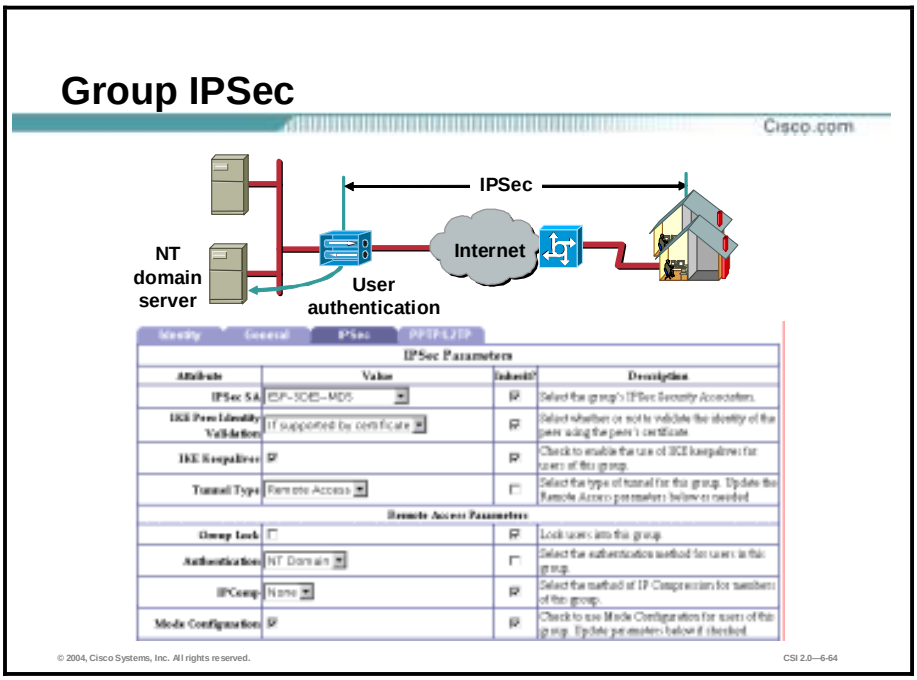

The IPSec tab enables you to configure IPSec parameters that apply to this group. The window can be divided into two sections: IPSec and remote-access parameters. IPSec parameters can be set as follows:

- IPSec SA drop-down menu—Click the drop-down menu button and choose the IPSec SA assigned to this group's IPSec clients. During tunnel establishment, the IPSec client and server negotiate an SA that governs authentication, encryption, encapsulation, key management, and so on. View or modify IPSec SAs on the Configuration>Policy Management>Traffic Management>Security Associations window.
- IKE Peer Identity Validation drop-down menu—This option applies only to tunnel negotiations based on digital certificates.
- IKE Keepalives check box—Select this check box to enable the feature. (IKE Keepalives is enabled by default.) This feature enables the VPN 3000 Concentrator to monitor the continued presence of a remote peer, and to report its own presence to that peer. If the peer becomes unresponsive, the VPN 3000 Concentrator initiates removal of the connection. Enabling IKE keepalives prevents hung connections when rebooting either the host or the peer. For this feature to work, both the VPN 3000 Concentrator and its remote peer must support IKE keepalives. The following peers support IKE keepalives:
  - Cisco VPN Client (Release 4.x)
  - Cisco VPN 3000 Client (Release 4.x)
  - Cisco VPN 3002 Hardware Client
  - Cisco VPN 3000 Concentrators (with IKE support)
  - Cisco IOS software
  - Cisco PIX Firewall
- Reauthentication on Rekey check box—By selecting the **Reauthentication on Rekey** check box, the VPN 3000 Concentrator prompts the user for an identification and password whenever a re-key occurs. The default is disabled.

- Tunnel Type drop-down menu—Click the drop-down menu and choose the remote access tunnel type. Choose **Remote access for IPSec client to LAN applications**.

Implementation—Layer 3 Switch

This topic describes the implementation of the Layer 3 switch in the SAFE SMR medium network campus module.

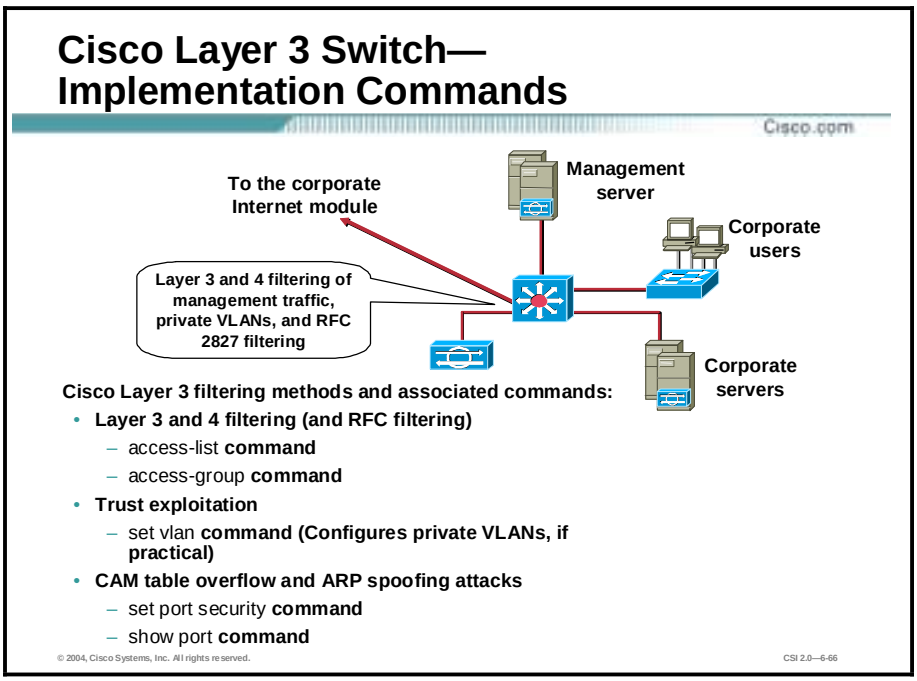

The following commands and features are used to implement SAFE SMR on a Layer 3 switch:

- Layer 3 and 4 filtering (and RFC filtering)
  - **access-list** command
  - **access-group** command
- Trust exploitation
  - **set vlan** command (Configures private VLANs, if practical)
- Content Addressable Memory (CAM) table overflow and Address Resolution Protocol (ARP) spoofing attacks
  - **set port security** command
  - **show port** command

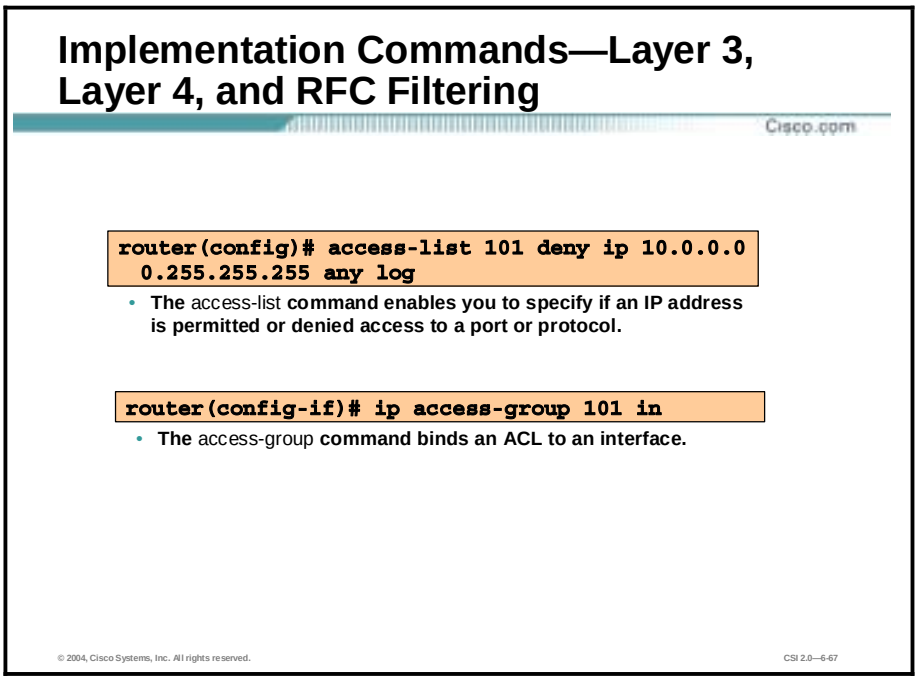

As with most devices using Cisco IOS or Cisco IOS-like commands, ACLs can be used to implement Layer 3 and 4 filtering.

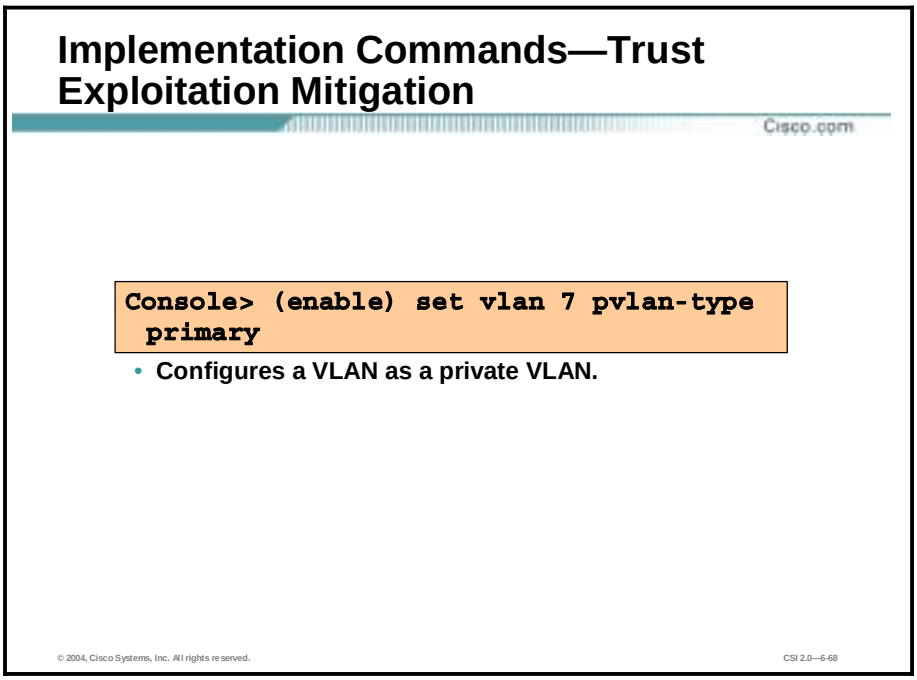

Private VLANs (PVLANs) are available on the Catalyst 6000 running the Cisco Catalyst Operating System (CatOS) 5.4 or later, and are available on the Catalyst 4000, 2980G, 2980G-A, 2948G, and 4912G running CatOS 6.2 or later.

PVLANs are tools that enable the segregating of traffic at Layer 2 and the turning of a broadcast segment into a nonbroadcast multiaccess-like segment. Traffic that comes to a switch from a promiscuous port (that is, a port that is capable of forwarding both primary and secondary VLANs) is able to go out on all the ports that belong to the same primary VLAN. Traffic that comes to a switch from a port mapped to a secondary VLAN (either an isolated, community, or two-way community VLAN) can be forwarded to a promiscuous port or a port belonging to the same community VLAN. Multiple ports mapped to the same isolated VLAN cannot exchange any traffic.

You must configure a private VLAN on the supervisor engine.

Valid values for **pvlan-type** are as follows:

- **primary**—Specifies the VLAN as the primary VLAN in a PVLAN
- **isolated**—Specifies the VLAN as the isolated VLAN in a PVLAN
- **community**—Specifies the VLAN as the community VLAN in a PVLAN
- **twoway-community**—Specifies the VLAN as a bidirectional community VLAN that carries the traffic among community ports, and to and from community ports to and from the Multilayer Switch Feature Card (MSFC)
- **none**—Specifies that the VLAN is a normal Ethernet VLAN, not a PVLAN

Note VLANs 1001, 1002, 1003, 1004, and 1005 cannot be used in PVLANs.

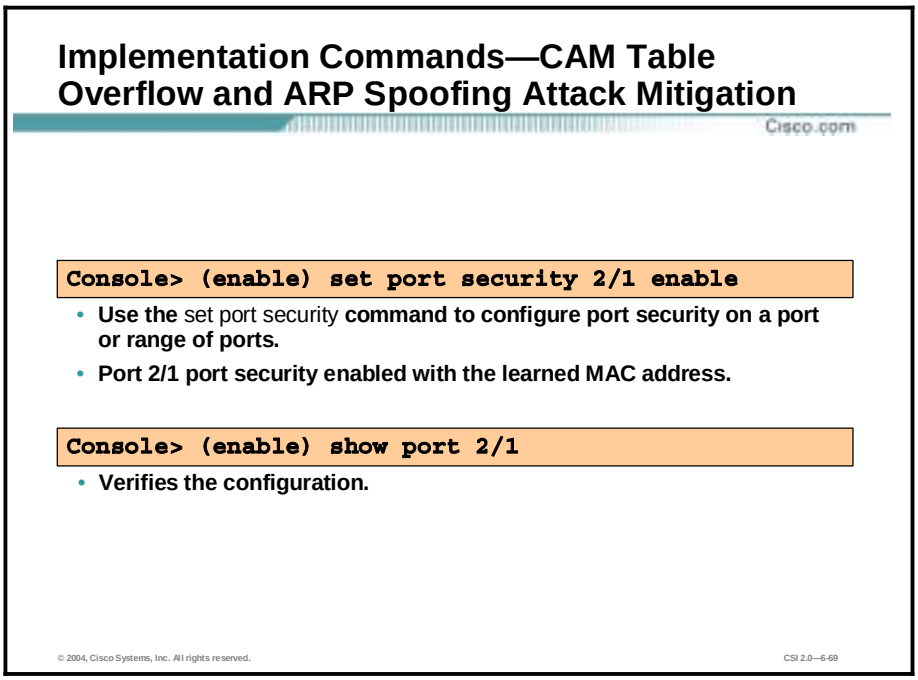

Configuring port security on the switch can mitigate these attacks. This option provides for either the specification of the MAC addresses on a particular switch port or the specification of the number of MAC addresses that can be learned by a switch port. When an invalid MAC is detected on the port, the switch can either block the offending MAC address or shut down the port.

Summary

This topic summarizes this lesson.

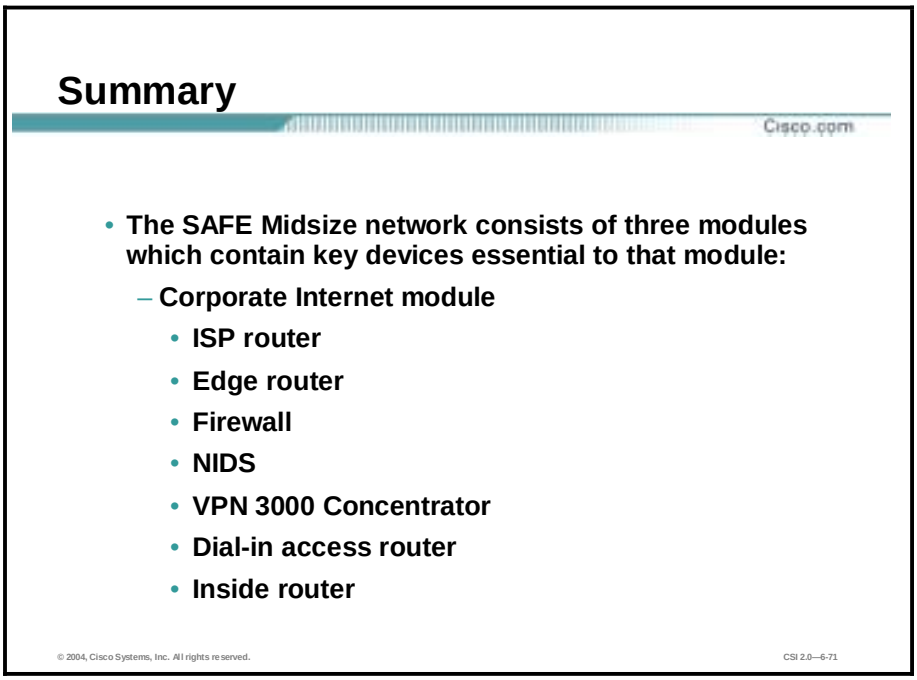

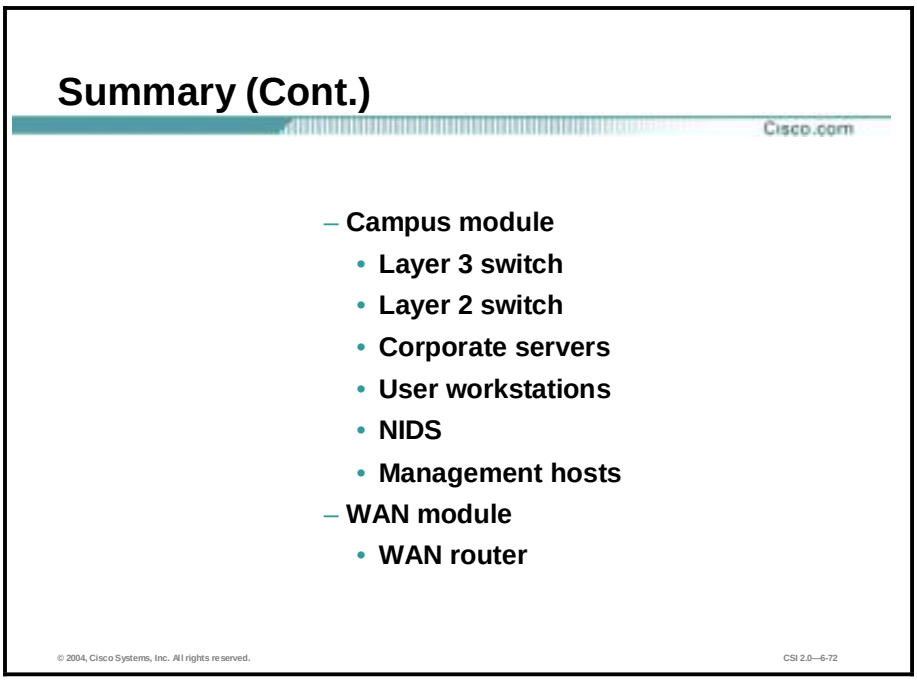

Summary (Cont.)

Cisco.com

- The mitigation roles identified for each threat in SAFE SMR are integral to a successful implementation.
- Specific configurations and commands are used to apply the mitigation roles identified for each threat.
- Alternative devices and configurations can be used in order to provide existing device integration, ease of implementation, and cost effectiveness.

7

Remote User Network Implementation

Overview

This lesson introduces the Cisco Remote User Network implementation strategy and components. It includes the following topics:

- Objectives
- Design overview
- Key devices and threat mitigation
- Software client option
- Remote site firewall option
- VPN 3002 hardware client option
- Remote site router option
- Summary
- Lab exercise

Objectives

This topic lists the lesson's objectives.

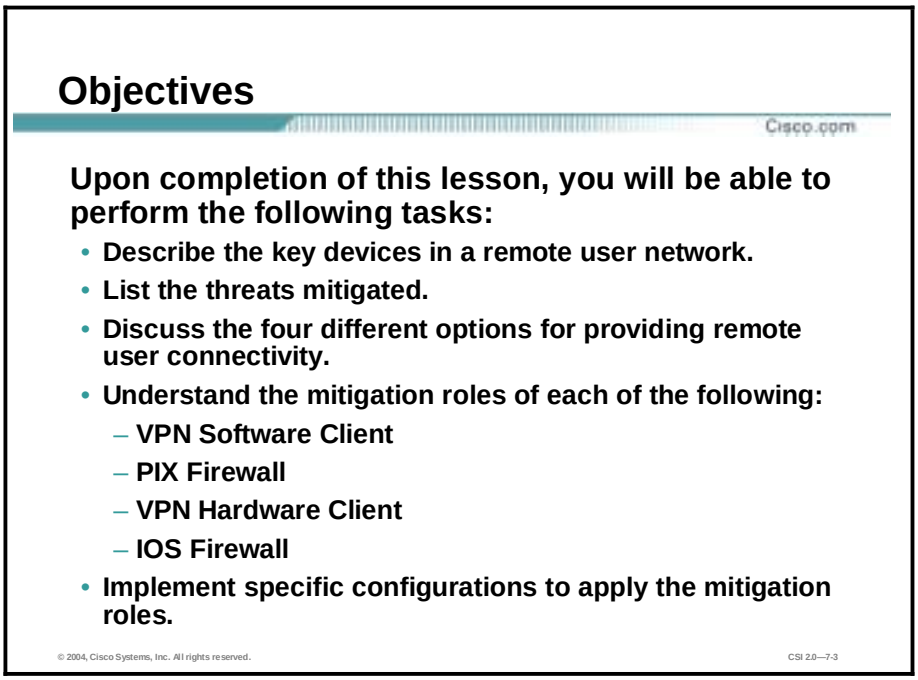

Design Overview

This topic gives a brief overview of the SAFE Extending the Security Blueprint to Small, Midsize, and Remote-User Networks (SAFE SMR) remote-user network implementation.

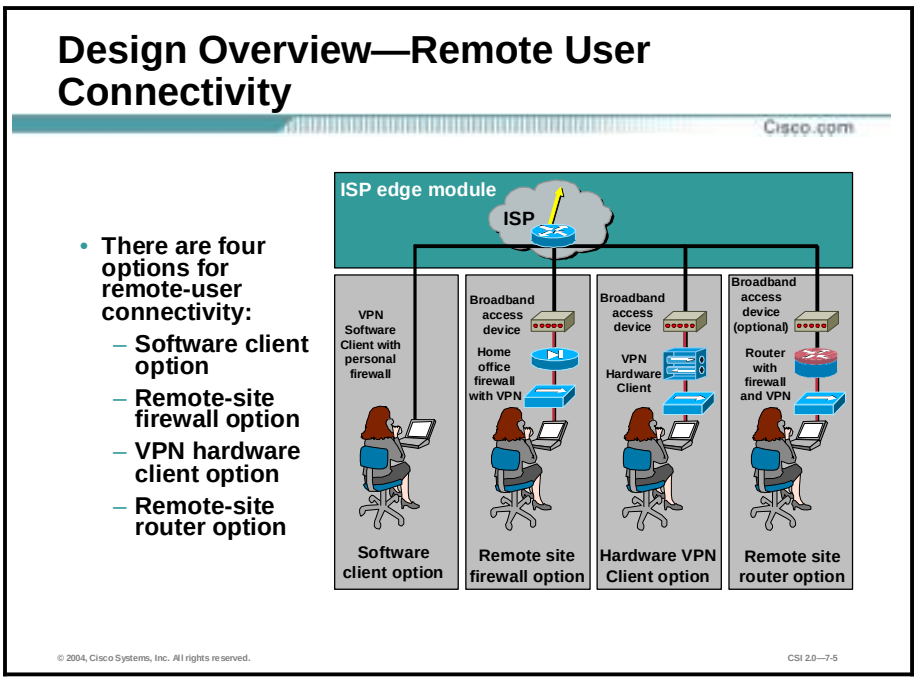

Remote connectivity applies to both mobile workers and home-office workers. The primary purpose of these designs is to provide connectivity from the remote site to the corporate headquarters and the Internet. The following four options include software-only, software-with-hardware, and hardware-only solutions:

- Software access option—Remote user with a Cisco Virtual Private Network (VPN) Software Client and a personal firewall software on the PC.
- Remote-site firewall option—The remote site is protected with a dedicated firewall that provides firewall protection and IPSec VPN connectivity to corporate headquarters. WAN connectivity is provided via an ISP-provided broadband access device (for example, a digital subscriber line [DSL] or cable modem).
- Cisco VPN 3002 Hardware Client option—Remote site using a dedicated Hardware Client that provides IPSec VPN connectivity to corporate headquarters. WAN connectivity is provided via an ISP-provided broadband access device.
- Remote-site router option—Remote site using a router that provides both firewall capabilities and IPSec VPN connectivity to corporate headquarters. This router can either provide direct broadband access or go through an ISP-provided broadband access device.

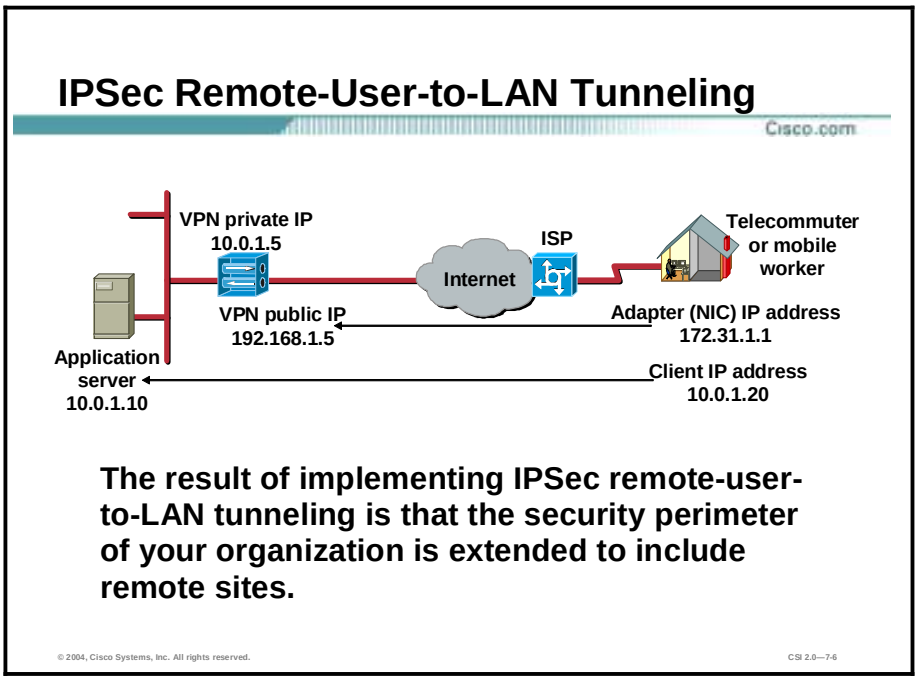

A VPN is defined as customer connectivity deployed on a shared infrastructure with the same policies as a private network. The shared infrastructure can augment a service provider IP network, Frame Relay, ATM backbone, or the Internet. The result is that the security perimeter of the organization is extended to the remote site.

The Cisco end-to-end hardware and Cisco IOS software networking products enable a complete access VPN solution by providing the following:

- Sophisticated security for sensitive private transmissions over the public infrastructure
- Quality of Service (QoS) through traffic differentiation
- Reliability for mission-critical applications
- Scalability for supporting large amount of data
- Comprehensive network management

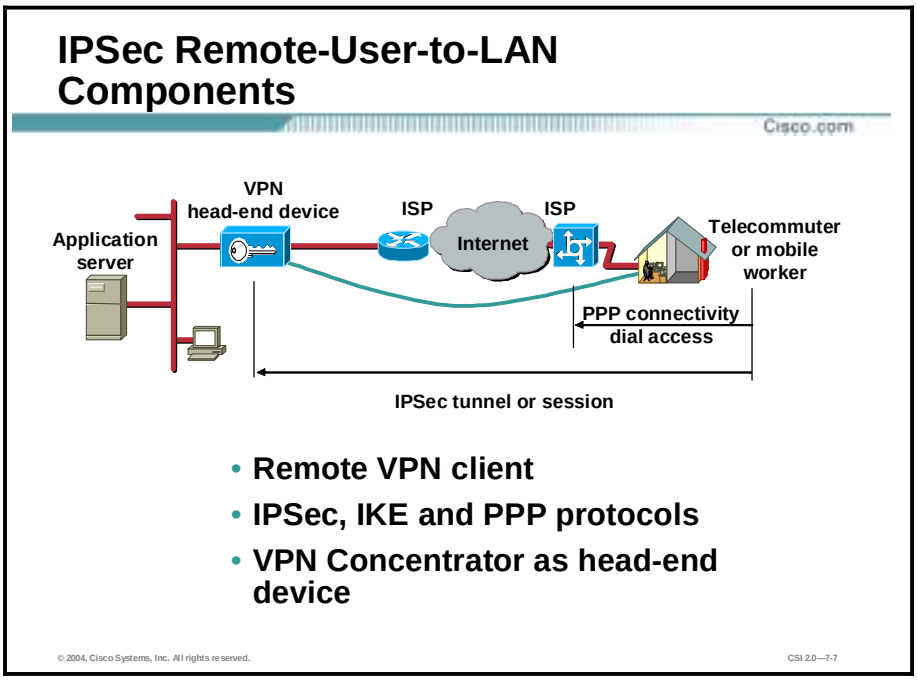

IPSec remote user-to-LAN components consist of the following:

- Remote VPN Client—Terminates one end of the tunnel
  - Software Client—Resides on a PC
  - Hardware Clients—Reside on a device
    - VPN 3002
    - PIX Firewall
    - Cisco IOS router
- IPSec and Internet Key Exchange (IKE)
  - Establishes a secure tunnel or session through the Internet to a Cisco VPN Concentrator
  - For dialup applications, IPSec relies on Point-to-Point Protocol (PPP) to establish the physical connection to the local ISP and Internet
- Concentrator
  - Terminates the tunnels and sessions
  - Encrypts, authenticates, and encapsulates data
  - May provide both authentication and data encryption options

Note	Cisco VPN Client (CVPN) v3.x, which is compliant with Cisco Easy VPN, has replaced the Cisco VPN Client (CSVPN) v1.0 and v1.1. Cisco VPN Client v3.x provides greater ease of use and customized update capabilities.

Note	The Cisco VPN 3002 Hardware Client Software package works with the Cisco VPN 3002 Hardware Client. This software incorporates management interfaces, and supports a broad variety of routing and tunneling protocols, encryption and authentication algorithms, and key management and system administration options.

Key Devices and Threat Mitigation

This topic provides details of the key devices for remote-user network implementation.

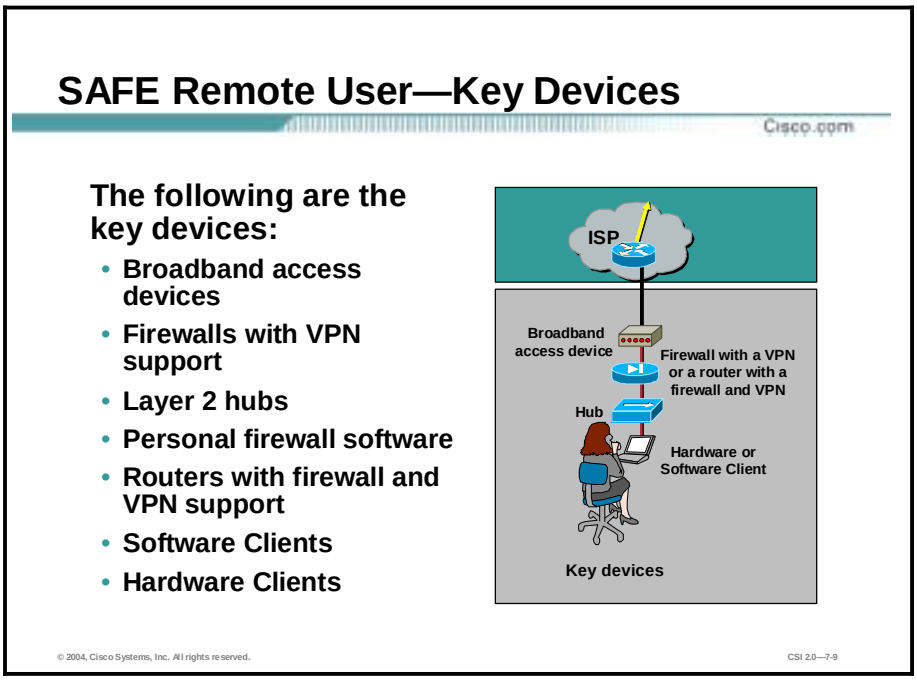

The following are the key devices in the SAFE remote-user configuration:

- Broadband access devices—Provides access to the broadband network (for example, DSL, cable, and so on)
- Firewalls with VPN support—Provides secure, end-to-end encrypted tunnels between the remote site and the corporate head-end, and provides network-level protection of remote-site resources and stateful filtering of traffic
- Layer 2 hubs—Provides connectivity for devices within the remote site, which can be integrated into the firewall or Hardware Client
- Personal firewall software—Provides device-level protection for individual PCs
- Routers with firewall and VPN support—Provides secure, end-to-end encrypted tunnels between the remote site and the corporate head-end, provides network-level protection of remote-site resources and stateful filtering of traffic, and can provide advanced services such as voice or QoS
- Software Clients—Provides secure, end-to-end encrypted tunnels between individual PCs and the corporate head-end
- Hardware Clients—Provides secure, end-to-end encrypted tunnels between the remote site and the corporate head-end

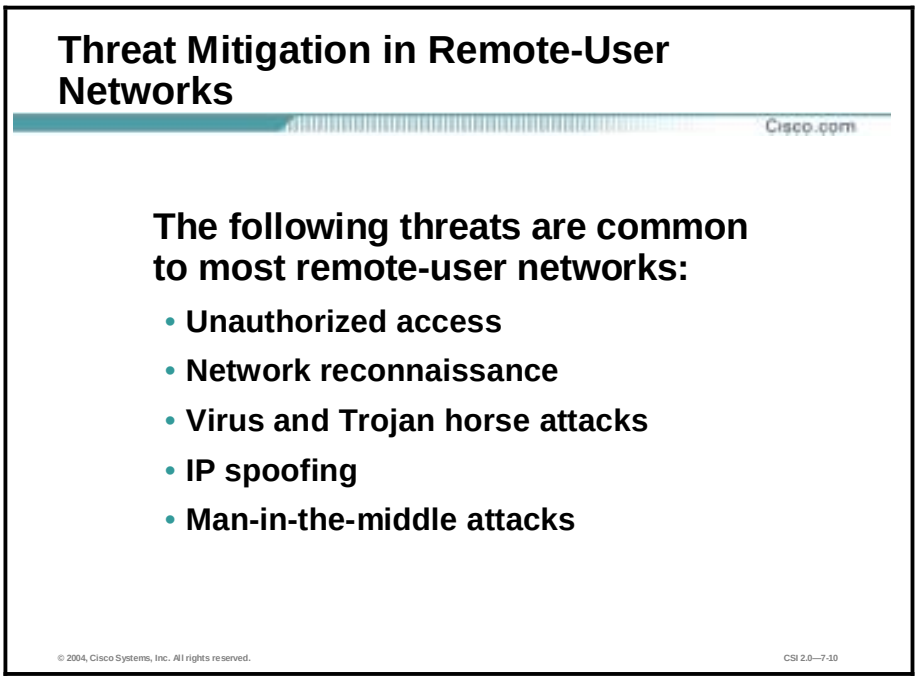

The following threats are expected in a remote-user environment:

- Unauthorized access—Any unauthorized data that traverses a network.
- Network reconnaissance—Information-gathering activities by which hackers collect data that is used to later compromise networks.
- Virus and Trojan horse attacks—Viruses are computer programs that are written by devious programmers and are designed to replicate themselves and infect computers when triggered by a specific event. Trojan horse attacks are software programs that appear to be harmless (for example, computer games), but are actually vehicles for destructive code.
- IP spoofing—Posing as an authorized party in the data transmission by using the IP address of one of the data recipients.
- Man-in-the-middle attacks—Require that the attacker has access to network packets that come across the networks. Such attacks are often implemented using network packet sniffers and routing and transport protocols. The possible uses of such attacks are theft of information, hijacking of an ongoing session to gain access to your internal network resources, traffic analysis to derive information about your network and its users, denial of service, corruption of transmitted data, and introduction of new information into network sessions.

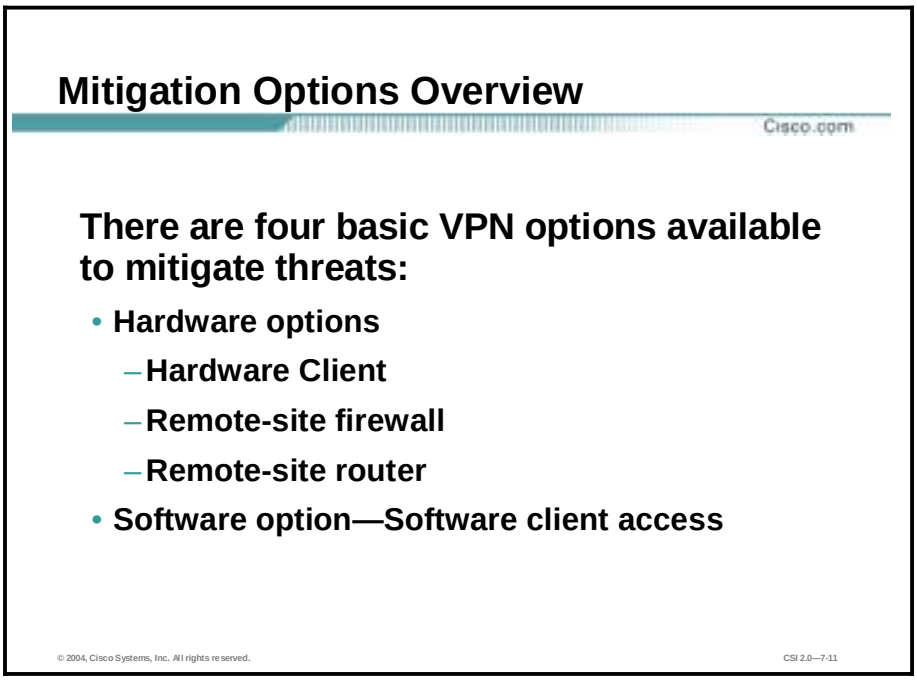

The following are mitigation solutions for remote-access network threats:

- Hardware Client—Use of a hardware client that is independent of the host provides a connection with a greater level of security than a software client.
- Remote-site firewall—Use of a remote-site firewall can provide a high level of security.
- Remote-site router—Use of a remote-site router can provide a medium level of security, and can establish an IPSec tunnel and basic security that is independent of the host.
- Software client access—Use of software for access provides security at a host level, which uses the host processor and memory.

Software Access Option

This topic provides details of the software access option for remote users.

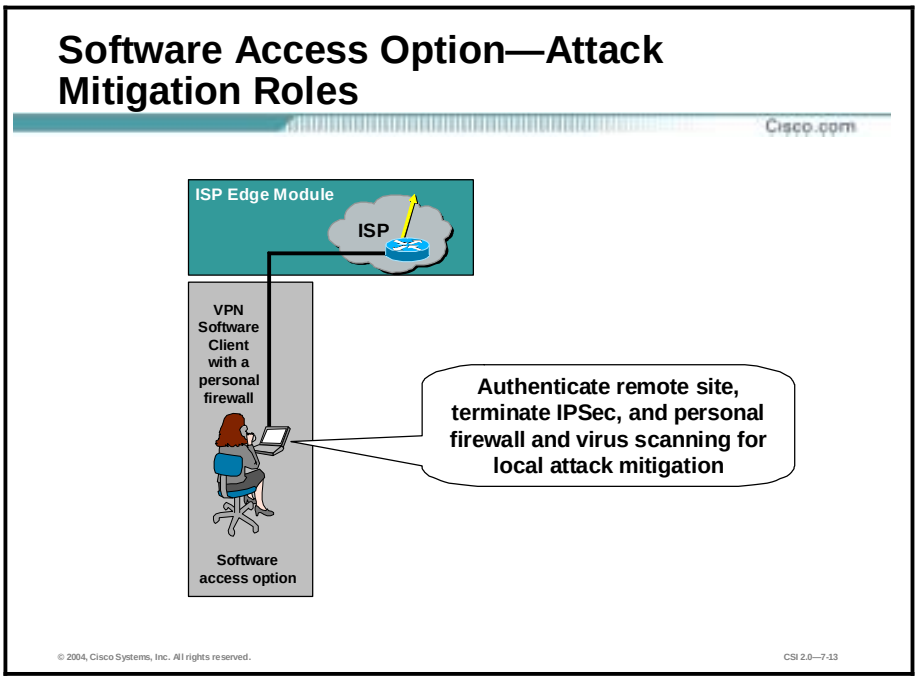

The following are the specific attack mitigation roles for the software access option:

- Authenticate remote site—Properly identify and verify a user or service
- Terminate IPSec—Successfully establish an IPSec tunnel between the remote site and the host site
- Personal firewall and virus scanning for local attack mitigation—Allay the risk of virus infection at the remote site

Software Access Option—Design Guidelines

Cisco.com

- Geared toward mobile- and home-office worker.
- Remote user needs VPN software and Internet access.
- Authentication and configuration are controlled from the headquarters.
- Split tunneling is disabled when the VPN tunnel is operational.
- Personal firewall software is recommended to protect the remote user when split tunneling is enabled or the VPN is not connected.
- Virus scanning software is recommended.

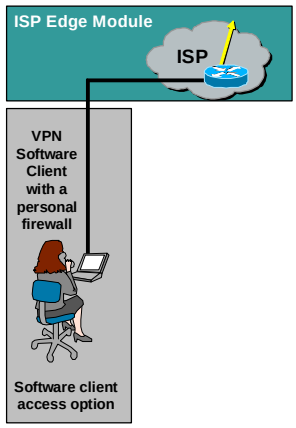

The software access option is geared toward the mobile worker as well as the home-office worker. All the remote user requires is a PC with Cisco VPN Client software and connectivity to the Internet or ISP network via a dial-in or Ethernet connection.

The primary function of the Software Client is to establish a secure, encrypted tunnel from the client device to a VPN head-end device. Access and authorization to the network are controlled from the headquarters location when filtering takes place on the firewall, and on the client itself if access rights are pushed down via policy. The remote user is first authenticated, and then receives IP parameters such as a virtual IP address, which is used for all VPN traffic, and the location of name servers (Domain Name System [DNS] and Windows Internet Name Service [WINS]).

Split tunneling can also be enabled or disabled via the central site. For the SAFE design, split tunneling was disabled, making it necessary for all remote users to access the Internet via the corporate connection when they have a VPN tunnel established. Because the remote user may not always want the VPN tunnel established when connected to the Internet or ISP network, personal firewall software is recommended to mitigate unauthorized access to the PC. Virus-scanning software is also recommended to mitigate viruses and Trojan horse programs infecting the PC.

Software Access Option—Implementation

Cisco.com

The Cisco VPN Client version 3.5 or higher is the recommended product for implementation of the software access option:

- **Provides integrated VPN and firewall functionality**
- **Simple install process**
- **Configured via the head-end VPN termination device**

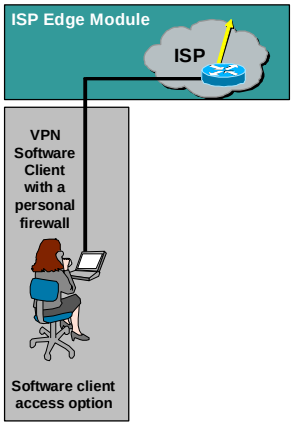

It is recommended that you use the latest level of the Software Client unless technical specifications or compatibility issues require an earlier version.

The Cisco VPN Client v3.x or higher is the recommended product for implementation of the software access option, for these reasons:

- It provides integrated VPN and firewall functionality.
- It has a simple install process.
- It is configured via the head-end VPN termination device.

Cisco VPN Client for Windows

Cisco.com

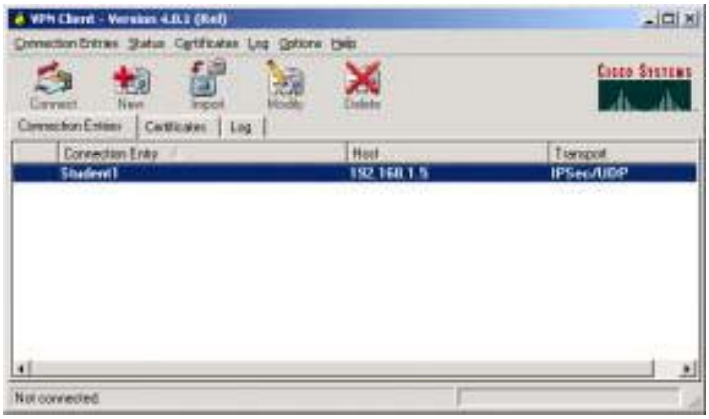

CSI 2.0—7-16

The Cisco VPN Client for Windows is a software program that runs on Windows 95, 98, ME, 2000, XP, and NT 4.0. The Software Client on a remote PC, communicating with a Concentrator at an enterprise or service provider, creates a secure connection over the Internet that lets you access a private network as if you were an on-site user.

The figure shows the Software Client window. From this window, you can launch the new connection wizard, change or set optional parameters, and launch the Software Client.

Cisco VPN Client for Windows Run Mode

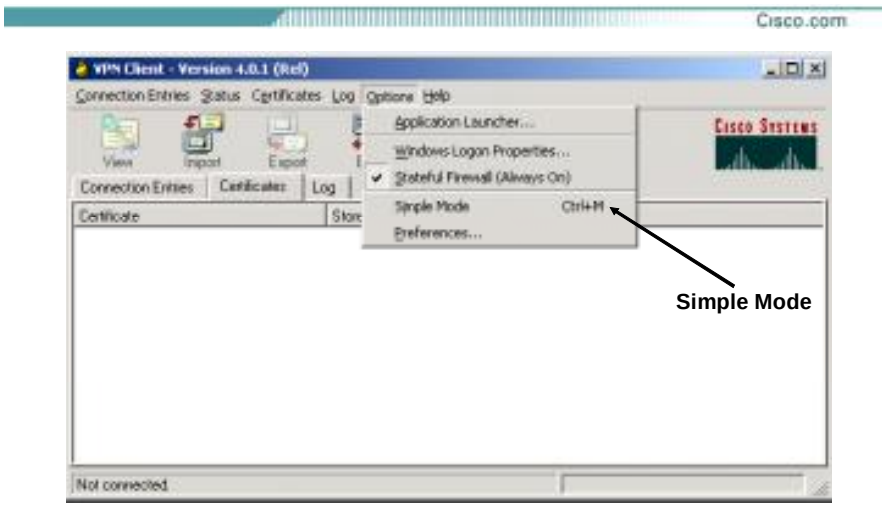

You can run the Software Client in simple mode or in advanced mode. The default is advanced mode, although your network administrator might have configured simple mode as the default.

Use simple mode if you want only to start the Software Client application and connect to a VPN device using the default connection entry.

Use advanced mode for the following tasks:

- Managing the Software Client
- Configuring connection entries
- Enrolling for and managing certificates
- Viewing and managing event logging
- Viewing tunnel routing data

To toggle between advanced mode and simple mode, press **Ctrl-M**. Alternatively, you can choose your mode from the Options menu.

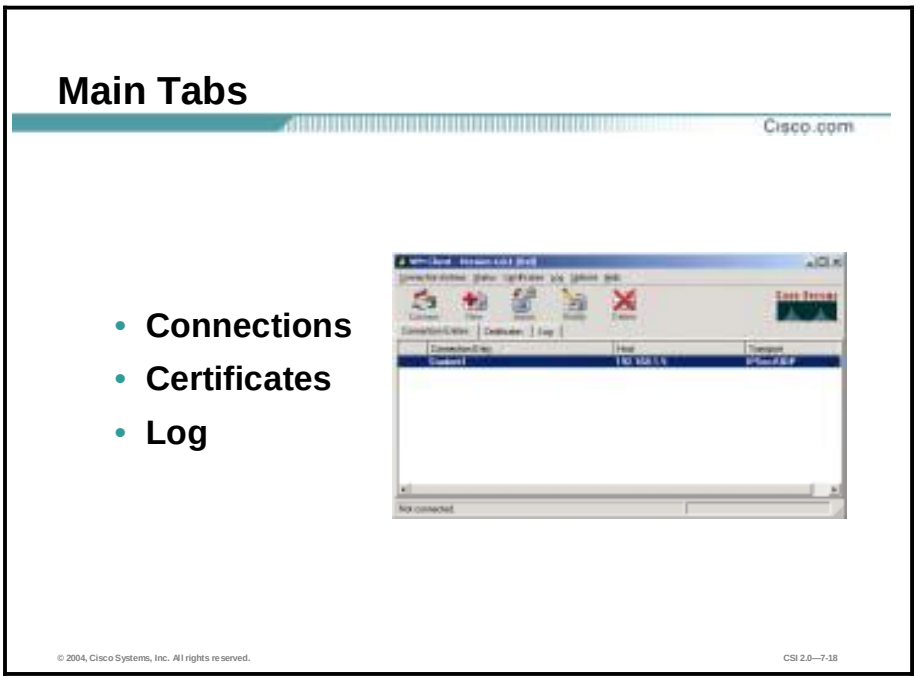

The following are the main tabs in the Cisco VPN Client run mode:

- Connection Entries tab—Displays the list of current connection entries, the host, which is the VPN device that each connection entry uses to gain access to the private network, and the transport properties that are set for each connection entry.
- Certificates tab—Displays the list of certificates in the VPN Client certificate store. Use this tab to manage certificates.
- Log tab—Displays event messages from all processes that contribute to the client-peer connection: enabling logging, clearing the event log, viewing the event log in an external window, and setting logging levels.

Menus—Connection Entries

Cisco.com

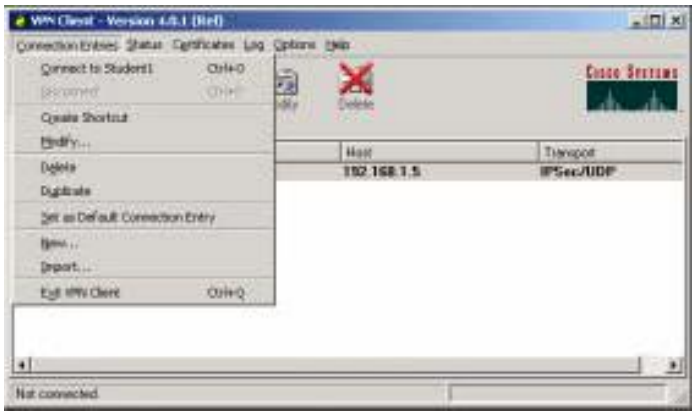

Use the Connection Entries menu as a shortcut to frequently used connection entry operations. The following actions are available:

- Connect to—Connect to a VPN device using the selected connection entry. If the Connections tab is not selected, a submenu is displayed, listing all available connection entries.
- Disconnect—Disconnect your current VPN session.
- Create Shortcut—Create a shortcut on your desktop for the current connection entry.
- Modify—Edit the current connection entry.
- Delete—Delete the current connection entry.
- Duplicate—Duplicate the selected connection entry. This menu choice lets you create a new connection entry using the configuration from a current connection entry as a template.
- Set as Default Connection Entry—Make the current connection entry the default.
- New—Create a new connection entry.
- Import—Bring in a new connection entry profile from a file.
- Exit VPN Client—Close the Cisco VPN Client application.

Menus—Status

Cisco.com

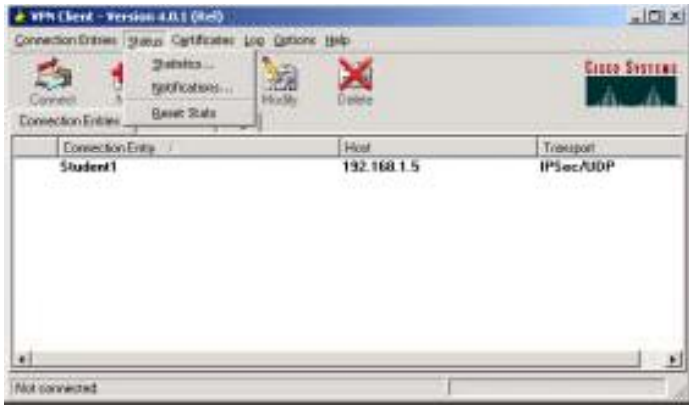

CSI 2.0—7-20

Use the Status menu to display routes and notifications and to reset the statistics display. The following actions are available:

- Statistics—View tunnel details, route details, and firewall information for the current VPN session.
- Notifications—View notices from the VPN device you are currently connected to.
- Reset Stats—Clear the statistics from the statistics displays and start over.

Menus—Certificates

Cisco.com

Use the Certificates menu to enroll and manage certificates. The following actions are available:

- View—Display the properties of the selected certificate.
- Import—Import a certificate file from a specified file location.
- Export—Export the selected certificate to a specified file location.
- Enroll—Enroll with a Certificate Authority (CA) to obtain a certificate.
- Verify—Verify that a certificate is still valid.
- Delete—Remove the selected certificate.
- Change Certificate Password—Change the password that protects the selected certificate in the Cisco VPN Client certificate store.
- Retry Certificate Enrollment—Retry a previously attempted certificate enrollment.
- Show CA/RA Certificates—Display digital certificates issued by either a CA or a Registration Authority (RA).

Menus—Log

Cisco.com

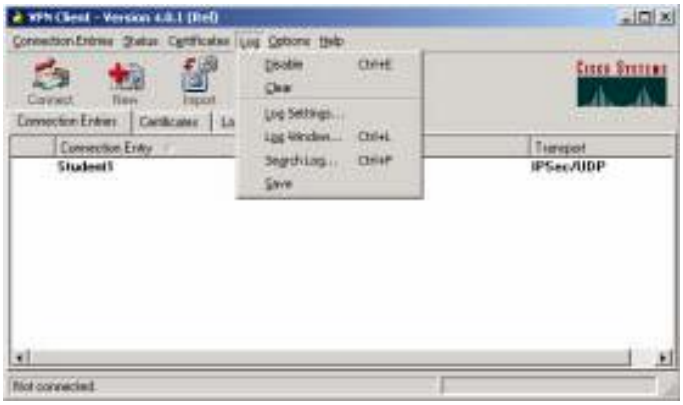

CSI 2.0—7-22

Use the Log menu to manage the log. The following actions are available:

- Enable/Disable—Start collecting events (Enable); stop collecting events (Disable).
- Clear—Erase the events displayed on the Log tab (and Log window).
- Log Settings—Change the logging levels of event classes.
- Log Window—Display a separate window that shows events. From this window you can save the display, edit logging levels by event class, and clear both log displays. This window shows more events than the display area of the main advanced mode window.
- Search Log—Display a dialog box into which you enter the exact string to be matched. The search string is not case sensitive, and wild cards are not supported. Matched instances are highlighted on the Log tab, not the Log window.
- Save—Store the current log in a specified log file.

Menus—Options

Cisco.com

Use the Options menu to perform actions such as launching an application. The following actions are available:

- Application Launcher—Start an application before connecting to a VPN device.
- Windows Logon Properties—Control logon features are applicable only to the Windows NT platform. The following logon features are available:
  - Ability to start a connection before logging on to a Windows NT system
  - Permission to launch a third-party application before logging on to a Windows NT system
  - Control of autodisconnect behavior when logging off
- Stateful Firewall (Always On)—Enable and disable the internal stateful firewall.
- Simple Mode—Switch to simple mode.
- Preferences—Set the following features:
  - Save window settings—Save any changes you make to the Cisco VPN Client window.
  - Hide upon connect—Place the Cisco VPN Client window in the dock when the VPN connection is established.
  - Enable tool tips—Enable tool tips for the toolbar action buttons.

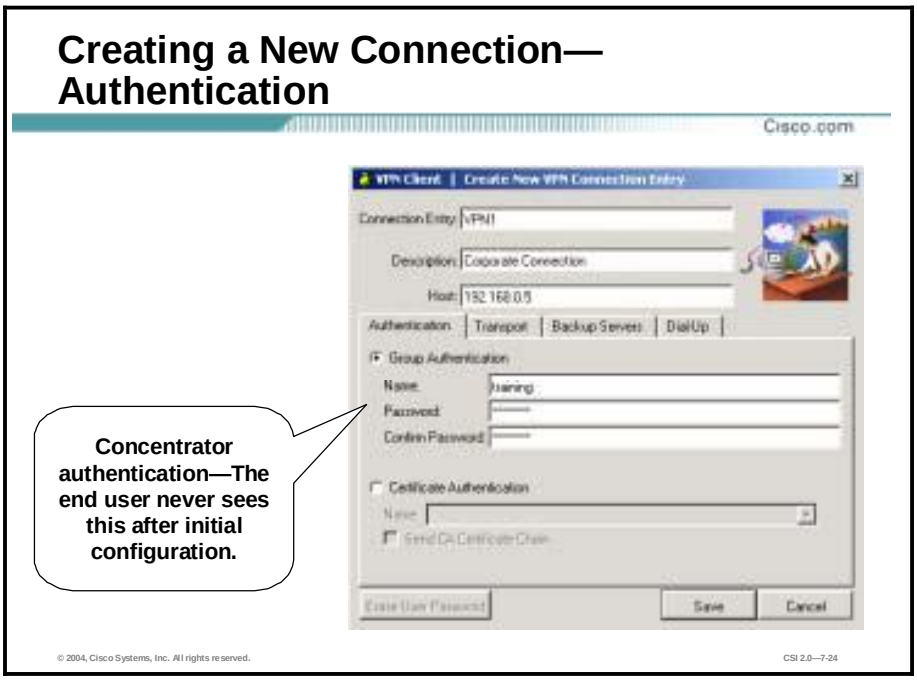

Clicking **New** from the toolbar or the Connection Entries menu displays the Create New VPN Connection Entry window. The following parameters need to be entered:

- Connection Entry—Enter a unique name for this new connection. You can use any name to identify this connection; for example, Engineering. This name can contain spaces, and it is not case sensitive.
- Description—Enter a description of this connection. This field is optional, but it helps further identify this connection. For example, Connection to Engineering remote server.
- Host—Enter the hostname or IP address of the remote VPN device you want to access.

Under the Authentication tab, you must choose whether you are going to be using group or certificate authentication and fill in the required fields as follows:

- Name—Enter the name of the IPSec group to which you belong. This entry is case sensitive.
- Password—Enter the password (which is also case sensitive) for your IPSec group. The field displays only asterisks.
- Confirm Password—Verify your password by entering it again.

For certificates to be exchanged, the Certificate Authentication radio button must be selected. In the Name drop-down menu, any personal certificates loaded on your PC are listed. Choose the certificate to be exchanged with the Concentrator during connection establishment. If no personal certificates are loaded in your PC, the drop-down menu is blank.

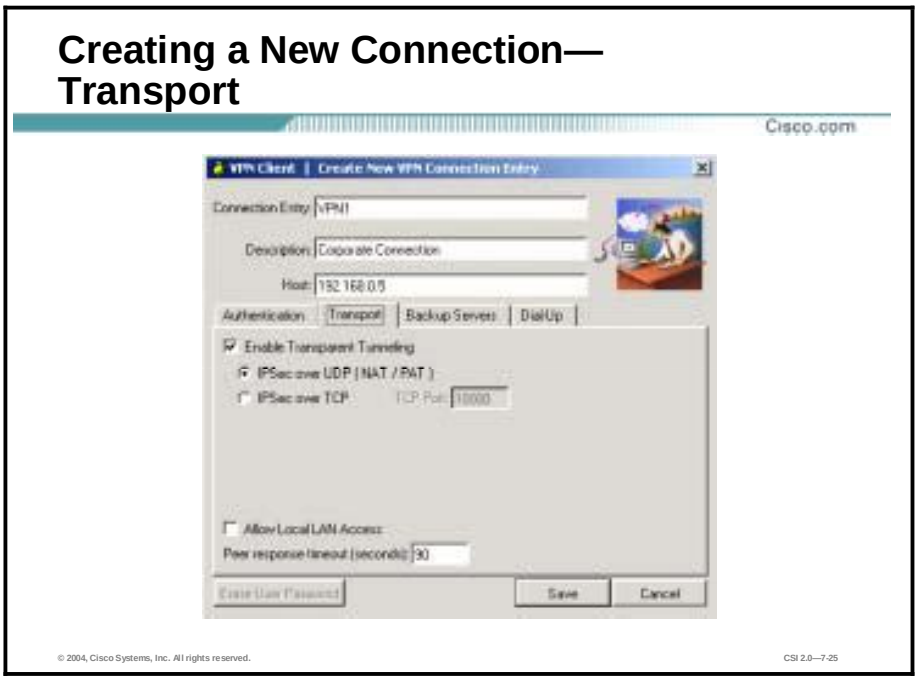

Transparent Tunneling

Transparent tunneling allows secure transmission between the Cisco VPN Client and a secure gateway through a router serving as a firewall, which may also be performing Network Address Translation (NAT) or Port Address Translation (PAT). Transparent tunneling encapsulates Protocol 50 (Encapsulating Security Payload [ESP]) traffic within UDP packets and can allow for both IKE (UDP 500) and Protocol 50 traffic to be encapsulated in TCP packets before it is sent through the NAT or PAT devices or firewalls. The most common application for transparent tunneling is behind a home router performing PAT. The central-site group in the Cisco VPN device must be configured to support transparent tunneling. This parameter is enabled by default. To disable transparent tunneling, deselect the **Enable Transparent Tunneling** check box under the Transport tab. It is recommended that you always keep this parameter selected.

You must choose a mode of transparent tunneling, either over UDP or over TCP. The mode you use must match that used by the secure gateway to which you are connecting. Either mode operates properly through a PAT device. Multiple simultaneous connections might work better with TCP. In an extranet environment, TCP mode is generally preferable, because UDP does not operate with stateful firewalls.

The following transport tunneling options are available:

- Using IPSec over UDP (NAT/PAT)—To enable IPSec over UDP (NAT/PAT), select the **IPSec over UDP (NAT/PAT)** radio button. With UDP, the port number is negotiated. UDP is the default mode.
- Using IPSec over TCP (NAT/PAT/Firewall)—To enable IPSec over TCP, select the **IPSec over TCP** radio button. When using TCP, you must also enter the port number for TCP in the TCP port field. The port number must match the port number configured on the secure gateway. The default port number is 10000.

Allowing Local LAN Access

In a multiple network interface card (NIC) configuration, local LAN access pertains only to network traffic on the interface on which the tunnel was established. The Allow Local LAN Access parameter gives you access to the resources on your local LAN (printer, fax, shared files, and other systems) when you are connected through a secure gateway to a central-site VPN device. When this parameter is enabled and your central site is configured to permit it, you can access local resources while connected. When this parameter is disabled, all traffic from your Cisco VPN Client system goes through the IPSec connection to the secure gateway.

To enable this feature, select the **Allow Local LAN Access** check box; to disable it, deselect the check box. If the local LAN you are using is not secure, you should disable this feature. For example, you should disable this feature when you are using a local LAN in a hotel or airport.

A network administrator at the central site configures a list of networks at the VPN Client side that you can access. You can access up to ten networks when this feature is enabled. When local LAN access is allowed and you are connected to a central site, all traffic from your system goes through the IPSec tunnel except traffic to the networks excluded from doing so (in the network list).

When the local LAN access feature is enabled and configured on the Cisco VPN Client and permitted on the central-site VPN device, you can see a list of the local LANs available by looking at the Routes table.

Adjusting the Peer Response Timeout Value

The Cisco VPN Client uses a keepalive mechanism called dead peer detection (DPD) to check the availability of the VPN device on the other side of an IPSec tunnel. If the network is unusually busy or unreliable, you might need to increase the number of seconds to wait before the Cisco VPN Client decides that the peer is no longer active. The default number of seconds to wait before terminating a connection is 90 seconds. The minimum number you can configure is 30 seconds, and the maximum is 480 seconds.

To adjust the setting, enter the number of seconds in the **Peer response timeout (seconds)** field. The Cisco VPN Client continues to send DPD requests every five seconds, until it reaches the number of seconds specified by the peer response timeout value.

Creating a New Connection—Backup Servers

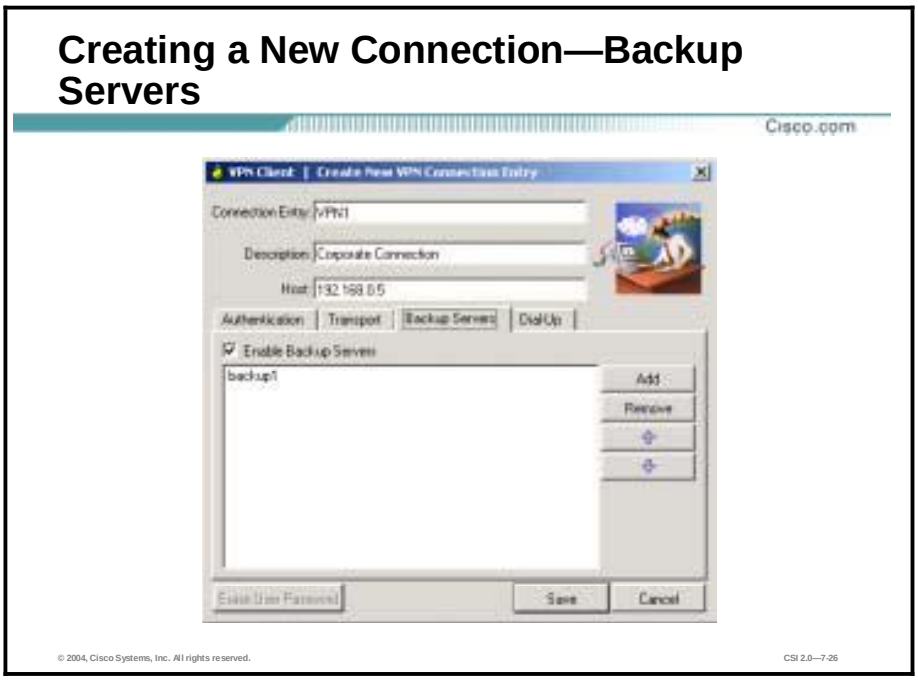

The private network may include one or more backup VPN servers to use if the primary server is not available. Your system administrator tells you whether to enable backup servers. Information on backup servers can download automatically from the Concentrator, or you can manually enter the information.

Creating a New Connection—Dial-Up

Cisco.com

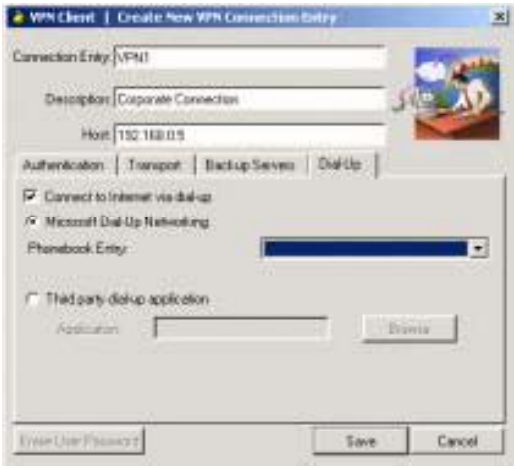

CSI 2.0—7-27

To enable and configure a connection to the Internet through dial-up networking, select the **Connect to Internet via dial-up** check box. This feature is not selected by default.

You can connect to the Internet using the Cisco VPN Client application in either of the following ways:

- Microsoft Dial-Up Networking (DUN)—If you have DUN phonebook entries and have enabled the Connect to Internet via dial-up feature, Microsoft DUN is enabled by default. To link your Cisco VPN Client connection entry to a DUN entry, click the **Phonebook Entry** drop-down arrow and choose an entry from the menu. The Cisco VPN Client then uses this DUN entry to automatically dial into the Microsoft network before making the VPN connection to the private network.
- Third-party dial-up program—If you have no DUN phonebook entries and have enabled the Connect to Internet via dial-up feature, then the third-party dial-up application is enabled by default. Click the **Browse** button to enter the name of the program in the Application field. This application launches the connection to the Internet. The string you choose or enter in this field is the path name to the command that starts the application and the name of the command; for example: c:\isp\ispdialer.exe dial Engineering. Your network administrator might have set the third-party dial-up program up for you. If not, consult your network administrator.

Remote Site Firewall Option

This topic provides details of the remote site firewall access option for remote users.

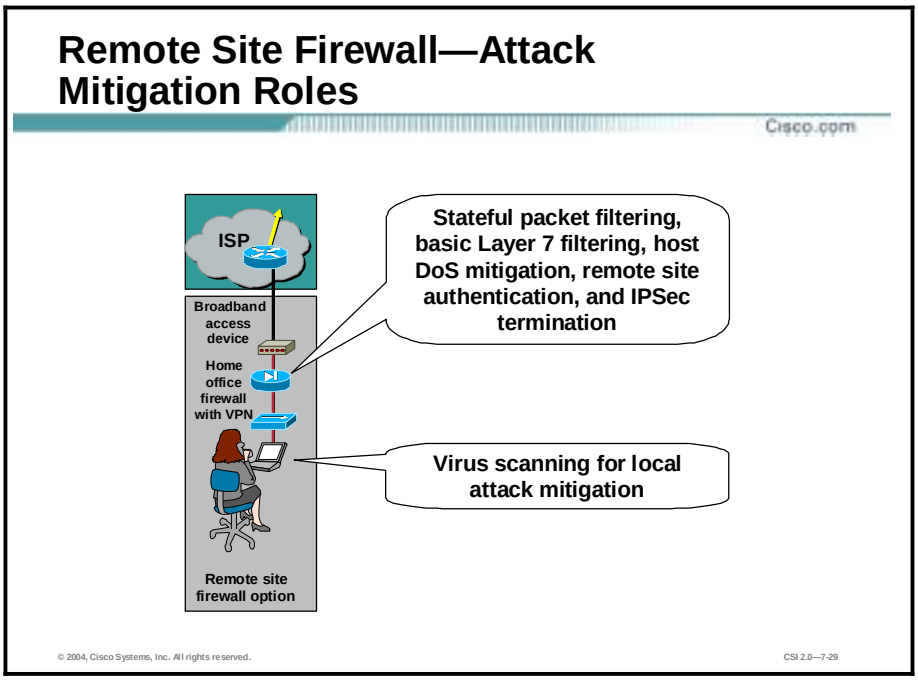

This figure provides attack mitigation roles for the remote site firewall option:

- Stateful packet filtering—Offers strong security by thoroughly inspecting data packets and maintaining critical addresses and port numbers in a lookup table
- Basic Layer 7 filtering—Offers basic filtering at the application layer of the OSI model
- Host DoS mitigation—Prevents host-based denial of service (DoS) attacks
- Remote site authentication—Properly identifies and verifies a user or service
- IPSec termination—Successfully establishes an IPSec tunnel between the remote site and host site
- Virus scanning for local attack mitigation—Allays the risk of virus infection at the remote site

Remote Site Firewall—Design Guidelines

Cisco.com

- Geared toward a home-office worker or a very small branch office
- Firewall provides connection-state enforcement
- Termination point for site-to-site IPSec
  - For remote management and production
  - Client does not need individual software (firewall or VPN)
  - NAT is not used in IPSec tunnel
  - Device authentication is used at head-end
  - Allows split tunneling
- Authentication is controlled from the headquarters
- Virus checking software is still recommended
- Personal firewall software can be used to protect the remote user when split tunneling is enabled
- You can use an IDS on a PIX Firewall

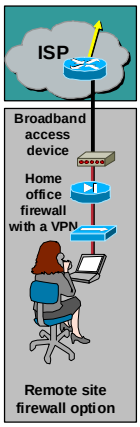

The remote-site firewall option is geared toward the home-office worker, or potentially a very small branch office. The remote site should have some form of broadband access available from a service provider. The firewall is installed behind the DSL or cable modem.

The primary function of the firewall is to establish the secure, encrypted tunnel between itself and a VPN head-end device, as well as providing connection-state enforcement and detailed filtering for sessions initiated through it. Individual PCs on a remote-site network do not need VPN Client software to access corporate resources. Additionally, because the stateful firewall protects access to the Internet, personal firewall software is not required on the individual PCs. However, to provide an additional level of security, a network administrator may implement personal firewall software on remote-site PCs. This setup may be useful if the home worker also travels and connects to the Internet directly over some public network. Because the corporate headquarters has a stateful firewall protecting the hosts, the remote site can have direct access to the Internet, rather than passing all traffic back through the corporate headquarters. Unless NAT is used when communicating with the headquarters, the IP addresses of the remote-site devices should be assigned in such a manner as to not overlap addressing space in the headquarters location or another remote site. Remote-site devices that require direct access to the Internet will require address translation to a registered address. This address translation can be achieved by translating all Internet-bound sessions to the public IP address of the firewall itself.

Access and authorization to the corporate network and the Internet are controlled by the configuration of both the remote-site firewall and the VPN head-end device. Configuration and security management of the remote-site firewall can be achieved via an IPSec tunnel from the public side of the firewall back to the corporate headquarters. This setup ensures that the remote-site user is not required to perform any configuration changes on the home-office firewall. Authentication should be set up on the firewall to prevent a local user from inadvertently modifying their firewall configuration and thereby compromising the security policy of that device. Individual users at the remote site who access the corporate network are not

authenticated with this option. Instead, the remote-site firewall and VPN head-end use device authentication.

Virus-scanning software is still recommended to mitigate against viruses and Trojan horse programs infecting individual PCs at the remote site—just like all the PCs in the entire corporation.

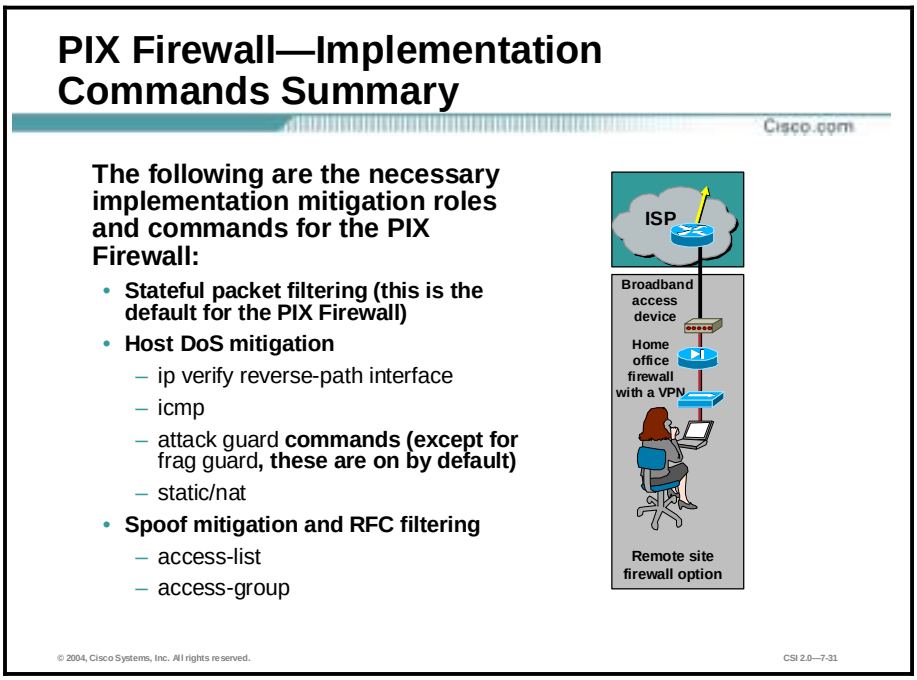

The following mitigation roles and commands are used to implement the policy on the Cisco PIX Firewall:

- Stateful packet filtering (this is the default for the PIX Firewall)
- Host DoS mitigation
  - **ip verify reverse-path interface**—Implements Unicast Reverse Path Forwarding (RPF) IP spoofing protection
  - **icmp**—Enables or disables pinging to an interface
  - **attack guard** commands—Enabled by default
  - **static/nat**—Implements static or dynamic NAT
- Spoof mitigation and RFC filtering
  - **access-list**—Creates an access control list (ACL)
  - **access-group**—Binds the ACL to an interface

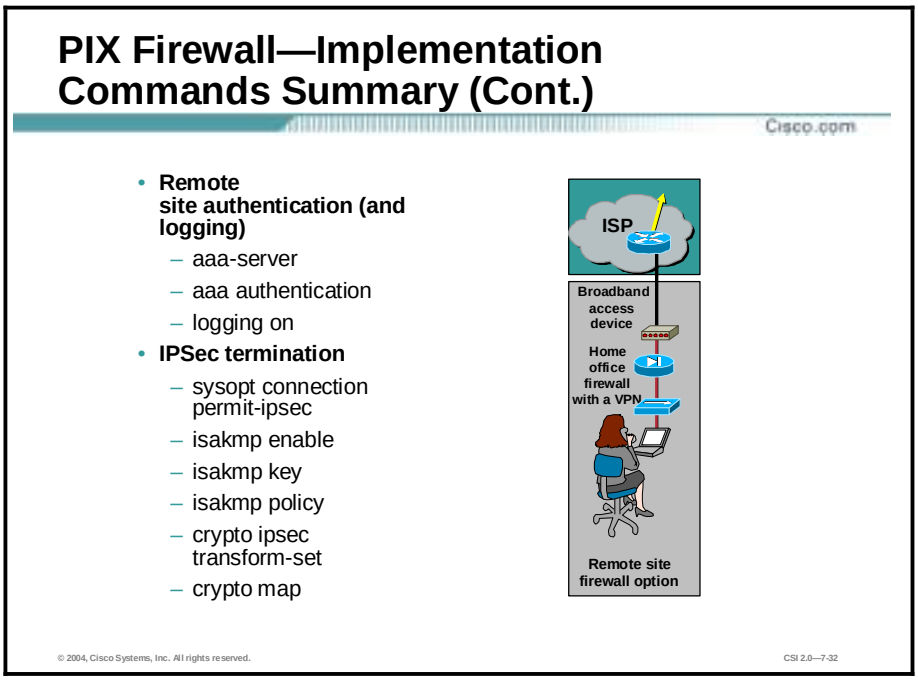

- Remote site authentication (and logging)
  - **aaa-server—**Specifies an authentication, authorization, and accounting (AAA) server
  - **aaa authentication**—Enables, disables, or views local, Terminal Access Controller Access Control System Plus (TACACS+), or Remote Access Dial-In User Service (RADIUS) user authentication
  - **logging on**—Enables or disables Syslog and Simple Network Management Protocol (SNMP) logging
- IPSec termination
  - **sysopt connection permit-ipsec**—Implicitly permits any packet that came from an IPSec tunnel
  - **isakmp enable**—Enables Internet Security Association Key Management Protocol (ISAKMP) negotiation on the interface on which the IPSec peer communicates with the PIX Firewall
- **isakmp key**—Specifies the authentication pre-shared key
  - **isakmp policy**—Uniquely identifies the IKE policy and assigns a priority to the policy
  - **crypto ipsec transform-set**—Creates, views, or deletes IPSec security associations (SAs), SA global lifetime values, and global transform sets
  - **crypto map**—Creates, modifies, views, or deletes a crypto map entry

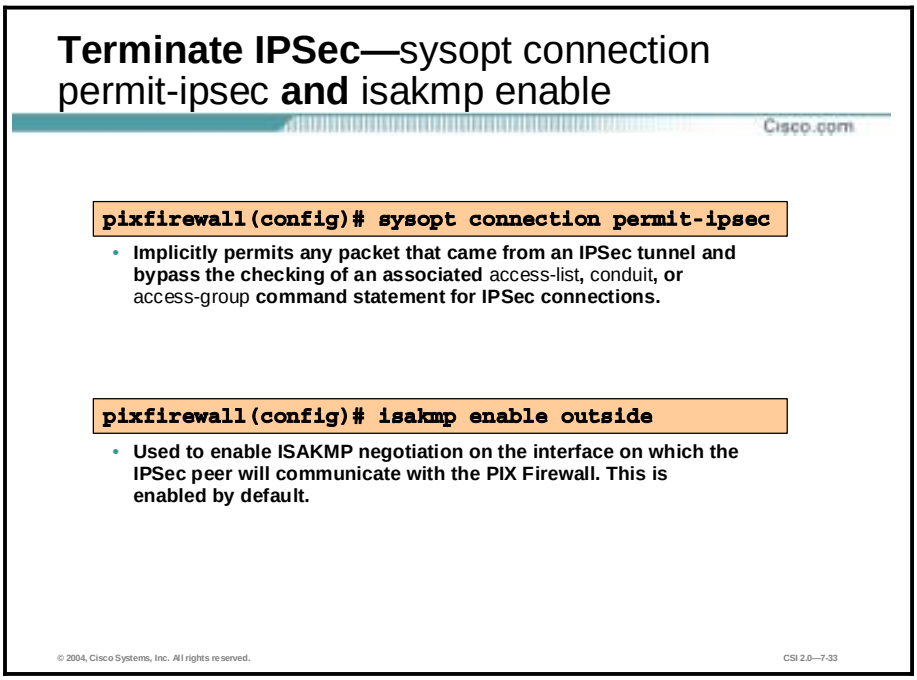

The **sysopt connection permit-ipsec** command implicitly permits any packet that came from an IPSec tunnel and bypasses the checking of an associated **access-list, conduit,** or **access-group** command statement for IPSec connections.

The **isakmp enable** command is used to enable ISAKMP negotiation on the interface on which the IPSec peer will communicate with the PIX Firewall. Use the **no isakmp enable** command to disable IKE.

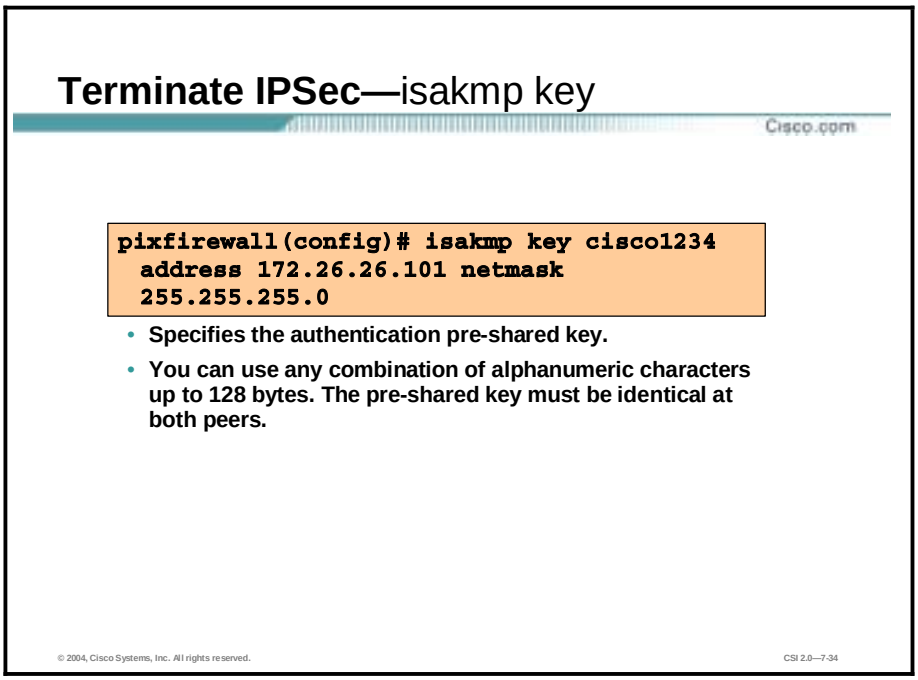

To configure a pre-shared authentication key and associate the key with an IPSec peer address or hostname, use the **isakmp key address** command. Use the **no isakmp key address** command to delete a pre-shared authentication key and its associated IPSec peer address.

You must configure the pre-shared key at both peers whenever you specify pre-shared key in an IKE policy. Otherwise, the policy cannot be used because it will not be submitted for matching by the IKE process.

A netmask of 0.0.0.0. can be entered as a wildcard indicating that any IPSec peer with a valid pre-shared key is a valid peer.

Note	The PIX Firewall or any IPSec peer can use the same authentication key with multiple peers, but this is not as secure as using a unique authentication key between each pair of peers.

The **no-xauth** or **no-config-mode** command options are to be used only if the following criteria are met:

- You are using the pre-shared key authentication method within your IKE policy.
- The security gateway and VPN Client peers terminate on the same interface.
- The Xauth or IKE Mode Configuration feature is enabled for VPN Client peers.

Both the Xauth and IKE Mode Configuration features are specifically designed for remote VPN Clients. The Xauth feature enables the PIX Firewall to challenge the peer for a username and password during IKE negotiation. The IKE Mode Configuration enables the PIX Firewall to download an IP address to the peer for dynamic IP address assignment. Most security gateways do not support the Xauth and IKE Mode Configuration features.

If you have the **no-xauth** command option configured, the PIX Firewall does not challenge the peer for a username and password. Similarly, if you have the **no-config-mode** command option configured, the PIX Firewall does not attempt to download an IP address to the peer for dynamic IP address assignment.

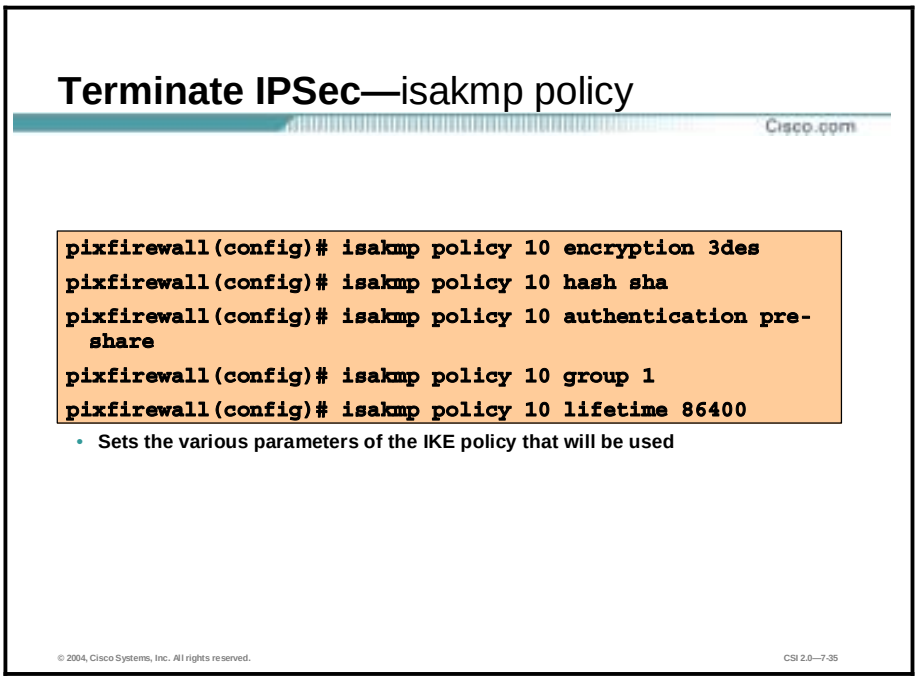

The **isakmp policy** command enables you to negotiate IPSec SAs and enable IPSec secure communications. Several parameters can be configured with this command as illustrated in the figure.

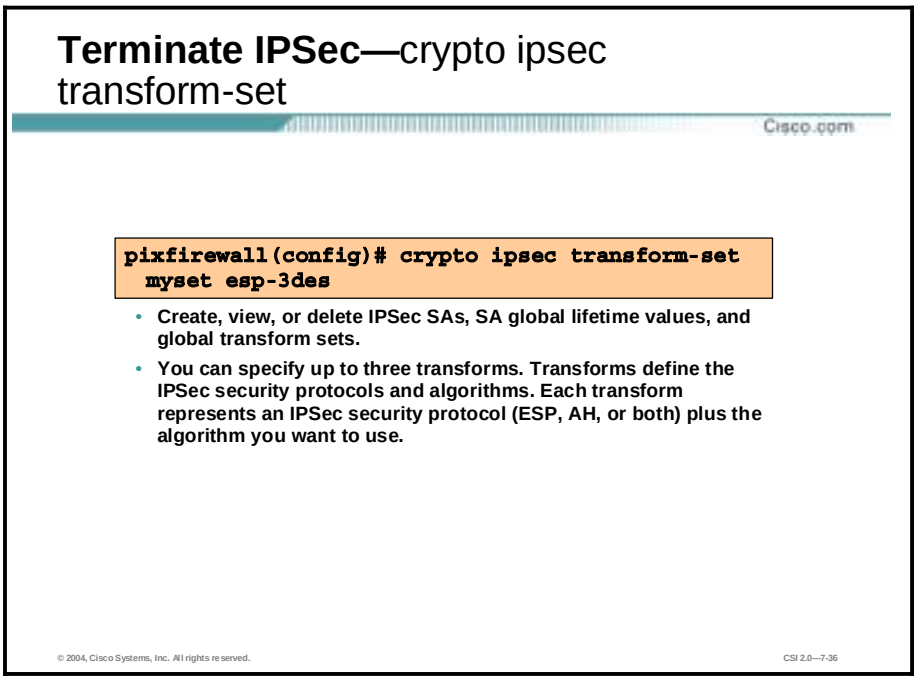

The **crypto ipsec transform-set** command defines a transform set. To delete a transform set, use the **no crypto ipsec transform-set** command. To view the configured transform sets, use the **show crypto ipsec transform-set** command.

A transform set specifies one or two IPSec security protocols (either ESP or Authentication Header [AH], or both) and specifies which algorithms to use with the selected security protocol. During the IPSec SA negotiation, the peers agree to use a particular transform set when protecting a particular data flow.

You can configure multiple transform sets, and then specify one or more of these transform sets in a crypto map entry. The transform set defined in the crypto map entry is used in the IPSec SA negotiation to protect the data flows specified by that crypto map entry's ACL. During the negotiation, the peers search for a transform set that is the same at both peers. When such a transform set is found, it is selected and is applied to the protected traffic as part of both peer's IPSec SAs.

When SAs are established manually, a single transform set must be used. The transform set is not negotiated.

Before a transform set can be included in a crypto map entry, it must be defined using the **crypto ipsec transform-set** command.

To define a transform set, you specify one to three "transforms"—each transform represents an IPSec security protocol (ESP or AH) plus the algorithm you want to use. When the particular transform set is used during negotiations for IPSec SAs, the entire transform set (the combination of protocols, algorithms, and other settings) must match a transform set at the remote peer.

In a transform set you can specify the AH protocol, the ESP protocol, or both. If you specify an ESP protocol in a transform set, you can specify just an ESP encryption transform or both an ESP encryption transform and an ESP authentication transform.

Examples of acceptable transform combinations are as follows:

- ah-md5-hmac
- esp-des
- esp-des and esp-md5-hmac
- ah-sha-hmac and esp-des and esp-sha-hmac

If one or more transforms are specified in the **crypto ipsec transform-set** command for an existing transform set, the specified transforms will replace the existing transforms for that transform set.

If you change a transform set definition, the change is only applied to crypto map entries that reference the transform set. The change will not be applied to existing SAs, but will be used in subsequent negotiations to establish new SAs. If you want the new settings to take effect sooner, you can clear all or part of the security association database by using the **clear [crypto] ipsec sa** command.

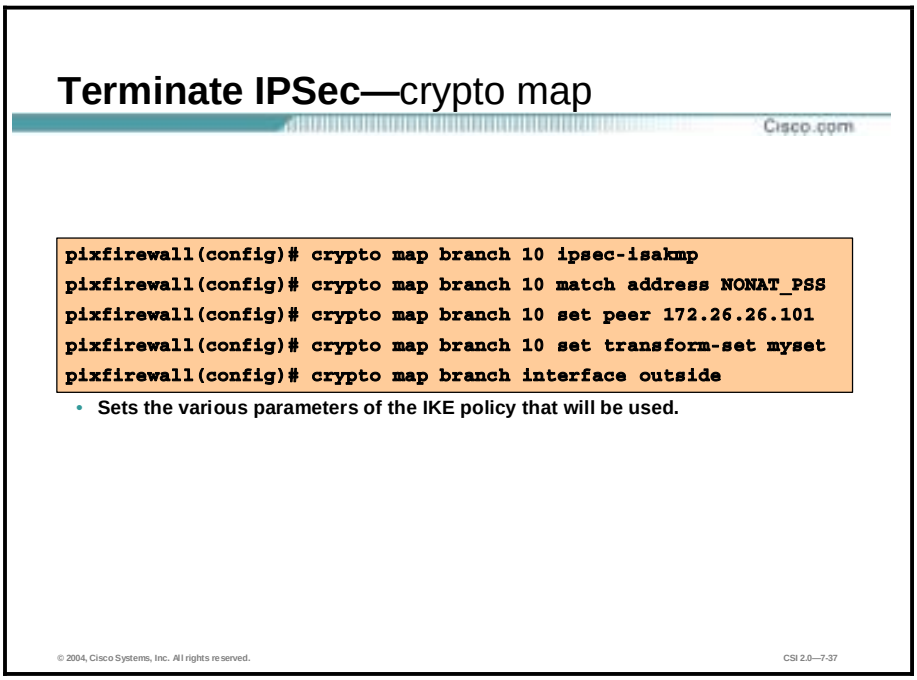

To create or modify a crypto map entry, use the **crypto map ipsec-manual** | **ipsec-isakmp** command. To create or modify an ipsec-manual crypto map entry, use the **ipsec-manual option** of the command. To create or modify an ipsec-isakmp crypto map entry, use the **ipsec-isakmp** option of the command. Use the **no crypto map** command to delete a crypto map entry or set.

Crypto maps provide two functions: filtering and classifying traffic to be protected, and defining the policy to be applied to that traffic. The first function affects the flow of traffic on an interface; the second affects the negotiation performed (via IKE) on behalf of that traffic.

IPSec crypto maps link together definitions of the following:

- What traffic should be protected?
- To which IPSec peers should the protected traffic be forwarded? These are the peers with which an SA can be established.
- Which transform sets are acceptable for use with the protected traffic?
- How should keys and SAs be used and managed? (Or what are the keys, if IKE is not used)

A crypto map set is a collection of crypto map entries, each with a different seq-num but the same map-name. Therefore, for a given interface, you could have certain traffic forwarded to one peer with specified security applied to that traffic, and other traffic forwarded to the same or a different peer with different IPSec security applied. To accomplish this you would create two crypto map entries, each with the same map-name, but each with a different seq-num.

The number you assign to the seq-num argument should not be arbitrary. This number is used to rank multiple crypto map entries within a crypto map set. Within a crypto map set, a crypto map entry with a lower seq-num is evaluated before a map entry with a higher seq-num; that is, the map entry with the lower number has a higher priority.

VPN Hardware Client Option

This topic covers the use of the Cisco VPN Hardware Client for remote access.

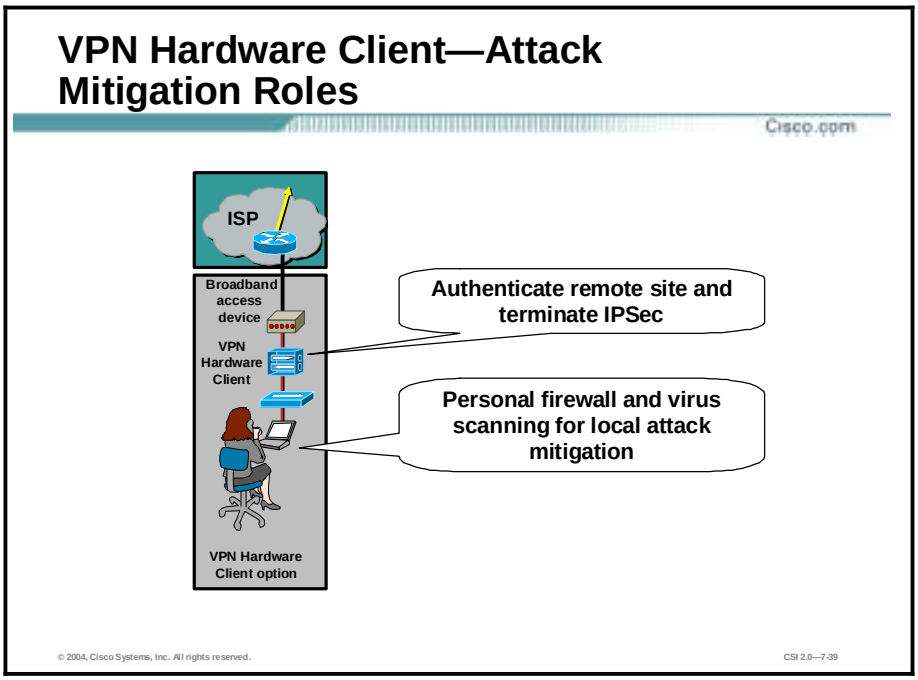

The following are the mitigation roles for the VPN Hardware Client:

- Authenticate remote site—Properly identifies and verifies a user or service
- Terminate IPSec—Successfully establishes an IPSec tunnel between the remote site and host site
- Personal firewall and virus scanning for local attack mitigation—Provides firewall inspection and allays the risk of virus infection at the remote site

VPN Hardware Client—Design Guidelines

Cisco.com

- Same guidelines as remote site firewall option except that the VPN Hardware Client does not have resident stateful firewall
- Use a personal firewall on individual hosts (if split tunneling will be used)
- If no personal firewall is in use, security behind the VPN device is dependent upon NAT (with split tunneling enabled)
- Access and authentication are controlled from the headquarters
- Configuration and security management is done from the headquarters
- VPN Client software is not needed

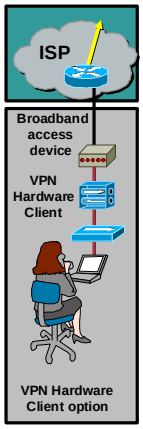

CSI 2.0—7-40

The VPN Hardware Client option is identical to the remote-site firewall option except that the VPN Hardware Client does not have a resident stateful firewall. This setup requires use of a personal firewall on the individual hosts, particularly when split tunneling is enabled. Without the personal firewall, the security of the individual hosts behind the VPN device is dependent upon the attacker being unable to circumvent NAT. This is because when split tunneling is enabled, connections to the Internet pass through a many-to-one NAT translation and do not undergo any filtering at Layer 4 and above. With split tunneling disabled, all access to the Internet must be through the corporate headquarters. This setup partially mitigates the requirement for personal firewalls on the end systems.

Using a VPN Hardware Client offers two primary advantages:

- As with the VPN Software Client, access and authorization to the corporate network and the Internet are controlled centrally from the headquarters location. Configuration and security management of the VPN Hardware Client device itself is done via a Secure Sockets Layer (SSL) connection from the central site. This setup ensures that the remote-site user is not required to perform any configuration changes on the VPN Hardware Client.
- Individual PCs on the remote-site network do not need VPN Client software to access corporate resources. However, individual users at the remote site who access the corporate network are not authenticated with this option. Instead, the VPN Hardware Client and VPN head-end Concentrator authenticate each other.

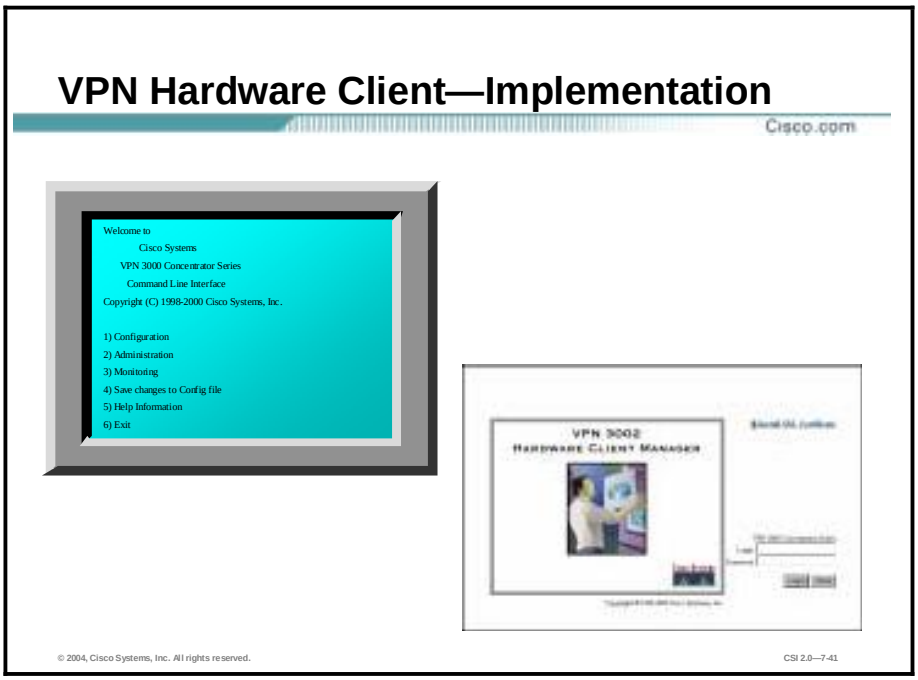

Once connected, the administrator must gain access to the VPN 3002 Hardware Client Manager. To gain access, complete the following steps:

Step 1 The VPN 3002 hardware client comes from the factory with a private interface IP address of 192.168.10.1. Hook up a PC to the private port and configure the TCP/IP address of the PC.

Step 2 Point the browser to the IP address of the private interface (for example, http://192.168.10.1).

Step 3 Log in using the login name and password *admin*. No command line interface (CLI) intervention is required.

However, if you would rather configure the VPN 3002 via CLI or if you need to change the default address on the private LAN interface, you can use the CLI. The default CLI setting is 9600 8N1.

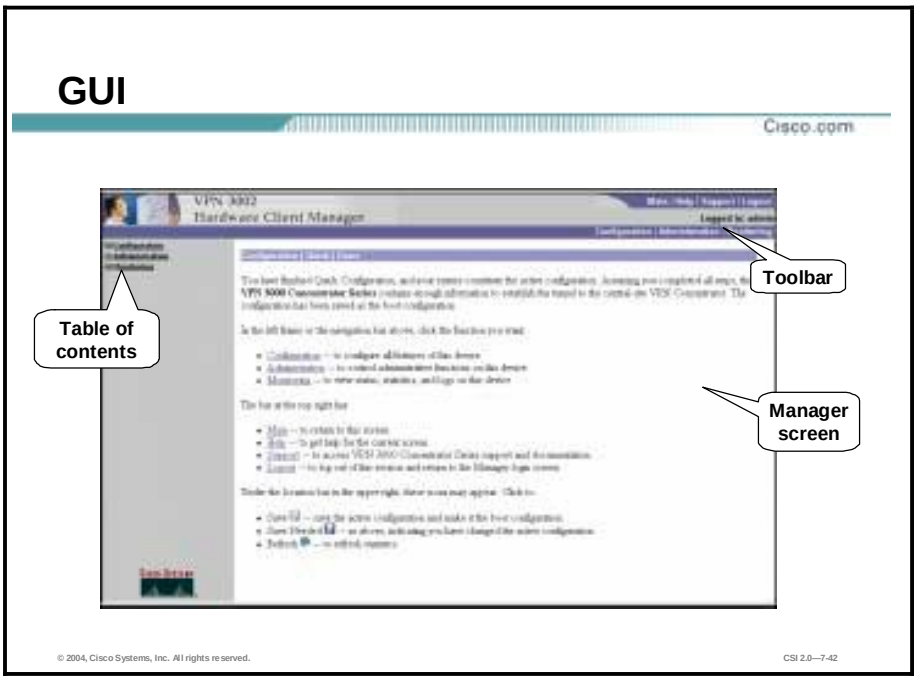

The main window of the VPN 3002 Hardware Client Manager after logging into the device consists of the following:

- The top frame (Manager toolbar) provides quick access to Manager functions: configuration, administration, and monitoring.
- The left frame provides the table of contents (TOC) to the Manager's windows.
- The main frame displays the current Manager window. From here you can navigate the Manager using either the TOC in the left frame or the Cisco toolbar in the top frame.

Under the location bar at the upper right, the **Save Needed** icons may appear. When finished with a configuration window, click **Apply**. Apply enables the configuration to take effect immediately. To save the changes to memory, click the **Save Needed** icon. If you reboot without saving, your configuration changes are lost.

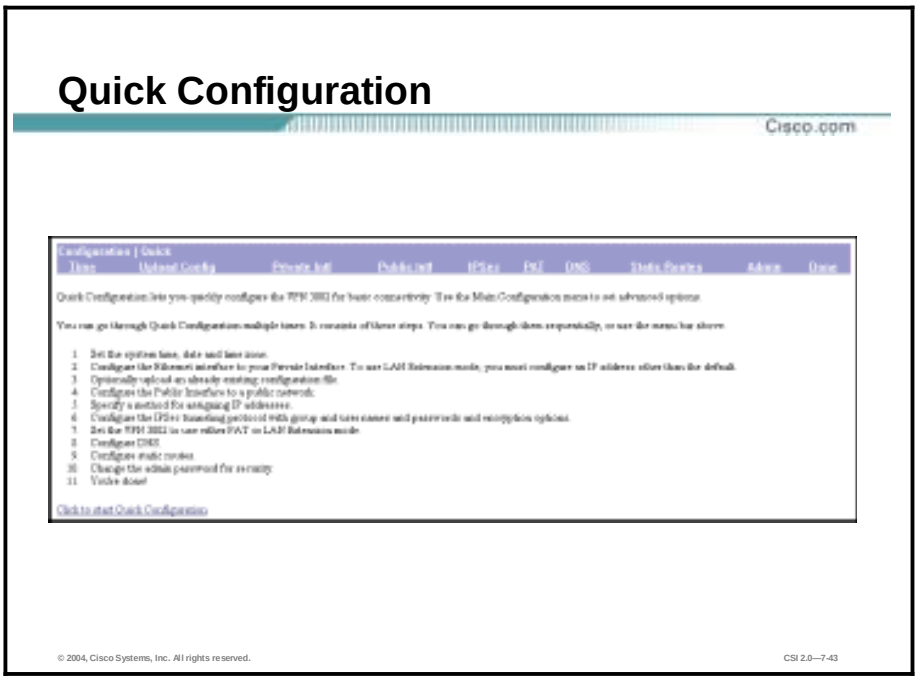

There are two ways to configure the Hardware Client: Quick Configuration and the main menu. The goal of Quick Configuration is to provide the *minimal* parameters needed for operation. You can access Quick Configuration from the Configuration>Quick window. The VPN 3002 Quick Configuration can be run multiple times.

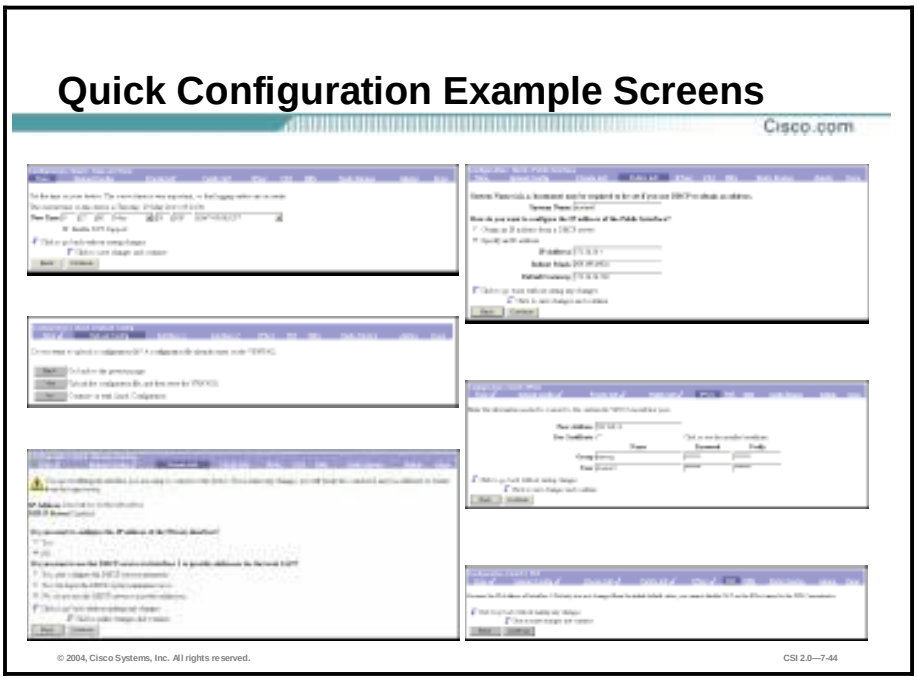

Quick Configuration guides you through the windows necessary to get a single tunnel up and running. Use the main menu to tune an application or configure features individually. The windows in the figure illustrate some of the IPSec remote access configuration screens using Quick Configuration.

Launching the Client

Cisco.com

CSI 2.0—7-45

The Hardware Client is configured, and now the tunnel needs to be established. In Client Mode, by default, there is no tunnel established. There are two ways to initiate a tunnel:

- By clicking **Connect Now**, under the Monitoring>System Status window
- By sending traffic to the Hardware Client destined for the remote end

You can verify the configuration by trying to ping an interface on the remote Concentrator. The Hardware Client recognizes the remote-bound traffic and attempts to establish a tunnel. If a tunnel is established, it is viewable on this screen. If the tunnel does not come up, check the event log of the Hardware Client and the Concentrator.

Remote Site Router Option

This topic summarizes the remote-site router option and provides configuration details.

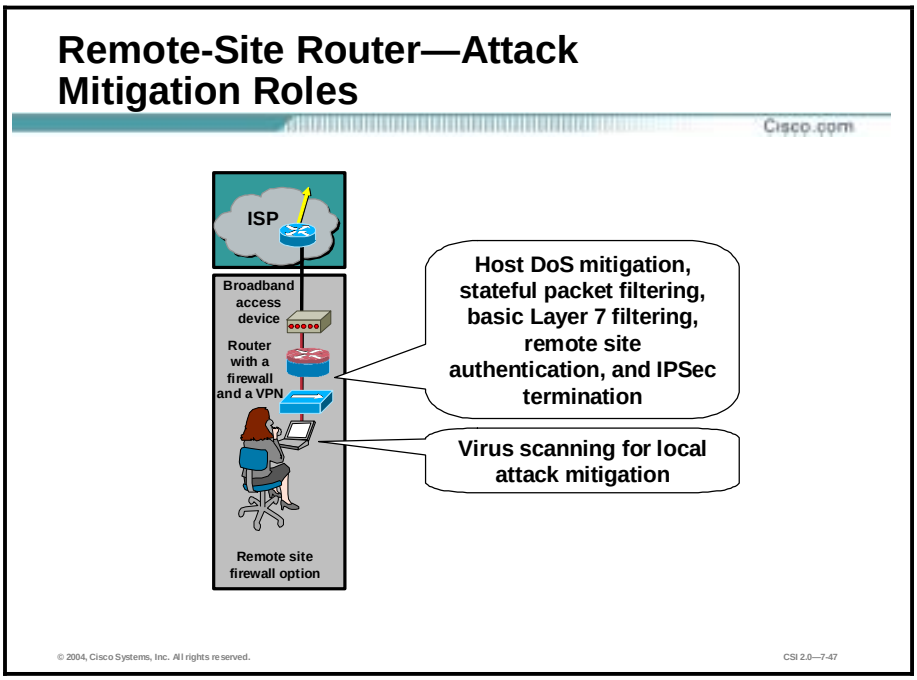

The following are the attack mitigation roles for the remote-site router option:

- Host DoS mitigation—Prevents host-based DoS attacks
- Stateful packet filtering—Offers strong security by thoroughly inspecting data packets and maintaining critical addresses and port numbers in a lookup table
- Basic Layer 7 filtering—Offers basic filtering at the application layer of the OSI model
- Authenticate remote site—Properly identifies and verifies a user or service
- Terminate IPSec—Successfully establishes an IPSec tunnel between the remote site and host site
- Virus scanning for local attack mitigation—Allays the risk of virus infection at the remote site

Remote Site Router—Design Guidelines

Cisco.com

- Uses the same guidelines as the remote site firewall option.
- The router can support QoS, routing, and more encapsulation options.
- Broadband capability can be integrated into the router:
  - This removes the need for a separate broadband access device.
  - This is typically managed by a service provider.
- An IDS on a router can be used (may not be available on all router platforms).

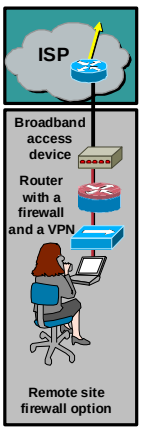

The remote-site router option is nearly identical to the remote-site firewall option with a few exceptions. When deployed behind a standalone broadband access device, the only difference is that the router can support advanced applications such as QoS, routing, and more encapsulation options. Additionally, if broadband capability is integrated into the router, a standalone broadband access device is not needed. This option requires that your ISP allows you to manage the broadband router, which is not a common scenario.

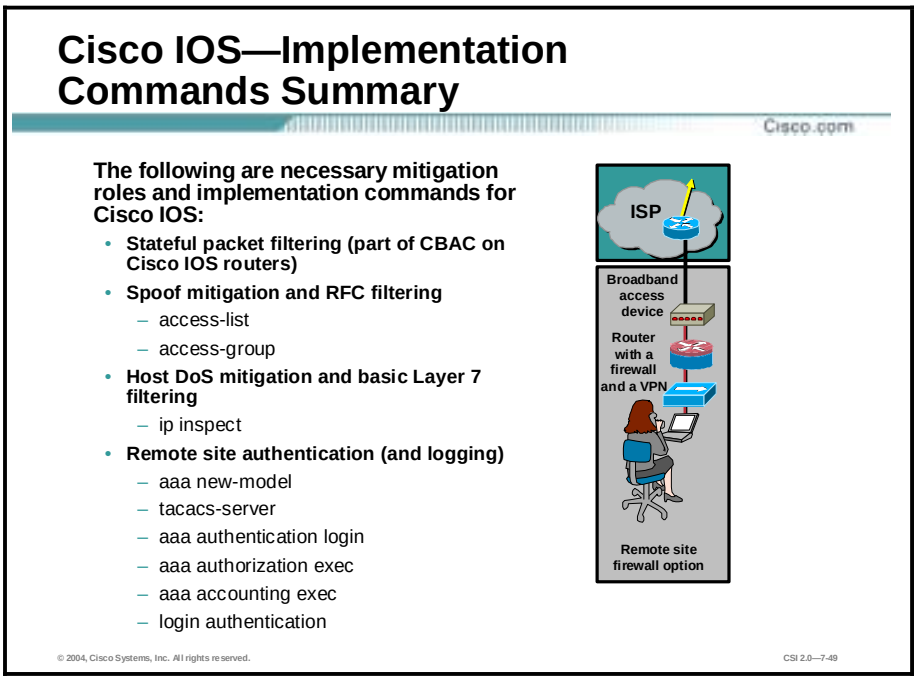

The following are necessary mitigation roles and implementation commands for the IOS Firewall:

- Stateful packet filtering—Part of Context-Based Access Control (CBAC) on Cisco IOS routers
- Spoof mitigation and RFC filtering
  - **access-list**—The **access-list** command enables you to specify if an IP address is permitted or denied access to a port or protocol.
  - **access-group**—Binds an ACL to an interface.
- Host DoS mitigation and basic Layer 7 filtering
  - **ip inspect**—Defines the application protocols to inspect.
- Remote site authentication (and logging)
  - **aaa new-model**—To define a set of inspection rules, enter this command for each protocol that you want to inspect, using the same inspection-name.
  - **tacacs-server**—Defines the TACACS server.
  - **aaa authentication login**—Enables AAA authentication at login.
  - **aaa authorization exec**—Restricts network access to a user.
  - **aaa accounting exec**—Runs accounting for EXEC shell session.
  - **login authentication**—Specifies the name of a list of AAA authentication methods to try at login.

Cisco IOS—Implementation Commands Summary (Cont.)

Cisco.com

- **IPSec commands—Provide for IPSec tunnel termination**
  - crypto isakmp policy
  - encryption
  - authentication
  - group
  - crypto isakmp key
  - crypto ipsec transform-set
  - crypto map
  - set peer
  - set tranform-set
  - match address

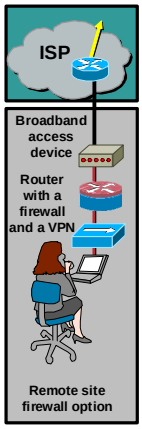

- IPSec commands—Provide for IPSec tunnel termination
  - **crypto isakmp policy**—Specifies the parameters to be used during an IKE negotiation.
  - **encryption**—Sets the algorithm to be negotiated.
  - **authentication**—Specifies the authentication method within an IKE policy.
  - **group**—Specifies the AAA server group to use for pre-authentication.
  - **crypto isakmp key**—Configures pre-shared authentication keys.
  - **crypto ipsec transform-set**—An acceptable combination of security protocols, algorithms and other settings to apply to IP security protected traffic.
  - **crypto map**—Configures filtering and classifying traffic to be protected and defines the policy to be applied to that traffic.
  - **set peer**—Specifies an IPSec peer for a crypto map.
  - **set transform-set**—Specifies which transform sets to include in a crypto map entry.
  - **match address**—Specifies an extended access list for a crypto map entry.

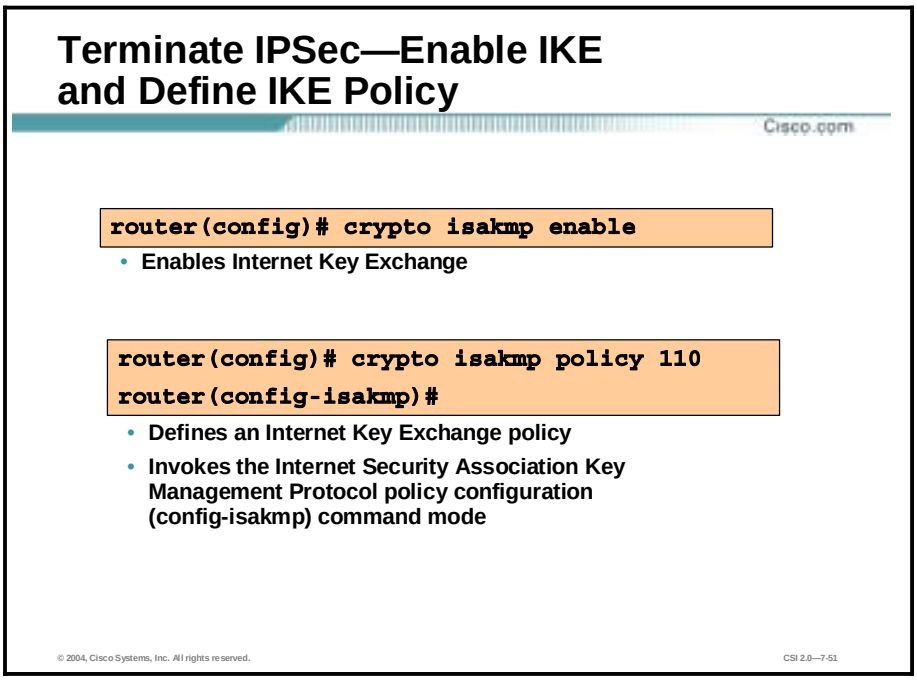

IKE is enabled by default. IKE does not have to be enabled for individual interfaces, but is enabled globally for all interfaces at the router.

If you do not want IKE to be used in your IPSec implementation, you can disable IKE at all your IPSec peers. If you disable IKE at one peer, you must disable it at all your IPSec peers.

If you disable IKE, you will have to make the following concessions at the peers:

- You must manually specify all the IPSec SAs in the crypto maps at the peers.
- The IPSec SAs of the peers never time out for a given IPSec session.
- During IPSec sessions between the peers, the encryption keys never change.
- Anti-replay services are not available between the peers.
- Certification Authority (CA) support cannot be used.

IKE negotiations must be protected, so each IKE negotiation begins by agreement of both peers on a common (shared) IKE policy. This policy states which security parameters will be used to protect subsequent IKE negotiations and mandates how the peers are authenticated.

After the two peers agree upon a policy, the security parameters of the policy are identified by an SA established at each peer, and these SAs apply to all subsequent IKE traffic during the negotiation.

You can create multiple, prioritized policies at each peer to ensure that at least one policy will match the policy of a remote peer.

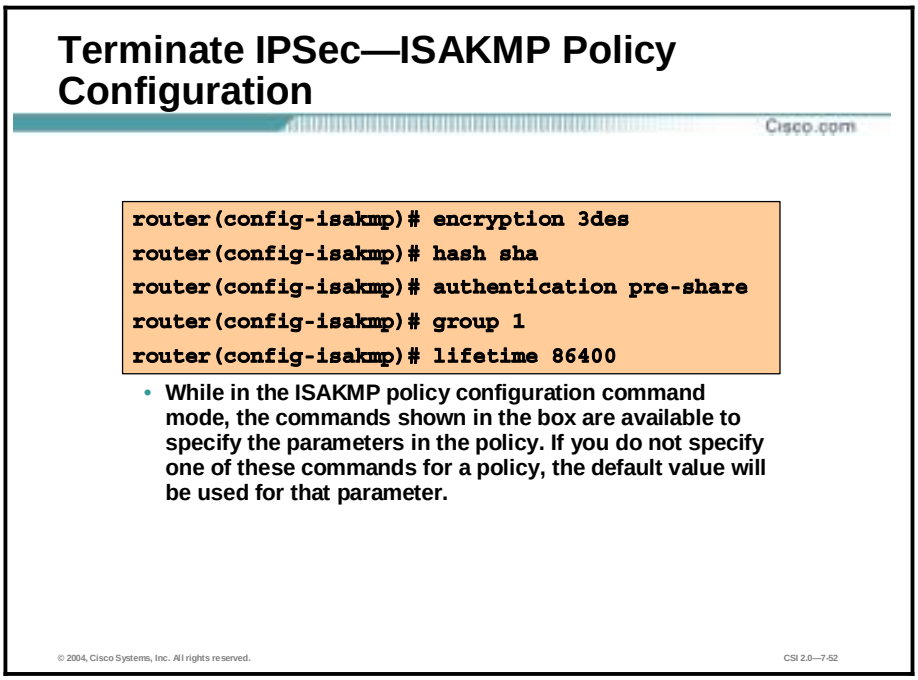

If you are interoperating with a device that supports only one of the values for a parameter, your choice is limited to the value supported by the other device. Aside from this, there is often a trade-off between security and performance, and many of these parameter values represent such a trade-off. You should evaluate the level of security risks for your network and your tolerance for these risks. After doing that, the following tips might help you select which value to specify for each parameter:

- The encryption algorithm has two options: 56-bit Data Encryption Standard-cipher block chaining (DES-CBC), and 168-bit DES.
- The hash algorithm has two options: Secure Hash Algorithm-1 (SHA-1), and Message Digest 5 (MD5). MD5 has a smaller digest and is considered to be slightly faster than SHA-1. There has been a demonstrated successful (but extremely difficult) attack against MD5; however, the hash-based message authentication code (HMAC) variant used by IKE prevents this attack.
- The authentication method has three options: Rivet, Shamir, and Adelman (RSA) signatures, RSA encrypted nonces, and pre-shared keys.
  - RSA signatures provide nonrepudiation for the IKE negotiation (you can prove to a third party after the fact that you did indeed have an IKE negotiation with the remote peer). RSA signatures allow the use of a CA. Using a CA can dramatically improve the manageability and scalability of your IPSec network. Additionally, RSA signature-based authentication uses only two public key operations, whereas RAS encryption uses four public key operations, making it costlier in terms of overall performance.
  - RSA encrypted nonces provide repudiation for the IKE negotiation. The RSA-encrypted nonces method uses the RSA encryption public key cryptography standard. It requires that each party generate a pseudo-random number (a nonce) and encrypt it in the other party's RSA public key. Authentication occurs when each party decrypts the other party's nonce with a local private key (and other publicly and privately available information) and then uses the decrypted nonce to compute a keyed hash.

This system provides for deniable transactions. That is, either side of the exchange can plausibly deny that it took part in the exchange.

— Pre-shared keys are clumsy to use if your secured network is large, and they do not scale well with a growing network. However, they do not require use of a CA, as do RSA signatures, and might be easier to set up in a small network with fewer than ten nodes. RSA signatures also can be considered more secure when compared with pre-shared key authentication.

- The Diffie-Hellman (DH) group identifier has two options: 768-bit and 1024-bit DH. The 1024-bit DH option is harder to crack, but requires more CPU time to execute.
- The lifetime of the SA can be set to any value. As a general rule, the shorter the lifetime (up to a point), the more secure your IKE negotiations will be. However, with longer lifetimes, future IPSec SAs can be set up more quickly.

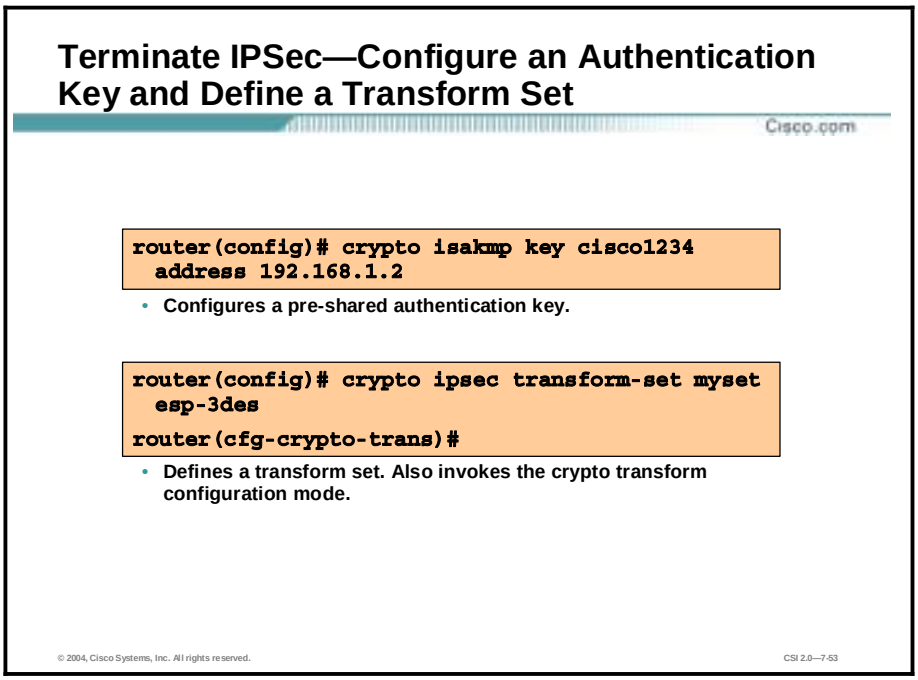

Complete the following steps at each peer that uses pre-shared keys in an IKE policy to configure pre-shared keys:

- Set the ISAKMP identity of each peer. Each peer's identity should be set to either its hostname or IP address. By default, a peer's identity is set to its IP address.
- Specify the shared keys at each peer. Note that a given pre-shared key is shared between two peers. At a given peer you could specify the same key to share with multiple remote peers; however, a more secure approach is to specify different keys to share between different pairs of peers.

A transform set is an acceptable combination of security protocols, algorithms, and other settings to apply to IPSec-protected traffic. During the IPSec SA negotiation, the peers agree to use a particular transform set when protecting a particular data flow.

You can configure multiple transform sets, and then specify one or more of these transform sets in a crypto map entry. The transform set defined in the crypto map entry is used in the IPSec SA negotiation to protect the data flows specified by the ACL of that crypto map entry. During the negotiation, the peers search for a transform set that is the same at both peers. When such a transform set is found, it is selected and will be applied to the protected traffic as part of the IPSec SAs of both peers.

To define a transform set, you specify one to three "transforms." Each transform represents an IPSec security protocol (ESP or AH) plus the algorithm you want to use. When the particular transform set is used during negotiations for IPSec SAs, the entire transform set (the combination of protocols, algorithms, and other settings) must match a transform set at the remote peer.

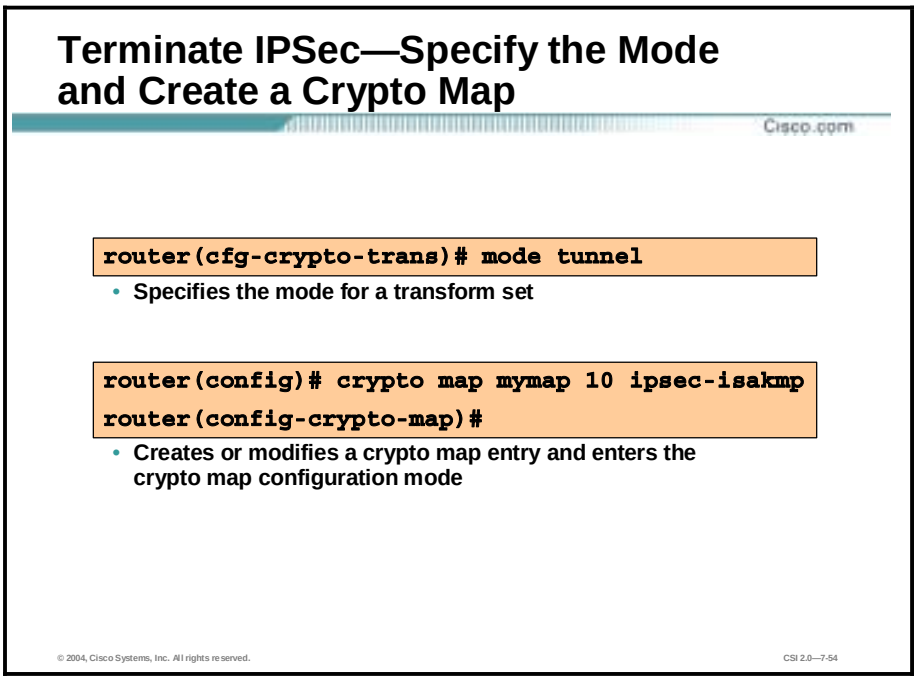

After you define a transform set, you are put into the crypto transform configuration mode. While in this mode you can change the mode to either tunnel or transport. This change applies only to the transform set just defined.

If the traffic to be protected has the same IP address as the IPSec peers and transport mode is specified, during negotiation the router requests transport mode but accepts either transport or tunnel mode. If tunnel mode is specified, the router requests tunnel mode and accepts only tunnel mode.

If you do not change the mode when you first define the transform set, but later decide you want to change the mode for the transform set, you must reenter the transform set (specifying the transform name and all its transforms) and then change the mode.

If you use this command to change the mode, the change only affects the negotiation of subsequent IPSec SAs via crypto map entries, which specify this transform set. (If you want the new settings to take effect sooner, you can clear all or part of the SA database.)

Crypto map entries created for IPSec pull together the various information used to set up IPSec SAs, including the following:

- Which traffic should be protected by IPSec (per a crypto ACL)
- The granularity of the flow to be protected by a set of SAs
- Where IPSec-protected traffic should be sent (who the remote IPSec peer is)
- The local address to be used for the IPSec traffic
- What IPSec security should be applied to this traffic (selecting from a list of one or more transform sets)
- Whether SAs are manually established or are established via IKE
- Other parameters that might be necessary to define an IPSec SA

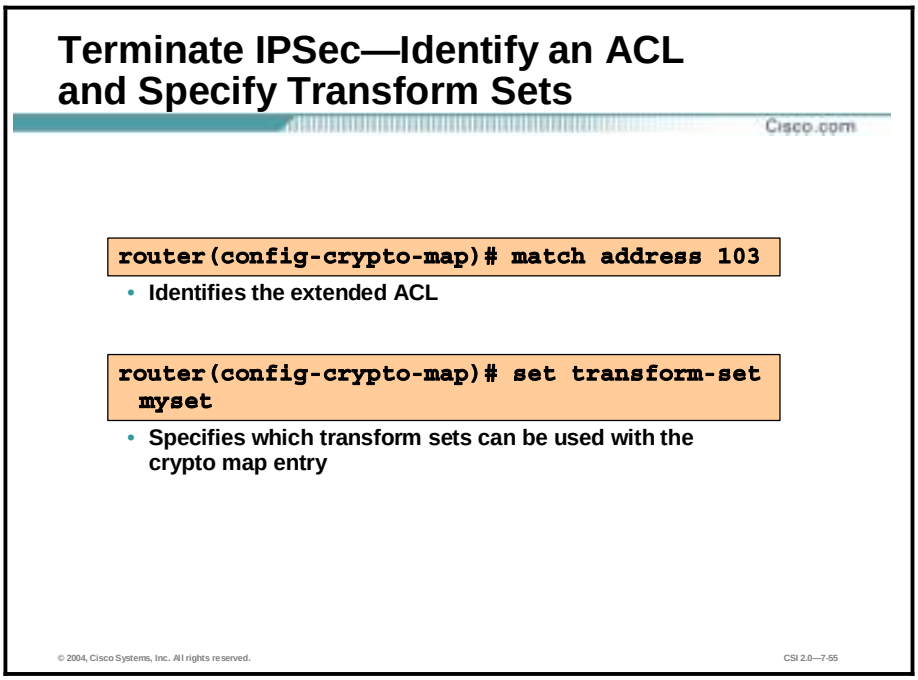

The **match address** command is required for all static crypto map entries. If you are defining a dynamic crypto map entry (with the **crypto dynamic-map** command), this command is not required but is strongly recommended.

Use this command to assign an extended ACL to a crypto map entry. You also need to define this ACL using the **access-list** or **ip access-list** extended commands.

The extended ACL specified with this command is used by IPSec to determine which traffic should be protected by crypto and which traffic does not need crypto protection. (Traffic that is permitted by the ACL is protected. Traffic that is denied by the ACL is not protected in the context of the corresponding crypto map entry.)

Note	The crypto ACL is not used to determine whether to permit or deny traffic through the interface. An ACL applied directly to the interface makes that determination.

The crypto ACL specified by the **match address** command is used when evaluating both inbound and outbound traffic. Outbound traffic is evaluated against the crypto ACLs specified by the interface's crypto map entries to determine if it should be protected by crypto and if so (if traffic matches a permit entry), which crypto policy applies. (If necessary, in the case of static IPSec crypto maps, new SAs are established using the data flow identity as specified in the permit entry; in the case of dynamic crypto map entries, if no SA exists, the packet is dropped.) After passing the regular ACLs at the interface, inbound traffic is evaluated against the crypto ACLs specified by the entries of the interface's crypto map set to determine if it should be protected by crypto, and, if so, which crypto policy applies. (In the case of IPSec, unprotected traffic is discarded because it should have been protected by IPSec.)

In the case of IPSec, the ACL is also used to identify the flow for which the IPSec SAs are established. In the outbound case, the permit entry is used as the data flow identity, in general,

while in the inbound case the data flow identity specified by the peer must be permitted by the crypto ACL.

The **set transform-set** command is required for all static and dynamic crypto map entries. Use this command to specify which transform sets to include in a crypto map entry.

For an ipsec-isakmp crypto map entry, you can list multiple transform sets with the **set transform-set** command. List the higher priority transform sets first.

If the local router initiates the negotiation, the transform sets are presented to the peer in the order specified in the crypto map entry. If the peer initiates the negotiation, the local router accepts the first transform set that matches one of the transform sets specified in the crypto map entry.

The first matching transform set that is found at both peers is used for the SA. If no match is found, IPSec does not establish an SA. The traffic is dropped because there is no SA to protect the traffic.

For an ipsec-manual crypto map entry, you can specify only one transform set. If the transform set does not match the transform set at the remote peer's crypto map, the two peers will fail to correctly communicate because the peers are using different rules to process the traffic.

If you want to change the list of transform sets, re-specify the new list of transform sets to replace the old list. This change is only applied to crypto map entries that reference this transform set. The change will not be applied to existing SAs, but will be used in subsequent negotiations to establish new SAs. If you want the new settings to take effect sooner, you can clear all or part of the SA database by using the **clear crypto sa** command.

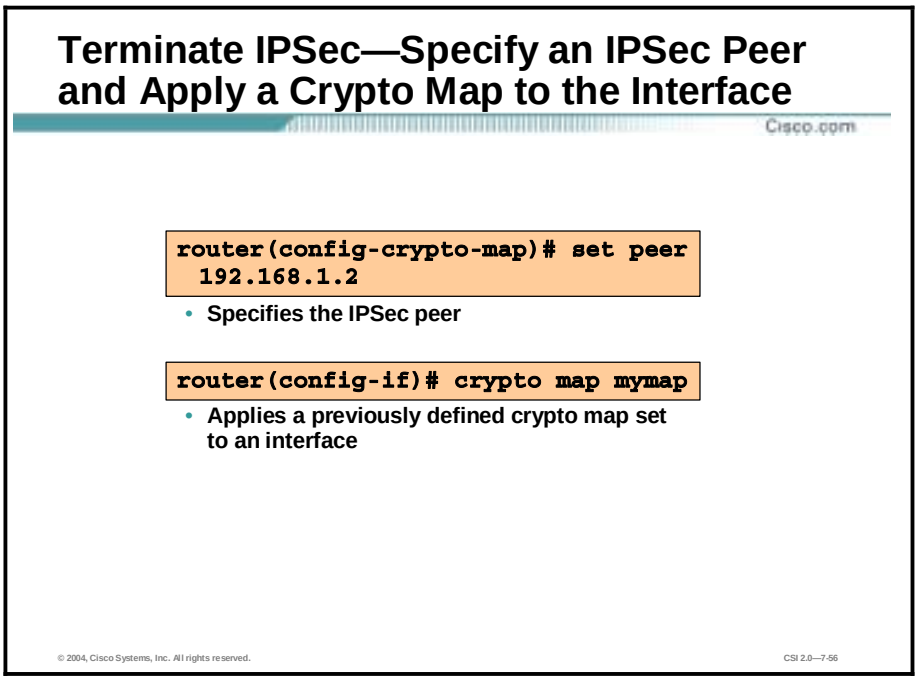

Use the **set peer** command to specify an IPSec peer for a crypto map.

This command is required for all static crypto maps. If you are defining a dynamic crypto map (with the **crypto dynamic-map** command), this command is not required, and in most cases is not used (because, in general, the peer is unknown).

For ipsec-isakmp crypto map entries, you can specify multiple peers by repeating this command. The peer that packets are actually sent to is determined by the last peer that the router heard from (received either traffic or a negotiation request from) for a given data flow. If the attempt fails with the first peer, IKE tries the next peer on the crypto map list.

For ipsec-manual crypto entries, you can specify only one IPSec peer per crypto map. If you want to change the peer, you must first delete the old peer and then specify the new peer.

You can specify the remote IPSec peer by its hostname only if the hostname is mapped to the peer's IP address in a DNS or if you manually map the hostname to the IP address with the **ip host** command.

Use the **crypto map** command to assign a crypto map set to an interface. You must assign a crypto map set to an interface before that interface can provide IPSec services. Only one crypto map set can be assigned to an interface. If multiple crypto map entries have the same map-name but a different seq-num, they are considered to be part of the same set and are all applied to the interface. The crypto map entry with the lowest seq-num is considered the highest priority and is evaluated first. A single crypto map set can contain a combination of cisco, ipsec-isakmp, and ipsec-manual crypto map entries.

Summary

This topic summarizes the information you learned in this lesson.

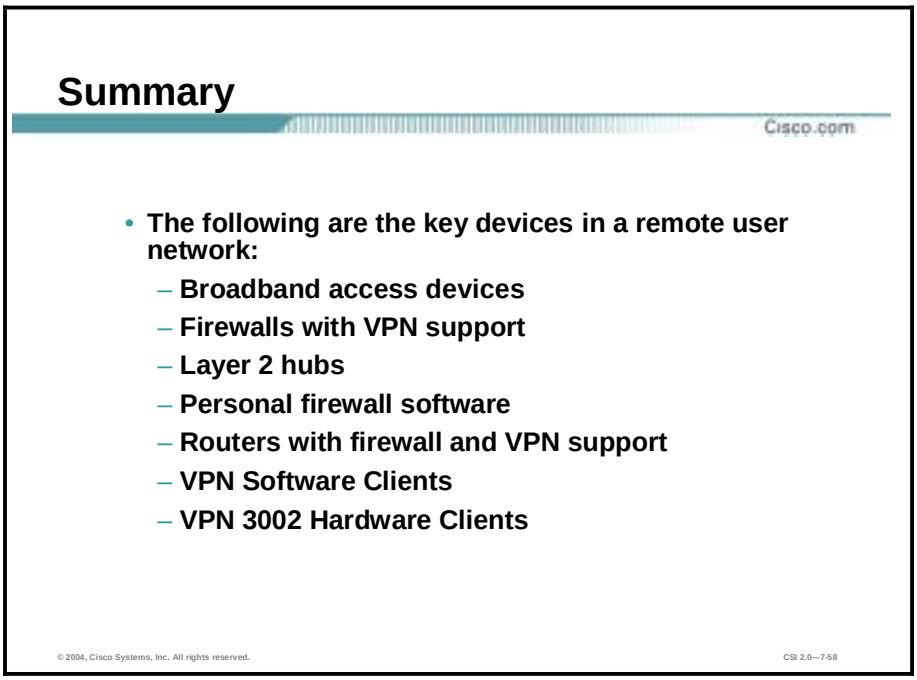

Summary (Cont.)

Cisco.com

- The following threats can be expected:
  - Unauthorized access
  - Network reconnaissance
  - Virus and Trojan horse attacks
  - IP spoofing
  - Man-in-the-middle attacks
- Four basic options are available to mitigate threats: one software-based and three hardware-based options.

8

SAFE Enterprise Network Design

Overview

This lesson describes SAFE Enterprise network design. It includes the following topics:

- Objectives
- Enterprise network design overview
- Enterprise Network Campus
- Enterprise Network Edge
- Summary

Objectives

This topic lists the lesson's objectives.

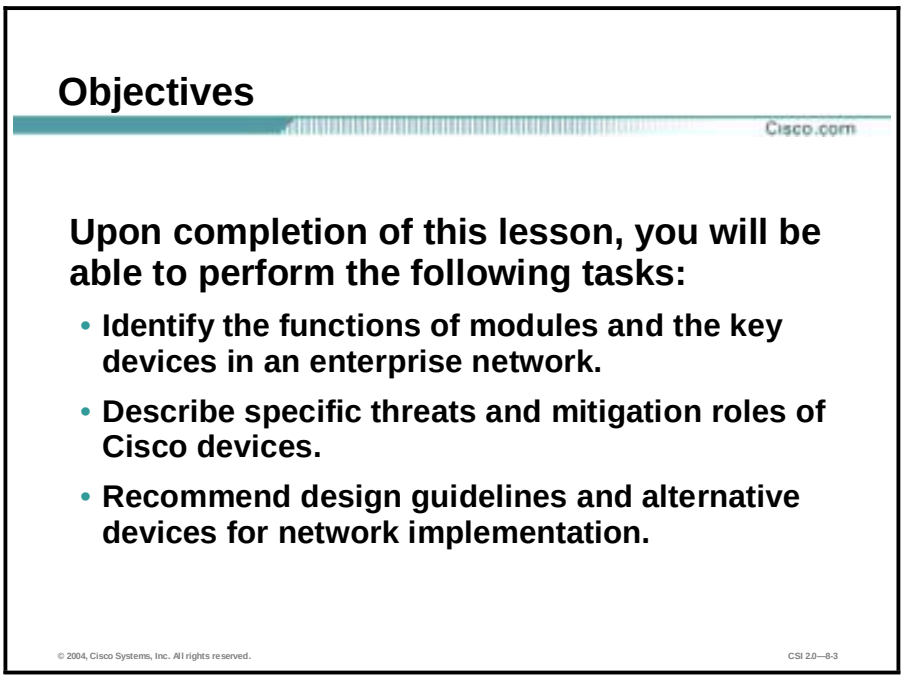

Enterprise Network Design Overview

This topic provides an overview of the SAFE Enterprise network design.

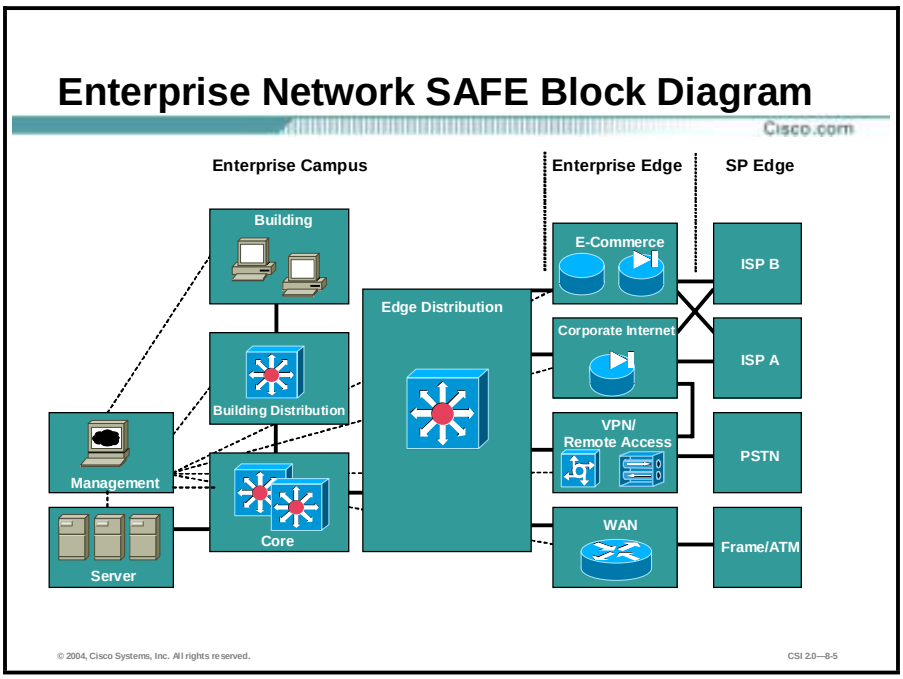

The enterprise comprises two functional areas: the Enterprise Network Campus and the Enterprise Network Edge. These two areas are further divided into modules that define the various functions of each area in detail.

Enterprise Network Campus

This topic describes the modules contained within the Enterprise Network Campus.

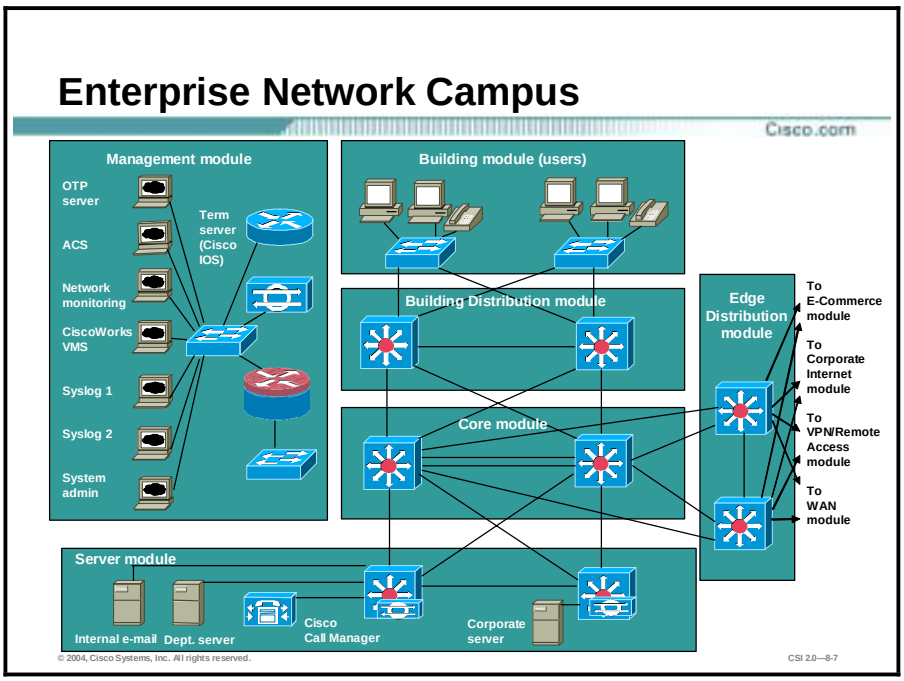

The SAFE Enterprise Network Campus is composed of six modules:

- Management module
- Core module
- Building Distribution module
- Building module
- Server module
- Edge Distribution module

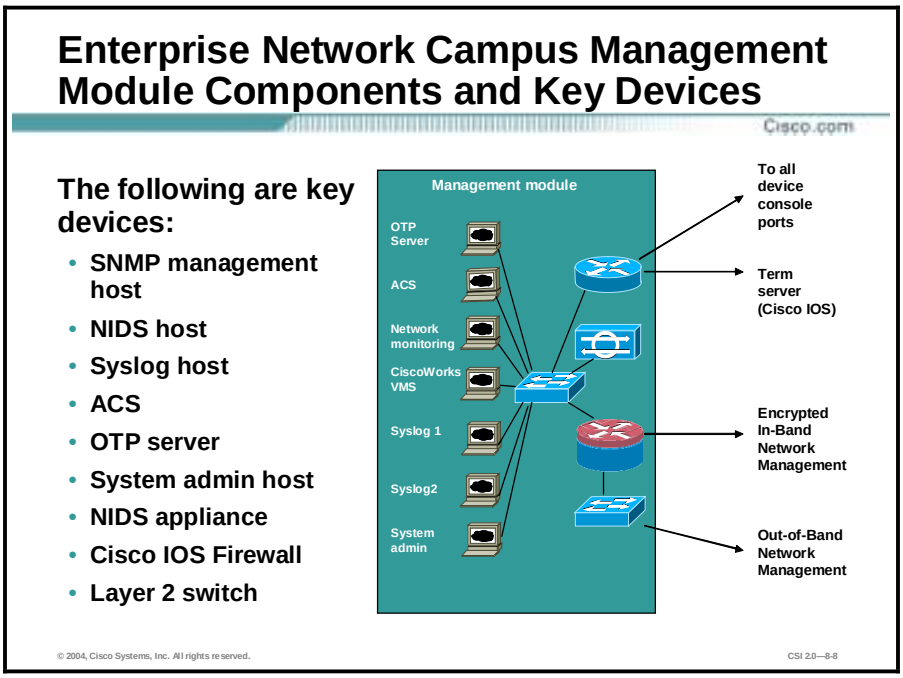

The primary goal of the Management module is to facilitate the secure management of all devices and hosts within the enterprise SAFE architecture. Logging and reporting information flow from the devices to the management hosts. Content, configurations, and new software flow from the management hosts to the devices.

The following are key devices in the Enterprise Network Campus Management module:

- SNMP management host—Provides Simple Network Management Protocol (SNMP) management for devices
- Syslog host—Aggregates log information for Firewall and NIDS hosts
- Access Control Server (ACS)—Delivers one-time, two-factor authentication services to the network devices
- One-time password (OTP) server—Authorizes one-time password information relayed from the access control server
- System administration host—Provides configuration, software, and content changes on devices
- NIDS appliance—Provides Layer 4 to Layer 7 monitoring of key network segments in the module
- NIDS host—Provides alarm aggregation for all network-based intrusion detection system (NIDS) devices in the network
- Cisco IOS Firewall—Allows granular control for traffic flows between the management hosts and the managed devices
- Layer 2 switch (with private VLAN support)—Ensures that data from managed devices can only cross directly to the Cisco IOS Firewall

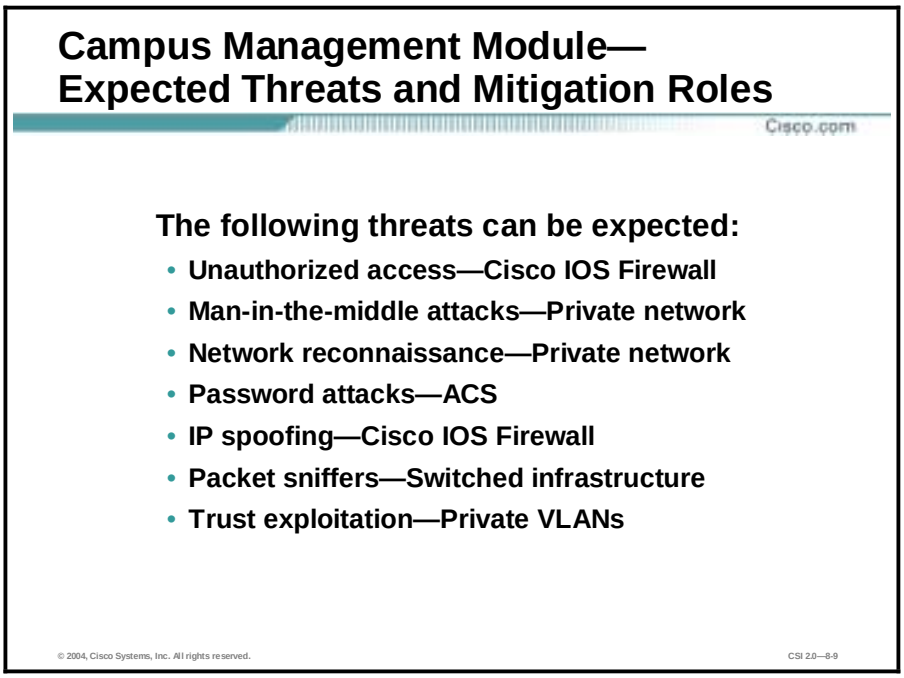

The following are expected threats to management hosts and servers in the Network Campus Management module:

- Unauthorized access—Filtering at the IOS firewall stops most unauthorized traffic in both directions.
- Man-in-the-middle attacks—Management data crossing a private network makes man-in-the-middle attacks difficult.
- Network reconnaissance—Because all management traffic crosses the network, it does not cross the production network where it could be intercepted.
- Password attacks—The ACS allows for strong two-factor authentication at each device.
- IP spoofing—Spoofed traffic is stopped in both directions at the IOS firewall.
- Packet sniffers—A switched infrastructure limits the effectiveness of sniffing.
- Trust exploitation—Private VLANs prevent a compromised device from masquerading as a management host.

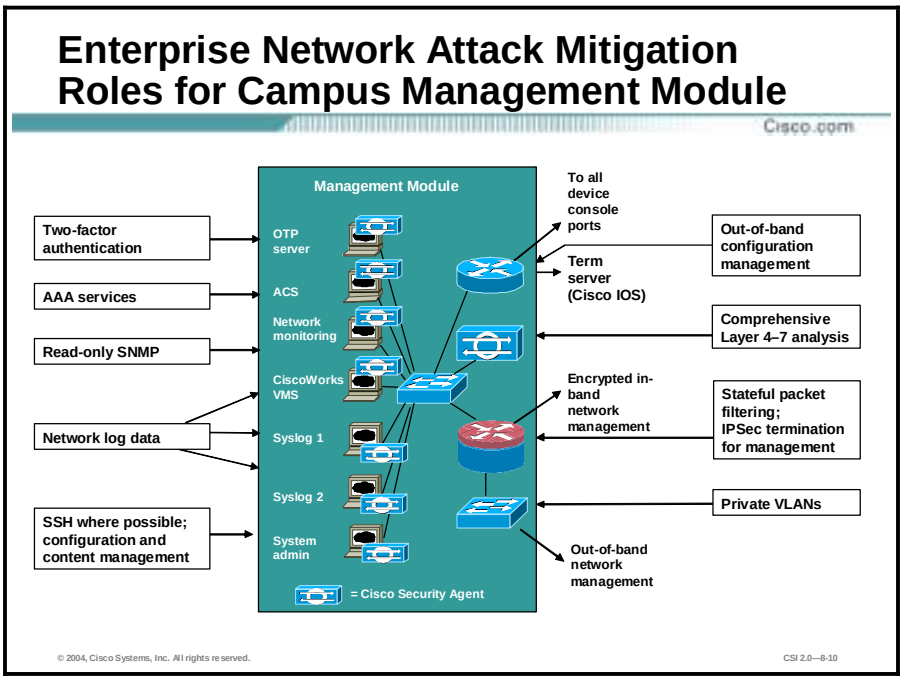

The attack mitigation roles for each device in the SAFE Enterprise Network Campus Management module are as follows:

- OTP server—Two-factor authentication
- ACS—Authentication, authorization, and accounting (AAA) services
- Network monitoring—Read-only SNMP
- CiscoWorks Virtual Private Network (VPN)/Security Management Solution (VMS)—Network log data
- Syslog server—Network log data
- System administration server
  - Secure Shell (SSH) where possible
  - Configuration and content management
- Router
  - Out-of-band configuration management
  - Stateful packet filtering
  - IPSec termination for management
- Layer 2 switch—Private VLANs

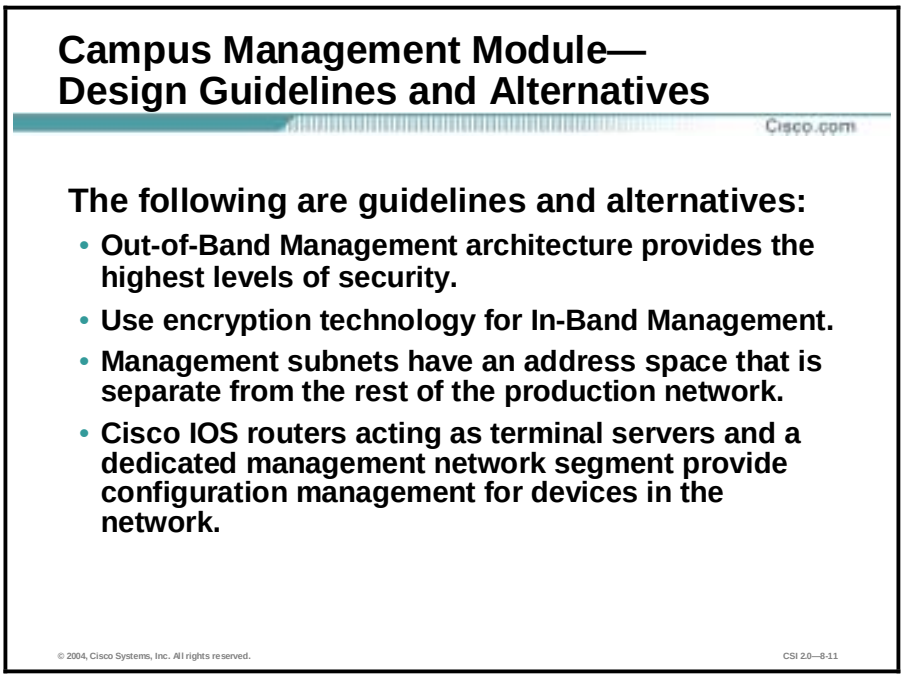

The SAFE enterprise management network has two network segments that are separated by an IOS router that acts as a firewall and a VPN termination device. The segment outside the firewall connects to all the devices that require management. The segment inside the firewall contains the management hosts themselves and the IOS routers that act as terminal servers. The remaining interface connects to the production network but only for selective Internet access, limited In-Band Management traffic, and IPSec-protected management traffic from predetermined hosts. The following SAFE guidelines and alternatives are recommended:

- Out-of-Band Management architecture is recommended as it provides the highest levels of security. The alternative is In-Band Management if the goal is a cost-effective security deployment. In the Out-of-Band Management environment, each network device and host has its own dedicated management interface, which connects to the private management network. This setup mitigates the risk of passing insecure management protocols such as Telnet, Trivial File Transfer Protocol (TFTP), SNMP, and Syslog over the production network where it could be intercepted or modified.
- In-Band Management only occurs when the application itself would not function out of band or if the Cisco device being managed did not physically have enough interfaces to support the normal management connection. It is this latter case that employs IPSec tunnels. The IOS firewall is configured to allow Syslog information into the management segment, as well as Telnet, SSH, and SNMP if these are first initiated by the inside network.
- Private addresses—Both management subnets operate under an address space that is completely separate from the rest of the production network. This ensures that the management network will not be advertised by any routing protocols.
- Cisco IOS routers act as terminal servers and a dedicated management network segment. They provide configuration management for nearly all devices in the network. The routers provide a reverse-Telnet function to the console ports on the Cisco devices throughout the enterprise. More extensive management features (software updates, content updates, log and alarm aggregation, SNMP management) are provided through the dedicated management network segment.

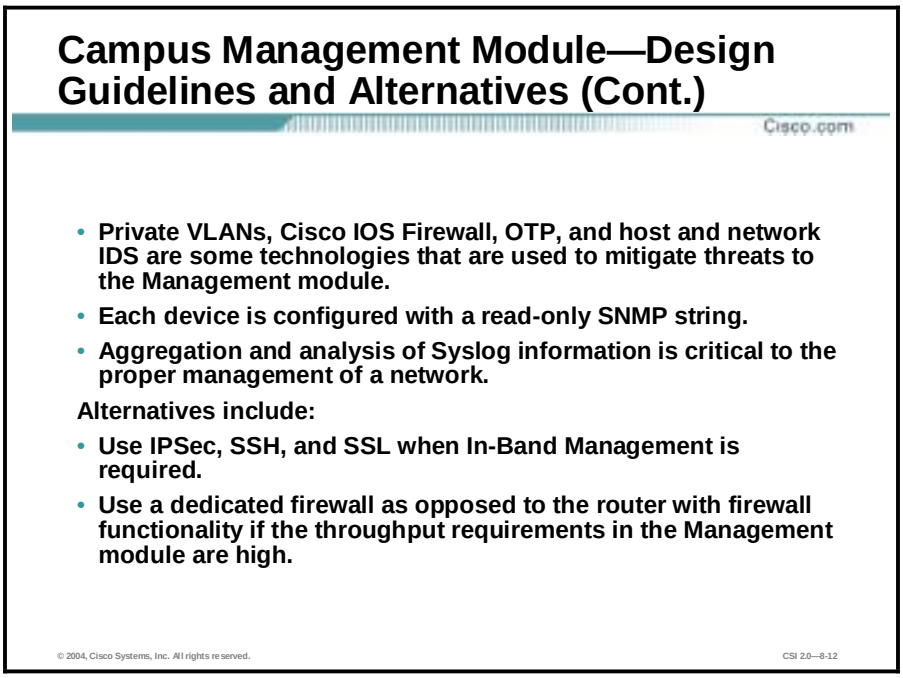

- Securing the management network module—The Management module has been built with the following technologies designed to mitigate risks:
  - To mitigate the threat of a compromised device, access control is implemented at the firewall and at every other possible device to prevent exploitation of the management channel. A compromised device cannot even communicate with other hosts on the same subnet because private VLANs on the management segment switch force all traffic from the managed devices directly to the IOS firewall where filtering takes place. Password sniffing only reveals useless information because of the OTP environment.
  - Host and network IDSs are implemented on the management subnet and are configured in a very restrictive level. Because the types of traffic on this network should be very limited, any signature match on this segment should be met with an immediate response.
- SNMP management—With SAFE, SNMP management pulls information from devices rather than allowing it to push changes. To ensure this, each device is configured with a "read-only" string.
- Proper aggregation and analysis of the Syslog information—Syslog provides important information regarding security violations and configuration changes. Depending on the device in question, different levels of Syslog information might be required. Having full logging with all messages sent might provide too much information for an individual or Syslog analysis algorithm to sort. SNMP "read-write" may be configured when using an out-of-band network, but be aware of the increased security risk due to a clear text string allowing modification of device configurations.
- The following are recommended alternatives for this module:
  - Complete Out-of-Band Management is not always possible. When In-Band Management is required, more emphasis needs to be placed on securing the transport of the management protocols. This can be through the use of IPSec, SSH, Secure Sockets Layer (SSL), or any other encrypted and authenticated transport that allows

management information to traverse it. When management happens on the same interface that a device uses for user data, importance needs to be placed on passwords, community strings, cryptographic keys, and the access lists that control communications to the management services.

— If the throughput requirements in the Management module are high, consider the use of a dedicated firewall as opposed to the router with firewall functionality. The router was chosen because of its flexibility in IPSec configuration and its routing options.

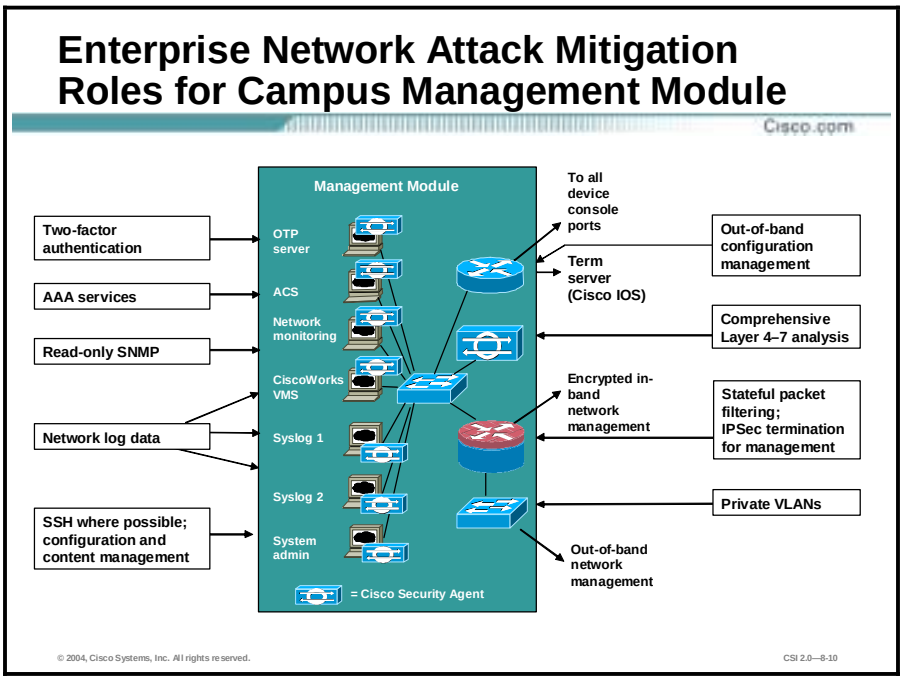

The Core module in the SAFE architecture is nearly identical to the Core module of any other network architecture. It merely routes and switches traffic as fast as possible from one network to another.

The Layer 3 switch is the key device in the SAFE Enterprise Network Campus Core module. This switch routes and switches production network data from one module to another.

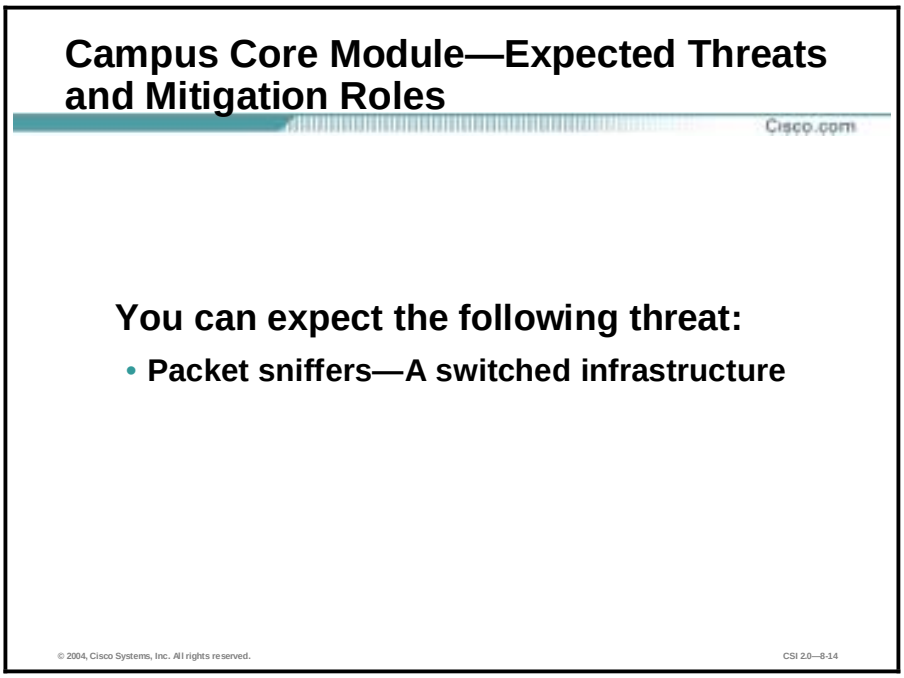

The threats that you can expect to the Enterprise Campus Core module are packet sniffers, and mitigation of this threat is a switched infrastructure that limits the effectiveness of sniffing.

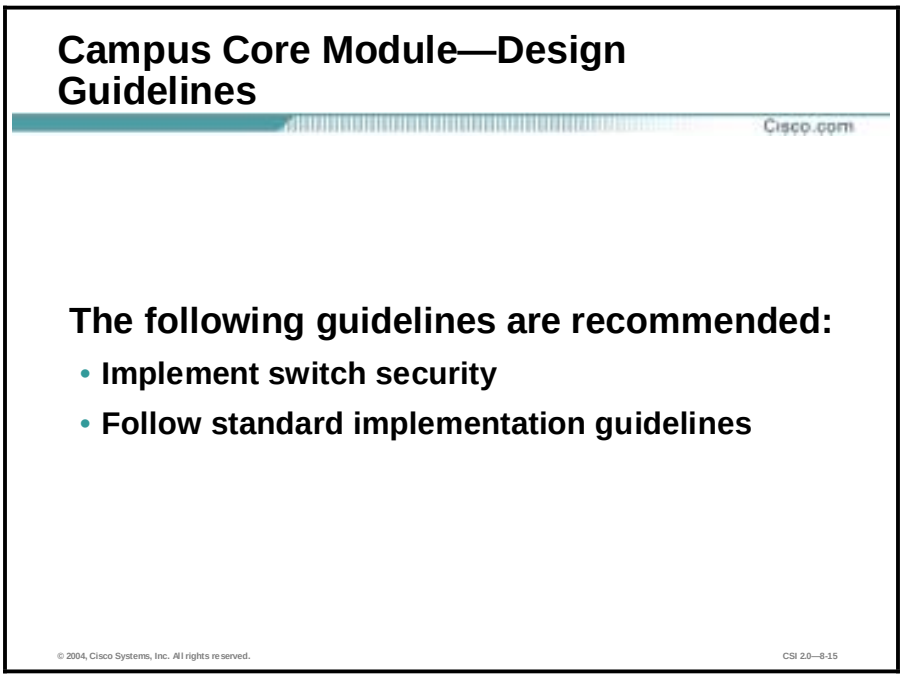

The following SAFE design guidelines are recommended for the Enterprise Campus Core module:

- Implement switch security in the core switches to ensure that they are well protected against direct attacks.
- Follow standard implementation guidelines in accordance with the core, distribution, and access layer deployments commonly seen in well-designed Cisco networks.

Enterprise Network Campus Core Module Key Devices

Cisco.com

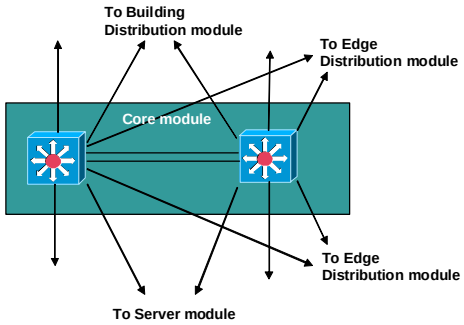

Key device: Layer 3 switch

The goal of this module is to provide distribution layer services to the building switches, including routing, Quality of Service (QoS), and access control. Requests for data flow into these switches and onto the core, and responses follow the identical path in reverse.

The key device for the SAFE Enterprise Network Campus Building Distribution module is a Layer 3 switch. The Layer 3 switch aggregates Layer 2 switches in the Building module and provides advanced services.

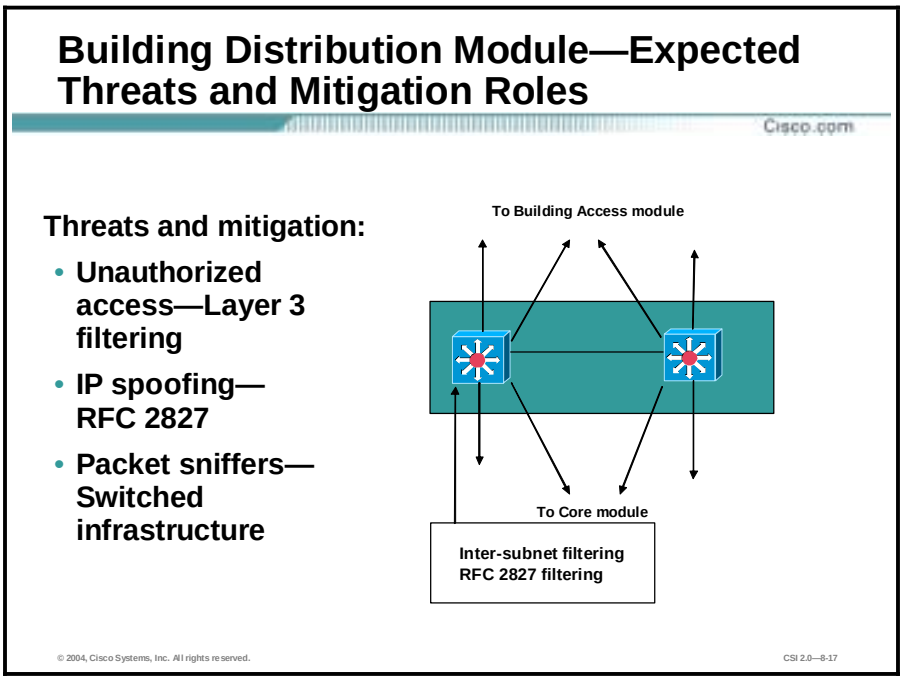

The following are expected threats and mitigation roles to the SAFE Enterprise Network Campus Building Distribution module:

- Unauthorized access—Attacks against server module resources are limited by Layer 3 filtering of specific subnets.
- IP spoofing—RFC 2827 filtering stops most spoofing attempts.
- Packet sniffers—A switched infrastructure limits the effectiveness of sniffing.

Note	Inter-subnet filtering uses ACLs or other filtering mechanisms to control data flow between subnets and VLANs on Layer 3 switches.

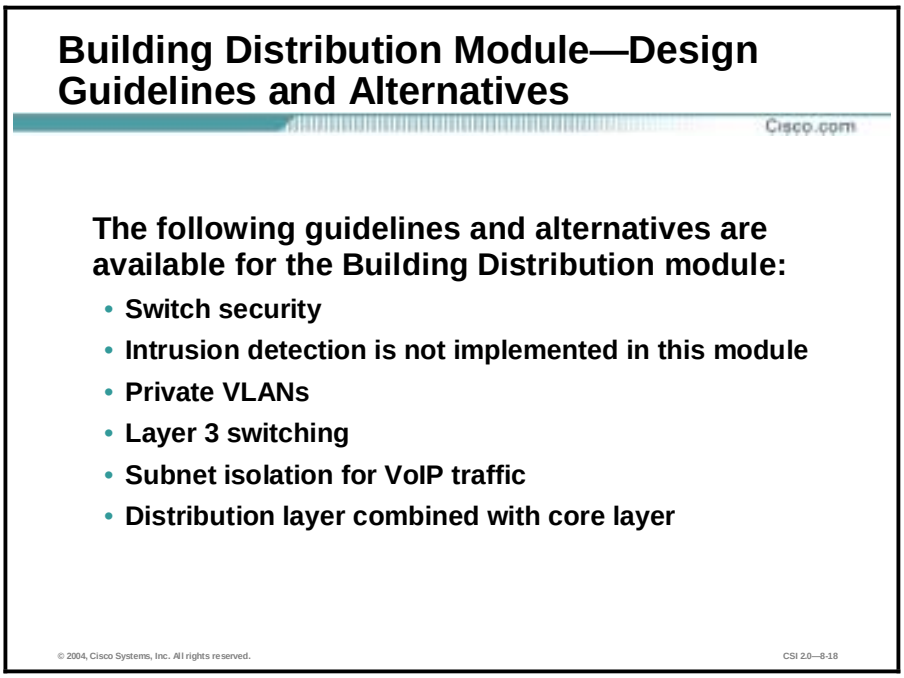

The following design guidelines and alternatives are recommended in the Building Distribution module:

- Switch security—In addition to standard network design fundamentals, optimize the switches to provide added security within the enterprise user community.
- Intrusion detection is not implemented at the Building Distribution module—Intrusion detection is implemented in the modules that contain the resources that are likely to be attacked for their content (server, remote access, Internet, and so forth).
- Private VLANs—The Building Distribution module provides the first line of defense and prevention against internally originated attacks. It can mitigate the chance of a department accessing confidential information on another department's server through the use of access control.
- Layer 3 switching—For performance reasons, it is important that access control be implemented on a hardware platform that can deliver filtered traffic at near wire rates. This generally dictates the use of Layer 3 switching as opposed to more traditional dedicated routing devices. This same access control can also prevent local source-address spoofing through the use of RFC 2827 filtering.
- Subnet isolation—Subnet isolation is used to route voice-over-IP (VoIP) traffic to the call manager and any associated gateways. This prevents VoIP traffic from crossing the same segments that all other data traffic crosses, reducing the likelihood of sniffing voice communications, and allows a smoother implementation of QoS. Complete secure IP telephony deployment details are outside the scope of this document.
- Combine distribution layer with the core layer—Depending on the size and performance requirements of the network, the distribution layer can be combined with the core layer to reduce the number of devices required in the environment.

Enterprise Network Campus Building Module

Cisco.com

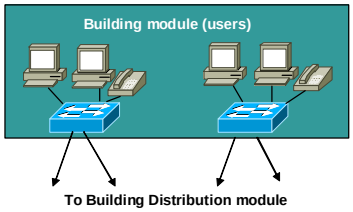

Key devices:

- **Layer 2 switch**
- **User workstation**
- **IP phone**

CSI 2.0—8-19

SAFE defines the Building module as the extensive network portion that contains end-user workstations, phones, and their associated Layer 2 access points. Its primary goal is to provide services to end users.

The following are key devices for the SAFE Enterprise Network Campus Building module:

- Layer 2 switch—Provides Layer 2 services to phones and user workstations
- User workstation—Provides data services to authorized users on the network
- IP phone—Provides IP telephony services to users on the network

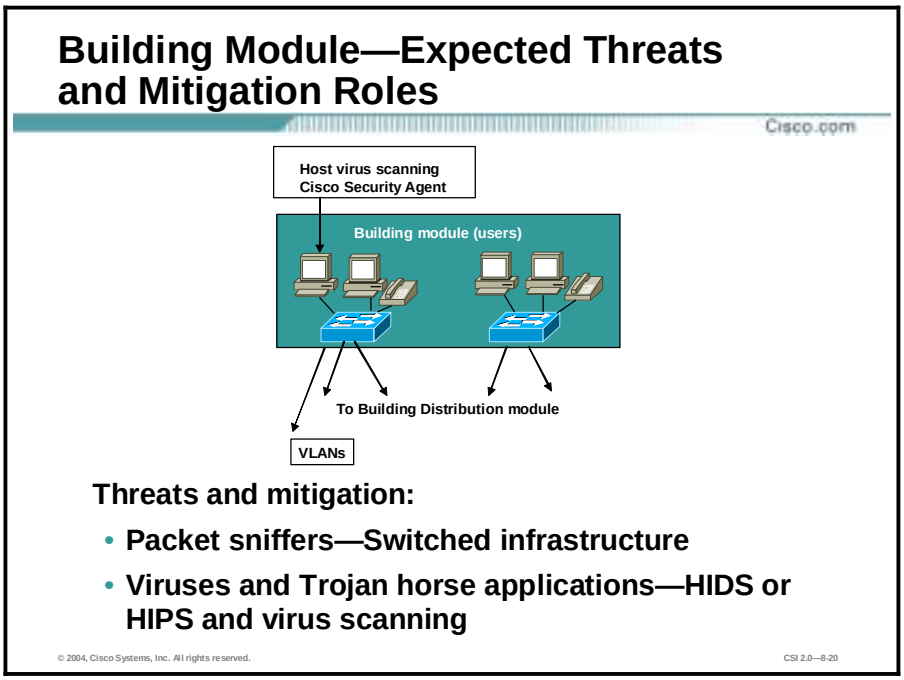

The following are expected threats and mitigation roles to the SAFE Enterprise Network Campus Building module:

- Packet sniffers—A switched infrastructure and default VLAN services limit the effectiveness of sniffing.
- Virus and Trojan horse applications—Host-based intrusion detection system (HIDS) or host-based intrusion prevention system (HIPS) and virus scanning prevents most viruses and many Trojan horses.

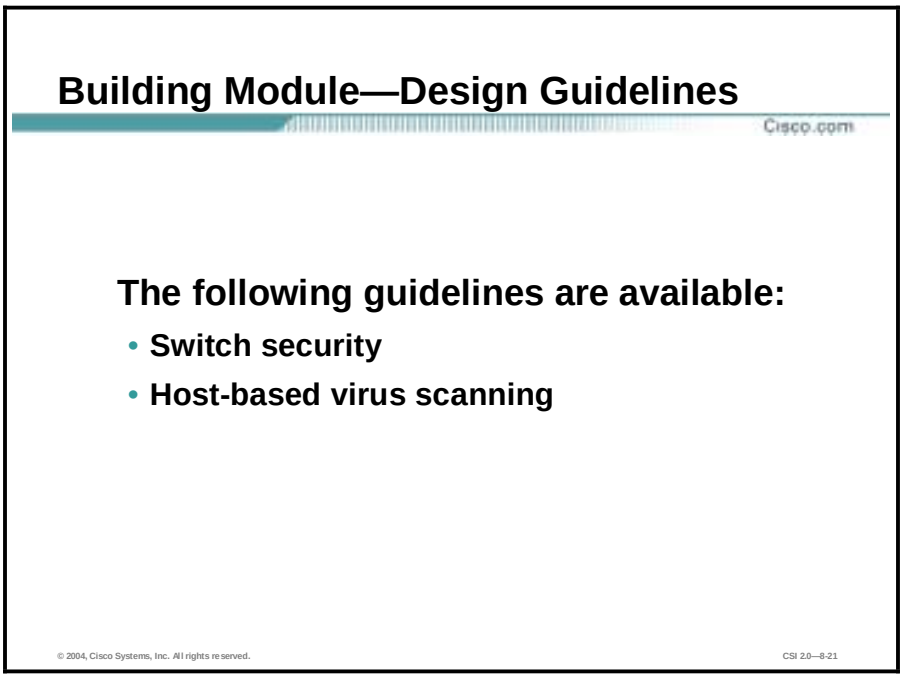

The following design guidelines are available for the Building module:

- Switch security—Implementing security in a concise and effective manner is challenging because user devices are generally the largest single element of the network. From a security perspective, the Building Distribution module, rather than anything in the Building module, provides most of the access control that is enforced at the end-user level. This is because the Layer 2 switch that the workstations and phones connect to has no capability for Layer 3 access control.
- Host-based virus scanning—This type of security is implemented at the workstation level.

Enterprise Network Campus Server Module

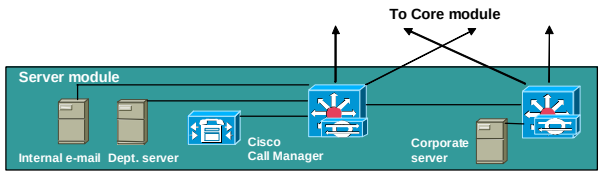

Key devices:

- **Layer 3 switch**
- **Call Manager**
- **Corporate and department servers**
- **E-mail server**

The Server module's primary goal is to provide application services to end users and devices. Traffic flows on the Server module are inspected by on-board intrusion detection within the Layer 3 switches.

The following are the key devices for the SAFE Enterprise Network Campus Server module:

- Layer 3 switch—Provides Layer 3 services to the servers and inspects data crossing the Server module with NIDS
- Call Manager—Performs call routing functions for IP telephony devices in the enterprise
- Corporate and department servers—Deliver file, print, and DNS services to workstations in the Building module
- E-mail server—Provides Simple Mail Transfer Protocol (SMTP) and Post Office Protocol Version 3 (POP3) services to internal users

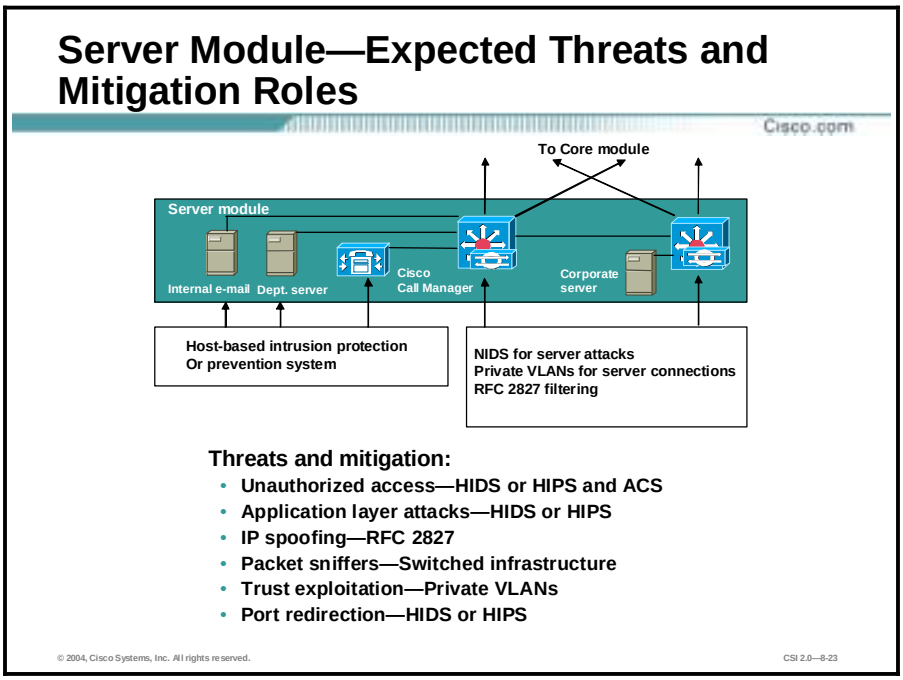

The following are expected threats and mitigation for the SAFE Enterprise Network Campus Server module:

- Unauthorized access—This is mitigated through the use of either HIDS or HIPS and ACS.
- Application layer attacks—Operating systems, devices, and applications are kept up-to-date with the latest security fixes and protected by either HIDS or HIPS.
- IP spoofing—RFC 2827 filtering prevents source address spoofing.
- Packet sniffers—A switched infrastructure limits the effectiveness of sniffing.
- Trust exploitation—Trust arrangements are very explicit. Private VLANs prevent hosts on the same subnet from communicating unless necessary.
- Port redirection—A host-based intrusion prevention system (HIPS) prevents port redirection agents from being installed.

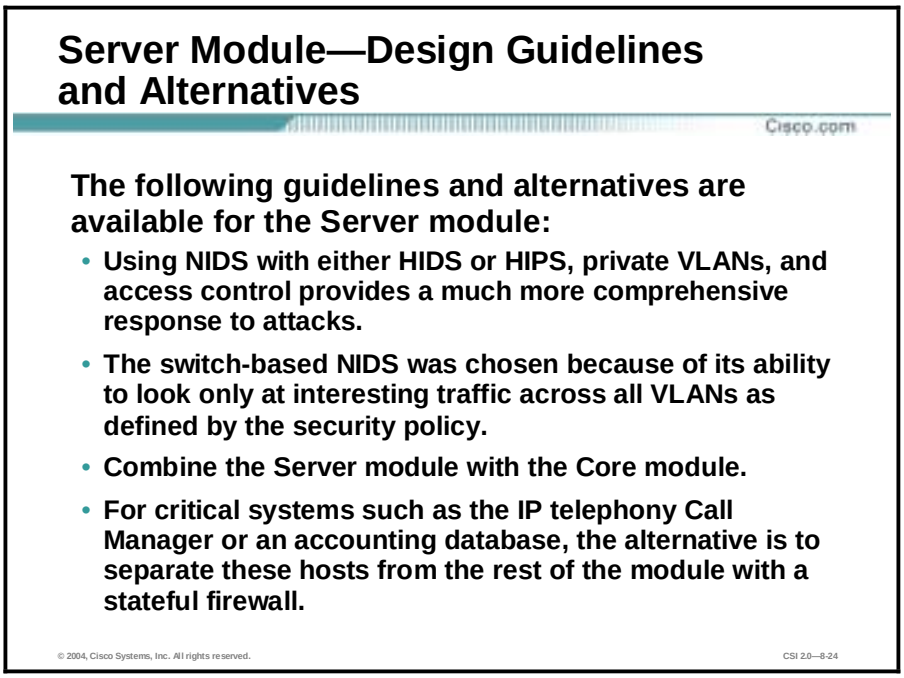

Q2DEV KMC: 1st bullet: “Using NIDS with either HIDS or HIPS”. Edit ok?-EDITJag – “Using NIDS, either HIDS or HIPS…..” will be better.

The Server module is often overlooked from a security perspective. When examining the levels of access most employees have to the servers to which they attach, the servers can often become the primary goal of internally originated attacks. Simply relying on effective passwords does not provide for a comprehensive attack mitigation strategy. The following guidelines and alternatives are available:

- Server security technologies—The use of host and network-based IDS, private VLANs, access control, and good system administration practices (such as keeping systems up-to-date with the latest patches), provides a much more comprehensive response to attacks.
- Switch-based NIDS—Because the NIDS system is limited in the amount of traffic it can analyze, it is important to send it attack-sensitive traffic only. This varies from network to network, but should likely include SMTP, Telnet, FTP, and the World Wide Web. The switch-based NIDS was chosen because of its ability to look only at interesting traffic across all VLANs as defined by the security policy. Once properly tuned, these IDSs can be set up in a restrictive manner because required traffic streams should be well known.
- Combine Server module with the Core module—Combine these modules if performance needs do not dictate separation. For very sensitive high-performance server environments, blades installing more than one NIDS blade and directing policy-matched traffic to specific blades can scale the NIDS capability in the Layer 3 switch.
- Separate domains—For critical systems such as the IP telephony Call Manager or an accounting database, consider separating these hosts from the rest of the module with a stateful firewall.

Enterprise Network Campus Edge Distribution Module

Cisco.com

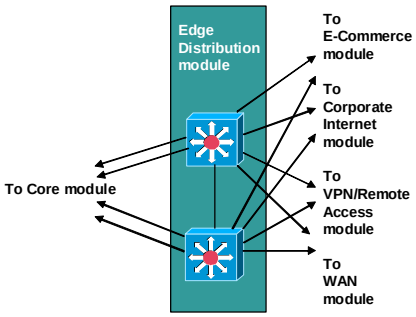

Key device: Layer 3 switch

This goal of the module is to aggregate the connectivity from the various elements at the edge. Traffic is filtered and routed from the edge modules and routed into the core.

The Layer 3 switch is the key device for the SAFE Enterprise Network Campus Edge Distribution module. The Layer 3 switch aggregates edge connectivity and provides advanced services.

Edge Distribution Module—Expected Threats and Mitigation Roles

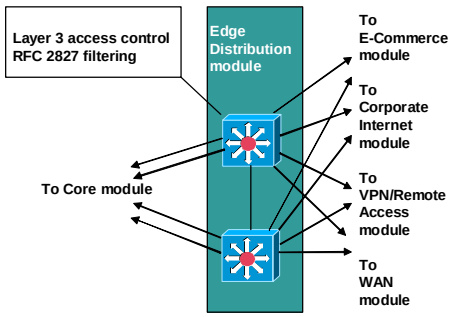

Threats and mitigation:

- **Unauthorized access—ACLs**
- **IP spoofing—RFC 2827**
- **Network reconnaissance—Filtering**
- **Packet sniffers—Switched infrastructure**

The following are expected threats and mitigation of those threats to the SAFE Enterprise Network Campus Edge Distribution module:

- Unauthorized access—Filtering provides granular control over specific edge subnets and their ability to reach areas within the campus.
- IP spoofing—RFC 2827 filtering limits locally initiated spoof attacks.
- Network reconnaissance—Filtering limits nonessential traffic from entering the campus, limiting a hackers ability to perform network reconnaissance.
- Packet sniffers—A switched infrastructure limits the effectiveness of sniffing.

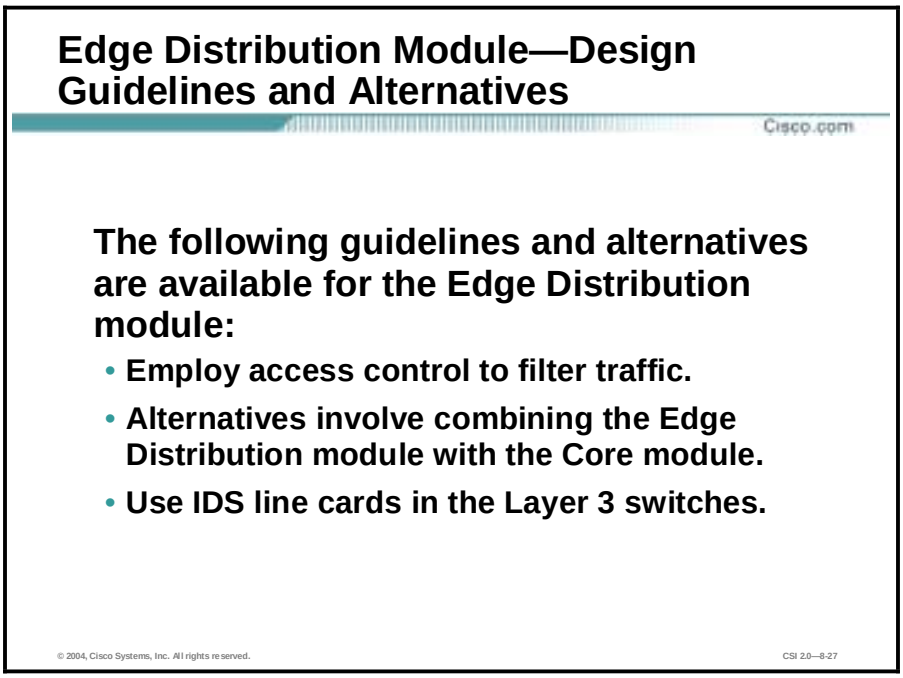

The Edge Distribution module is similar in some respects to the Building Distribution module in terms of overall function. Both modules employ access control to filter traffic, although the Edge Distribution module can rely somewhat on the entire edge functional area to perform additional security functions. Both modules use Layer 3 switching to achieve high performance, but the Edge Distribution module can add additional security functions because the performance requirements are not as great. The following SAFE guidelines and alternatives are available:

- Access control—The Edge Distribution module provides the last line of defense for all traffic destined to the campus module from the edge module. This includes mitigation of spoofed packets, erroneous routing updates, and provisions for network layer access control.
- Combine Edge Distribution module with the Core module—Combine these modules if performance requirements do not require separation.
- Switch-based NIDS—NIDS is not present in this module, but could be placed here through the use of IDS line cards in the Layer 3 switches. It would then reduce the need for NIDS appliances at the exit from the critical edge modules as they connect to the campus. However, performance reasons may dictate that dedicated intrusion detection be placed in the various edge modules as opposed to the Edge Distribution module.

Enterprise Network Edge

This topic discusses all the modules contained within the Enterprise Network Edge.

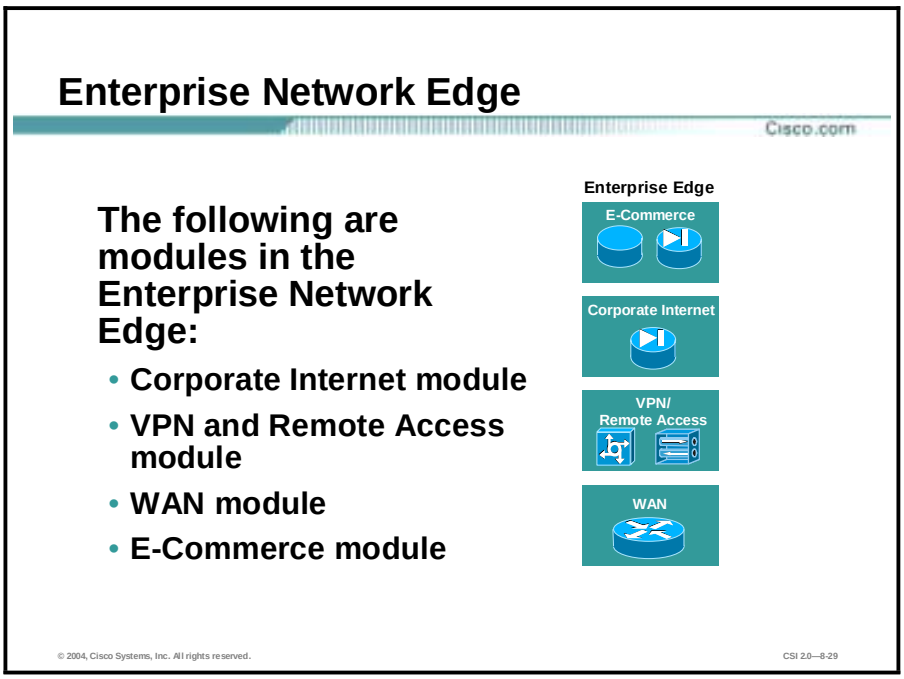

The SAFE Enterprise Network Edge is composed of four modules:

- E-commerce module
- Corporate Internet module
- VPN and Remote Access module
- WAN module

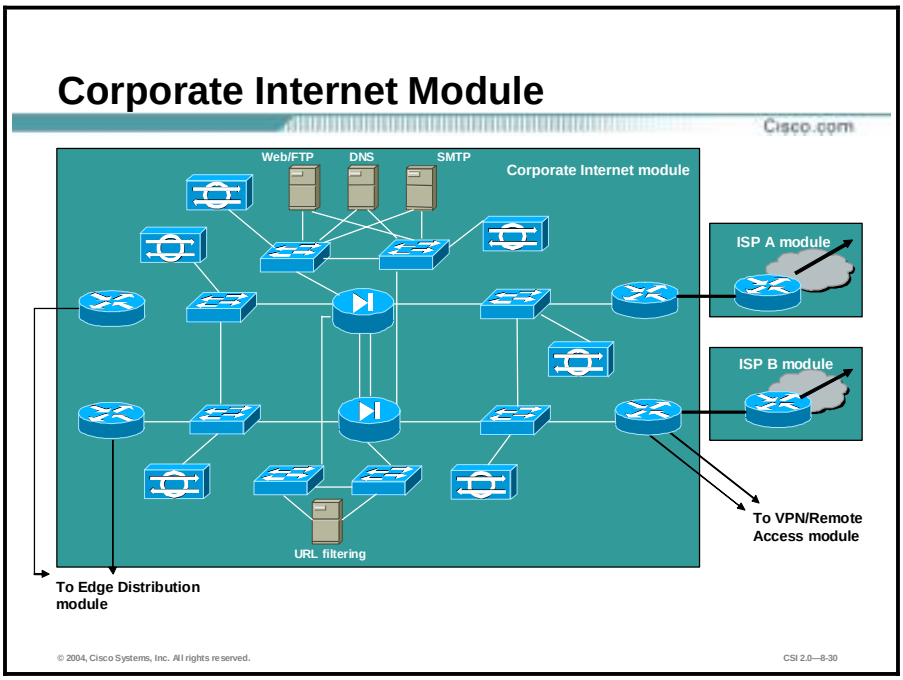

The Corporate Internet module provides internal users with connectivity to Internet services and Internet users access to information on public servers. Traffic also flows from this module to the VPN and Remote Access module where VPN termination takes place. This module is not designed to serve e-commerce type applications.

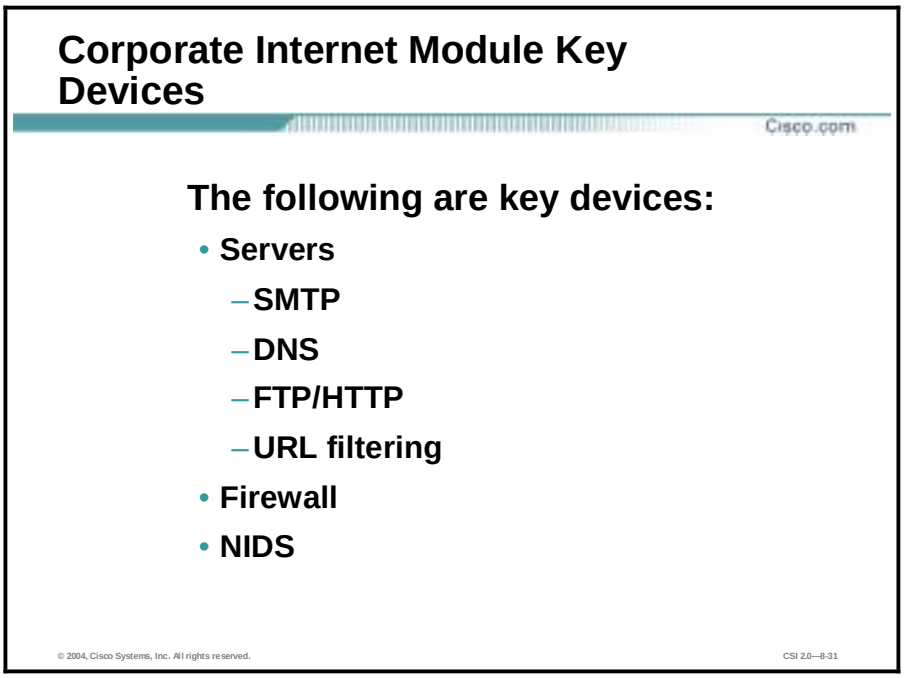

The following are key devices for the SAFE Enterprise Network Edge Corporate Internet module:

- SMTP server—Acts as a relay between the Internet and the Internet mail servers, and inspects content
- DNS server—Serves as authoritative external DNS server for the enterprise, and relays internal requests to the Internet
- FTP/HTTP server—Provides public information about the organization
- URL filtering server—Filters unauthorized URL requests from the enterprise
- Firewall—Provides network-level protection of resources and stateful filtering of traffic
- NIDS appliance—Provides Layer 4 to Layer 7 monitoring of key network segments in the module

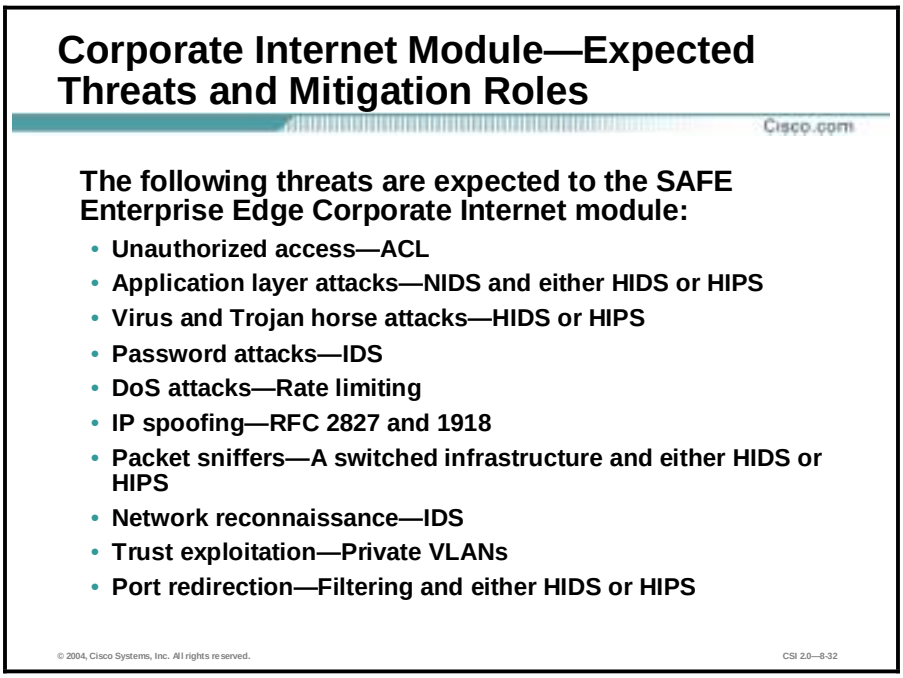

The following are expected threats and mitigation of those threats to SAFE Enterprise Network Edge Corporate Internet module:

- Unauthorized access—Mitigated through filtering at the ISP, edge router, and corporate firewall
- Application layer attacks—Mitigated through IDS at the host and network levels
- Virus and Trojan horse attacks—Mitigated through e-mail content filtering and either HIDS or HIPS
- Password attacks—Limited services available against brute-force attacks, but operating systems and IDS can detect the threat
- Denial of service (DoS) attacks—Rate limiting at ISP edge and TCP setup controls at firewall
- IP spoofing—RFC 2827 and 1918 filtering at ISP edge and enterprise edge router
- Packet sniffers—Switched infrastructure and either HIDS or HIPS to limit exposure
- Network reconnaissance—IDS detects recon, protocols filtered to limit effectiveness
- Trust exploitation—Restrictive trust model and private VLANs to limit trust-based attacks
- Port redirection—Restrictive filtering and either HIDS or HIPS to limit attacks

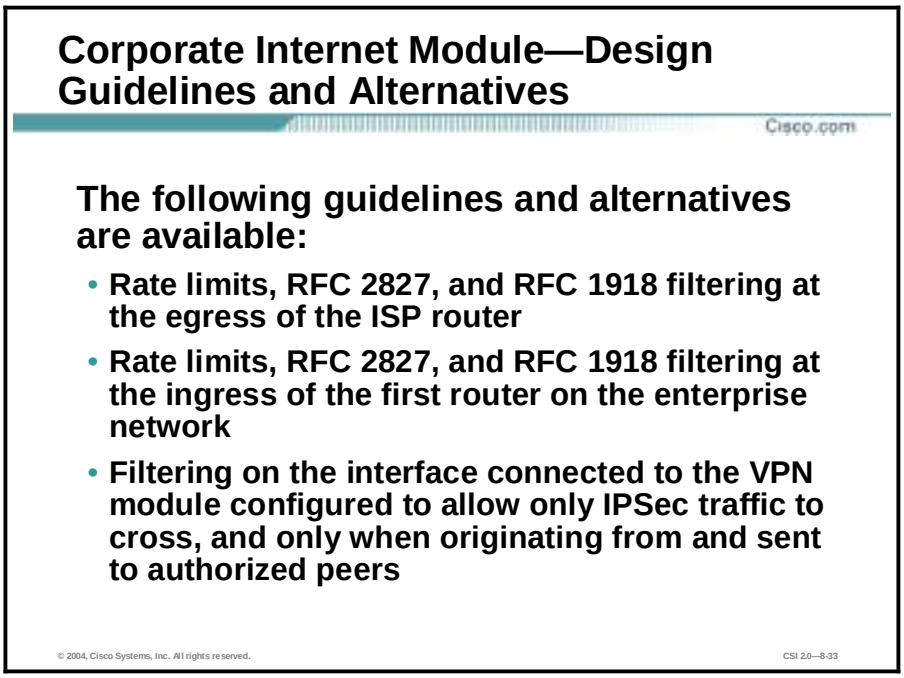

The heart of the Corporate Internet module is a pair of resilient firewalls, which provide protection for the Internet public services and internal users. Stateful inspection examines traffic in all directions, ensuring that only legitimate traffic crosses the firewall. Aside from the Layer 2 and Layer 3 resilience built into the module and the stateful failover capability of the firewall, all other design considerations center around security and attack mitigation. The following guidelines and alternatives are available:

- Rate limits and access control at egress—Starting at the customer-edge router in the ISP, the egress out of the ISP rate-limits nonessential traffic that exceeds specified thresholds in order to mitigate against distributed denial of service (DDoS) attacks. Also at the egress of the ISP router, RFC 2827 and RFC 1918 filtering mitigate against source-address spoofing of local networks and private address ranges.
- Rate limits and access control at the ingress—On the first router in the Enterprise network, basic filtering limits the traffic to the expected traffic (addresses and IP services), providing a coarse filter for the most basic attacks. RFC 1918 and 2827 filtering is also provided here as a supplement to ISP's filtering. In addition, the router is configured to drop most fragmented packets because of the enormous security threat that they create. Any legitimate traffic lost because of this filtering is considered an acceptable risk when compared to the risk of allowing such traffic.
- IPSec traffic filters—Any IPSec traffic destined for the VPN and Remote Access module is routed appropriately. Filtering on the interface connected to the VPN module is configured to allow only IPSec traffic to cross and only when originated from and sent to authorized peers. With remote-access VPNs, you generally do not know the IP address of the system coming in, so filtering can be specific only to the head-end peers with which the remote users are communicating.

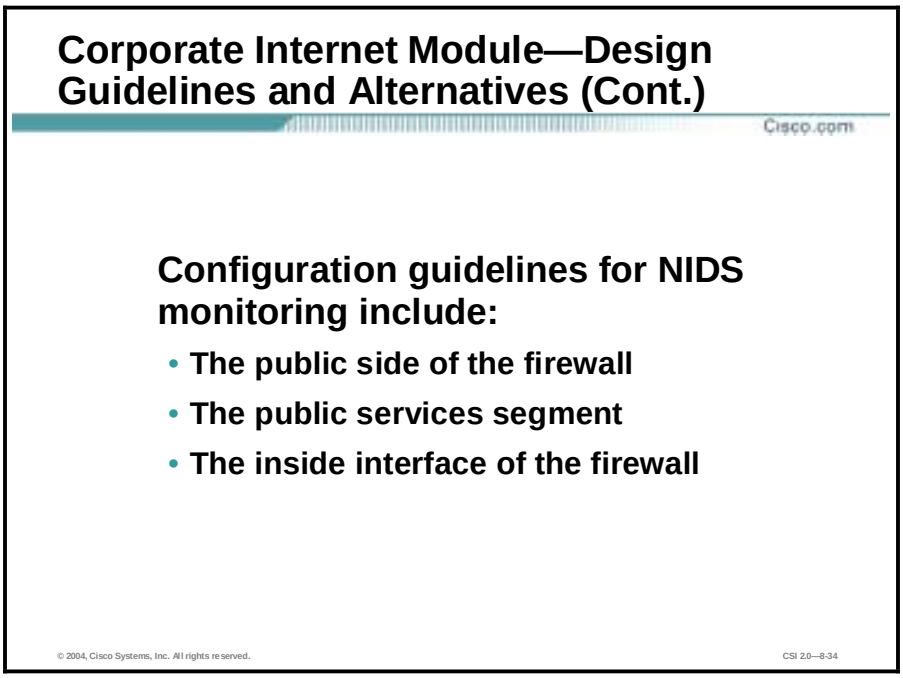

Configuration guidelines for NIDS monitoring in the Corporate Internet module include the following:

- NIDS on the public side of the firewall—NIDS is monitoring for attacks based on Layer 4 to Layer 7 analysis and comparisons against known signatures. Because the ISP and enterprise edge router are filtering certain address ranges and ports, this allows the NIDS appliance to focus on some of the more complex attacks. Still, the NIDS should have alarms set to a lower level than appliances on the inside of the firewall because alarms seen here do not represent actual breaches, but merely attempts.
- NIDS on the public services segment—This NIDS detects attacks on ports that the firewall is configured to permit. These most often are application-layer attacks against a specific service or a password attack against a protected service. You need to set this NIDS to a more restrictive level than the NIDS on the outside of the firewall, because signatures matched here have successfully passed through the firewall.
- NIDS on the inside interface of the firewall—This NIDS provides a final analysis of attacks. Very few attacks should be detected on this segment because only responses to initiated requests, and a few select ports from the public services segment, are allowed on the inside. Only sophisticated attacks should be seen on this segment because the appearance of an attack on this segment generally means that a system on the public services segment has been compromised and the hacker is attempting to leverage this foothold to attack the internal network. If attacks are seen on this segment, the responses to those attacks should be more severe than responses on other segments, because they probably indicate that a compromise has already occurred. The use of a TCP reset to thwart, for example, an SMTP attack, should be seriously considered.

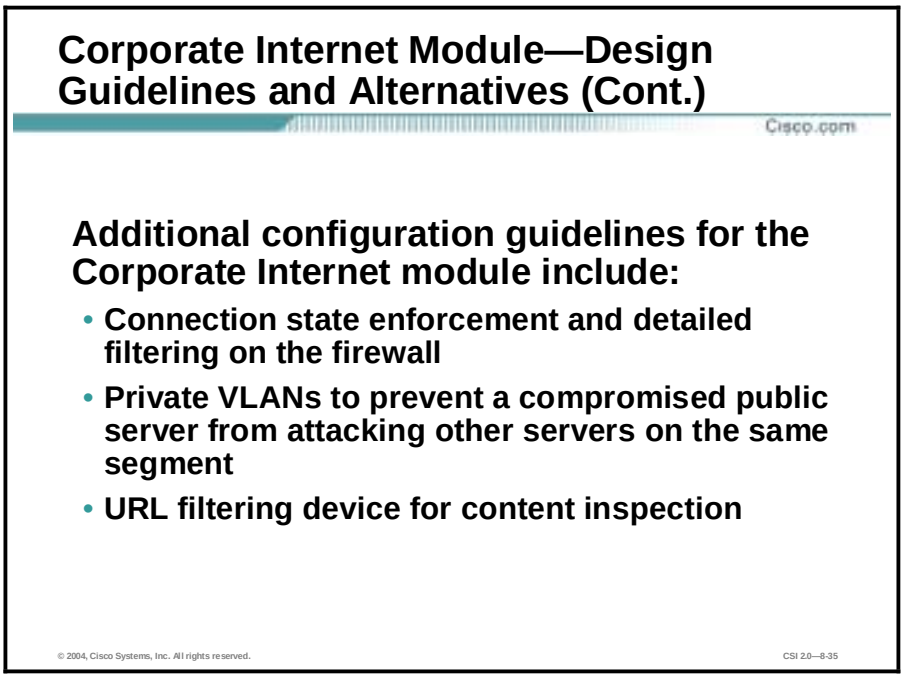

Additional configuration guidelines for the Corporate Internet module include the following:

- Connection state enforcement and detailed filtering—Publicly addressable servers have some protection against TCP SYN floods through the use of half-open connection limits on the firewall. In addition to limiting traffic on the public services segment to relevant addresses and ports, filtering in the opposite direction also takes place.

  Specific filtering prevents unauthorized requests from being generated by the public servers to mitigate an attack that compromises one of the public servers (by circumventing the firewall, HIPS, and NIDS). Such a server should not be able to further attack the network.
- Private VLANs—Private VLANs prevent a compromised public server from attacking other servers on the same segment. This traffic is not even detected by the firewall, which is why private VLANs are critical.
- URL filtering—Traffic on the content inspection segment is limited to URL filtering requests from the firewall to the URL filtering device. In addition, authenticated requests are allowed from the Enterprise URL filtering device out to a master server for database updates. The URL filtering device inspects outbound traffic for unauthorized web requests. It communicates directly with the firewall and approves or rejects URL requests sent to its URL inspection engine by the firewall. Its decision is based on a policy managed by the enterprise using classification information of the web provided by a third-party service. URL inspection was preferred over standard access filtering because IP addresses often change for unauthorized web sites, and such filters can grow to be very large. With URL filtering you sacrifice performance of your HTTP traffic for the greater control this inspection provides.

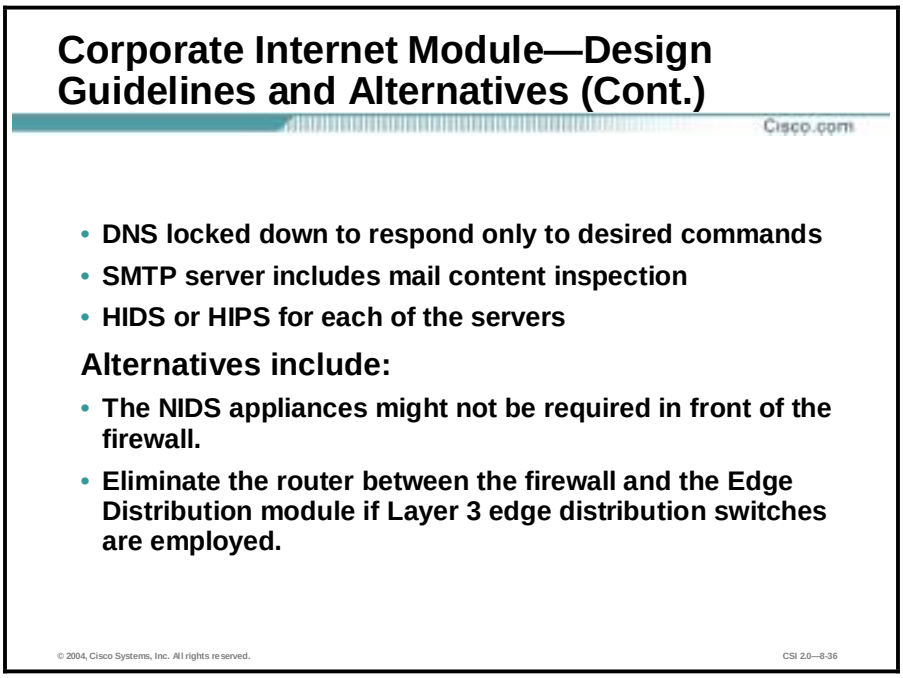

Additional configuration guidelines for the Corporate Internet module include the following:

- DNS host—Should be locked down to respond only to desired commands and eliminate any unnecessary responses that might assist hackers in network reconnaissance. This includes preventing zone transfers from anywhere but the internal DNS servers.
- SMTP server—Should include mail content inspection services that mitigate against virus and Trojan-type attacks generated against the internal network, which are usually introduced through the mail system. The firewall itself filters SMTP messages at Layer 7 to allow only necessary commands to the mail server.
- HIDS or HIPS—Each of the servers has host intrusion detection software on them to monitor against any rogue activity at the operating system level, as well as activity in common server applications (HTTP, FTP, SMTP, and so forth).

The following are alternative designs for this module:

- Eliminate NIDS appliances in front of the firewall—In fact, without basic filtering on the access router, this type of monitoring is not recommended. With the appropriate basic filters, the IDS outside the firewall can provide important alarm information that would otherwise be dropped by the firewall. Because the number of alarms generated on this segment is probably large, alarms generated here should have a lower severity than alarms generated behind a firewall. You should consider logging alarms from this segment to a separate management station to ensure that legitimate alarms from other segments get the appropriate attention. With the visibility that NIDS outside the firewall provides, you can better evaluate the attack types that your organization is attracting. In addition, you can evaluate the effectiveness of ISP and enterprise edge filters.
- Eliminate the router between the firewall and the Edge Distribution module—Though its functions can be integrated into the Edge Distribution module, the functional separation between modules would be lost because the edge distribution switches would need to be aware of the entire topology of the Corporate Internet module to ensure proper routing. In addition, eliminating the router limits your ability to deploy this architecture in a modular

fashion. If an enterprise's current core is Layer 2, for example, the routing provided in the Corporate Internet module would be required.

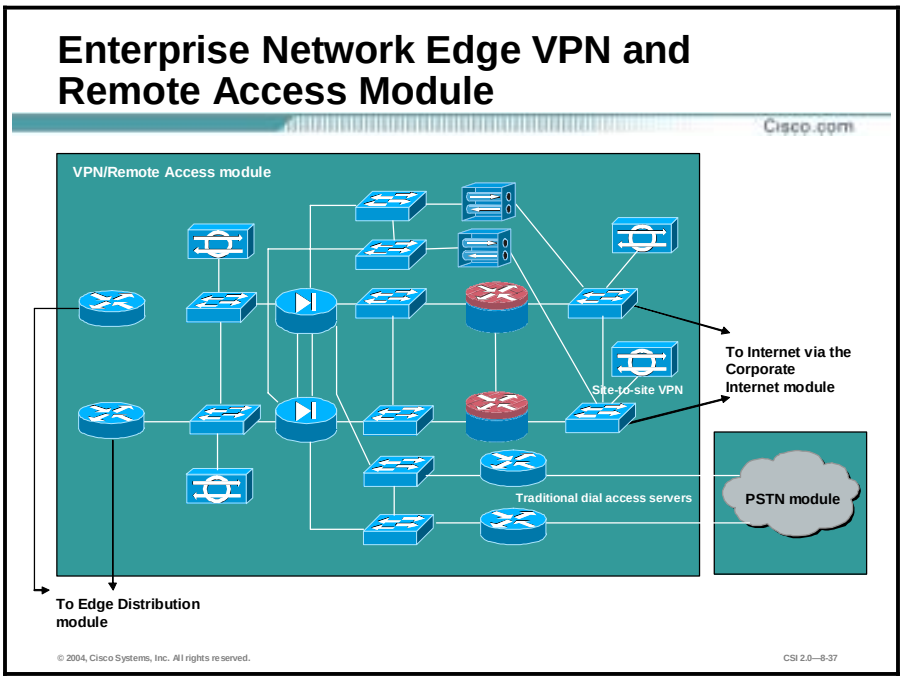

The primary objective of the VPN and Remote Access module is threefold: terminate the VPN traffic from remote users, provide a hub for terminating VPN traffic from remote sites, and terminate traditional dial-in users. All the traffic forwarded to the Edge Distribution module is from remote corporate users that are authenticated in some fashion before being allowed through the firewall.

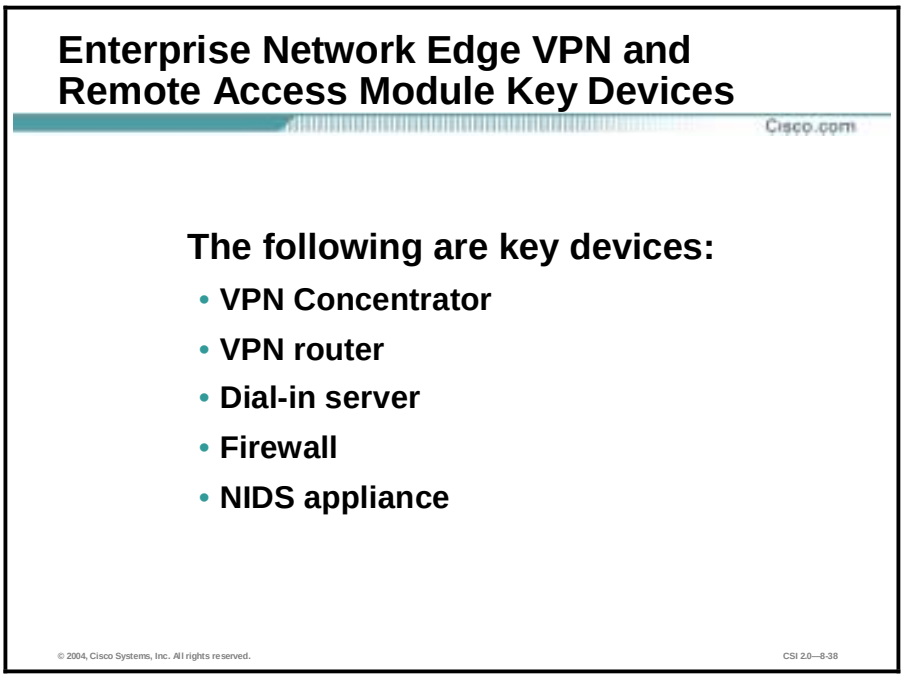

The following are key devices in SAFE Enterprise Edge VPN and Remote Access module:

- VPN Concentrator—Authenticates individual remote users using Extended Authentication (XAUTH) and terminates their IPSec tunnels
- VPN router—Authenticates trusted remote sites and provides connectivity using Generic Routing Encapsulation (GRE) and IPSec tunnels
- Dial-in server—Authenticates individual remote users using TACACS+ and terminates their analog connections
- Firewall—Provides differentiated security for the three different types of remote access
- NIDS appliance—Provides Layer 4 to Layer 7 monitoring of key network segments in the module

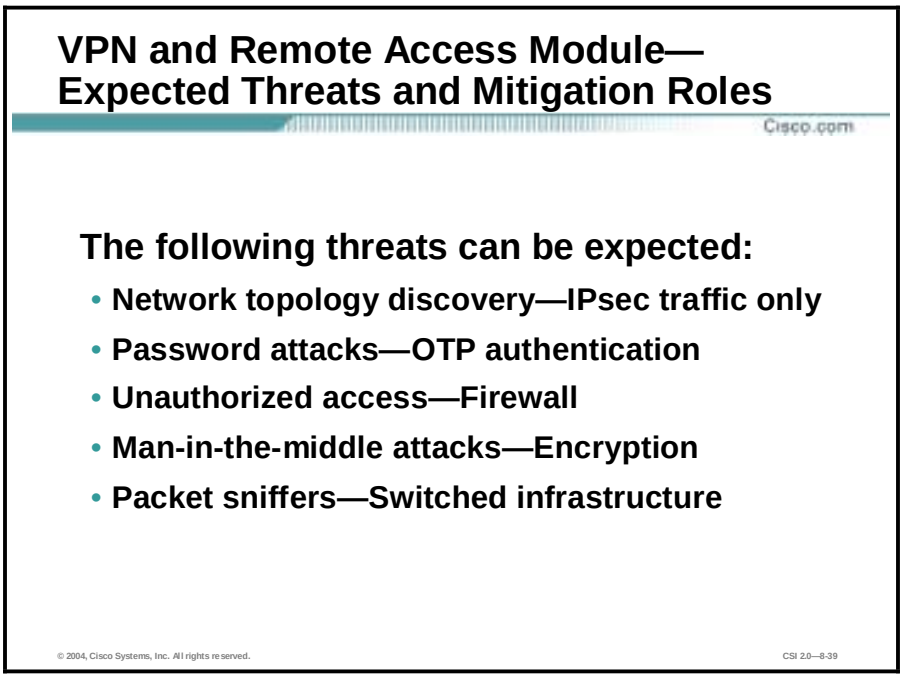

The following are expected threats and mitigation of those threats to the SAFE Enterprise Edge VPN and Remote Access module:

- Network topology discovery—Only Internet Key Exchange (IKE) and Encapsulating Security Payload (ESP) are allowed into this segment from the Internet.
- Password attacks—OTP authentication reduces the likelihood of a successful password attack.
- Unauthorized access—Firewall services after packet decryption prevent traffic on unauthorized ports.
- Man-in-the-middle attacks—Mitigated through encrypted remote traffic.
- Packet sniffers—A switched infrastructure limits the effectiveness of sniffing.

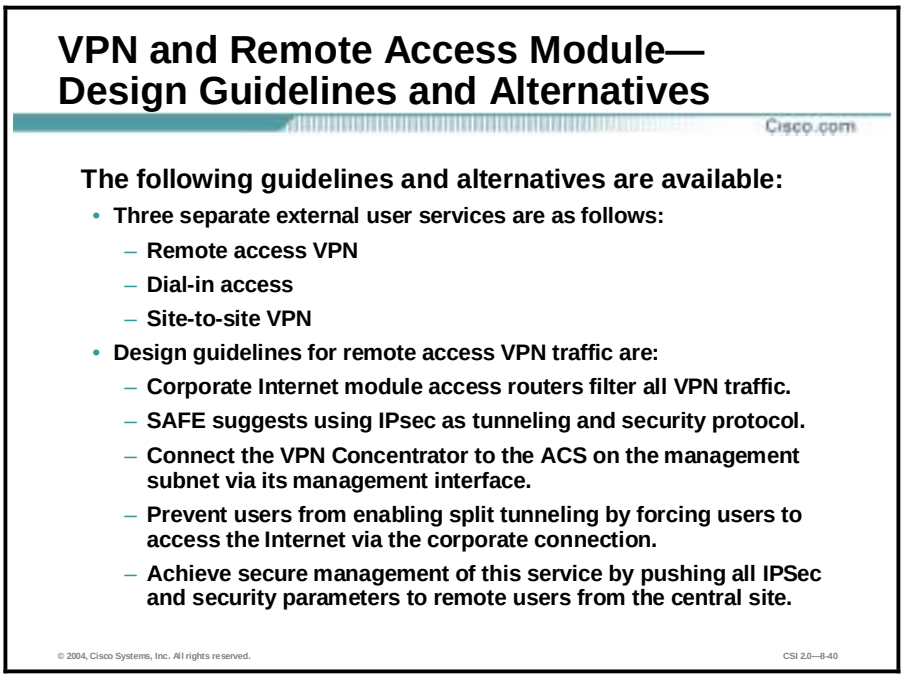

The main requirements of this module are resilience and to have three separate external user services authenticate and terminate: remote-access VPN, dial-in access, and site-to-site VPN. A separate interface is provided on the firewall for each of these three services. Design guidelines for remote-access VPN traffic are as follows:

- The VPN traffic is forwarded from the Corporate Internet module access routers, where it is first filtered at the egress point to the specific IP addresses and protocols that are part of the VPN services. SAFE recommends IPSec protocol because the clients require minimal configuration and at the same time provide good security. The IPSec parameters that are being used are Triple-Data Encryption Standard (3DES) for encryption and Secure Hash Algorithm (SHA) and hash-based message authentication code (HMAC) for data integrity.
- The remote-access VPN traffic will be addressed to one specific public address. IKE extensions (XAUTH and ModeCFG) provide an additional user authentication and authorization mechanism to control the access of the remote user.
- The VPN Concentrator is connected to the ACS on the management subnet via its management interface. Strong passwords are provided via the OTP server.
- The users are prevented from enabling split tunneling, thereby forcing the user to access the Internet via the corporate connection.
- Following termination of the VPN tunnel, traffic is sent through a firewall to ensure that VPN users are appropriately filtered.
- Secure management of this service is achieved by pushing all IPSec and security parameters to the remote users from the central site. Additionally, connections to all management functions are on a dedicated management interface.

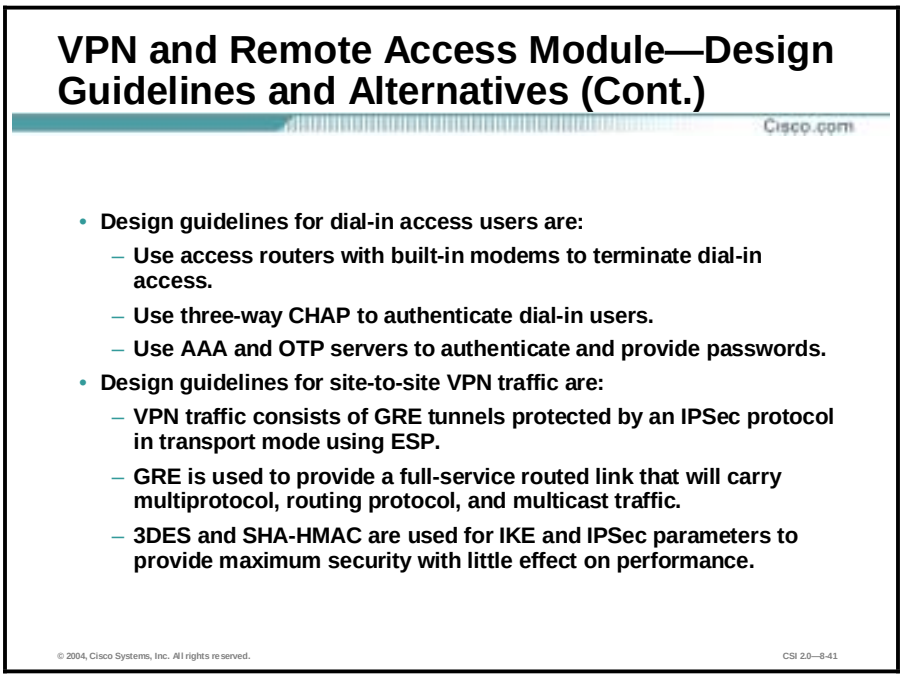

- Design guidelines for dial-in access users are as follows:
  - Use access routers with built-in modems terminate dial-in connections.
  - Three-way CHAP is used to authenticate the user.
  - The AAA and OTP servers are used to authenticate and provide passwords. Once authenticated, the users are provided with IP addresses from an IP pool through PPP.
- Design guidelines for site-to-site VPN traffic are as follows:
  - The VPN traffic associated with site-to-site connections consists of GRE tunnels protected by an IPSec protocol in transport mode using ESP. The traffic that is forwarded from the Corporate Internet module can be limited to the specific destination addresses on the two VPN routers and the source addresses expected from the remote sites. The ESP protocol and the IKE protocol will be the only two expected on this link. 3DES and SHA-HMAC are used for IKE and IPSec parameters to provide the maximum security with little effect on performance.
  - GRE is used to provide a full-service routed link that will carry multiprotocol, routing protocol, and multicast traffic. Because routing protocols (Enhanced Interior Gateway Routing Protocol [EIGRP] is being used between remote sites) can detect link failure, the GRE tunnel provides a resilience mechanism for the remote sites if they build two GRE connections, one to each of the central VPN routers.
  - IPSec hardware accelerators are used in the VPN routers.

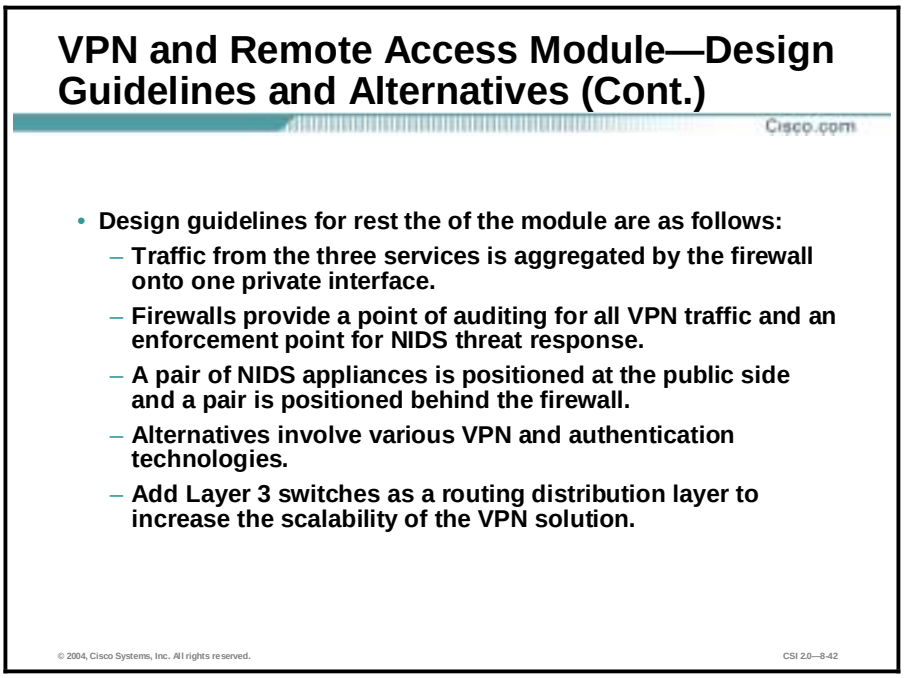

Design guidelines for rest of the module are as follows:

- The traffic from the three services is aggregated by the firewall onto one private interface before being sent to the Edge Distribution module via a pair of routers.
- The firewall must be configured with the right type of constraining access control to allow only the appropriate traffic through to the inside interface of the firewall from each of the services. In addition to access control, the firewalls provide a point of auditing for all VPN traffic and an enforcement point for NIDS threat response.
- A pair of NIDS appliances is positioned at the public side of the module to detect any network reconnaissance activity targeted at the VPN termination devices. On this segment, only IPSec (IKE/ESP) traffic should be seen. Because the NIDS system cannot see inside the IPSec packets, any alarm on this network indicates a failure or compromise of the surrounding devices. As such, these alarms should be set to high severity levels.
- A second pair of NIDS appliances is positioned behind the firewall to detect any attacks that made it through the rest of the module. All users crossing this segment should be bound to, or come from a remote location, so that any shunning or TCP resets will only affect those users. This allows a more restrictive level for the NIDS as opposed to, say, the Corporate Internet module where some of the NIDS devices have the potential to shut out legitimate users if too loosely configured.
- In VPN and authentication technology, there are many alternatives available depending on the requirements of the network. These alternatives are listed below for reference:
  - Smartcard and/or biometric authentication
  - Layer 2 Tunneling Protocol (L2TP) and/or Point-to-Point Tunneling Protocol (PPTP) remote-access VPN tunnels
  - Certificate Authorities (CAs)
  - IKE keep-alive resilience mechanism
  - Multiprotocol Label Switching (MPLS) VPNs

- An alternative VPN design is to add Layer 3 switches as a routing distribution layer before the clear-text traffic is sent through the firewall. This design significantly increases the scalability of the VPN solution.

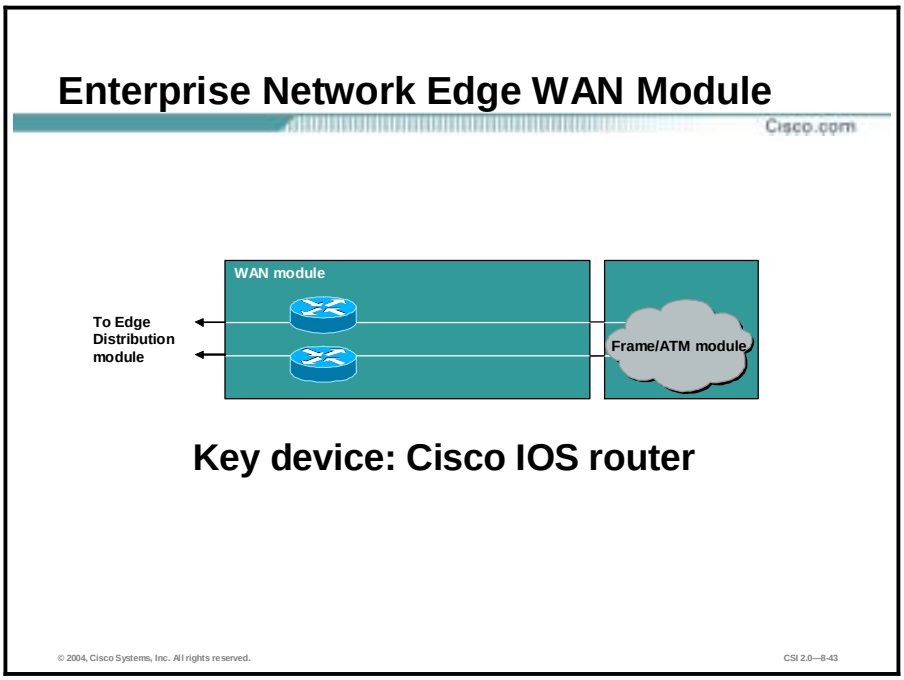

The WAN module shows resilience and security for WAN termination, rather than including all potential WAN designs. Using Frame Relay encapsulation, traffic is routed between remote sites and the central site.

The IOS router is the key device for the Enterprise Network Edge WAN module, using routing, access control, and QoS mechanisms.

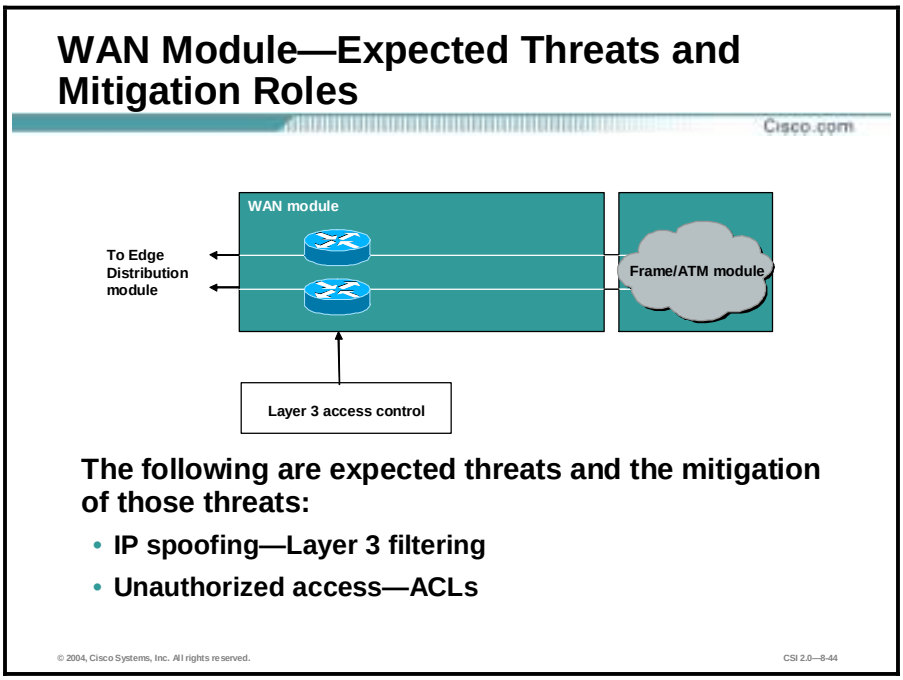

The following are expected threats and mitigation to those threats to the SAFE Enterprise Edge WAN module:

- IP spoofing—IP spoofing is mitigated through Layer 3 filtering.
- Unauthorized access—Simple access control on the router can limit the types of protocols to which branches have access.

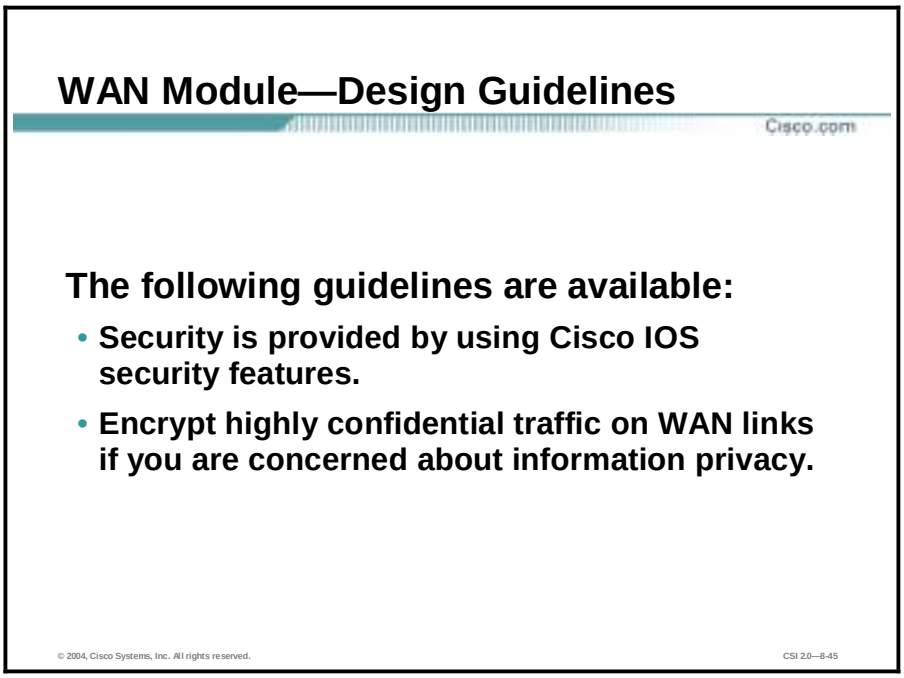

The resilience is provided by the dual connection from the service provider, through the routers, and to the Edge Distribution module. Following are the design guidelines:

- Security is provided by using IOS security features. Input access lists are used to block all unwanted traffic from the remote branch.
- Some organizations that are very concerned about information privacy encrypt highly confidential traffic on their WAN links. Similarly to site-to-site VPNs, you can use IPSec to achieve this information privacy.

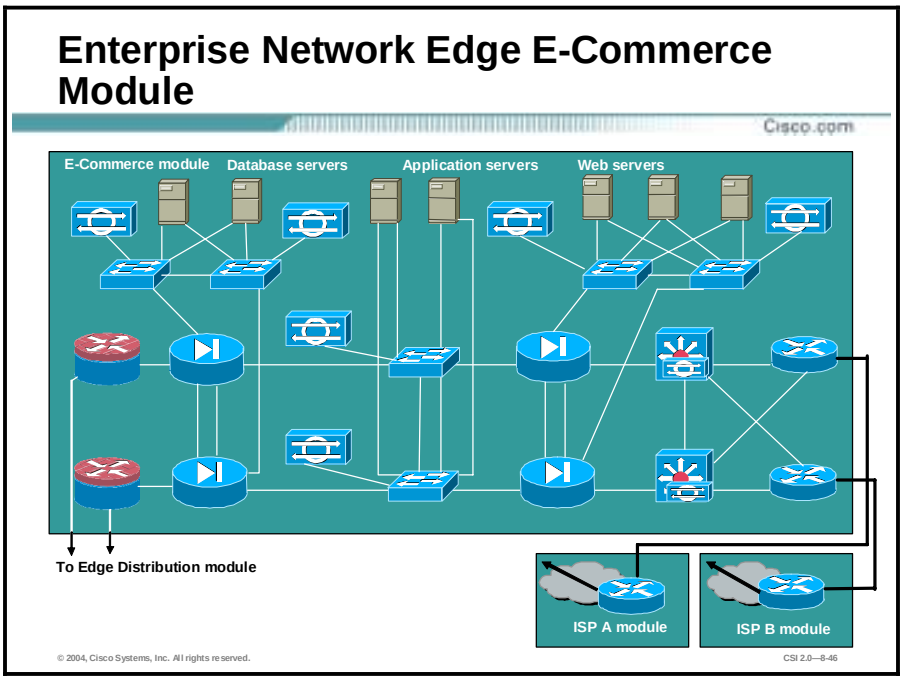

Because e-commerce is the primary objective of this module, the balance between access and security must be carefully weighed. Splitting the e-commerce transaction into three components (Database, Application and Web servers) allows the architecture to provide various levels of security without impeding access.

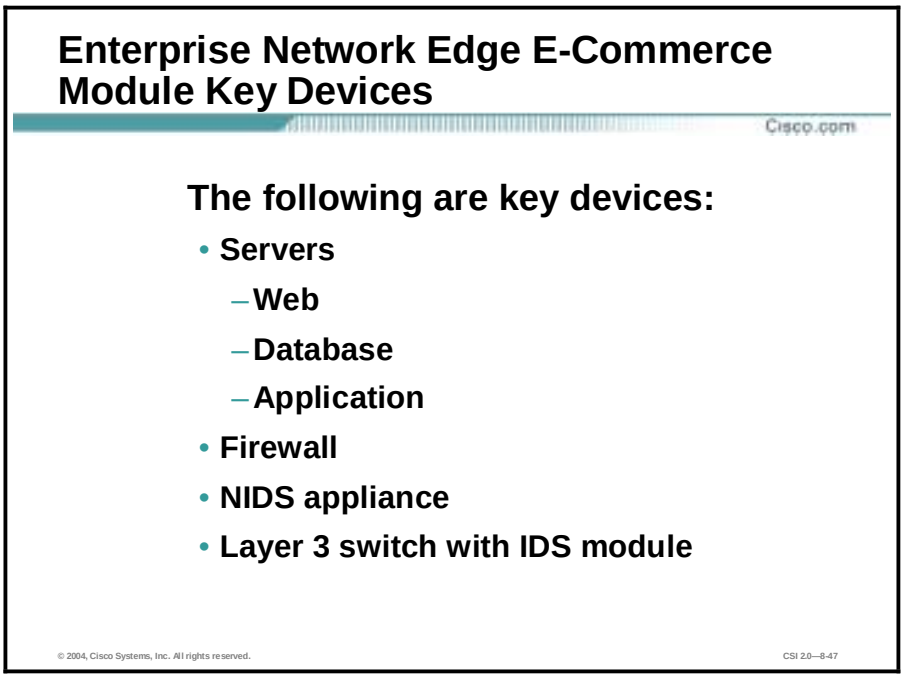

The following are key devices for the SAFE Enterprise Edge E-Commerce module:

- Web server—Acts as the primary user interface for the navigation of the e-commerce store
- Application server—The platform for the various applications required by the web server
- Database server—The critical information that is the heart of the e-commerce business implementation
- Firewall—Governs communication between the various levels of security and trust in the system
- NIDS appliance—Provides monitoring of key network segments in the module
- Layer 3 switch with IDS module—The scalable e-commerce input device with integrated security monitoring

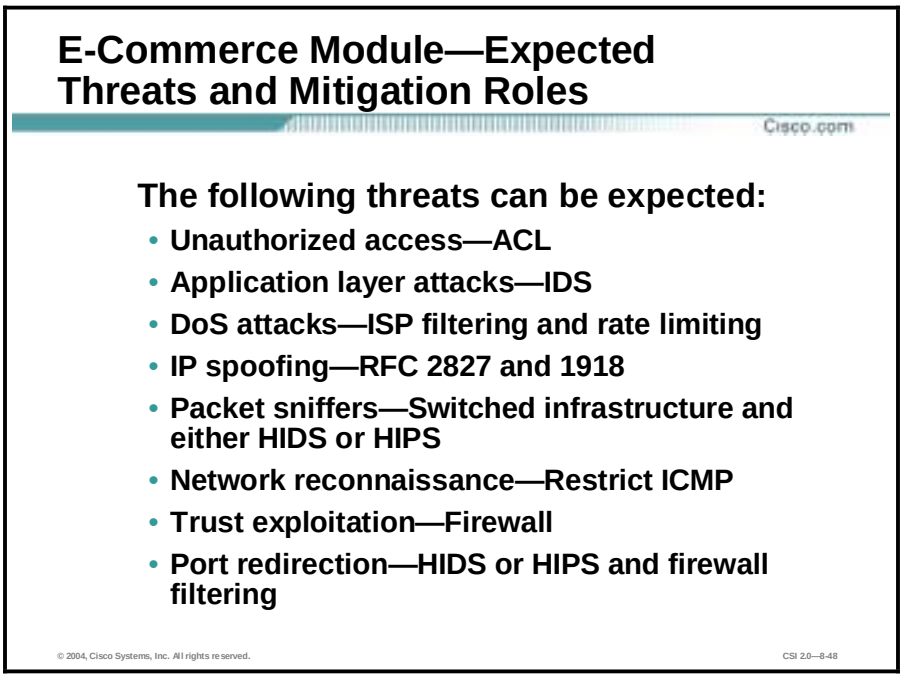

The following are expected threats and mitigation of those threats to the SAFE Enterprise Edge E-Commerce module:

- Unauthorized access—Stateful firewalls and ACLs limit exposure to specific protocols.
- Application layer attacks—Attacks are mitigated through the use of IDS.
- DoS—ISP filtering and rate-limiting reduce DDoS and DoS potential
- IP spoofing—RFC 2827 and 1918 prevent locally originated spoofed packets and limit remote spoof attempts.
- Packet sniffers—A switched infrastructure and either HIDS or HIPS to limit the effectiveness of sniffing.
- Network reconnaissance—Ports are limited to only what is necessary. Internet Control Message Protocol (ICMP) is restricted.
- Trust exploitation—Firewalls ensure communication flows only in the proper direction on the proper service.
- Port redirection—HIDS or HIPS and firewall filtering limit exposure to these attacks.

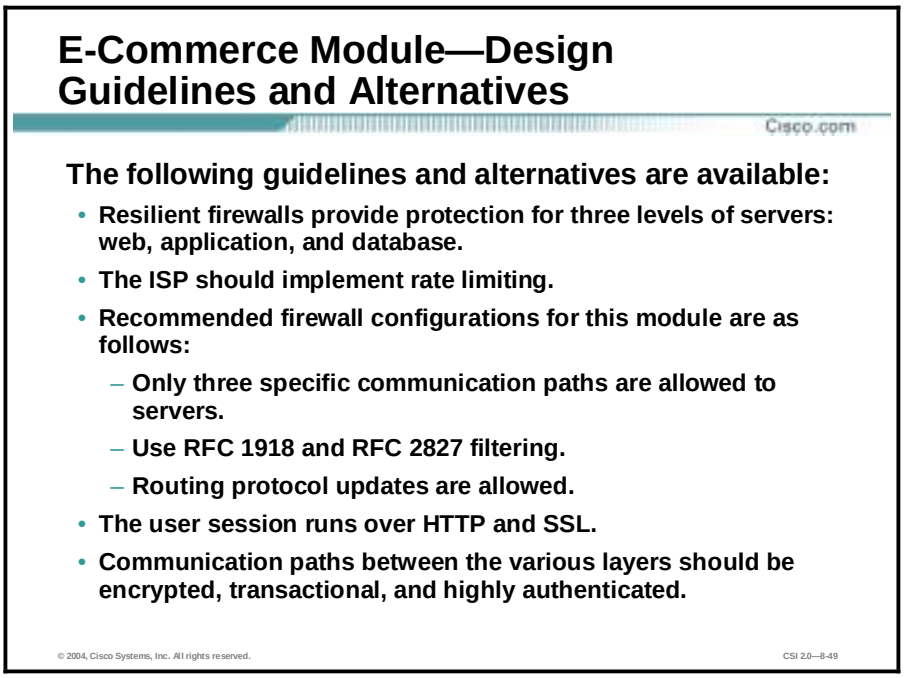

The heart of the module is two pairs of resilient firewalls that provide protection for the three levels of servers: web, application, and database. Some added protection is provided by the ISP edge routers at the ISP and the enterprise. The following design guidelines and alternatives are available:

- The ISP should implement rate limiting in order to mitigate DDoS and DoS attacks.
- Recommended configurations for the firewall in this module are as follows:
  - Firewalls must allow only three specific communication paths, each with its own protocol, and block all other communication unless it is the return path packets that are associated with the three original paths.
  - The e-commerce firewalls are initially protected by the customer edge router at the ISP. At the router egress point, towards the enterprise, the ISP can limit the traffic to the small number of protocols required for e-commerce with a destination address of the web servers only.
  - Routing protocol updates (generally Border Gateway Protocol [BGP]) are required by the edge routers, and all other traffic should be blocked.
  - Filtering according to RFC 1918 and RFC 2827 should also be implemented by the ISP.
- The DNS is hosted on a different network to reduce the amount of protocols required by the e-commerce application.
- User's entire session runs over HTTP and SSL with no ability to communicate directly with the application server or the database server.
- Communications paths between the various layers (web, applications, and database) should be encrypted, transactional, and highly authenticated.

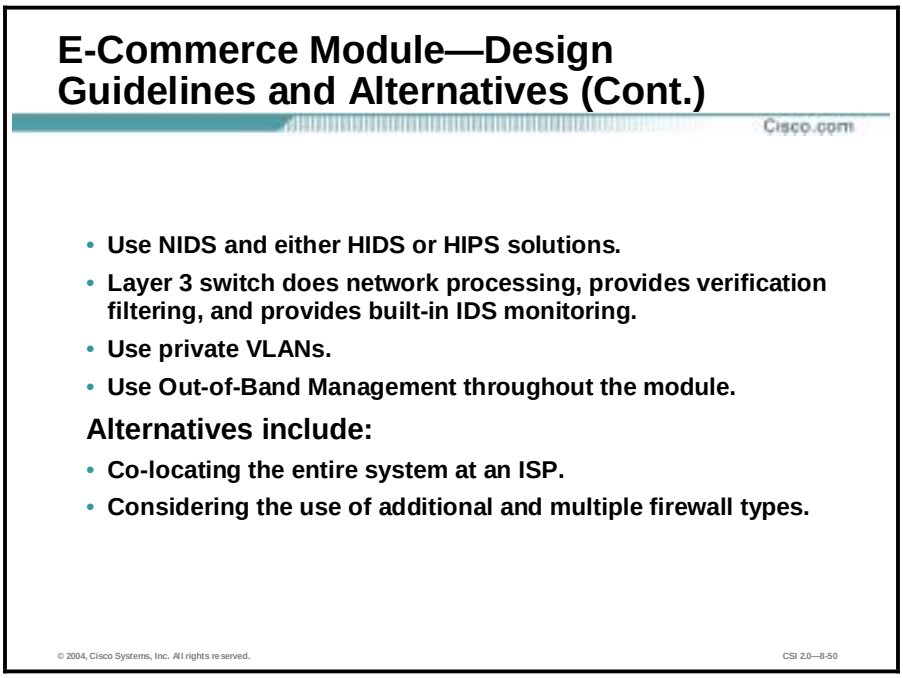

Additional design guidelines and alternatives for the E-Commerce module are as follows:

- HIDS or HIPS—All servers must be fully protected, especially the web server, which is a publicly-addressable host. The operating system and server applications must be patched to the latest versions and monitored by the host intrusion detection software.
- NIDS appliances—NIDS appliances should be behind the various interfaces of the firewall, to monitor the segments for any attacks that might have penetrated the first line of defense. The false-positives must be removed so that all true attack detections are treated with the correct level of priority. In fact, because only certain types of traffic exist on certain segments, you can tune NIDS very tightly.
- Layer 3 switches—These switches do the entire network processing because they have features off-loaded to hardware processors. The Layer 3 switches participate in the full BGP routing decision in order to decide which ISP has the better route to the particular user. The Layer 3 switches also provide verification filtering in keeping with the ISP filtering described above, which provides overlapping security. Finally, the Layer 3 switches provide built-in IDS monitoring.
- Private VLANs—Use private VLANs to implement a trust model that matches the desired traffic communication on a particular segment and eliminates all others.
- Out-of-band management—The management of the entire module is done completely out of band as in the rest of the architecture.
- The alternatives to this deployment are as follows:
  - Co-locating the entire system at an ISP.
  - The use of additional firewalls. This allows each firewall to only control communications for one primary system.
  - For very high security requirements, the use of multiple firewall types may be considered.

Summary

This topic summarizes the information you learned in this lesson.

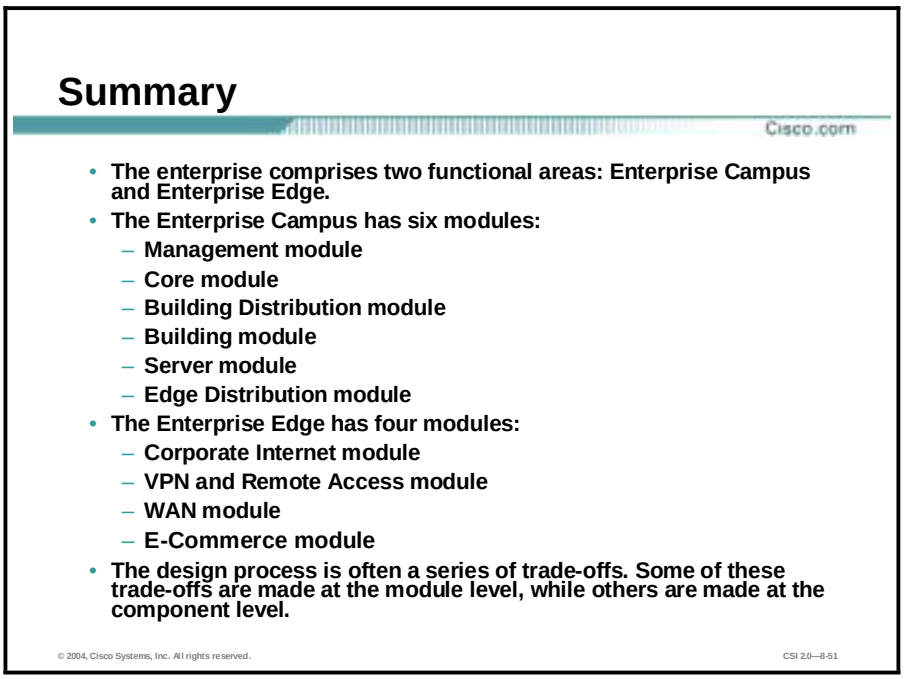

9

SAFE: IP Telephony Security in Depth

Overview

This lesson includes the following topics:

- Objectives
- IP telephony concepts
- IP telephony caveats
- IP telephony axioms
- Cisco IP telephony product portfolio
- SAFE IP telephony design considerations
- Small network IP telephony design
- Medium network IP telephony design
- Large network IP telephony design
- Summary
- Review questions

Objectives

This topic lists the lesson's objectives.

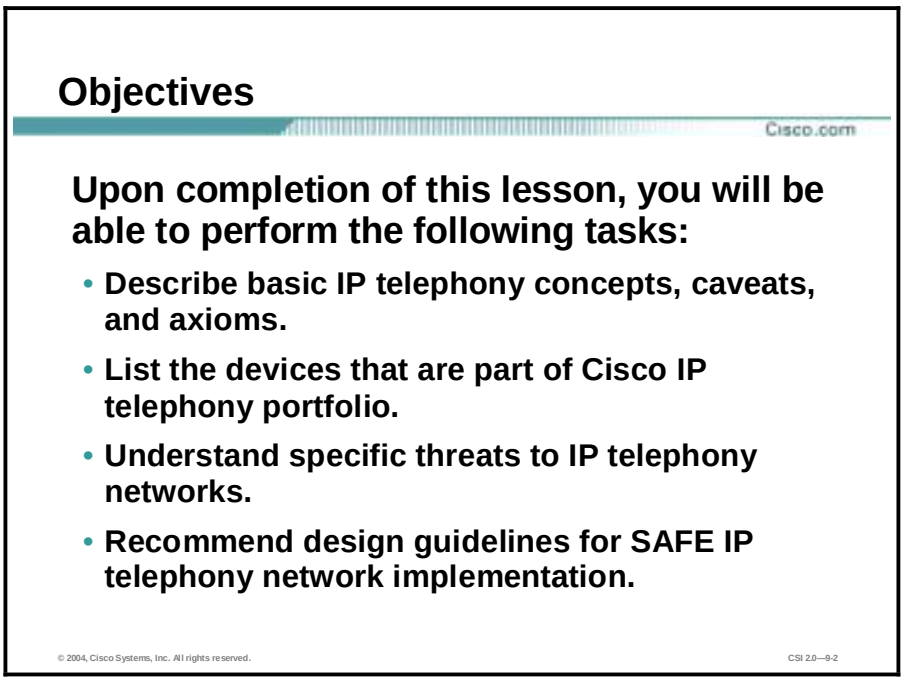

IP Telephony Concepts

This topic provides an overview of basic IP telephony concepts.

The convergence of voice and data traffic on a single IP network is revolutionizing communications. Voice can now be transported as high-priority data, lowering the cost of network ownership and enhancing business communications via the enablement of value-added applications. Many of the new business applications that are now deployed on converged networks provide immediate ways to increase personal and workgroup productivity while enhancing customer care and responsiveness.

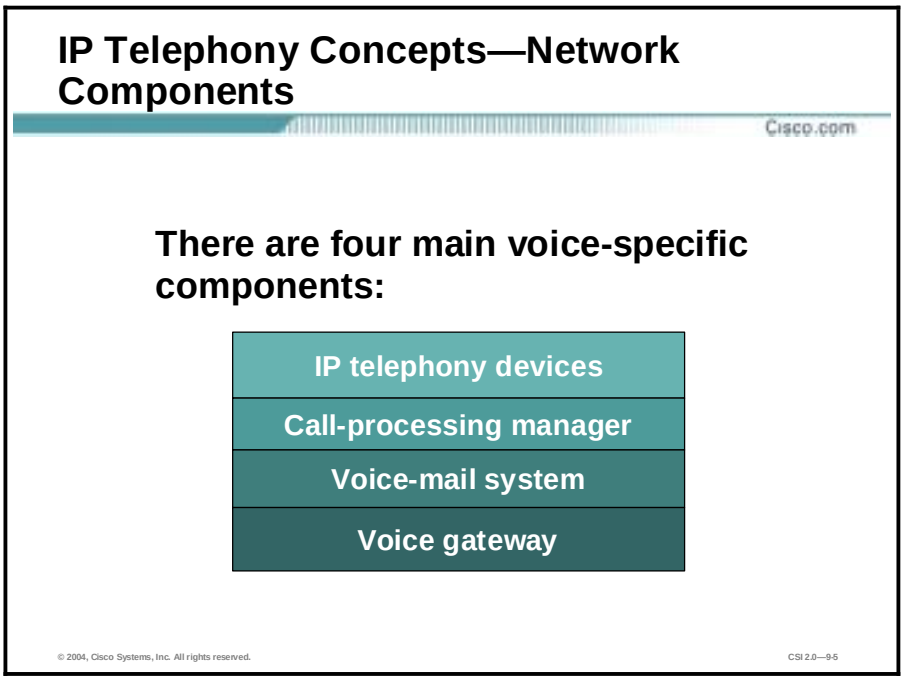

There are four main voice-specific components of an IP telephony network:

- IP telephony devices—This refers to any device that supports placing calls in an IP telephony network. IP phones are included as well as applications installed on user systems with speakers and microphones. IP phones offer services such as user directory lookups and Internet access for stock quotes; these are referred to as IP phone services and are accessed via a proxy server.
- Call-processing manager—This server provides call control and configuration management for IP telephony devices. It is also known as “IP PBX.” This device provides the core functionality to bootstrap IP telephony devices, provide call setup, and route calls throughout the network to other voice devices, including voice gateways and voice-mail systems.
- Voice-mail system—Provides IP-based voice-mail storage and an auto attendant (an automated attendant providing voice services) for services such as user directory lookup and call forwarding.
- Voice gateway—A general term used to refer to any gateway that provides voice services, including such features as Public Switched Telephone Network (PSTN) access, IP packet routing, backup call-processing, and voice services. This is the device that provides access to legacy voice systems for local calls, toll bypass, and WAN backup in case of failure. Backup call processing allows for the voice gateway to take over call processing in case the primary call-processing manager goes offline for any reason. Typically the voice gateway supports a subset of the call-processing functionality supported by the call-processing manager.

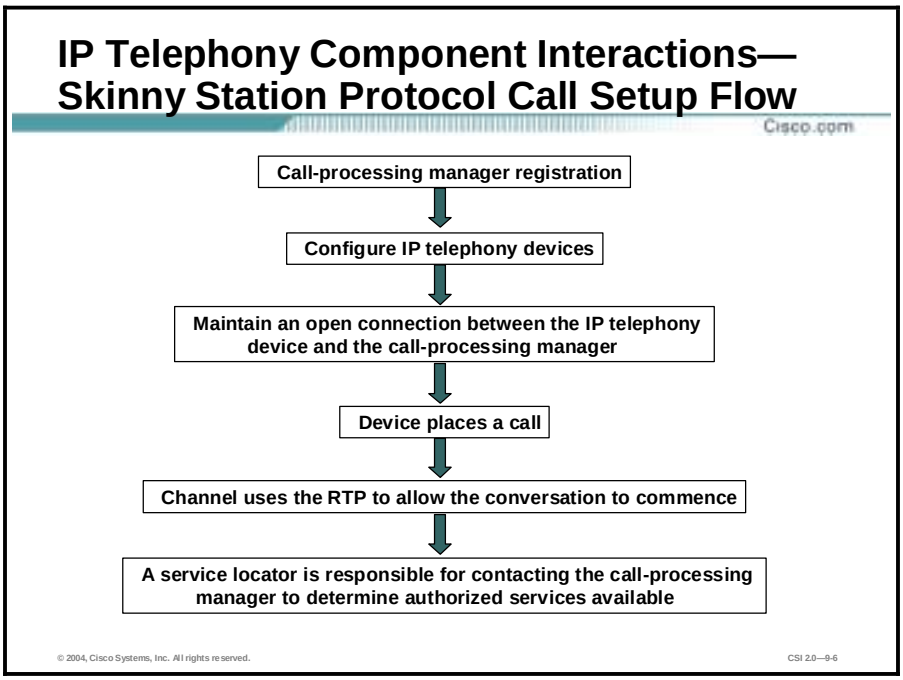

These IP telephony component interactions model that of the Skinny Station Protocol and are provided as a reference call setup flow:

- Call-processing manager registration—All IP telephony devices must complete call-processing manager registration before placing a call. For the Cisco Skinny Protocol, this process occurs over TCP port 2748 (control channel) and UDP port 69 (TFTP) for phone configuration and firmware updates.
- Configure IP telephony devices—The IP telephony devices are configured with access to voice-mail, data services, time-of-day, speed-dials, any other custom configurations, and an extension. If voice-mail is pending, the message-waiting indicator light will illuminate. If the device is an IP phone and there is new operating system code available, it will download the code, install it, reboot, and restart the registration process.
- Maintain an open connection at all times between the IP telephony device and the call-processing manager—This can be used to notify the IP telephony device if a voice-mail was received, and to ensure that a call-processing manager is always available for call processing.
- Device places a call—Once the user dials the extension of the remote user that they want to talk to, the extension is sent to the call-processing manager, which then in turn notifies the destination device that a call is incoming. Because the destination device went through the registration process, the call-processing manager will be able to fulfill the call. Once the remote user takes the device off-hook, the remote device notifies the call-processing manager that it is willing to accept the call. At this point the call-processing manager notifies both devices that a channel is now available for them to converse over.
- Channel uses the RTP to allow the conversation to commence—Real-Time Transport Protocol (RTP) running on top of UDP/IP is used by the channel to allow the conversation to commence. The UDP stream originates on the calling IP phone and terminates on the target IP phone; the call-processing manager does not access the stream. Once one party hangs up, the call-processing manager notifies the other side and the UDP session is torn down. If the remote user had been unavailable, the call-processing manager would have

instead established an RTP session with the voice-mail server. At that time the user could leave a voice-mail message.

- The services available to the user are displayed on the phone—IP phones support dynamic content services via an HTTP client and Extensible Markup Language (XML) support. The IP phone contains a service locator that is responsible for contacting the call-processing manager to determine what authorized services are available. When a service is accessed on the IP phone, it creates an HTTP connection to the proxy server. Once the list of services is determined, the phone then creates another HTTP connection to the call-processing manager. The phone's identification (Media Access Control [MAC] address) is included in this connection request. This allows the call-processing manager to check what services the phone has enabled and to then notify the phone. At this point the services available to the user are displayed on the phone. If the user chooses, for example, a stock quote, the request is sent to the proxy server over HTTP to fulfill the request. The proxy server then establishes its own HTTP connection outbound to retrieve the stock's price, records the result, and finally sends the result back to the phone and it is displayed to the user.

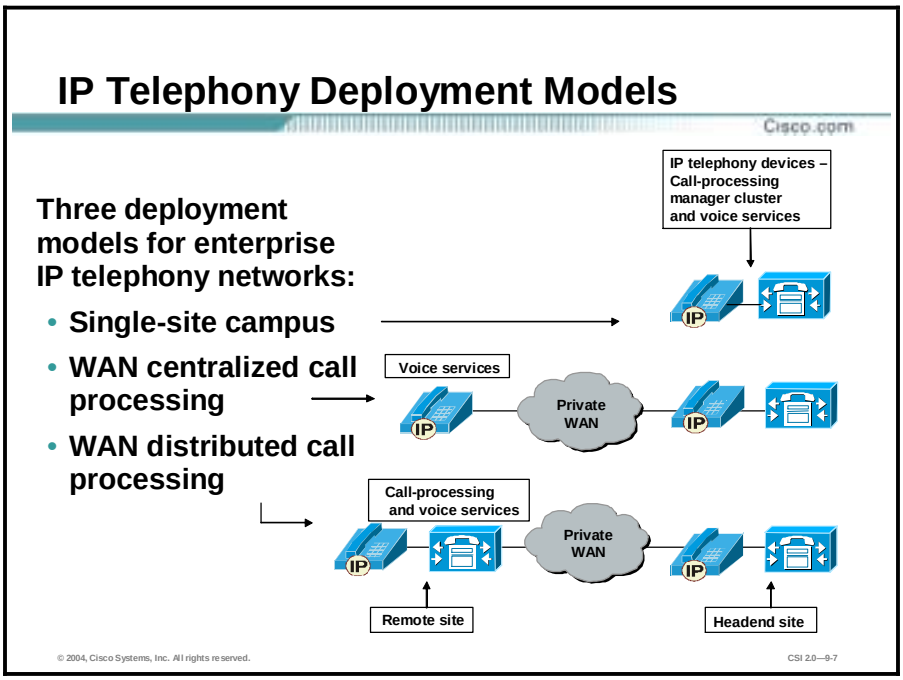

There are three models for the deployment of enterprise IP telephony networks:

- Single-site campus—This is the most basic deployment scenario in which all IP telephony devices reside in a single campus.
- WAN centralized call processing—This is a moderately complex scenario in which multiple sites are connected over a private WAN and the headend site contains the only call-processing manager cluster. Remote sites may have voice services such as voice mail.
- WAN distributed call-processing—This is the most complex scenario in which multiple sites are connected over a private WAN and one or more of the sites contains a call-processing manager cluster. Many of the sites will also have voice services such as voice mail (not covered in this lesson).

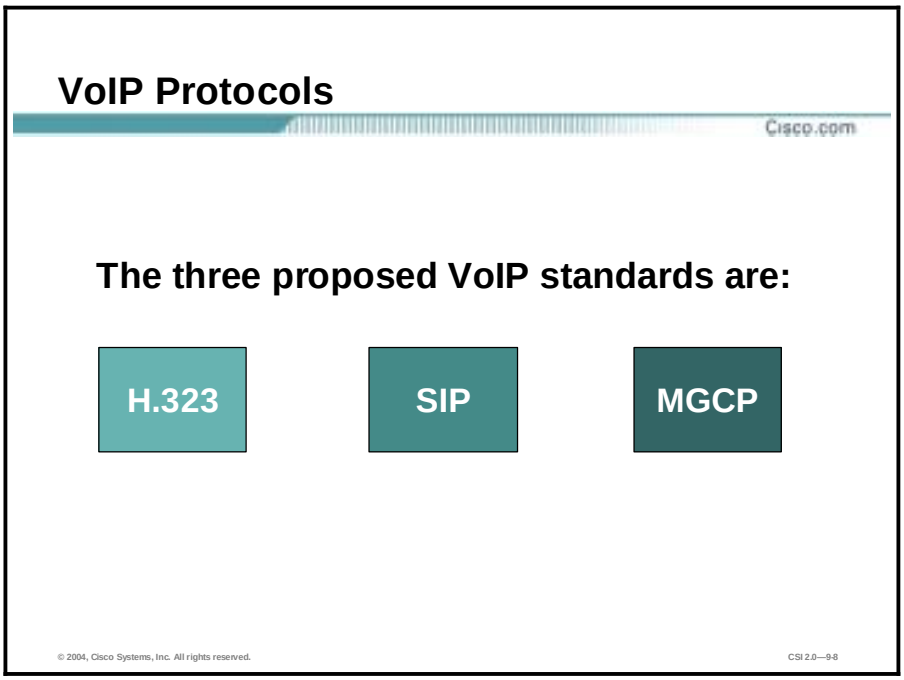

The three front-running proposed Voice over IP (VoIP) standards are H.323, Session Initiation Protocol (SIP), and Media Gateway Control Protocol (MGCP). They are described here for reference:

- H.323—The H.323 standard is a logical extension of the H.320 standard to enable corporate intranets and packet-switched networks to transport multimedia and conference traffic. The following are some important features of this protocol:
  - H.323 recommendations cover the IP devices that participate and control H.323 sessions and the elements that interact with the Switched Circuit Network (SCN).
  - H.323 standards do not include the LAN itself, nor the transport layer that interconnects LAN segments. In common with other International Telecommunication Union (ITU) multimedia teleconferencing standards, H.323 implementation applies to either point-to-point or multipoint sessions.
  - The H.323 recommendation allows multipoint conferences through a variety of methods or configurations.
  - Ports or sockets used for H.245 signaling, audio, video, or data channels are dynamically negotiated between endpoints.
  - The use of dynamic sockets makes it difficult to implement security, policy, and traffic shaping.
  - H.323 data conferencing uses both “reliable” and “unreliable” communications. Reliable transport is for control signals and data, because signals must be received in proper order and cannot be lost. Unreliable transport is used for audio and video streams, which are time sensitive. Delayed audio and video packets are dropped.
  - TCP is applied to the H.245 control channel, the T.120 data channels, and the call-signaling channel, whereas UDP applies to audio, video, and registration, admission, and status (RAS) channels.
  - Because H.323-compliant applications use dynamically allocated sockets for audio, video, and data channels, a firewall must be able to allow H.323 traffic through on

an intelligent basis. The firewall must be either H.323-enabled with an H.323 proxy, or it must be able to "snoop" the control channel to determine which dynamic sockets are in use for H.323 sessions, and allow traffic as long as the control channel is active.

- SIP—Session Initiation Protocol (SIP) is the Internet Engineering Task Force's (IETF's) standard for multimedia conferencing over IP. SIP is an ASCII-encoded, application-layer control protocol (defined in RFC 2543) that can be used to establish, maintain, and terminate calls between two or more endpoints. Like other VoIP protocols, SIP is designed to address the functions of signaling and session management within a packet telephony network. Signaling allows call information to be carried across network boundaries. Session management provides the ability to control the attributes of an end-to-end call. Users in a SIP network are identified by unique SIP addresses. A SIP address is similar to an e-mail address and is in the format sip:userID@gateway.com. The user ID can be either a username or an E.164 address. Users register with a registrar server using their assigned SIP addresses. The registrar server provides this information to the location server upon request. When a user initiates a call, a SIP request is sent to a SIP server (either a proxy or a redirect server). The request includes the address of the caller (in the From header field) and the address of the intended addressee (in the To header field). SIP provides the capabilities to do the following:
  - Determine the location of the target endpoint—SIP supports address resolution, name mapping, and call redirection.
  - Determine the media capabilities of the target endpoint—Via Session Description Protocol (SDP), SIP determines the lowest level of common services between the endpoints. Conferences are established using only the media capabilities that can be supported by all endpoints.
  - Determine the availability of the target endpoint—If a call cannot be completed because the target endpoint is unavailable, SIP determines whether the called party is already on the phone or did not answer in the allotted number of rings. It then returns a message indicating why the target endpoint was unavailable.
  - Establish a session between the originating and target endpoint—If the call can be completed, SIP establishes a session between the endpoints. SIP also supports midcall changes, such as the addition of another endpoint to the conference or the changing of a media characteristic or codec.
  - Handle the transfer and termination of calls—SIP supports the transfer of calls from one endpoint to another. During a call transfer, SIP simply establishes a session between the transferee and a new endpoint (specified by the transferring party) and terminates the session between the transferee and the transferring party. At the end of a call, SIP terminates the sessions between all parties.
  - Handle conferences consisting of two or more users which can be established using multicast or multiple unicast sessions.
- MGCP—In Media Gateway Control Protocol (MGCP), media gateway controllers or call agents provide control, signaling, and the processing skills to control the telephony gateways. One of the goals of MGCP is simplicity. A telephony gateway is a network device that provides conversion between the audio signals carried on telephone circuits and data packets carried on packet networks. MGCP assumes a call-control architecture wherein the call-control "intelligence" is outside the gateways and is handled by external call-control elements. MGCP is a master/slave protocol, wherein the gateways execute the commands sent by the call agents. The call agent implements the signaling layers of H.323 and appears to H.323 devices as an H.323 gatekeeper or one or more H.323 endpoints.

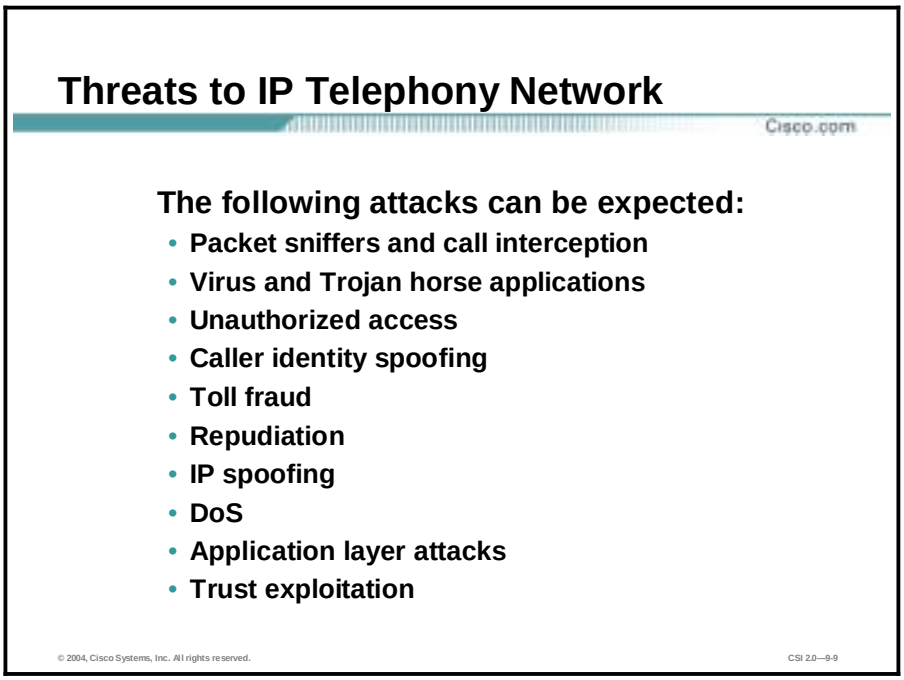

You can expect the following threats in an IP telephony network:

- Packet sniffers and call interception—A packet sniffer is a software application that uses a network adapter card in promiscuous mode to capture all network packets that are sent across a particular collision domain. Sniffers are used legitimately in networks today to aid in troubleshooting and traffic analysis. Since voice transport mechanisms generally don't use encryption, the voice streams can be saved and reassembled for listening. The tool "voice over misconfigured Internet telephones" (also known as "vomit") takes an IP phone conversation trace captured by the UNIX tool tcpdump, and reassembles it into a wave file for easy listening. The phones are not actually misconfigured. Rather, if someone was able to obtain access to the IP data stream at any point in the network, they could eavesdrop.
- Virus and Trojan horse applications—The primary vulnerabilities for end-user workstations are viruses and Trojan horse attacks. Viruses refer to malicious software that is attached to another program to execute a particular unwanted function on a user's workstation. A Trojan horse is different in that the entire application was written to look like something else, when in fact it is an attack tool. You should be concerned about viruses and Trojan horses in the data segment because they infect the PC-based IP Phone hosts that connect to the voice segment.
- Unauthorized access—An unauthorized-access attack, although not a specific type of attack, refers to most attacks executed in networks today. In order for someone to brute-force attack a Telnet login, they must first get the Telnet prompt on a system. These kinds of attacks can be initiated on both the outside and inside of a network. Hackers may attempt to gain unauthorized access to voice services with malicious intent in mind.
- Caller identity spoofing—This type of attack occurs when a hacker is able to trick a remote user into believing they are talking to a particular person when in fact they are really talking to the hacker. This type of attack typically occurs with the hacker assuming the identity of someone who is not familiar to the target. A complex attack would be to first place a rogue IP phone in the network and then via a secondary exploit assume the identity of a valid IP phone (assuming the identity that you want your target to see). On the other

hand, it may be as simple as bypassing physical building security and using an unattended IP phone.

- Toll fraud—This attack constitutes theft of service, namely phone calls. There are numerous methods the hacker could use to accomplish this task. The basic case of toll fraud involves an unauthorized user accessing an unattended IP phone in order to place calls. A more complex attack might include placing a rogue IP phone or gateway on the network in order to place unauthorized calls.
- Repudiation—If two parties talk over the phone and later one party denies that the conversation ever happened, what proof would the other party have that it ever occurred? This type of attack is not common and is difficult to mitigate. Call logging is available. Without strong user authentication (and users who log out when not using their IP phones), it is not possible to validate who actually placed a call.
- IP spoofing—An IP spoofing attack occurs when a hacker inside or outside a network impersonates the conversations of a trusted computer. A hacker can do this in one of two ways: Using an IP address that is within the range of trusted IP addresses for a network, or using an authorized external IP address that is trusted and to which access is provided to specified resources on a network. IP spoofing attacks are often launching points for other attacks. The classic example is the launching of a denial of service (DoS) attack using spoofed source addresses to hide the hacker's identity. Without spoof mitigation filters, a hacker might be able to spoof the address of the call-processing manager and UDP flood the entire voice segment.
- Denial of service—DoS attacks are the most publicized forms of attack, and are also among the most difficult to completely eliminate. Even among the hacker community, DoS attacks are regarded as trivial and considered bad form because they require so little effort to execute. Because of the ease of implementation and potentially significant damage, DoS attacks deserve special attention from security administrators. If not properly mitigated, the following DoS attacks could render a voice segment unusable:
  - TCP SYN flood
  - Ping of death
  - UDP fragment flood
  - Internet Control Message Protocol (ICMP) fragment flood
- Application layer attacks—Application layer attacks can be implemented using several different methods. One of the most common methods is exploiting well-known weaknesses in software that are commonly found on servers, such as sendmail, HTTP, and FTP. By exploiting software weaknesses, hackers can gain access to a computer with the permissions of the account running the application, which is usually a privileged system-level account. Application layer attacks are often widely publicized in an effort to allow administrators to rectify the problem with a patch. Unfortunately, many hackers also learn about the attacks. The primary problem with application layer attacks is that they often use ports that are allowed through a firewall. For example, a hacker executing a known vulnerability against a web server often uses TCP port 80 in the attack. Because the web server serves pages to users, a firewall needs to allow access on that port. From the perspective of the firewall, it is merely standard port 80 traffic. For this reason, either a host-based intrusion detection system (HIDS) or a host-based intrustion prevention system (HIPS) is also used on call-processing managers even though they are protected by a stateful firewall. Application layer attacks are the most common methods of gaining access to devices.
- Trust exploitation—Although not an attack as such, trust exploitation refers to a situation where an individual takes advantage of a trust relationship within a network. The classic

example is a perimeter network connection from a corporation. These network segments often house Domain Name System (DNS), Simple Mail Transfer Protocol (SMTP), and HTTP servers. Because they all reside on the same segment, a compromise of one system can lead to the compromise of other systems, because they might trust other systems attached to their same network. Another example is a system on the outside of a firewall that has a trust relationship with a system on the inside of a firewall. When the outside system is compromised, it can take advantage of that trust relationship to attack the inside network. This type of attack rarely surfaces in a secure IP telephony design. However, if multiple voice and data servers reside in the same segment, and the data server is compromised, it may then be possible to compromise the voice server as well, even though they are both protected by a stateful firewall.

IP Telephony Caveats

This topic describes IP telephony caveats.

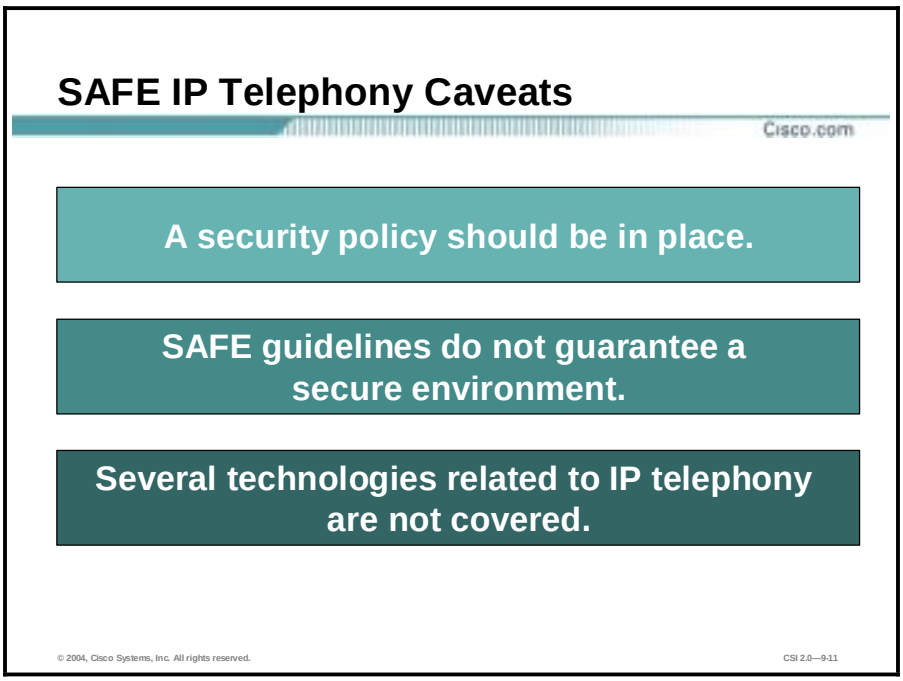

The following are the SAFE IP telephony caveats:

- A security policy should be in place—It is presumed that you already have a security policy in place. Cisco Systems does not recommend deploying any technology without an associated security policy. It presumes you are aware of what data is sensitive in your network so that it can be properly protected when transported throughout the network.
- SAFE guidelines do not guarantee a secure environment—Following the guidelines in the SAFE security documents does not guarantee a secure environment. SAFE practices help you make an informed choice about the risks and benefits of the technology. Only after you have weighed the risks and benefits should you deploy any technology. You can achieve reasonable security by establishing a good security policy, following the guidelines in this and the SAFE security documents, staying up-to-date on the latest developments in the hacker and security communities, and maintaining and monitoring all systems with sound system administration practices. SAFE attempts to build secure IP telephony networks. However, there are no widely adopted protocols for call control or voice transport that support integrated security features such as confidentiality and strong authentication. Until the standards progress and are widely adopted, you must rely on securing the surrounding network and its components.
- Several technologies related to IP telephony are not covered—Though the SAFE white paper contains details on many aspects of IP telephony technologies, the discussion is not exhaustive. The following technologies that relate to IP telephony are not covered:
  - Legacy systems—SAFE focuses on voice security as it relates to IP, not legacy systems.

- Best practices for implementing quality of service (QoS) in IP telephony-enabled networks—A viable IP telephony-enabled network requires implementation of QoS throughout the infrastructure. However, discussion of such implementation is out of the scope of this document. If the network does not support QoS, do not attempt to add IP telephony.
- Use of IPSec virtual private networks (VPNs) to provide secure transport for IP telephony—For this reason, the SAFE VPN-centric remote designs are not covered.
- Distributed call processing—SAFE addresses centralized but not distributed call processing. It is assumed that all remote sites have a redundant link to the headend or local call-processing backup in case of headend failure.
- The interaction between Network Address Translation (NAT) and IP telephony—It is assumed that all networks are private in order to guarantee QoS, in which case non-overlapping address ranges should be used.

IP Telephony Axioms

This topic provides an overview of SAFE IP telephony axioms.

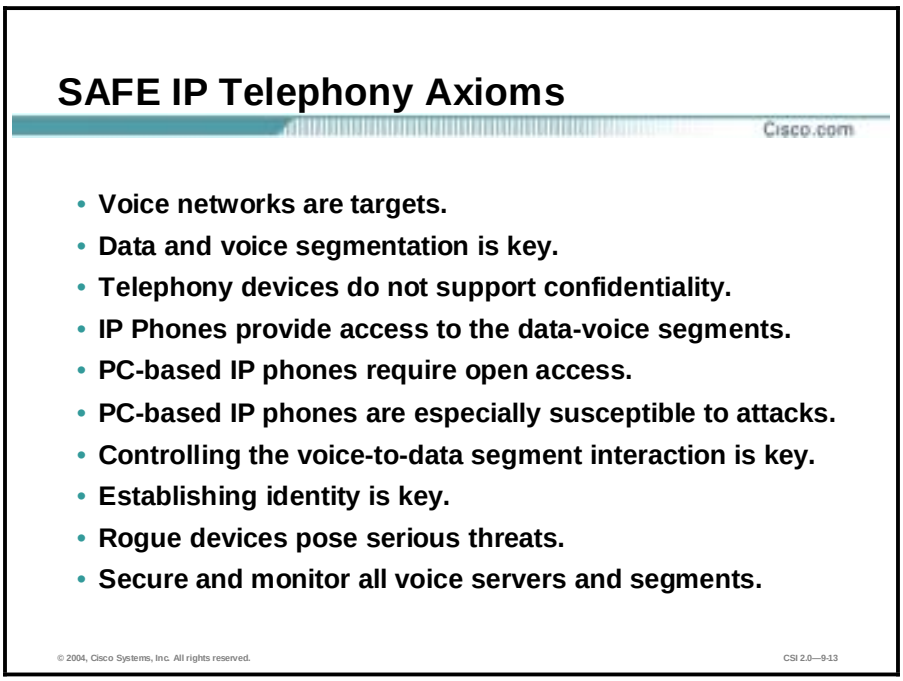

Although IP telephony design differs greatly with the size of enterprises, the underlying best practices remain virtually the same. The following axioms represent important design considerations that affect nearly every design within SAFE IP telephony:

- Voice networks are targets.
- Data and voice segmentation is key.
- Telephony devices don’t support confidentiality.
- IP phones provide access to the data-voice segments.
- PC-based IP phones require open access.
- PC-based IP phones are especially susceptible to attacks.
- Controlling the voice-to-data segment interaction is key.
- Establishing identity is key.
- Rogue devices pose serious threats.
- Secure and monitor all voice servers and segments.

SAFE IP Telephony Axioms—Voice Networks Are Targets

Cisco.com

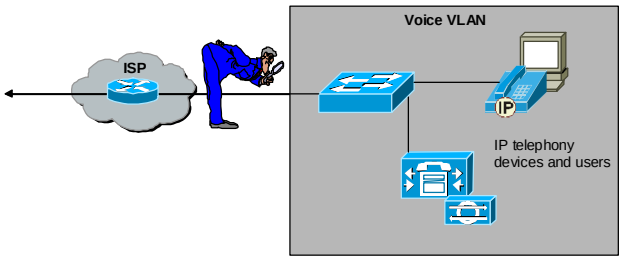

The main issue with voice networks today is that they are generally wide open and require little or no authentication to gain access.

Voice networks are targets. A hacker can eavesdrop on phone conversations between you and your customers, and even forward your calls to your competitors. The main issue with voice networks today is that they are generally wide open and require little or no authentication to gain access. The reason for this is that the model chosen for IP voice networks parallels that chosen for legacy voice systems. It is expected that in the future traditional security features such as strong authentication and encryption will integrate with IP telephony standards. Regardless of the future need for a standards-based and integrated approach to secure IP telephony, there are a number of existing datacentric security technologies you can use today for increased voice security.

SAFE IP Telephony Axioms—Segment Data and Voice Traffic

Cisco.com

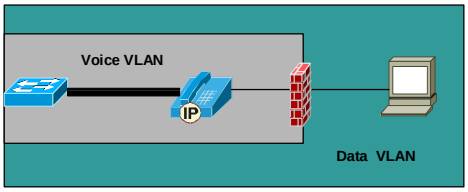

The following technologies provide voice and data segmentation:

- **VLANs**
- **ACLs**
- **Stateful firewalls**

CSI 2.0—9-15

Data and voice segmentation is key. IP-based telephony provides a means of providing telephony over the existing IP data network. However, deployment of IP telephony devices and IP data devices should occur on two logically disparate segments, for reasons including QoS, scalability, manageability, and security. Segmenting IP voice from the traditional IP data network greatly increases attack mitigation capability and allows use of the same access, core, and distribution layers. Although the segments should be disparate, SAFE by no means recommends deploying two IP infrastructures. Technologies such as VLANs, access control lists, and stateful firewalls provide the Layer 3 segmentation necessary to keep the voice and data segments separate at the access layer.

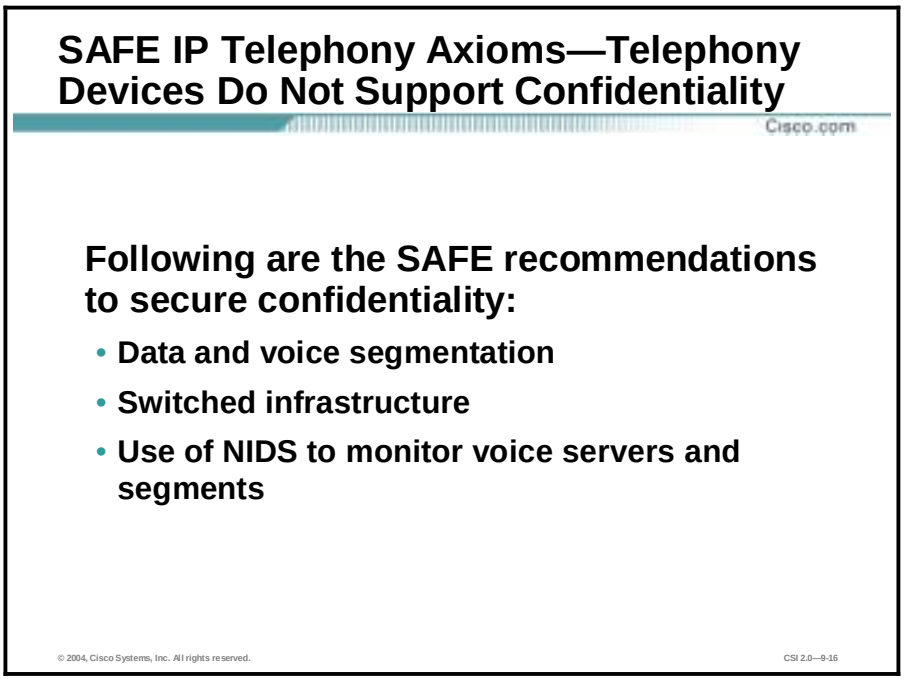

Telephony devices don’t support confidentiality. The combination of data and voice segmentation and a switched infrastructure strongly mitigates call eavesdropping attacks. If someone is able to obtain access to the IP data stream at any point in the network, they can eavesdrop. To a degree, keeping the segments separate prevents devices in the data segment from listening to calls in the voice segment. An obvious way around this segmentation is to unplug an IP phone in the voice segment and plug in another device such as a workstation. Using a switched infrastructure should thwart a device even in the same segment from call monitoring. However, tools such as **dsniff** effectively turn the switched medium into a shared medium. Thus segmentation provides minimal attack mitigation by itself. The true value of segmentation is the ability to tune network-based intrusion detection systems (NIDSs), as outlined in the “Secure and monitor all voice servers and segments” axiom. If a hacker has access to the local switched segment, the hacker might be able to insert a phone into the voice segment with a spoofed MAC address, assume the target phone’s identity, and intercept a call.

Note	dsniff is a collection of tools for network auditing and penetration testing. This collection is freely available and is used at the backbone to allow the monitoring of "Our Network," (ON-2).

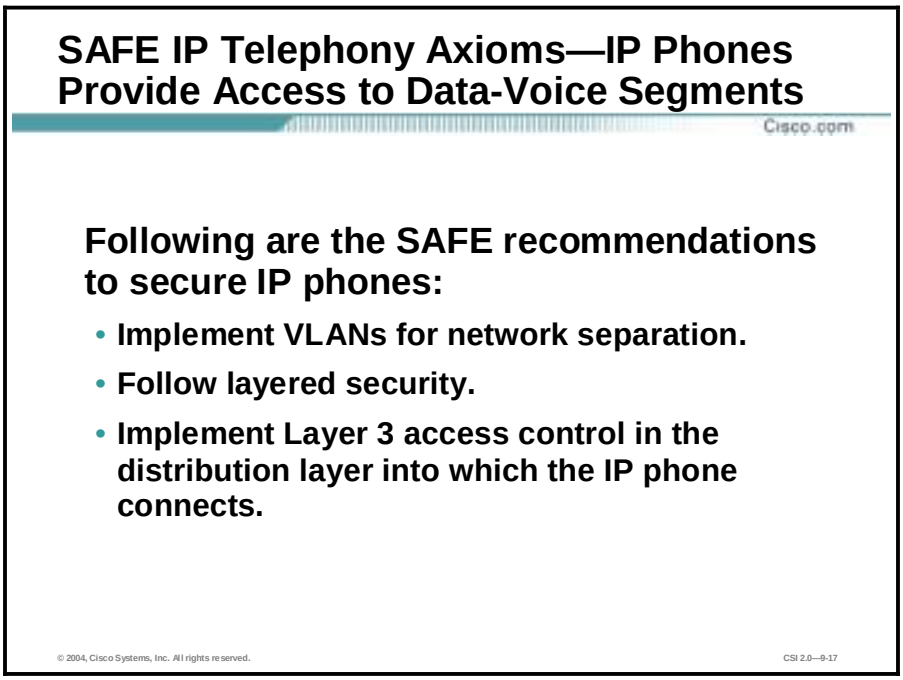

IP phones provide access to the data-voice segments. Many IP phones support a data port to allow the connection of a PC to the phone, so that only a single cable is necessary to provide data and voice connectivity to the user's workspace. When this occurs, follow the data/voice segmentation principle. Some IP phones only provide basic Layer 2 connectivity where the IP phone essentially acts as a hub combining the data and voice segments. Some IP phones provide enhanced Layer 2 connectivity with the option to use VLAN technology, such as 802.1q, to place the phone and the data port in two different VLANs. This architecture assumes that the IP phones support VLANs to keep the data and voice segments separate. Security designs should not rely solely on VLANs for network separation. Rather, follow layered security best practices and also rely on Layer 3 access control in the distribution layer into which the IP phone connects.

SAFE IP Telephony Axioms—PC-Based IP Phones Require Open Access

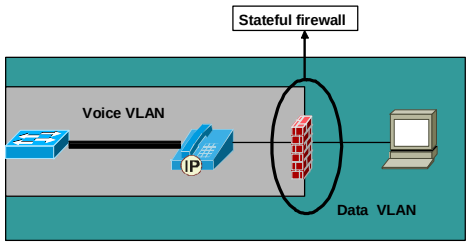

For PC-based IP phones, SAFE recommends deploying a stateful firewall to broker data-voice interaction.

PC-based IP phones require open access. Because the deployment of PC-based IP phones provides a path for attacks against the voice segment, SAFE doesn't recommend their usage unless a stateful firewall brokers the data-voice interaction. PC-based IP phones reside in the data segment and require access to the voice segment in order to access call control, place calls to IP phones, and leave voice messages. Calls placed between IP telephony devices generally use dynamically assigned UDP port numbers greater than 16384, thus requiring a stateful inspection device to allow pinpoint access between the segments. Without a stateful firewall brokering all connections between the data and voice networks, you would have to allow wide UDP port ranges. In most networks it will not be possible to secure all connections between the data and voice segments with a stateful firewall. Consider that in an enterprise, multiple data and voice segments will exist, most likely on the same switch. Here stateful firewall segmentation would not be feasible, nor would Layer 3 stateless filtering suffice. Without a stateful firewall present, a UDP flood DoS attack launched from the data segment could easily overwhelm the voice segment. For this reason, SAFE does not recommend placing IP telephony-capable devices on the data segment unless a stateful firewall is present.

SAFE IP Telephony Axioms—PC-Based IP Phones Are Susceptible to Attacks

Cisco.com

PC-based IP Phones are not as resilient under attack as their IP phone counterparts due to the following reasons:

- **Operating system vulnerabilities**
- **Application vulnerabilities**
- **Service vulnerabilities**
- **Viruses**

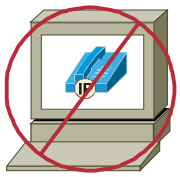

PC-based IP phones are especially susceptible to attacks. PC-based IP phones are not as resilient under attack as their IP phone counterparts. In comparison, PC-based IP phone hosts are more susceptible to attacks due to the number of vectors into the system. These include operating system vulnerabilities, application vulnerabilities, service vulnerabilities, worms, viruses, and so on. IP phones run custom operating systems with limited service support and are less likely to have vulnerabilities. Another issue is that because the PC-based IP phone resides in the data segment, it is susceptible to any attack against that entire segment and not just the host itself. The Code Red and Nimda worms and viruses, for instance, bogged down PC-based IP Phone user systems and the segments they resided in to such a point that they were unusable. No amount of QoS will prioritize voice traffic over data traffic in the data segment if the end system placing the call is unusable.

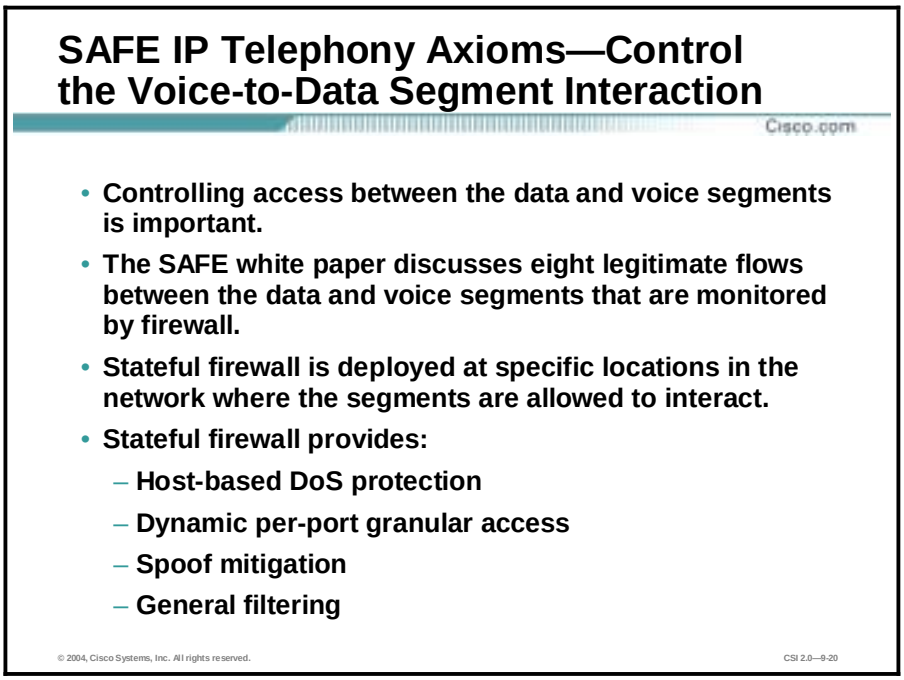

Controlling the voice-to-data segment interaction is key. By methodically controlling access between the data and voice segments, it is possible to deploy a secured IP telephony network. To accomplish this task use a stateful firewall, since it provides host-based DoS protection against connection starvation and fragmentation attacks, dynamic per-port granular access through the firewall when legitimate and necessary, spoof mitigation, and general filtering. SAFE advocates placing a stateful firewall between all segments. SAFE requires a stateful firewall at specific locations in the network where the segments are allowed to interact. Anywhere else in the network, no communication between the segments is allowed and to carry this out stateless Layer 3 filtering is sufficient. There are multiple legitimate flows between the data and voice segments that should be allowed and SAFE discusses each flow's requirements and security issues. In general, the firewall will broker the following connections:

- The voice-mail system, when placed in the data segment connecting to the call-processing manager in the voice segment—Unified voice-mail systems use the traditional e-mail store in the data segment for voice-message storage and require communication with the call-processing manager to notify users of voice mail. Generally this service runs over well-known TCP ports. The firewall mitigates connection starvation, DoS, and protocol attacks against the call-processing manager. The voice-mail system may reside in the data or voice segments, depending on the scalability requirements and the location of the existing e-mail system if unified messaging is deployed.
- IP phones in a voice segment connecting to the call-processing manager in another voice segment for call establishment control and configuration—Generally these services run over a combination of well-known TCP ports and UDP. The firewall mitigates connection starvation, DoS, and protocol attacks against the call-processing manager. It also opens port-level granular access for UDP between the segments. To mitigate attacks, it is recommended that the call-processing manager and IP phones reside in separate voice segments. This is normally done for increased scalability and ease of management.
- IP phones in the voice segment connecting to the voice-mail system when placed in the data segment—Users need to be able to leave voice messages on the voice-mail server not only locally, but also remotely in the case of branch offices. The firewall opens port-level

granular access for UDP between the segments and mitigates general DoS attacks against the IP phones.

- IP phones in the voice segment browsing resources via the proxy server in the voice segment—This might include employee user directories or even Internet access for news. Generally this service runs over well-known TCP ports. The firewall mitigates connection starvation, DoS, and protocol attacks against the proxy server.
- Users in the data segment browsing the call-processing manager in the voice segment—IP phone users will need to be able to modify custom configuration settings of their phone. Generally this service runs over well-known TCP ports. The firewall mitigates connection starvation, DoS, and protocol attacks against the call-processing manager.
- Proxy server in the voice segment accessing resources in the data segment—The server proxies all requests by the IP phone services. Generally these services run over well-known TCP ports. The firewall mitigates connection starvation, DoS, and protocol attacks against the proxy server.

The following two connections will exist only if PC-based IP phones are deployed:

- PC-based IP phone in the data segment accessing the call-processing manager in the voice segment for call establishment. Generally this service runs over well-known TCP ports. The firewall mitigates connection starvation, DoS, and protocol attacks against the call-processing manager.
- PC-based IP phones in the data segment accessing the voice-mail system when placed in the voice segment. This may not be possible in many networks, as discussed previously. Users need to be able to leave voice messages on the voice-mail server, not only locally, but also remotely in the case of branch offices. The firewall opens port-level granular access for UDP between the segments and mitigates general DoS attacks against the voice-mail system.

Note Use private address space (RFC 1918) for all IP telephony devices to reduce the likelihood that traffic could traverse outside of your network. This provides the added benefit that hackers outside of your network will not be able to scan the voice segment for vulnerabilities unless NAT is misconfigured. If possible, use different RFC 1918 address space in the data and voice segments to facilitate filtering and recognition. Although stateful firewalls are used in all of the designs to front-end the call-processing manager, NAT is not in effect for traffic routed within the voice segments. NAT is used between the data and voice segments in all designs to support IP phone services via the proxy server. All Layer 3 devices mitigate voice segment IP address spoofing via filtering as outlined in RFC 2827.

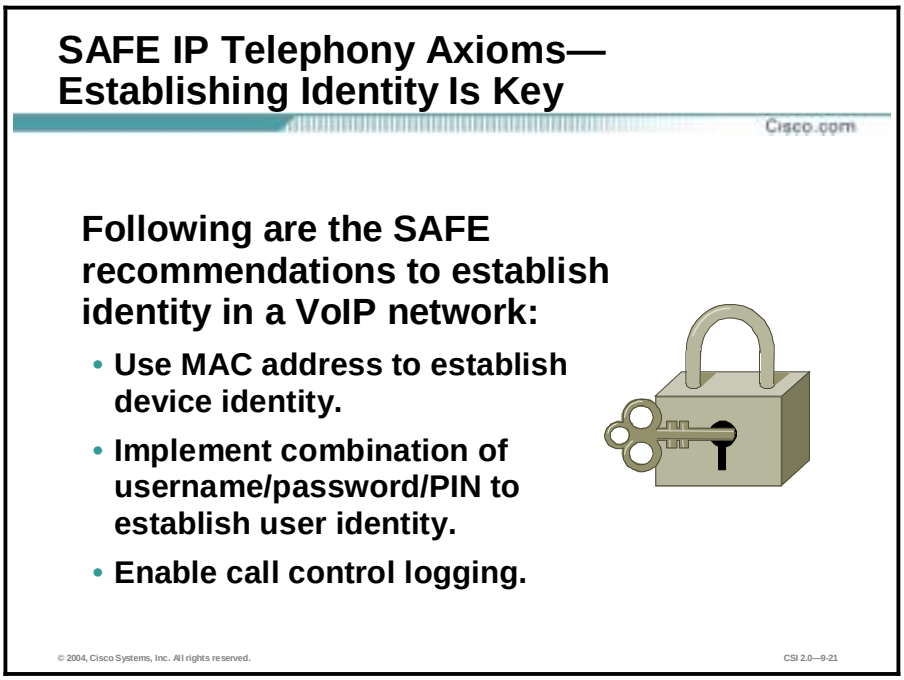

Establishing identity is key. Whenever possible, use user and device authentication, because this mitigates many attacks against the IP telephony network. The SAFE white paper discusses the following recommendations to establish user and device identity in a VoIP network:

- The primary method for the device authentication of IP phones is the MAC address. If a phone with an unknown MAC address attempts to download a network configuration from the call-processing manager and it has no knowledge of the IP phone's MAC address, then that IP phone will not receive a configuration assuming automatic registration has been disabled. This setup prevents someone from placing a rogue phone into the network and then placing a call, unless of course the person spoofs the MAC address in hopes of intercepting calls.
- Some IP phones support basic user authentication, which provides a facility for a user to log in to a phone. By providing either a valid password or personal identification number (PIN), the attacker is granted access to the phone and a custom configuration. User authentication occurs after successful device authentication.
- Some IP telephony systems also support legacy features such as requiring a user to enter an access code before placing calls to restricted locations. Typically these codes are fixed, changed infrequently, and are sent over the IP network in the clear. User authentication mitigates more effectively attacks in which a device spoofs a MAC address and attempts to assume the identity of its target. Requiring user authentication also provides some level of non-repudiation in that if both parties are successfully authenticated, it provides some level of certainty that you can trust the party on the other side of the call.
- Enable call-control logging on the call-processing manager to provide records of placed calls, because this aids in non-repudiation. Some PC-based IP phones provide Windows-based authentication while others use a username/password/PIN combination. In any case, where PINs or passwords are used, they should be aged and changed frequently.
- A username/password/PIN combination may also be used to identify the user to the call-processing manager. This feature allows users to access their custom configuration settings after successfully authenticating and is recommended. Some voice-mail systems support

two-factor authentication. In this scenario, users must undergo strong authentication in order to change their custom settings (for example, greeting message) or to listen to voice mail. SAFE does not recommend deploying this feature unless the sensitivity of the voice messages mandates it.

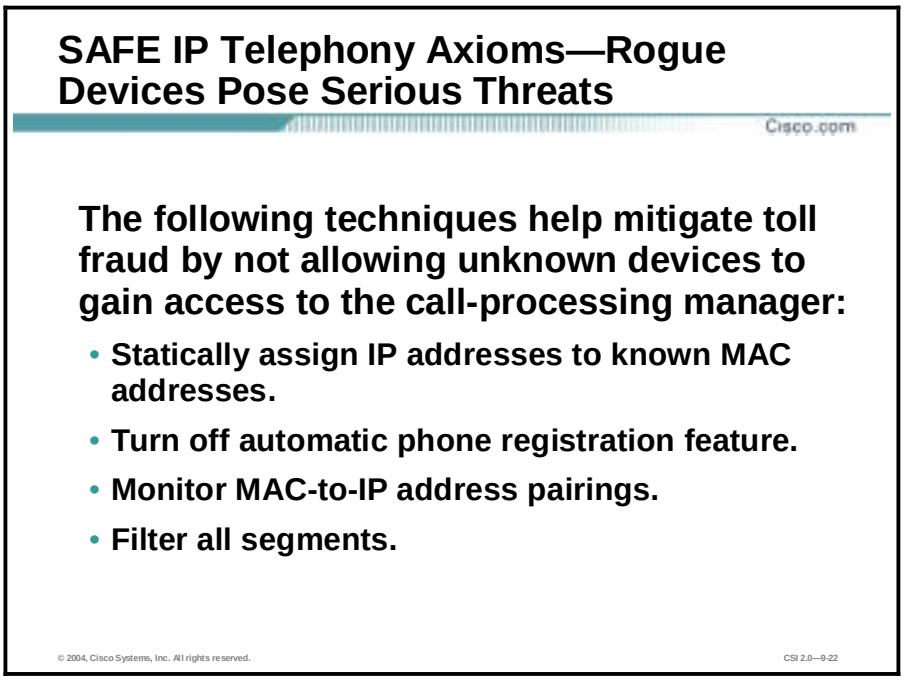

Rogue devices pose serious threats. Locking down the switched ports, segments, and services in the network will provide attack mitigation for rogue devices. As in any IP network, it is important to mitigate the capabilities of a rogue device plugged into the network. As the original SAFE papers addressed, mitigation methods include best practices such as disabling unused ports and deploying a switched environment. All of these best practices can also be used in the voice segment, but some additional steps should be taken. The following four best practices provide mitigation details specific to IP telephony:

- Statically assigning IP addresses to known MAC addresses—Since Dynamic Host Configuration Protocol (DHCP) is typically used for a scalable IP phone deployment, consider statically assigning IP addresses to known MAC addresses. This way, the IP phone always has the same address, and if an unknown device is plugged into the network, it does not receive an address. Although a hacker could still statically assign both an IP and MAC address on a device to subvert this practice, doing so is not easy. In case automatic registration is left enabled on the call-processing manager, with this practice in place it is less likely that a device could register on the network. By using separate DHCP servers for the voice and data segments, a DoS attack launched against the DHCP server in the data segment would not interfere with IP phone address allocation in the voice segment. Fixing IP address allocation to known MAC addresses greatly reduces the likelihood of a successful IP address DoS starvation attack against the voice DHCP server.
- Turning off the automatic phone registration feature—Many call-processing managers provide an automatic phone registration feature that bootstraps an unknown phone with a temporary configuration and then allows it to interact with the network. Turn this functionality off for normal day-to-day use. You should only use it temporarily for bulk deployment of phones. Configure the call-processing manager to deny unknown PC-based IP phone access. This will provide mitigation for unknown devices that should not be allowed to use the call-processing manager to register on the network.
- Monitoring MAC-to-IP address pairings—Consider using a utility such as Arpwatch to monitor the MAC addresses in your voice segment. In comparison to the data segment, MAC addresses are more likely to be static and Arpwatch will track the MAC addresses of

all devices in the voice segments. Arpwatch logs any changes in MAC-to-IP address pairings. For more information on Arpwatch, refer to http://www-nrg.ee.lbl.gov/nrg.html.

- Filtering all segments—This should limit devices in unknown segments from connecting to the call-processing manager. If a rogue device is placed in a segment not approved for IP telephony use, filtering should prevent the device from registering on the network via the call-processing manager. If a rogue call-processing manager is placed in the network, filtering should prevent the redirection of IP telephony devices to it. If a rogue voice gateway is placed in the network, filtering should not allow it to connect to the call-processing manager.

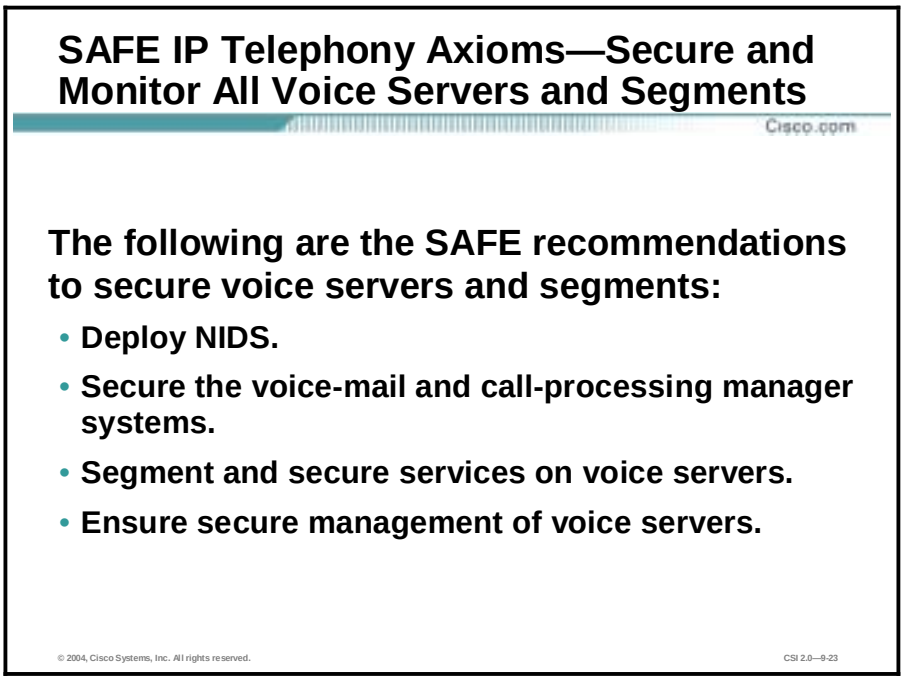

Secure and Monitor All Voice Servers and Segments. It is no surprise that the same attacks that have the potential to cripple key production servers in the data segment can also have the same effect on the voice servers in the voice segment. For this reason, many of the same precautions should be taken. The following are SAFE recommendations for securing and monitoring voice servers and segments:

- Deploy NIDS—NIDS is a powerful tool that should be used to trigger an alarm, log, and in some cases react to attack signatures detected in the network. As of now, NIDS does not provide voice control protocol attack signatures. It does, however, provide signatures for UDP DoS attacks and HTTP exploits that are applicable in voice environments. The SAFE white paper has the following suggestions for deploying NIDS in VoIP networks:
  - Deploy NIDS in front of the call-processing manager to detect attacks sourced from the data segment against its HTTP user service.
  - Deploy NIDS between the voice and data segments in order to detect DoS attacks against the voice segment.
  - Deploy NIDS in switches, which increases the effectiveness of the tool. An NIDS provides the most value when it is tuned to the environment in which it is deployed. Depending on the voice control protocols used in your environment, the tuning characteristics will be different. In general, the number of traffic flows in the voice segment is much more easily classified than the number in the data segment.
  - In controlling the voice-to-data segment interaction axiom, a total of eight flows were identified. The benefit of such a small set of traffic flows is twofold. First, any traffic other than those flows should trigger a severe NIDS alarm. This provides the defense-in-depth approach to attack mitigation. Second, given the small number of flows, filtering out false positives is relatively straightforward. In comparison, this is far different than in the data segment where multiple flows exist and false positives are commonplace. Since attack signature detection in the voice segment is most likely not a false positive, you should consider enabling NIDS reactive features.

- Once an attack signature is detected on a segment, shunning can be used to dynamically change the Layer 3 filtering configuration of a network device to drop all additional traffic from the source on that segment. Resets can be used to tear down a TCP session that triggers an attack signature. For example, NIDS resets will provide attack mitigation for attacks against the call-processing manager web and IP phone services.
- NIDS shunning could be used to block UDP flood attacks sourced from the data segment against the voice segment, but care should be taken to not block legitimate usage.

- Secure the voice-mail and call-processing manager systems—These security precautions include: turning off all unneeded services, patching the operating system and services with the latest security patches, hardening the operating system configuration, disabling any features on the voice servers that are not in use, and, finally, not running any unnecessary applications on the server (for example, an e-mail client). Doing all these tasks reduces the number of vectors into the system that an attacker might use. Installing either HIDS or HIPS is also recommended. Given the high target value of a voice server and the lag time involved in validating an application or operating system security patch for production use, either HIDS or HIPS will provide significant and immediate attack mitigation. Some call-processing managers do not support either HIDS or HIPS. Installation of either HIDS or HIPS on the mail server in the data segment should also occur if this system is being used as the voice-mail store in addition to e-mail content filtering or virus scanning. These recommendations are consistent with the SAFE security papers which recommend either HIDS or HIPS on all critical servers.
- Segment and secure services on voice servers—Voice servers may run multiple services that can be distributed across multiple devices in order to increase scalability and manageability. You should also use this feature to increase the level of security. For instance, call-processing managers typically support call control, web configuration, IP phone browsing services, conference calling, and device configuration services. Most configuration services do not support strong authentication. By segmenting these services the number of entry points into a system is reduced. Also make sure that the services use user or service accounts with only the privileges absolutely necessary to run normally. A compromised service shouldn't provide root or administrative access.
- Secure management of voice servers—Voice servers also support a variety of methods for management. These include protocols such as HTTP, Secure Sockets Layer (SSL), and Simple Network Management Protocol (SNMP). Follow the guidelines provided in the original SAFE paper regarding management operations.

Cisco IP Telephony Product Portfolio

This topic describes Cisco's IP telephony product line.

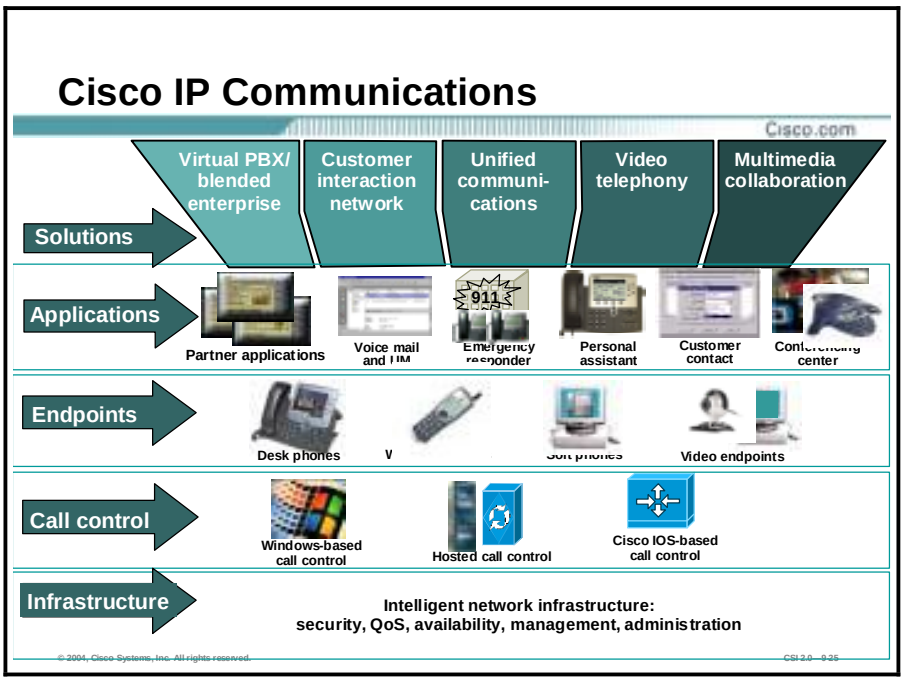

Cisco IP Communications is a comprehensive system of enterprise-class solutions—including IP telephony, unified communications, IP video/audio conferencing, and contact center—that facilitate more engaging and efficient interactions among employees, partners, and customers, and provide the foundation for a collaborative workforce. Enabled by Cisco Architecture for Voice, Video and Integrated Data (AVVID), IP Communications solutions dramatically improve operational efficiencies, increase organizational productivity, and enhance customer satisfaction to create an empowered, effective work environment. By promoting greater levels of workforce collaboration, Cisco IP Communications solutions help enterprises exceed customer expectations, outpace the competition, and realize a measurable return on their investments.

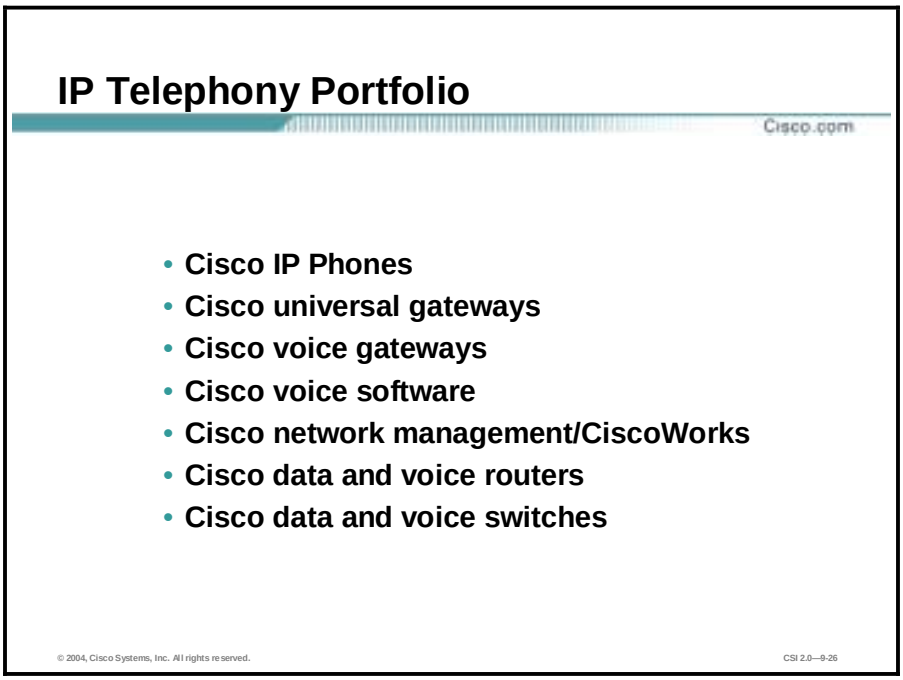

The Cisco IP Communications system includes the following primary components:

- Cisco IP Phones
- Cisco universal gateways
- Cisco voice gateways
- Cisco voice software
- Cisco network management/CiscoWorks
- Cisco data and voice routers
- Cisco data and voice switches

Cisco IP Phones are next-generation intelligent communication devices, delivering essential business communications at a touch. Fully programmable, the growing family of Cisco IP Phones provides the most frequently used business features. Cisco IP Phones as full-featured telephones can plug directly into your IP network. H.323 clients and computer telephony integration (CTI) ports comprise software-based devices that you configure similarly to the Cisco IP Phones. The Cisco CallManager allows you to configure phone features such as call forwarding and call waiting for your phone devices. You can also create phone button templates to assign a common button configuration to a large number of phones. Once you have added the phones, you can associate users with them. By associating a user with a phone, you give that user control over that device.

The newest Cisco IP Phones have the following enhancements:

- Display-based
- Straightforward user customization based on changing needs
- In-line power accepted from integrated Catalyst switch card or the Catalyst in-line power patch panel
- Two-port 10/100BASE-T Ethernet switch interface to ensure QoS
- Each Cisco IP Phone provides toll-quality audio and doesn't require a companion PC. Because it is an IP-based phone, it can be installed anywhere on a corporate, local, or wide-area IP network.

Cisco IP Phone Series are:

- Cisco 7900 Series IP Phones

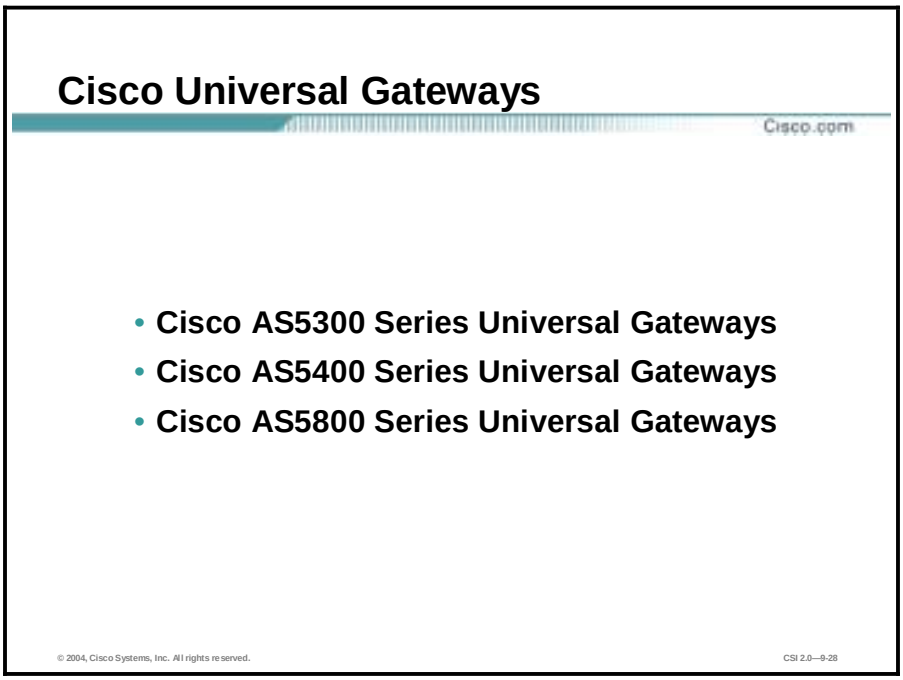

The Cisco Universal Gateways support a wide range of voice and data services and multiple architectures simultaneously allowing service providers to generate new revenue streams and quick return on investment. Following universal gateway models are available:

- Cisco AS5300 Series Universal Gateway—Provides reliable, scalable data and voice gateway functions, and supports PSTN signaling, gateway signaling, voice codecs, fax, VoiceXML, Remote Access Dial-In User Service (RADIUS), Tool Command Language (Tcl), and interactive voice response. Ideal for Tier 2 and Tier 3 ISPs, small points of presence, and enterprises, the AS5300 Series Universal Gateway also supports service provider data and voice applications, including:
  - Long distance
  - Prepaid calling
  - Local access
  - Hosted IP telephony
  - ASP hosting and termination
  - Unified communications
  - Access VPN
  - Dial access

Note	The AS5300 Series includes the Cisco AS5300 Access Server/Voice Gateway, an award-winning dialup remote access server and VoIP gateway, and the Cisco AS5350 Universal Gateway, a universal port-ready, one-rack unit, dual T1/E1 gateway that provides carrier-class reliability in a modular design.

- Cisco AS5400 Series Universal Gateway—A reliable, scalable data and voice gateway in a two-rack unit. It supports PSTN signaling, gateway signaling, voice codecs, fax,

VoiceXML, RADIUS, Tcl, and interactive voice response. Service provider data and voice applications are also supported, including:

- Long distance
- Prepaid calling
- Local access
- Hosted IP telephony
- ASP hosting and termination
- Unified communications
- Access VPN
- Dial access

Note The Cisco AS5400 Series includes the Cisco AS5400HPX Universal Gateway, which provides enhanced performance for processor-intensive voice and fax applications.

- Cisco AS5800 Series Universal Gateway—Provides reliable, scalable data and voice gateway functions for large service providers. The Cisco AS5800 Series Universal Gateway supports public switched telephone network signaling, gateway signaling, voice codecs, fax, RADIUS, Tool Command Language, and interactive voice response. Service provider data and voice applications are also supported, including:
  - Long distance
  - Prepaid calling
  - Local access
  - Hosted IP telephony
  - Application service provider (ASP) hosting and termination
  - Unified communications
  - Access VPN
  - Dial access

Note The Cisco AS5800 Series includes the Cisco AS5800 Access Server/Voice Gateway, which helps with deploying data or toll-quality voice and fax services over packet networks, and the Cisco AS5850 Universal Gateway, a high-density, carrier-class gateway supporting up to 5 x CT3, 96 T1, or 86 E1 of data, voice, and fax services simultaneously.

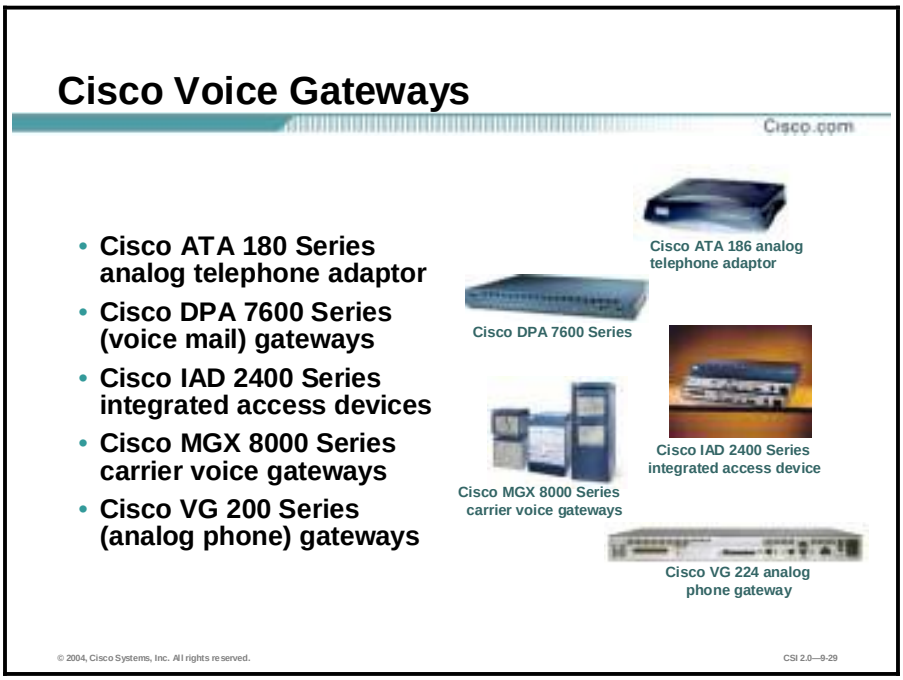

Cisco voice gateway solutions include:

- Cisco ATA 180 Series analog telephone adaptor—A handset-to-Ethernet adaptor that turns traditional telephone devices into IP devices. The Cisco ATA 186 supports two voice ports, each with its own independent telephone number, and a single 10BASE-T Ethernet port. This adaptor can make use of existing Ethernet LANs, in addition to broadband pipes such as DSL, fixed wireless, and cable modem deployments.
- Cisco DPA 7600 Series of voice mail gateways—A family of VoIP gateways that enables legacy voice-mail equipment to connect to a Cisco IP telephony solution network.
- Cisco IAD 2400 Series—Cisco's premier line of Smart Integrated Access Devices. The IAD 2400 Series offers the best-in-class integration of data and toll-quality analog or digital voice services for a customer premises solution. Designed to bridge the current and emerging multiservice needs of competitive local exchange carriers (CLECs) and other service providers, the Cisco IAD 2400 Series operates in both Class 5 switch-access and Class 5 switch-bypass architectures.
- Cisco MGX 8000 Series carrier voice gateways—High capacity carrier-class voice gateways that offer standards-based support for VoIP and voice over ATM (VoATM) services. These gateways include the Voice Interworking Service Module (VISM) deployed in the Cisco MGX 8230, 8830, 8250, and 8850 advanced ATM multiservice switches. The Cisco MGX 8000 Series carrier voice gateways combine the industry's highest quality packet voice with the extensible architecture and the proven network availability of the Cisco MGX 8000 Advanced ATM Multiservice Portfolio (AAMP).
- Cisco VG 200 Series analog phone gateway—The Cisco VG 224 is the latest model in the VG 200 Series analog phone gateways. The Cisco VG 224 melds the high density RJ-21 analog interface with IOS manageability to deliver a cost-effective platform to leverage existing analog phone equipment with Cisco AVVID. The Cisco VG 224 is a high-density twenty-four-port gateway for analog phones, fax machines, modems, and speaker phones. It enhances enterprise voice system architecture with Cisco CallManager or Cisco

CallManager Express on a Full Service Branch (FSB) router in a very compact 19-inch rack-mounted chassis.

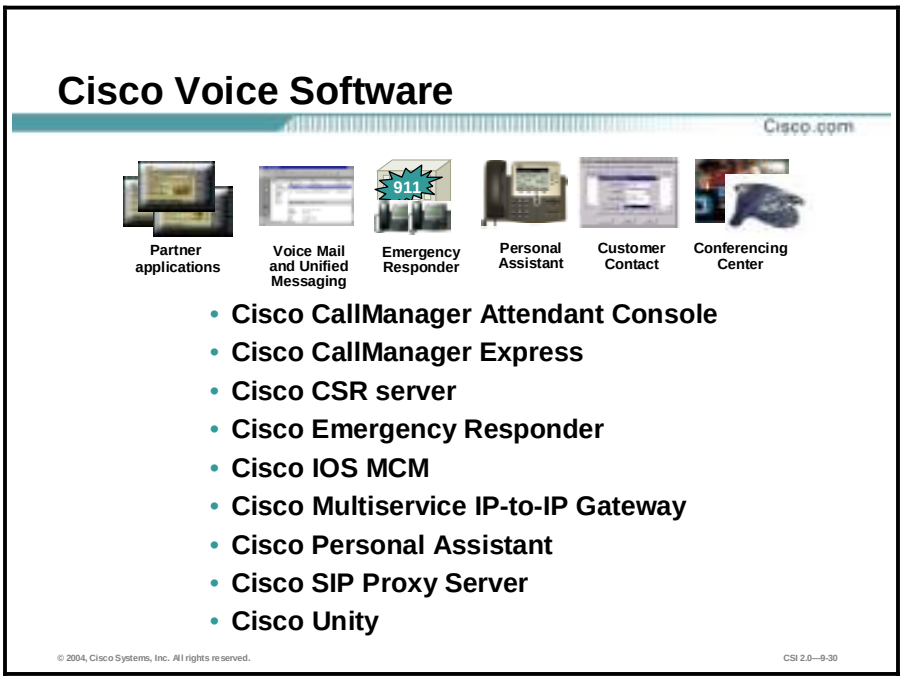

- Cisco CallManager Attendant Console—An affordable and scalable IP-based solution that replaces the traditional private branch exchange (PBX) manual attendant console. Associated with a Cisco IP Phone, the Cisco CallManager Attendant Console allows the attendant to quickly accept and dispatch calls to enterprise users. An integrated directory service provides traditional busy lamp field (BLF) and direct station select (DSS) functions for any line in the system. One of the primary benefits over traditional attendant console systems is the Cisco CallManager Attendant Console's ability to monitor the state of every line in the system and to efficiently dispatch calls. The absence of a hardware-based line monitor device offers a much more affordable and distributable manual attendant solution than traditional consoles.
- Cisco CallManager Express—A solution embedded in Cisco IOS software that provides call processing for Cisco IP Phones. This solution enables the large portfolio of Cisco access routers to deliver a robust set of features commonly used by business users, thereby enabling deployment of a cost-effective, highly reliable, IP Communications solution for the small office. Customers can now scale IP telephony to a small site or branch office with a solution that is very simple to deploy, administer, and maintain. The Cisco CallManager Express solution provides customers with a low-cost, reliable, feature-rich solution for a deployment of up to 100 users.
- The Cisco Carrier Sensitive Route (CSR) server—A software application that provides advanced, rules-based routing logic within a Cisco H.323 VoIP network. Targeted at the highly competitive wholesale voice marketplace, the Cisco CSR server allows service providers to optimize the routing decisions within their network to fit their business needs.
- Cisco Emergency Responder—Dynamically addresses the need to identify the location of 911 callers in an emergency, with no administration required when phones or people move from one location to another. This capability enhances the existing E9-1-1 functionality of Cisco CallManager, with a real-time location-tracking database and enhanced routing capabilities that direct emergency calls to the appropriate Public Safety Answering Point (PSAP) based on the location of the caller. Coupled with Cisco CallManager, Cisco Emergency Responder surpasses traditional PBX capabilities by introducing zero-cost user

or phone moves and changes, and dynamic tracking of user and phone locations for E9-1-1 safety and security purposes.

- Cisco Systems IOS Multimedia Conference Manager (MCM)—Provides network administrators with a mechanism to support H.323 applications without impacting mission-critical applications that run on today's networks. Cisco MCM is implemented on Cisco IOS and provides a network manager with the ability to limit the H.323 traffic on the LAN and WAN, provide user accounting for records based on the service utilization, inject QoS for the H.323 traffic generated by applications such as VoIP, data conferencing, and video conferencing, and provides the mechanism to implement security for H.323 communications. The MCM consists of an H.323 gatekeeper and an H.323 proxy.
- Cisco Multiservice IP-to-IP Gateway—A special IOS software load that runs on the Cisco 2600XM, 3660, 3725, and 3745 series of multiservice gateway platforms. It is designed to meet Enterprise and Internet Telephony Service Provider (ITSP) VoIP-to-VoIP interconnection needs. One set of images is available for VoIP-to-VoIP interconnection; a second set is available for interconnection via an Open Settlement Protocol (OSP) provider. The Multiservice IP-to-IP Gateway facilitates easy and cost-effective connectivity between independent VoIP service provider networks. The IP-to-IP Gateway provides a network-to-network interface point for billing, security, Cisco CallManager interconnectivity, call admission control, and signaling interworking. It will perform most of the same functions of a PSTN-to-IP gateway, but will join two VoIP call legs. Media packets can either flow through the gateway and hide the networks from each other, or flow around the IP-to-IP gateway if network security is not of primary importance.
- Cisco Personal Assistant—A telephony application that streamlines communications by helping users manage how and where they want to be reached. In addition to providing dedicated, customizable administrative assistance, Cisco Personal Assistant offers a variety of interfaces that are easy to use and modify, without special training or system administrator assistance. Cisco Personal Assistant gives speech-enabled access to your Cisco Unity voice messages, the corporate directory, conference calling features, and personal contact lists from any telephone. Its web-based and telephone user administration interfaces make it easy to forward or screen calls in advance or in real time. IP phone productivity services give you the option to check your calendar, personal contact information, e-mail, and voice messages using the pixel-based liquid crystal display (LCD) and interactive soft keys on the Cisco 7940 or 7960 model IP phones.
- Cisco SIP Proxy Server—A call-control software package that enables service providers to build scalable, reliable VoIP networks. The Cisco SIP Proxy provides a full array of call routing capabilities to maximize network performance in both small and large packet voice networks.
- Cisco Unity—A powerful unified communications solution that provides advanced, convergence-based communication services on a platform that offers the utmost in reliability, scalability, and performance. Cisco Unity integrates with the desktop applications to improve communications, boost productivity, and enhance customer service capabilities across your organization. With Cisco Unity, you can listen to your e-mail over the telephone, check voice messages from the Internet, and (when integrated with a supported third-party fax server) forward faxes to any local fax machine.

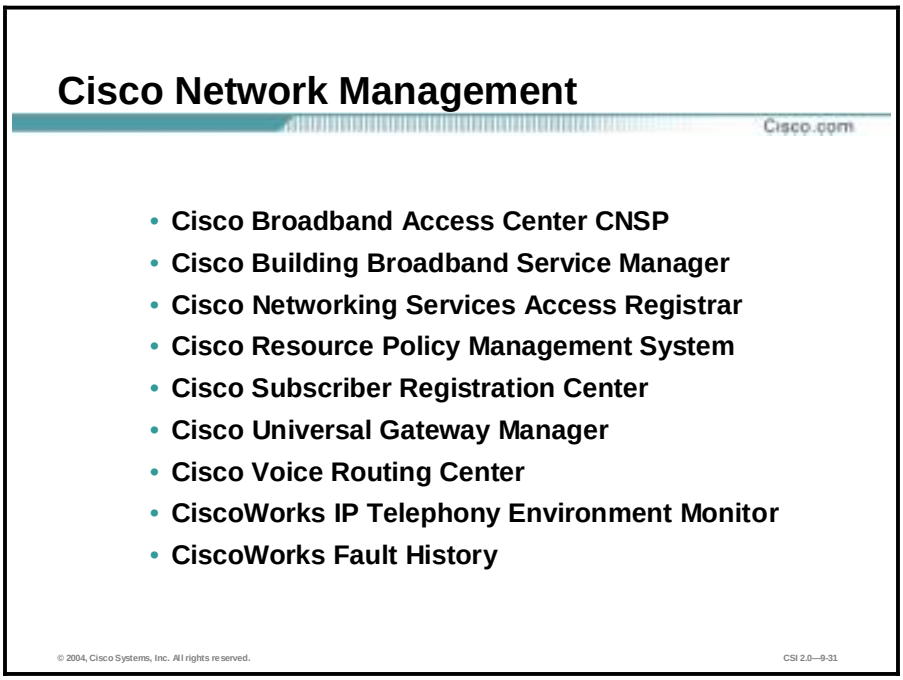

The following are Cisco network management solutions that support data and voice deployments:

- Cisco Broadband Access Center CNSP—Cisco Convergent Network Solution Center: Subscriber Provisioning CCNSC (SP) helps you provision subscriber Customer Premises Equipment (CPE) and devices (for example: 3660, 7200, and ESR10K) through a GUI. It allows a service provider to automate the provisioning process of T1 CPE and gateways. In addition, it allows a service provider to manage devices that have been externally provisioned, thus providing a single, consistent management system for T1 CPEs and 3660, 7200, and ESR10K gateways, regardless of the origins of their provisioning.
- Cisco Building Broadband Service Manager—BBSM (v5.x) is an access gateway for public-access networks that enables simple, plug-and-play access, end-user self-provisioning of services, multiple authentication and billing options, and web-based management, reporting and configuration. Cisco BBSM works with Cisco access-layer LAN products to provide a complete solution that enables service providers to create a market and operate broadband access services in vertical markets such as hospitality, public spaces, higher education, guest access to the enterprise businesses, and much more.
- Cisco Networking Services Access Registrar—A RADIUS-compliant, access policy server designed to support the delivery of dial, ISDN, and new services including DSL, cable with telco-return, wireless, and VoIP. Designed to meet the needs of service provider operations, Cisco Networking Services Access Registrar provides carrier-class performance and scalability as well as the extensibility required for integration with evolving service management systems.
- Cisco Resource Policy Management System—The Cisco Resource Policy Management System (RPMS) is a software tool that provides policy management of platform and gateway resources. RPMS 2.0 provides many enhancements including RADIUS-based port policy management, distributed architecture for greater scalability, reliability enhancements, and port reservation guarantees with oversubscription.

- Cisco Subscriber Registration Center—SRC Device Provisioning Registrar (DPR) 2.0 makes it easier than ever for service providers to deploy high-speed data and VoIP services over Data-over-Cable Service Interface Specification (DOCSIS) cable modems, Digital Video Broadcasting/Digital Audio and Visual Council (DVB/DAVIC) digital set-top-boxes (DSTBs), and fixed wireless devices. Cisco SRC DPR builds intelligence on top of the Cisco Network Registrar Protocol servers to allow service providers to automate the subscriber-provisioning process. Performance, scalability, and reliability were the design requirements behind this second-generation Cisco SRC product. In addition, Cisco SRC DPR includes a Java-based provisioning application-programming interface (API), ensuring quick and seamless integration with customers' existing and next-generation operations support systems (OSSs).
- Cisco Universal Gateway Manager—Network management applications and tools are critical for the successful deployment and operations of voice or data services. The Cisco Universal Gateway Manager (UGM) is an element management system for Cisco AS5000 Series Universal Gateways. The Cisco UGM enables network operators and administrators to efficiently deploy, manage, and maintain Cisco AS5000 Series Universal Gateways that support VoIP, managed voice, PSTN gateway, and dial-access services.
- Cisco Voice Routing Center—VRC is a graphical tool for designing, configuring, and managing dial plans in Cisco VoIP networks. The Cisco VRC centralizes and coordinates call routing configuration of up to 3000 Cisco H.323 VoIP gateways and gatekeepers in a single application. The Cisco VRC aids in building and managing any Cisco H.323 VoIP network, implemented for toll bypass between PBXs at different offices in an enterprise private voice network, or for national or international long distance in the Cisco Voice Infrastructure and Applications (VIA) solution.
- CiscoWorks IP Telephony Environment Monitor—ITEM is a powerful suite of applications and tools that continuously evaluates and reports the operational health of your Cisco IP telephony implementation. CiscoWorks ITEM is used to manage Cisco AVVID and Cisco IOS software-based IP telephony environments. It provides specialized operations and security tools beneficial to large and small IP telephony implementations.
- CiscoWorks Fault History—A browser-based management utility that provides detailed information about alerts and faults that are detected and processed by CiscoWorks Device Fault Manager (DFM). DFM is part of the CiscoWorks LAN Management Solution (LMS). CiscoWorks Fault History is intended for operations, network administration, and help-desk personnel who need access to historical information about the devices in their IP network.

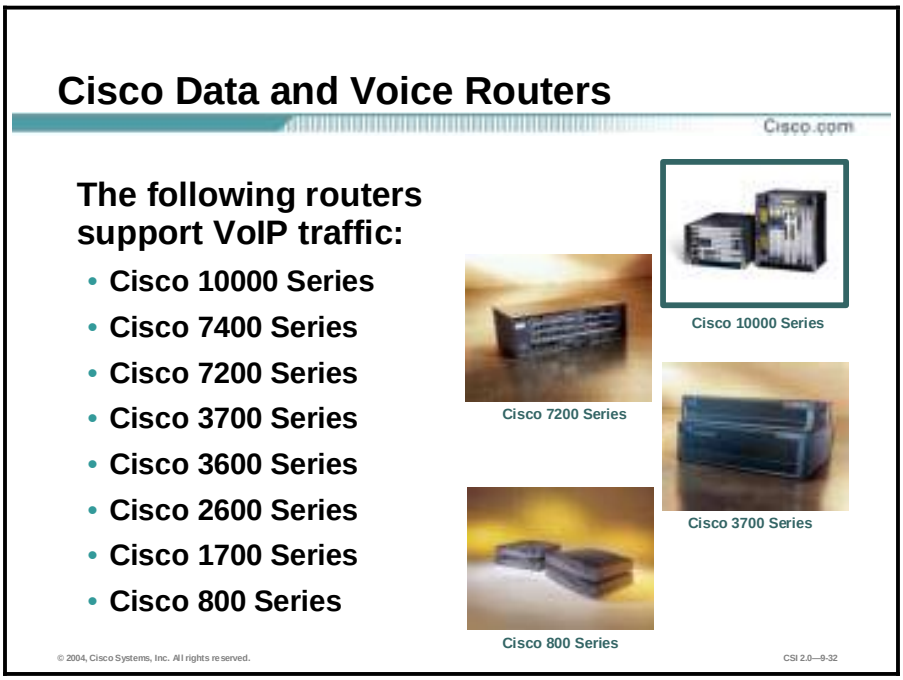

The following routers support data and voice traffic:

- Cisco 10000 Series
- Cisco 7400 Series
- Cisco 7200 Series
- Cisco 3700 Series
- Cisco 3600 Series
- Cisco 2600 Series
- Cisco 1700 Series
- Cisco 800 Series

Cisco Data and Voice Switches

Cisco.com

The following switches support data and voice traffic:

- **Catalyst 3500 Series XL switches**
- **Catalyst 4000 Series XL switches**
- **Catalyst 4500 Series XL switches**
- **Cisco MGX 8800 Series switches**

Catalyst 3500 Series

Catalyst 4000 Series

Catalyst 4500 Series

Cisco MGX 8800 Series

The following switches support data and voice traffic:

- Catalyst 3500 Series XL switches
- Catalyst 4000 Series XL switches
- Catalyst 4500 Series XL switches
- Cisco MGX 8800 Series switches

SAFE IP Telephony Design Considerations

This topic describes the SAFE IP telephony design considerations.

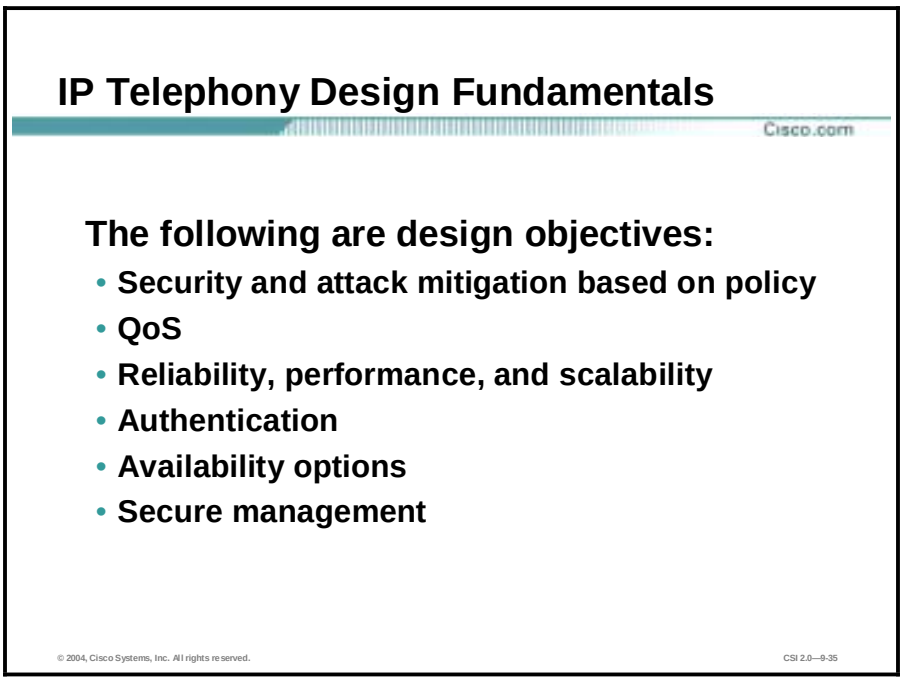

SAFE IP telephony emulates as closely as possible the functional requirements of today's networks. Implementation decisions vary, depending on the network functionality required. The following design objectives, listed in order of priority, guide the decision-making process.

- Security and attack mitigation based on policy
- Quality of service
- Reliability, performance, and scalability
- Authentication of users and devices
- Options for high availability (some designs)
- Secure management

First and foremost, SAFE IP telephony must provide ubiquitous IP telephony services to the locations and users that require it. It must maintain as many of the characteristics of traditional telephony as possible while doing so in a secure manner. Finally, it must integrate with existing network designs based on the SAFE security architecture and not interfere with existing functions.

Note The term "call-processing manager" is used throughout SAFE. Some readers may be more familiar with the term "IP PBX." The term "voice-mail system" is used to refer to an IP-based voice-mail storage device.

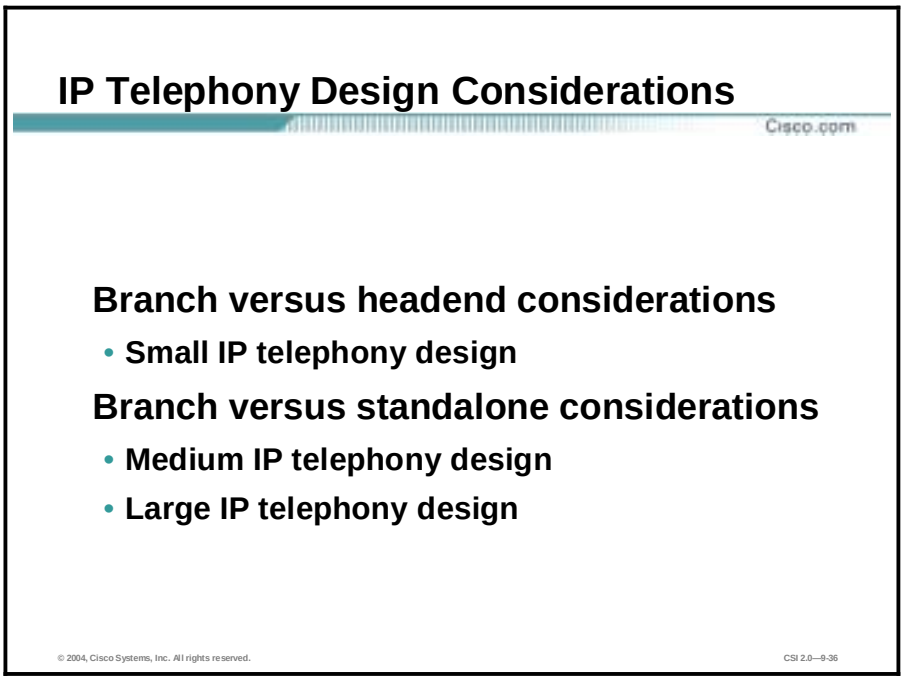

Small and medium designs can be used in the following two possible configurations:

- The design acts like a branch of a larger organization, built in the configuration described in SAFE Enterprise.
- The design is the headend of an organization's network. This headend may have connections to other offices of the same organization. For example, a large law office may use the medium network design for its headend, and several small network designs for its other locations.

Small Network IP Telephony Design

This topic describes small network IP telephony design.

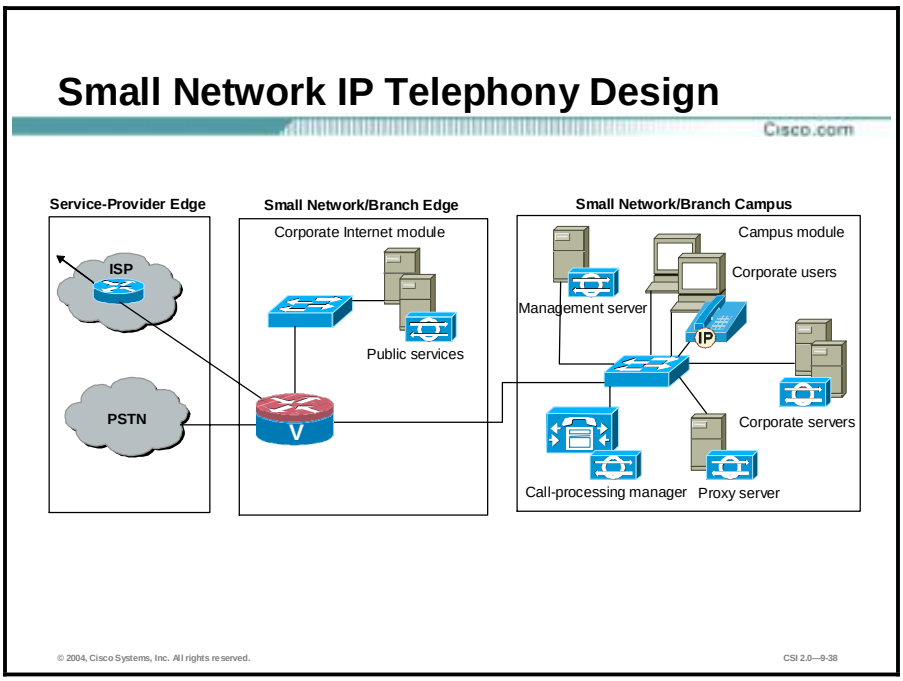

The small network IP telephony design uses the small network design from the SAFE security papers. The corporate Internet module has been modified to support voice services, including PSTN access for WAN backup and local calls, and VLANs for data and voice segmentation. The campus has been modified to support IP phones, PC-based IP Phones, proxy services, and VLANs. Most of the discussion of this design is based on this design operating as the headend for a corporation. Specific design changes when used as a branch are also included.

Small Network Corporate Internet Module—Key Device

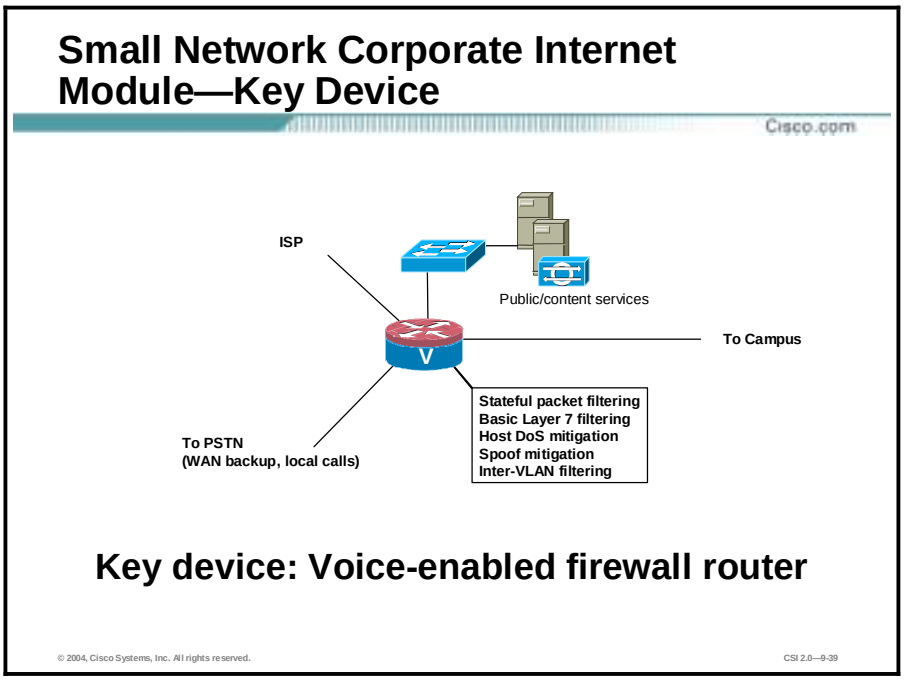

- The corporate Internet module provides internal users with connectivity to Internet services, Internet users access to information on public servers, and segmentation between the data and voice segments (located in the campus module).
- The key IP telephony device is a voice-enabled firewall router. It provides network-level protection of resources, stateful filtering of traffic, and voice services.

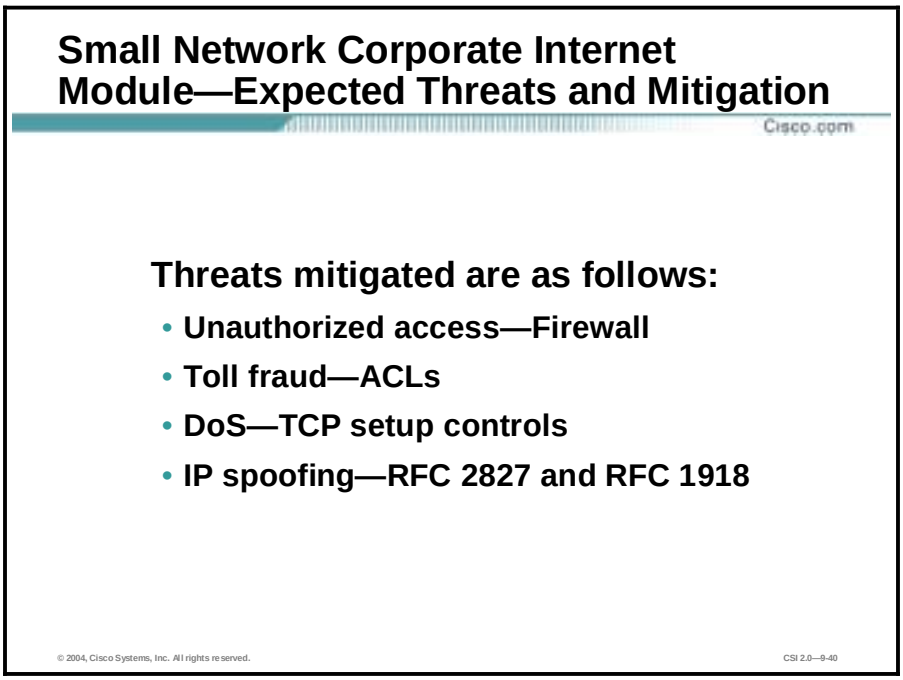

The following voice threats are expected and mitigated for the small network corporate Internet module:

- Unauthorized access—This type of access is mitigated through filtering at the firewall.
- Toll fraud—ACLs permit only known telephony devices to communicate with each other.
- Denial of service—TCP setup controls limit exposure of the call-processing manager.
- IP spoofing—RFC 2827 and RFC 1918 filters are placed at the ISP edge and the local firewall router.

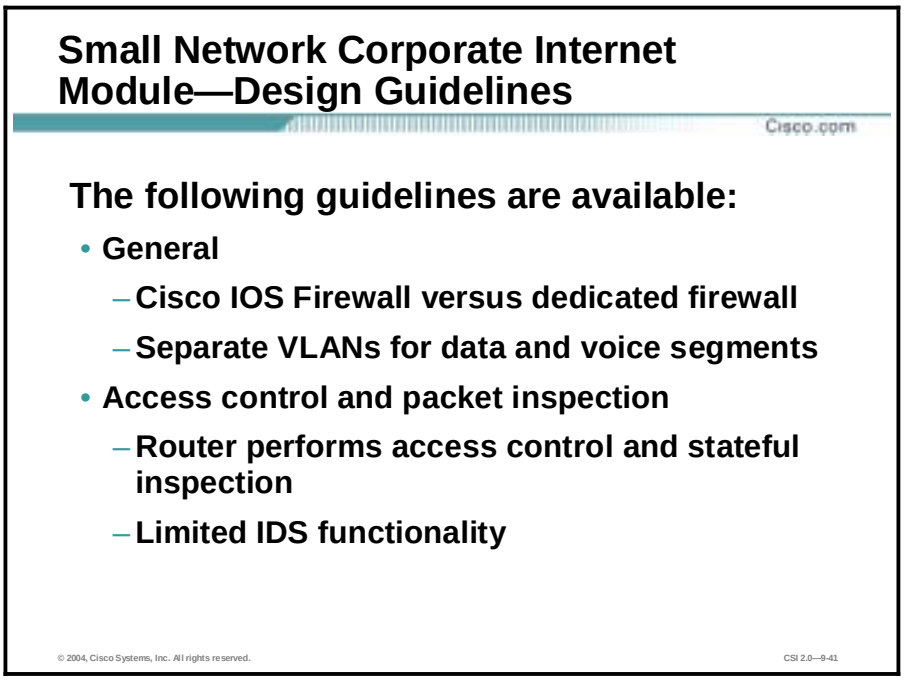

The following design guidelines are available for the small network corporate Internet module:

- General design considerations are as follows:
  - IOS firewall versus dedicated firewall—The small network corporate Internet module represents the ultimate in scaled-down network design, with all the features compressed into a single box. These features include routing, NAT, IDS, VLAN, voice services, VPN, and stateful firewall. Two principal alternatives were considered in the original SAFE security papers. The first was to use a router. This setup yielded the greatest flexibility for the small network because the router supports all the advanced services that may be necessary in today's networks. As an alternative, a dedicated firewall was also considered. However, the dedicated firewall places numerous general and IP telephony-specific restrictions on deployment:
    - Firewalls are generally Ethernet only, requiring some conversion to access PSTN and the WAN. This access would then most likely occur through the use of an additional router, a setup that would nullify the choice of using a single dedicated firewall in the first place.
    - Firewalls in this small scale of a design generally do not support enough interfaces or VLANs to provide segmentation between the Internet edge, public services, data, and voice segments.
    - For the branch mode of operation, firewalls do not support the same backup voice services for local call processing that routers do in case of headend failure. In a small network design, in which redundant links are not feasible, local backup call processing provides significant value.
  - Separate voice and data VLANs exist. The call-processing manager, proxy server, and IP phones reside in the voice segment. All other devices including management, user, and the voice-mail/mail system reside in the data segment.
- Access control and packet inspection guidelines are as follows:

- The router controls access between the data and voice segments via access control and stateful inspection. Because the call-processing center and IP phones reside in the same voice segment, access control is not possible between the two. The router controls access to all other flows listed in the “Controlling the voice-to-data segment interaction is key” axiom. Any other flows are denied and logged. Integrated IDS will trigger an alarm on any attack signatures detected in any of the above connections.
- In addition it should be noted that integrated IDS in routers and firewalls does not provide the full signature or feature set that a standalone NIDS appliance provides.

Small Standalone Network—Campus Module

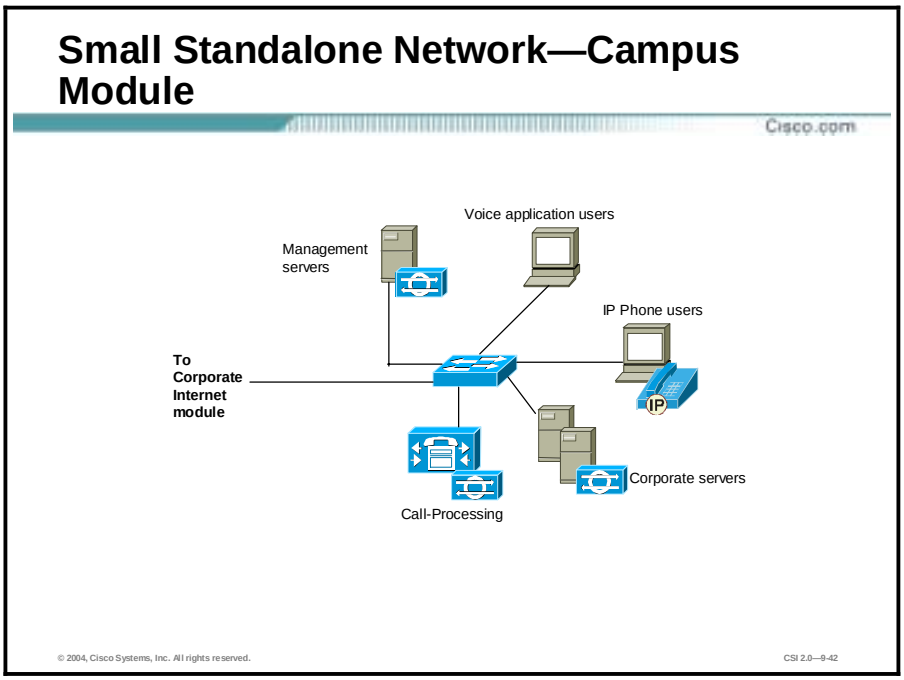

The campus module contains end-user workstations, corporate intranet servers, management servers, IP phones, and the associated Layer 2 infrastructure required to support the devices. Within the small network design, this Layer 2 functionality has been combined into a single switch.

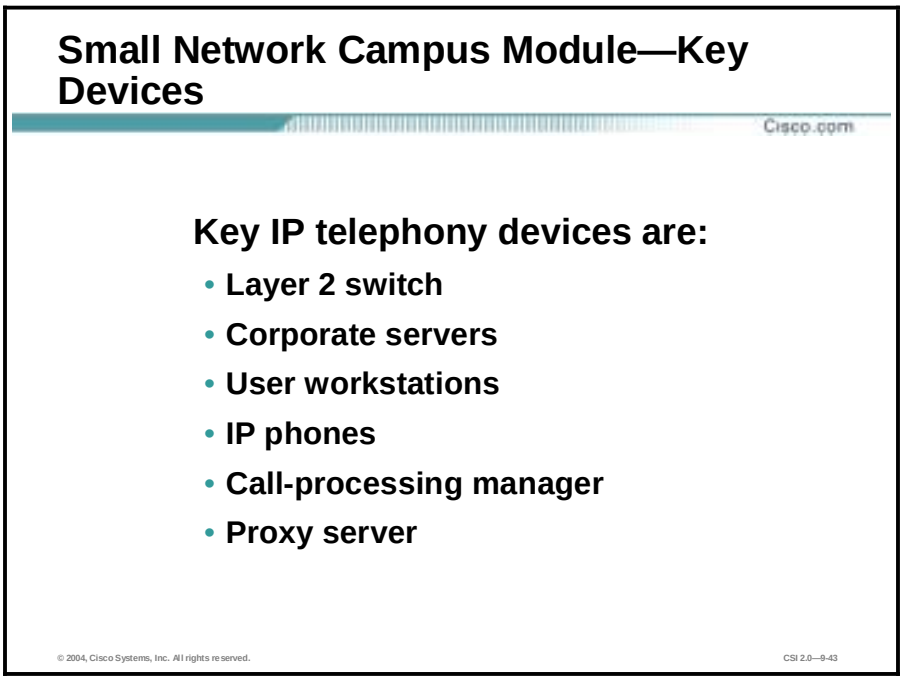

Key IP telephony devices for the small network campus module are the following:

- Layer 2 switch (with VLAN support)—Provides Layer 2 services to data and voice devices
- Corporate servers—Provide e-mail and voice-mail services to internal users, as well as delivering file, print, and DNS services to workstations
- User workstations—Provide data services and voice services via PC-based IP Phones to authorized users on the network
- IP phones—Provide voice services to users on the network
- Call-processing manager—Provides voice services to IP telephony devices in the network
- Proxy server—Provides data services to IP phones

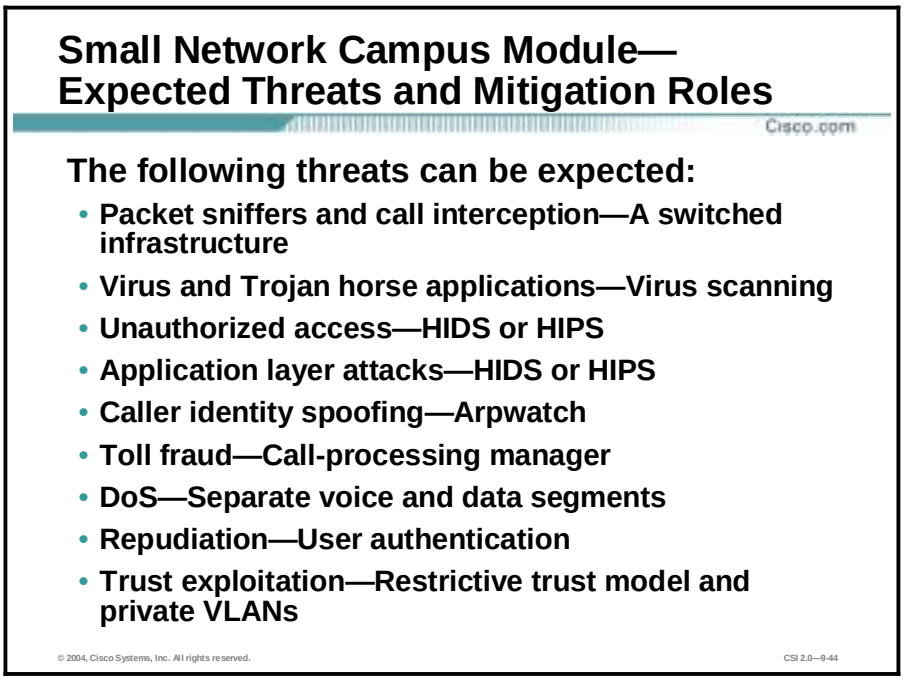

You can expect the following threats and mitigation roles in the small network campus module:

- Packet sniffers and call interception—A switched infrastructure limits the effectiveness of sniffing.
- Virus and Trojan horse applications—Host-based virus scanning prevents most viruses and many Trojan horses.
- Unauthorized access—This type of access is mitigated through the use of either HIDS or HIPS and application access control.
- Application layer attacks—Operating systems, devices, and applications are kept up-to-date with the latest security fixes, and most servers are additionally protected by either HIDS or HIPS.
- Caller identity spoofing—Arpwatch notifies the administrator of the unknown device.
- Toll fraud—The call-processing manager will not allow unknown phones to be configured.
- Denial of service—Separation of the voice and data segments significantly reduces the likelihood of an attack.
- Repudiation—Users are authenticated before accessing the telephony device, thus reducing the likelihood of a later denial that a call ever occurred.
- Trust exploitation—Restrictive trust model and private VLANs limit trust-based attacks.

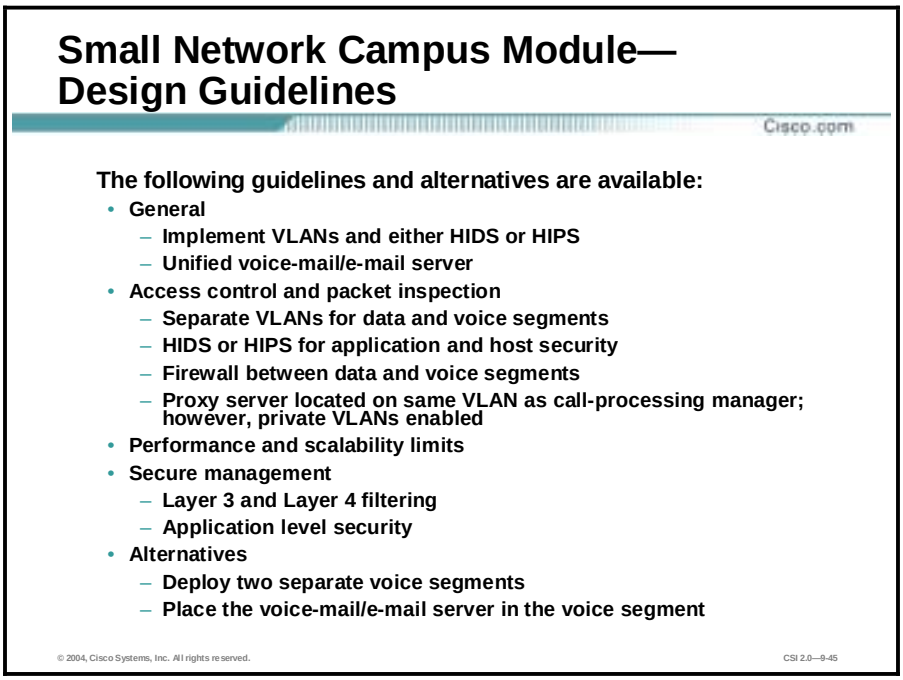

The following guidelines and alternatives are available in the small network campus module:

- General—The primary functions of the campus switch are to switch data, management, and voice traffic and to provide connectivity for the corporate, voice, and management servers and users. These functions are carried out via VLAN support and reliance on either HIDS or HIPS and virus scanning on key systems. A unified voice-mail or e-mail server can be placed in the data segment, given the small scale and the limited number of voice segments. This placement further restricts the number of hosts accessing the sole voice segment.
- Access control and packet inspection guidelines are as follows:
  - Within the Layer 2 switch, VLANs are enabled in order to mitigate attacks sourced from the data segment against the voice segment.
  - Because there are no Layer 3 services within the campus module, it is important to note that this design places an increased emphasis on application and host security because of the open nature of the internal network. Therefore, either HIDS or HIPS can be installed on key systems within the campus, including the corporate servers, call-processing manager, and management systems.
  - Because this network is so small and only a single voice and single data segment exist, it is actually more feasible to deploy a firewall between the two segments than it is in a larger design. Thus, deployment of PC-based IP phones in the data segment is more feasible as the stateful firewall functionality of the router can broker the data-voice interaction. However, the remaining issues covered in the “PC-based IP phones require open access” axiom still hold.
  - Virus scanning was installed on user systems. This provided virus or worm attack mitigation for the PC-based IP phone hosts against attacks on the data segment. If the hosts were infected, the virus or worm could potentially spread and attack the voice segment. Either HIDS or HIPS can be used to detect any anomalies on the mail, voice-mail, or call-processing devices. Access control is not carried out in this module.

- The proxy server is located on the same VLAN as the call-processing manager. However, private VLANs are enabled to mitigate local trust-exploitation attacks.

- Performance and Scalability—For the standalone network design, the scalability limit is the number of IP telephony devices supported by the local call-processing manager and voice-mail system. Performance in this design is not an issue because all necessary services are available locally on a Fast Ethernet switched network.
- Secure Management—Layer 3 and Layer 4 filtering was put in place to limit the administration of the voice servers by known management systems. Application-level security was used to provide confidentiality and user authentication for the configuration and monitoring of management traffic. IP phones download the latest configurations and operating system versions from the call-processing manager at regular intervals.
- Alternatives are as follows:
  - The first logical modification to the design would be to deploy two separate voice segments: one for IP phones and the other for the call-processing manager. This configuration would provide additional connection management between user and telephony service systems than that provided by the single segment. The design would then be very similar to that of the medium campus IP telephony design.
  - You might also consider placing the voice-mail/e-mail server in the voice segment. However, this is not recommended. Most likely the mail server is running multiple services. For instance, the mail server might also operate as the domain controller and DNS server. The associated risk is great enough that you should keep it separate rather than on the same segment as the IP phones and call-processing manager.

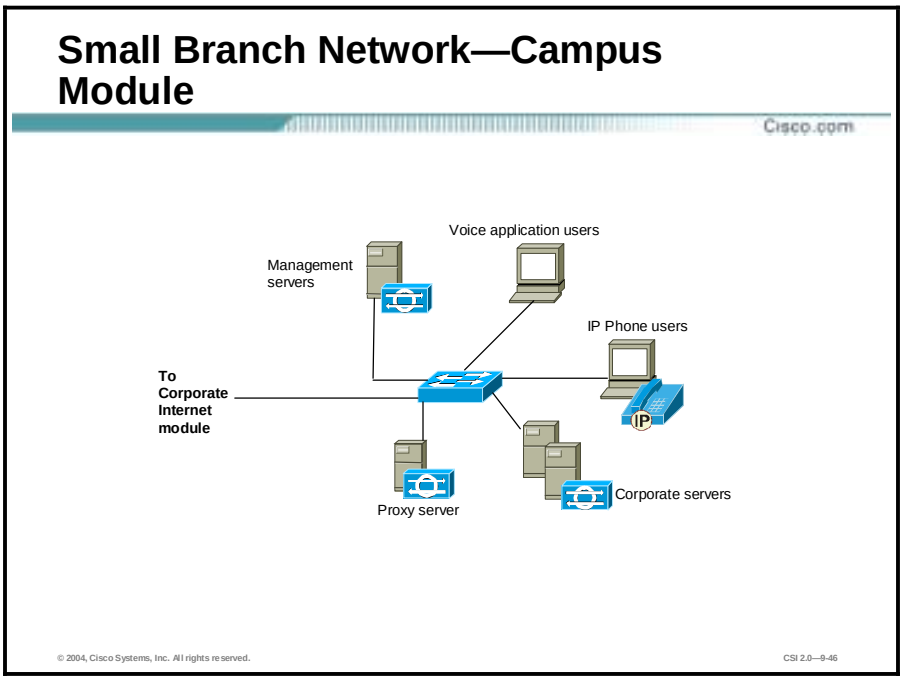

The primary consideration between the two designs is the location of the voice services. Because the voice services are located remotely in the branch design, configuration of the firewall router will be more complex to allow the connections as outlined in the “Controlling the voice-to-data segment interaction is key” axiom. This setup increases the likelihood of configuration errors. The branch design requires a different software image supporting call processing on the router in case of headend failure. This requirement constitutes the only difference in the corporate Internet module for both designs. The number of telephony devices supported will be limited first by the headend call-processing manager and the voice-mail system. The second limiting factor is the number of devices supported by the voice-enabled router in the corporate Internet module in case of headend failure. Performance may also be limited by the slower response time of the call-processing manager and voice-mail systems located at the headend.

Medium Network IP Telephony Design

This topic describes medium IP telephony design.

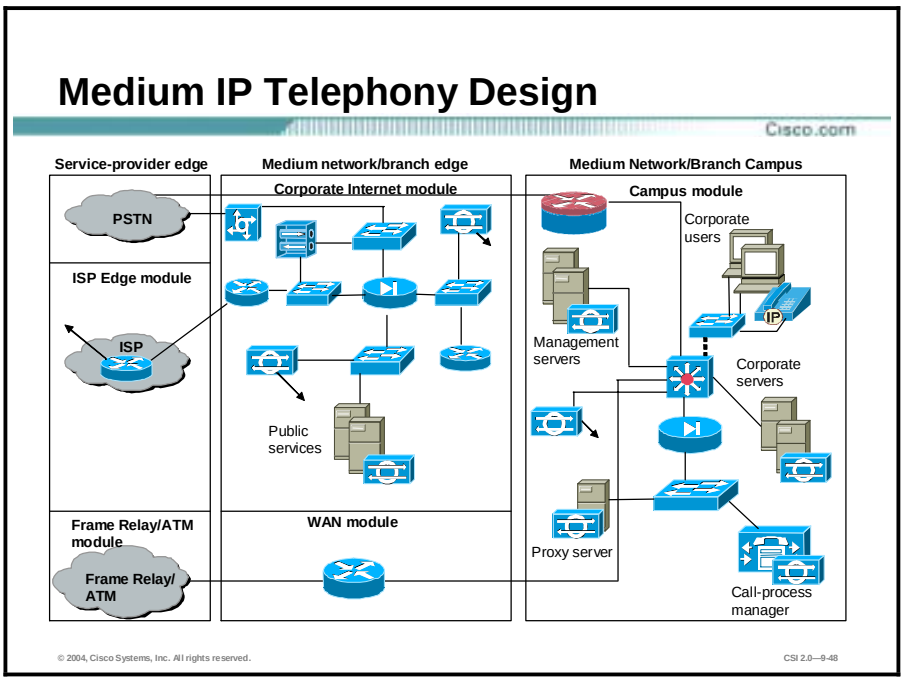

The medium IP telephony design uses the medium network design from the SAFE security papers. The corporate Internet module has not been modified. The campus module has been modified to support IP phones, PC-based IP Phones, voice services, proxy services, PSTN for WAN backup and local calls, and VLANs for data/voice segmentation. Most of the discussion about the medium IP telephony design is based on this design operating as the headend for a corporation. Specific design changes when used as a branch are also included.

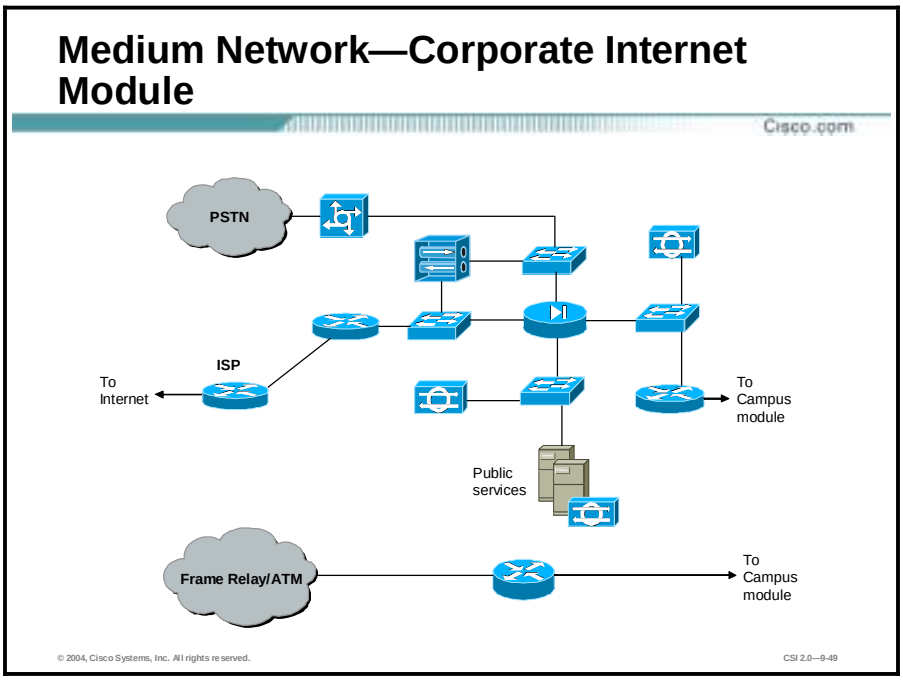

The corporate Internet module provides internal users with connectivity to Internet services and Internet users access to information about the public servers (HTTP, FTP, SMTP, and DNS). The WAN module provides connections to remote locations over a private network.

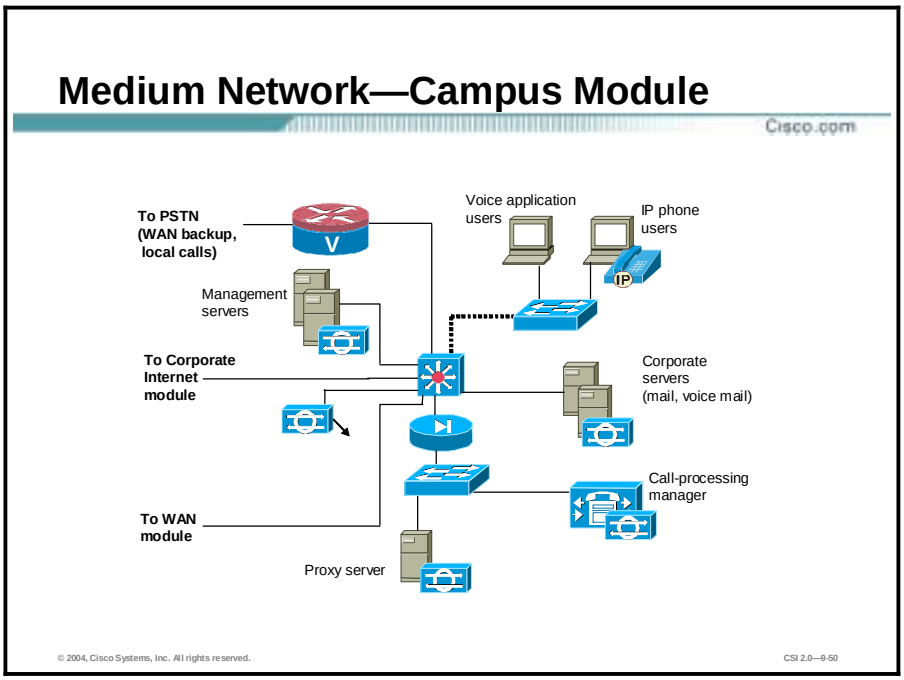

The original campus module contains end-user workstations, corporate servers, management servers, and the associated Layer 2 and Layer 3 infrastructure required to support the devices. Its primary purpose is to switch production and management traffic and to provide connectivity for the corporate and management servers and users. To support IP telephony, Cisco has added IP phones, voice servers, a proxy server, additional voice VLANs, and a call-processing manager with a stateful firewall for protection.

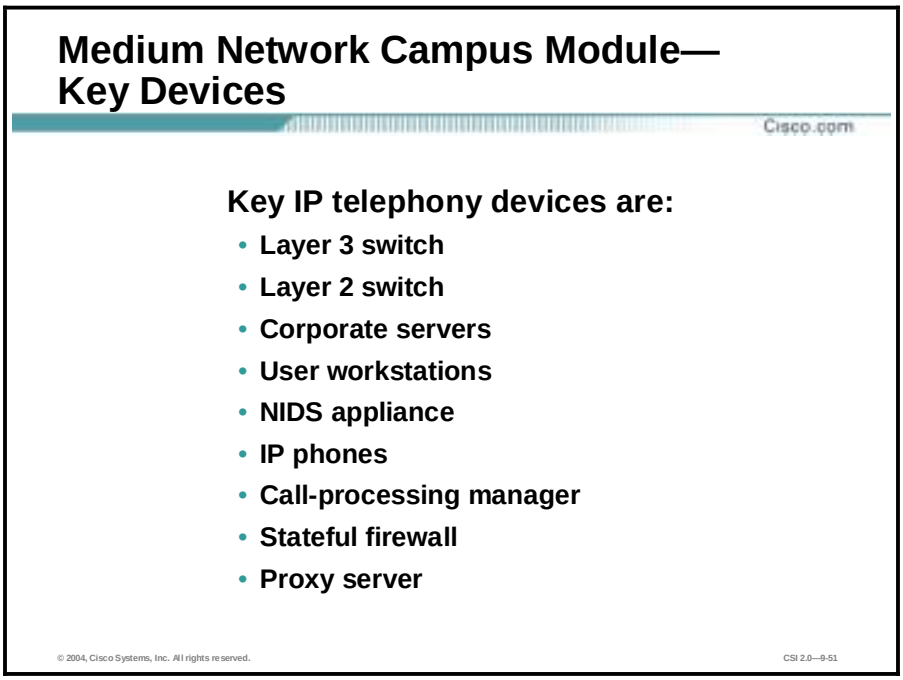

The following are the key IP telephony devices:

- Layer 3 switch—The Layer 3 switch routes and switches data, voice, and management traffic within the campus module; provides distribution layer services to the building switches; and supports advanced services such as traffic filtering.
- Layer 2 switch (with VLAN support)—The Layer 2 switch provides Layer 2 services to data and voice devices.
- Corporate servers—Corporate servers provide e-mail and voice-mail services to internal users, as well as delivering file, print, and DNS services to workstations.
- User workstations—User workstations provide data services and voice services via PC-based IP Phones to authorized users on the network.
- NIDS appliance—A NIDS appliance provides Layer 4-to-Layer 7 monitoring of key segments in the module.
- IP phones—IP phones provide voice services to users on the network.
- Call-processing manager—This feature provides voice services to IP telephony devices in the network.
- Stateful firewall—The stateful firewall provides network-level protection for the call-processing manager, including stateful filtering of traffic, DoS mitigation, and spoof mitigation.
- Proxy Server—Provides data services to IP phones.

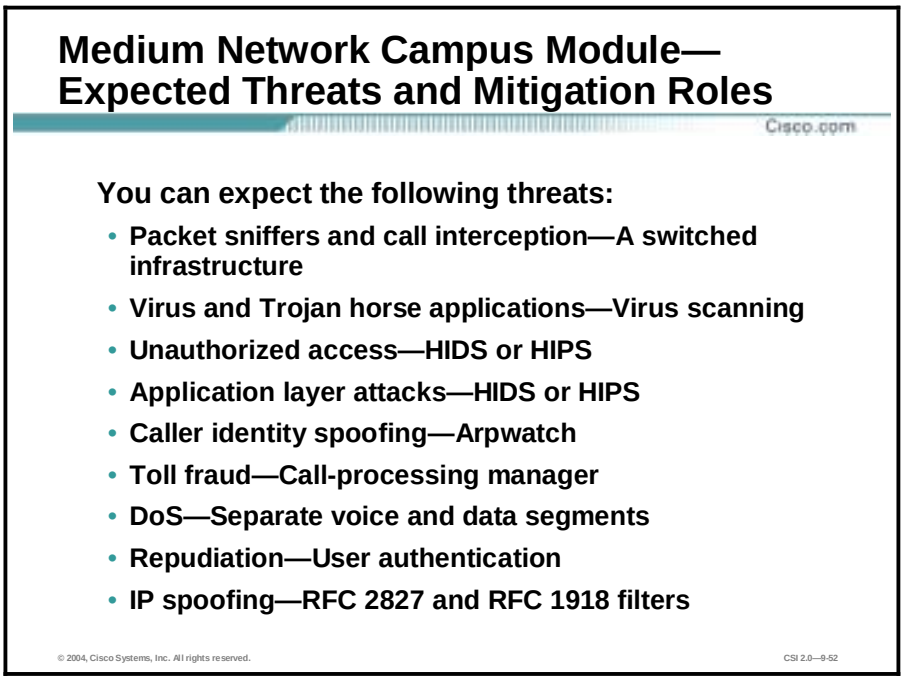

You can expect the following threats and mitigation roles in the medium network campus module:

- Packet sniffers and call interception—A switched infrastructure limits the effectiveness of sniffing.
- Virus and Trojan horse applications—Host-based virus scanning prevents most viruses and many Trojan horses.
- Unauthorized access—This type of access is mitigated through the use of either HIDS or HIPS and application access control.
- Application layer attacks—Operating systems, devices, and applications are kept up-to-date with the latest security fixes, and most servers are additionally protected by either HIDS or HIPS.
- Caller identity spoofing—Arpwatch notifies the administrator of the unknown device.
- Toll fraud—The call-processing manager will not allow unknown phones to be configured. Access control permits only known telephony networks to communicate with each other.
- Denial of service—Separation of the voice and data segments significantly reduces the likelihood of an attack. The stateful firewall TCP setup controls limit exposure to the call-processing manager and proxy server.
- Repudiation—Call-processing manager call-setup logs provide some level of non-repudiation.
- IP spoofing—RFC 2827 and RFC 1918 filters are placed at the ISP edge and local firewall.

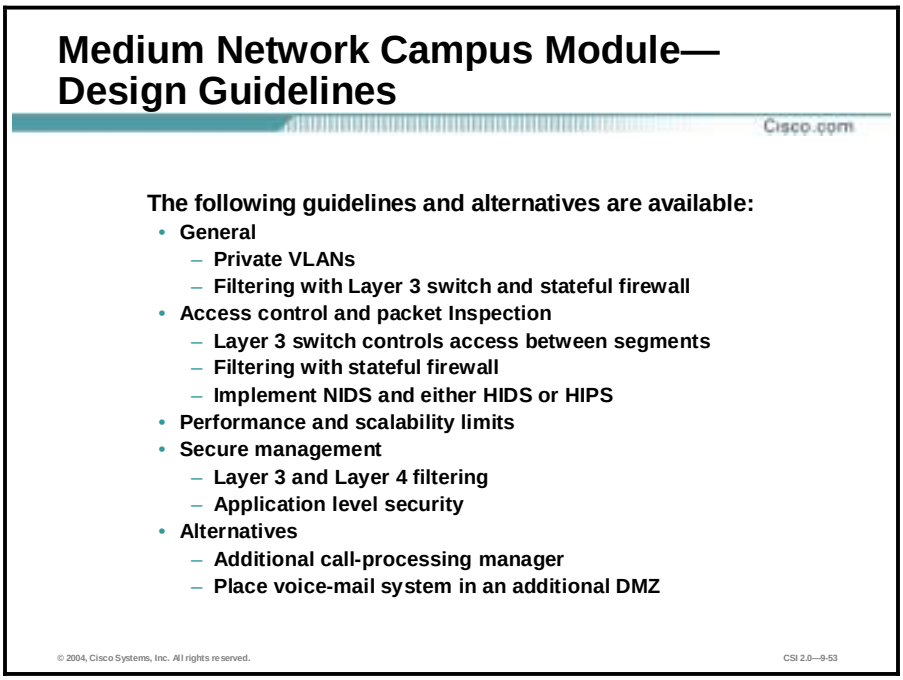

The following guidelines and alternatives are available in the medium network campus module:

- General—The primary function of the campus module is to switch data, voice, and management traffic and at the same time enforce the data network and voice segmentation. These functions are carried out by VLANs and filtering on both the Layer 3 switch and stateful firewall. Virus scanning protects PC-based IP Phone hosts on the data segment. Either HIDS or HIPS can be used to protect key voice services.
- Access control and packet inspection guidelines are as follows:
  - The Layer 3 switch controls access between the data and voice segments via access control and stateless filtering. The Layer 3 switch filters all flows listed in the "Controlling the voice-to-data segment interaction is key" axiom.
  - These connections are brokered by the stateful firewall. Any other flows between the voice and data segments are denied and logged by filtering on either the Layer 3 switch or the stateful firewall.
  - NIDS triggers an alarm on any attack signatures detected in any of these connections with the caveats outlined in the axioms section. Either HIDS or HIPS can be used to detect any anomalies on the mail, voice-mail, or call-processing devices. Virus scanning is installed on user systems to provide attack mitigation for the PC-based IP Phone against attacks on the data segment that could then traverse into the voice segment. The proxy server is located on the same VLAN as the call-processing manager. However, private VLANs are enabled to mitigate local trust-exploitation attacks.
- Performance and Scalability—For the standalone network, the limit of this design is the number of IP telephony devices supported by the local call-processing manager and voice-mail system. Performance in this design is not an issue because all necessary services are available locally on a Fast Ethernet switched network.
- Secure Management—Layer 3 and Layer 4 filtering is put in place to limit the administration of the voice servers by known management systems. Application-level

security is used to provide confidentiality and user authentication for the configuration and monitoring of management traffic. IP phones download the latest configurations and operating system versions from the call-processing manager at regular intervals.

The following alternatives are available:

- The addition of high availability for IP telephony—An administrator might consider adding a resilient firewall pair and an additional call-processing manager to accomplish this. This design would then be similar to the enterprise IP telephony design.
- To place the voice-mail system in an additional DMZ segment on the stateful firewall—This setup would provide for stateful inspection and filtering between the telephony devices and the voice-mail system instead of the current stateless filtering. This option would also provide DoS mitigation for the voice-mail system and stateful inspection between it and the mail server in the data segment. The only real drawback to this option is the increasing complexity of the configuration.

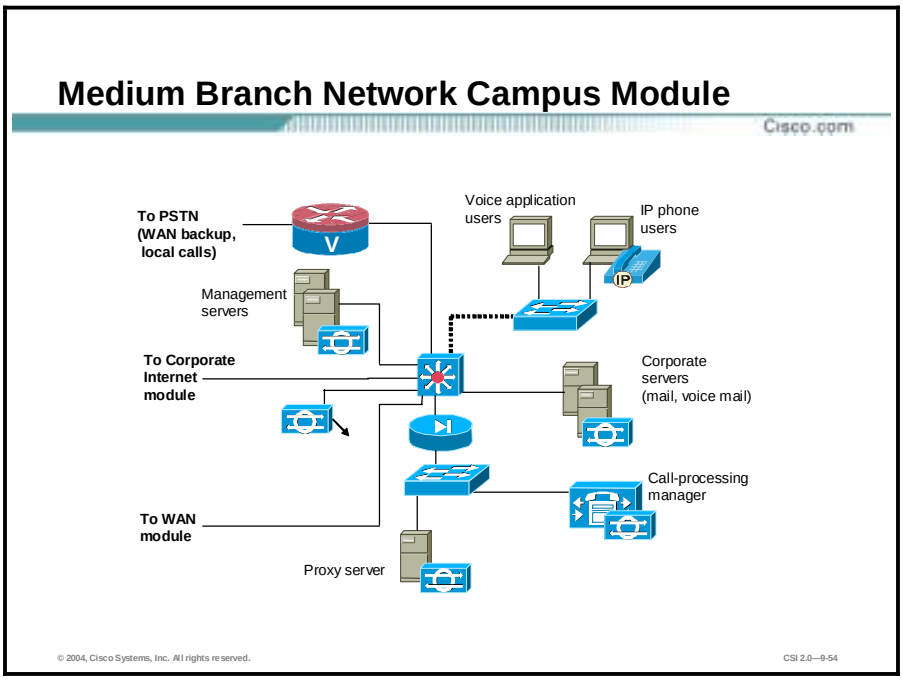

The primary consideration when using this design in a branch configuration is the location of the call-processing manager. Because the call-processing manager is located remotely in the branch design, configuration of the Layer 3 switch will be more complex to allow the connections as outlined in the “Controlling the voice-to-data segment interaction is key” axiom. This scenario increases the likelihood of configuration errors. In the branch design, performance may be limited by the slower response time of the call-processing manager located at the headend. Calls are transported to the headend via the WAN module.

Large Network IP Telephony Design

This topic describes the large network IP telephony design.

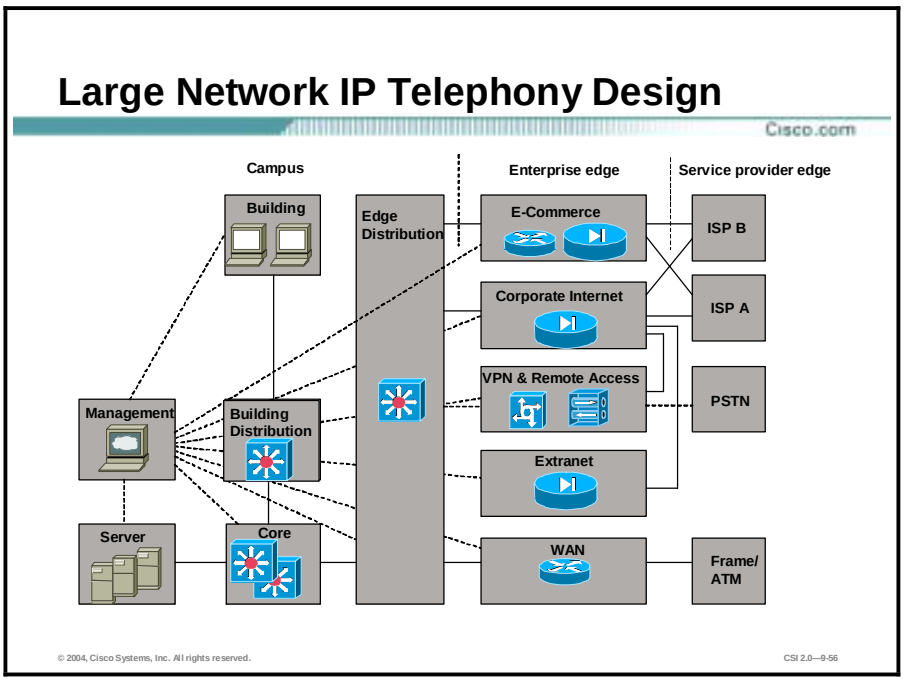

The large network IP telephony design uses the large network design from the SAFE: A security Blueprint for Enterprise Networks white paper. IP telephony is covered in the initial release of the paper but not discussed in depth. Some changes have been made to the design, including:

- PC-based IP Phones were added to the data segments of the research and development and marketing user groups.
- An additional voice segment was added for the voice-mail system.
- PSTN for local calls was added to the edge distribution module.
- The call-processing segment in the server module was made highly available and front ended with a pair of stateful firewalls.
- Either HIDS or HIPS was installed on all voice-related services.
- NIDS was tuned to the correct flows in the voice and related segments.

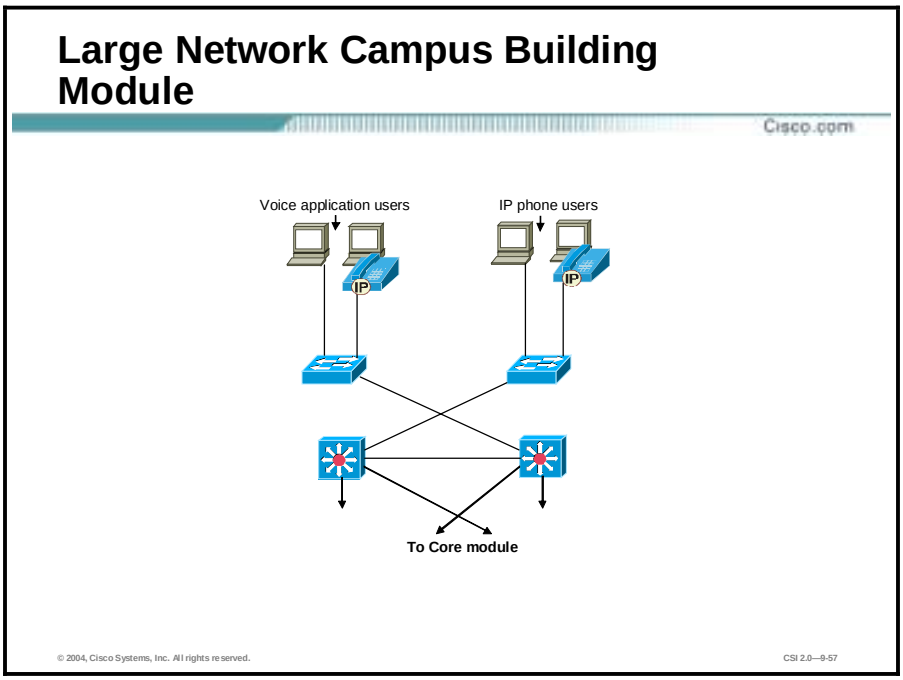

The building module contains end-user workstations, IP phones, and their associated Layer 2 access points. Its primary goal is to provide services to end-users. The topology of the building module remains unchanged from the original SAFE paper, but some of the mitigation factors have changed.

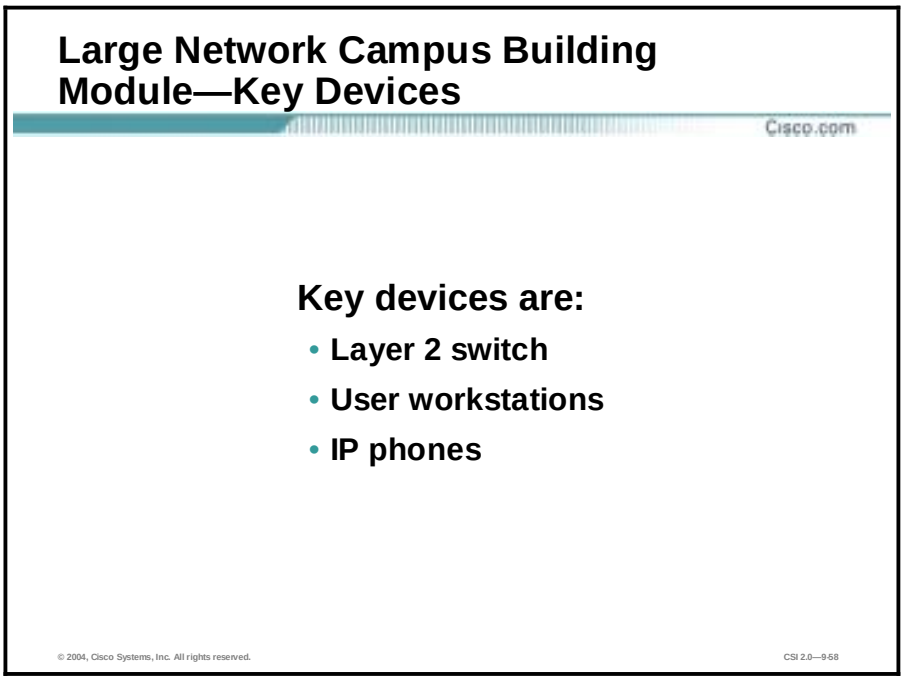

The following are the key IP telephony devices in the large network campus building module:

- Layer 2 switch (with VLAN support)—This switch provides Layer 2 services to data and voice devices.
- User workstations—User workstations provide data services and voice services via PC-based IP Phones to authorized users on the network.
- IP phones—IP phones provide voice services to users on the network.

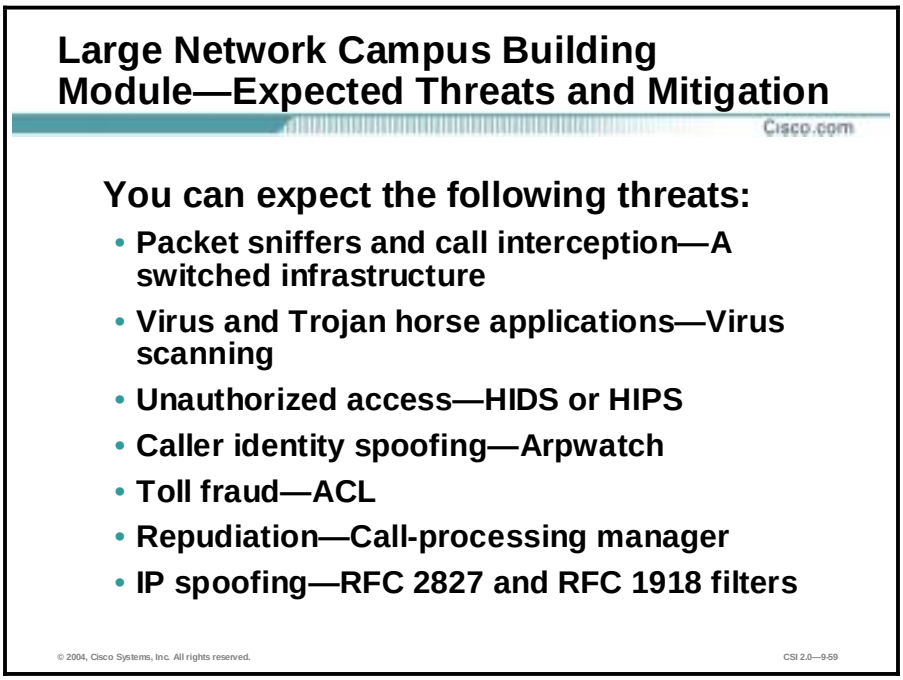

You can expect the following threats and mitigation roles in the large network campus building module:

- Packet sniffers and call interception—A switched infrastructure limits the effectiveness of sniffing.
- Virus and Trojan horse applications—Host-based virus scanning prevents most viruses and many Trojan horses.
- Unauthorized access—This type of access is mitigated through the use of either HIDS or HIPS and application access control.
- Caller identity spoofing—Arpwatch notifies the administrator of the unknown device.
- Toll fraud—Access control permits only known telephony networks to communicate with each other.
- Repudiation—Call-processing manager call-setup logs provide some level of nonrepudiation.
- IP spoofing—RFC 2827 and RFC 1918 filters are placed on the Layer 3 switches.

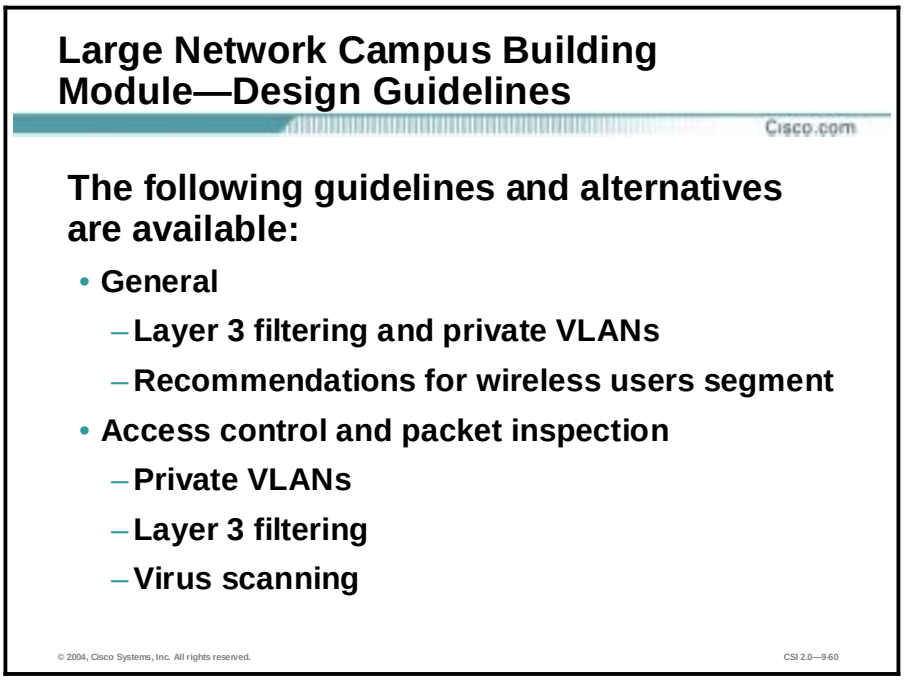

The following guidelines and alternatives are recommended in the large network campus building module:

- General—The primary function of the building module is to switch data and voice traffic and at the same time enforce the data and voice segmentation. These functions are carried out by stateless Layer 3 filtering and VLANs. Virus scanning protects user systems on the data segment. The SAFE wireless LAN security discusses the fact that user-group differentiation is not possible with today's wireless technology unless IPSec is used. A recommendation is made to disallow users on a wireless segment from accessing a controlled group segment, with the understanding that the risks outweigh the benefits.
- Access Control and Packet Inspection guidelines are as follows:
  - Within the switches, VLANs are enabled in order to mitigate attacks sourced from the data segment against the voice segment.
  - Stateless Layer 3 filtering controls the flows outlined in the "Controlling the voice-to-data segment interaction is key" axiom. Any other flows are denied and logged.
  - Virus scanning is installed to mitigate local attacks on the user systems. Since they are now running the PC-based IP Phone and have access to the voice segment, attack mitigation on these hosts is key to guard against attacks on the data segment from traversing into the voice segment.

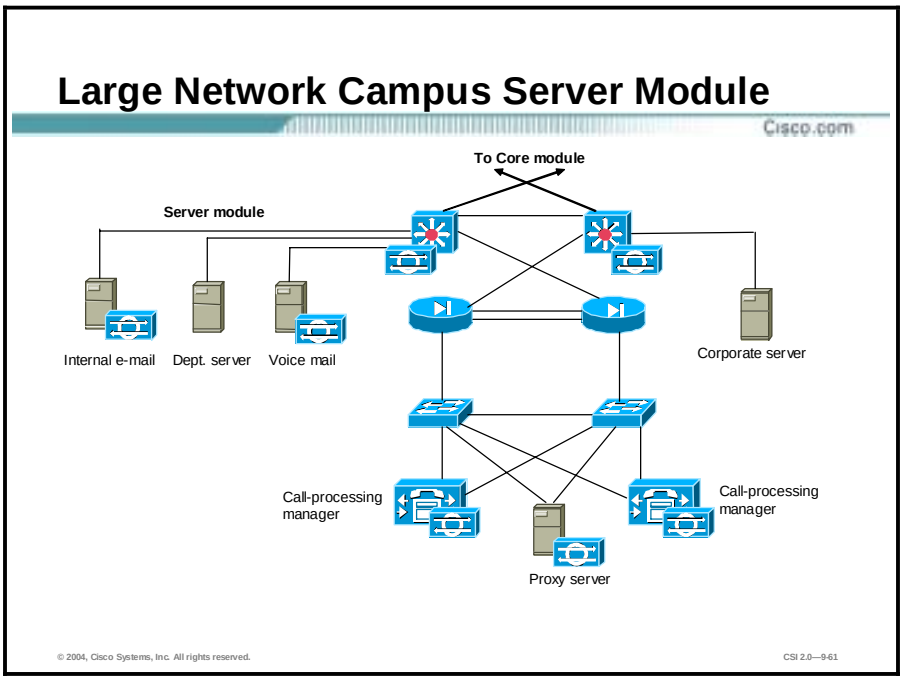

The primary goal of the server module is to provide application and voice services to end users and devices.

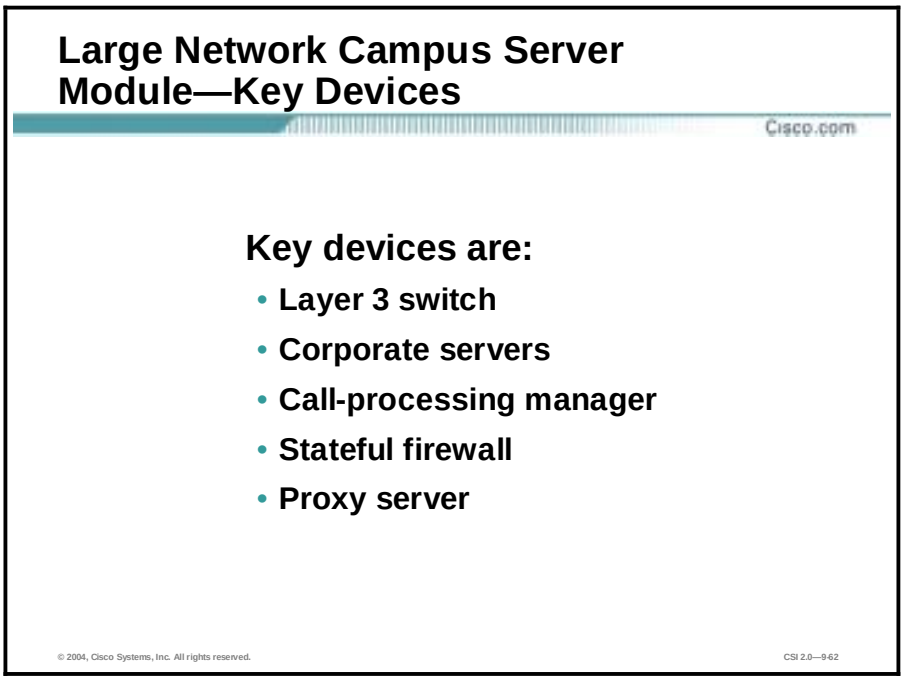

The following are the key IP telephony devices in the large network campus server module:

- Layer 3 switch—The Layer 3 switch routes and switches data, voice, and management traffic within the server module, and supports advanced services such as traffic filtering and NIDS.
- Corporate servers—Corporate servers provide e-mail and voice-mail services to internal users, as well as delivering file, print, and DNS services to workstations.
- Call-processing manager—The call-processing manager provides voice services to IP telephony devices in the network.
- Stateful firewall—The stateful firewall provides network-level protection for the call-processing manager, including stateful filtering of traffic, DoS mitigation, and spoof mitigation.
- Proxy Server—Provides data services to IP phones.

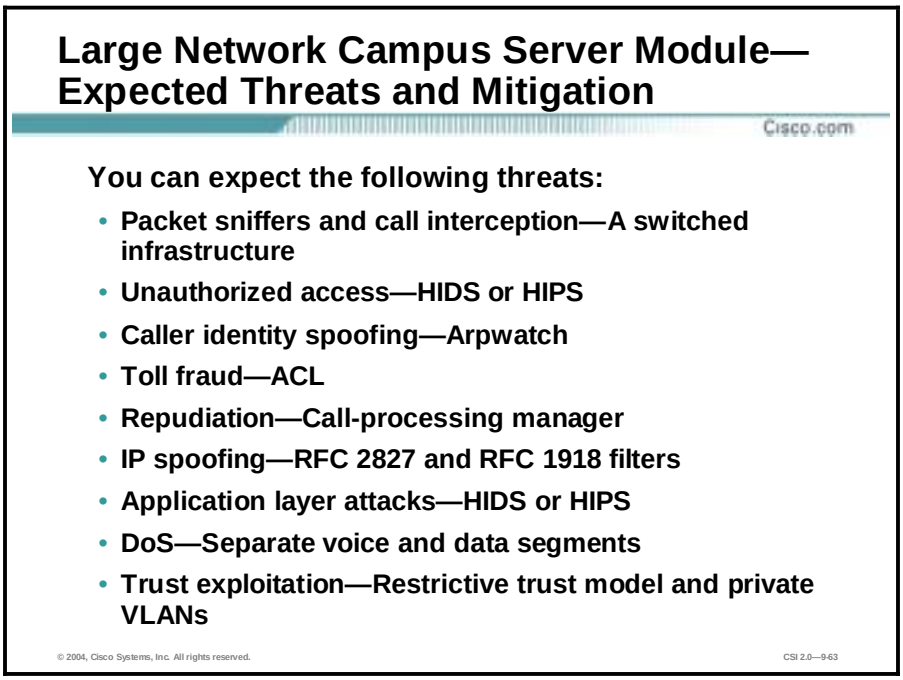

You can expect the following threats and mitigation roles in the large network campus server module:

- Packet sniffers and call interception—A switched infrastructure limits the effectiveness of sniffing.
- Unauthorized access—This type of access is mitigated through the use of either HIDS or HIPS and application access control.
- Caller identity spoofing—Arpwatch notifies the administrator of the unknown device.
- Toll fraud—The call-processing manager will not allow unknown phones to be configured. Access control permits only known telephony networks to communicate with each other.
- Repudiation—Call-processing manager call-setup logs provide some level of nonrepudiation.
- IP spoofing—IP spoofing provides RFC 2827 and RFC 1918 filtering on Layer 3 switches and stateful firewall.
- Application layer attacks—Operating systems, devices, and applications are kept up-to-date with the latest security fixes, and most servers are additionally protected by either HIDS or HIPS.
- Denial of service—Separation of the voice and data segments significantly reduces the likelihood of an attack. The stateful firewall TCP setup controls limit exposure to the call-processing manager and proxy server.
- Trust exploitation—Restrictive trust model and private VLANs to limit trust-based attacks.

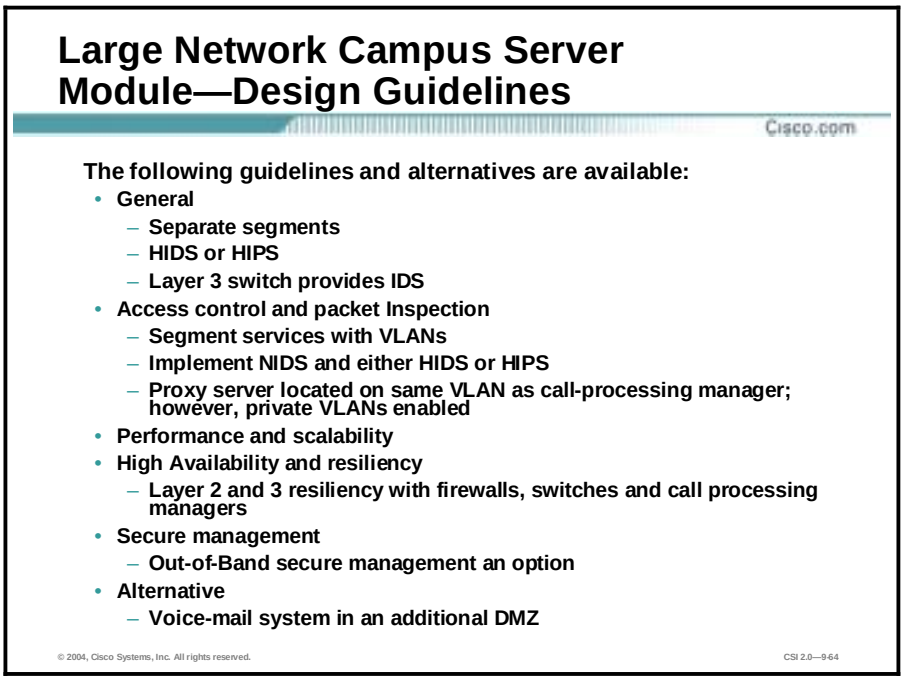

The following guidelines and alternatives are available:

- General—The server module contains all the voice services necessary for IP telephony. The call-processing manager, proxy server, voice-mail, and mail systems reside in separate segments, not only to model the scale of a large enterprise but also to provide layered security. All services have either HIDS or HIPS agents installed. All traffic flows in the server module are inspected by onboard IDSs within the Layer 3 switches. The call-processing segment is highly available and is protected by a stateful firewall pair.
- Access control and packet inspection recommendations are as follows:
  - Segmenting the services offers tremendous value in both scalability and security. It is easier to implement the required flows as outlined in the “Controlling the voice-to-data segment interaction is key” axiom, reducing the likelihood of configuration errors. Any other flows are denied and logged.
  - If NIDS detects a signature, the system will trigger an alarm with the caveats outlined in the axioms section. Either HIDS or HIPS detects anomalies on the mail, voice-mail, or call-processing devices.
  - The proxy server is located on the same VLAN as the call-processing manager. However, private VLANs are enabled to mitigate local trust-exploitation attacks.
- Performance and scalability—The scalability limit is the number of IP telephony devices supported by the call-processing manager and voice-mail system. Performance in this design is not an issue because all necessary services are available locally on a Fast Ethernet switched network, with the exception of some remote sites that use these local services over the WAN.
- High Availability and resiliency—This module provides Layer 2 and Layer 3 resiliency. With the addition of voice services, high availability is maintained. Two stateful firewalls connect the secured call processing manager segment to the dual Layer 3 switches in the server module. The internal segment is Layer 2 resilient, not only between the firewalls’ internal interfaces and the dual Layer 2 switches, but also on the dual-interfaced call-

processing managers as well. In this configuration, each call-processing manager operates with two network interface cards, both in the same network, with one connected to each switch.

- Secure management—The enterprise SAFE security design supports out-of-band secure management as one potential management option. This means that all network devices and all key servers are dual-homed to provide a dedicated and secure interface for management. In comparison, devices can be managed in-band over existing segments, but this mixes the production and management traffic and thus makes it more difficult to secure the device. All voice servers should support more than one interface to support this best practice. The voice services are critical components of the network and restricting management access to them is key in safeguarding them from attack. Layer 3 and Layer 4 filtering is put in place to limit the administration of the voice servers to known management systems. Application-level security is used to provide confidentiality and user authentication for the management traffic.

The following alternative is recommended:

- Place the voice-mail system on an additional DMZ segment on the stateful firewall. This setup provides for stateful inspection and filtering between the telephony devices and the voice-mail system, instead of the current stateless filtering. This option also provides DoS mitigation for the voice-mail system and stateful inspection between it and the mail server in the data segment. The only real drawback to this option is the increased complexity of the configuration.

Summary

Cisco.com

- There are four main voice-specific components:
  - IP telephony devices
  - Call-processing manager
  - Voice-mail system
  - Voice gateway
- Enterprise IP telephony networks can be deployed in three ways:
  - Single-site campus
  - WAN centralized call-processing
  - WAN distributed call-processing
- There are numerous attacks against the IP telephony network.

Summary (Cont.)

Cisco.com

- Discussed branch versus headend considerations and branch versus standalone considerations for small, medium, and large IP telephony networks.
- The mitigation roles identified for each threat in the SAFE white paper are integral to a successful VoIP network implementation.
- The design process is often a series of trade-offs. Some of these trade-offs are made at the module level, while others are made at the component level.

Review Questions

Q 1) What are four main voice-specific components of an IP telephony network?

A) Four main voice-specific components are IP telephony devices, call-processing manager, voice-mail system, and voice gateway.

Q 2) What protocol is used in call setup to allow the IP telephony conversation to commence?

A) Real-Time Transport Protocol (RTP) running on top of UDP/IP allows the conversation to commence.

Q 3) What is the main difference between WAN centralized call-processing deployment and WAN distributed call-processing deployment?

A) In WAN centralized call-processing deployment, only the headend site contains the call-processing manager cluster. WAN distributed call-processing deployment has one or more sites that contain a call-processing manager cluster.

Q 4) Name the three VoIP standards?

A) Proposed IP telephony standards are H.323, Session Initiation Protocol (SIP), and Media Gateway Control Protocol (MGCP).

Q 5) What function does MGCP perform?

A) In MGCP, media gateway controllers or call agents provide control, signaling, and the processing skills to control the telephony gateways.

Q 6) Why are PC-based IP Phones not as resilient under attack as their IP phone counterparts?

A) PC-based IP Phone hosts are more susceptible to attacks due to the number of vectors into the system. These include operating system vulnerabilities, application vulnerabilities, service vulnerabilities, worms, viruses, and so on.

Q 7) What two voice-to-data connections exist only if PC-based IP Phones are deployed?

A) A PC-based IP Phone in the data segment accessing the call-processing manager in the voice segment for call establishment, and a PC-based IP Phone in the data segment accessing the voice-mail system when placed in the voice segment.

Q 8) What is the primary method for device authentication in a VoIP network?

A) The primary method for device authentication of IP phones is the MAC address.

Q 9) What threat is mitigated by enabling call control on the call-processing manager?

A) Enabling call-control logging on the call-processing manager provides a record of placed calls, which aids in nonrepudiation.

Q 10) What are SAFE recommendations to secure voice-mail and call-processing manager systems?

A) Security precautions include: turning off all unneeded services, patching the operating system and services with the latest security patches, hardening the operating system configuration, disabling any features on the voice servers that are not in use, and finally, not running any unnecessary applications on the server. Installing either HIDS or HIPS is also recommended.

Q 11) What are the key benefits of voice and data segmentation?

A) Reasons to have two logically disparate segments for voice and data include QoS, scalability, manageability, and security. Segmenting IP voice from the traditional IP data network greatly increases attack mitigation capability and allows use of the same access, core, and distribution layers.

Q 12) What port numbers are used by calls placed between IP telephony devices?

A) Calls placed between IP telephony devices generally use dynamically assigned UDP port numbers greater than 16384.

10

SAFE: Wireless LAN Security in Depth

Overview

This lesson includes the following topics:

- Objectives
- Wireless LAN Security Concepts
- SAFE Wireless LAN Caveats and Axioms
- Wireless LAN Security Extensions
- Cisco Wireless LAN Product Portfolio
- Wireless LAN Design Approach
- Standard WLAN Design
- Enterprise Wireless LAN Design
- Medium Wireless LAN Design
- Small Wireless LAN Design
- Remote Wireless LAN Design
- SAFE WLAN Implementation
- Summary
- Lab Exercise

Objectives

This topic lists the lesson's objectives.

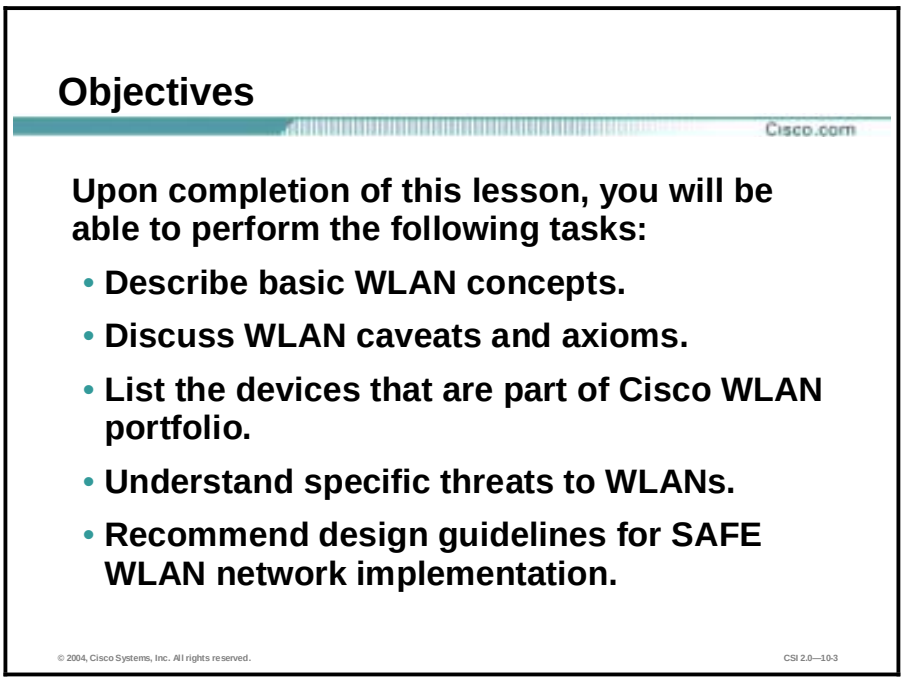

Wireless LAN Security Concepts

This topic provides an overview of basic wireless LAN (WLAN) concepts.

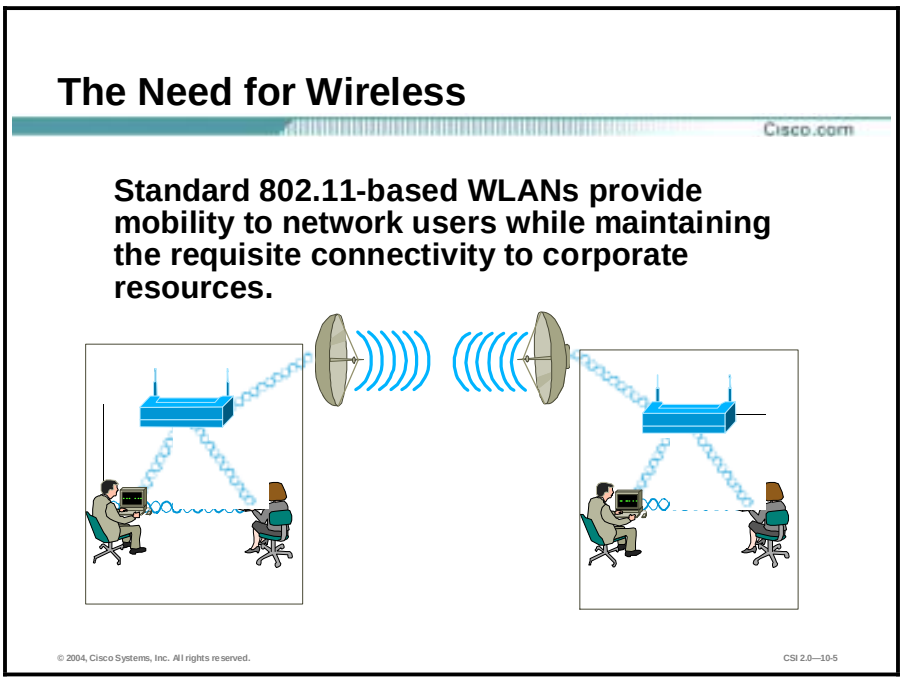

As laptops become more pervasive in the workplace, users are more prone to use them as their primary computing device, allowing greater portability in meetings and conferences and during business travel. WLANs offer organizations greater productivity per employee by providing constant connectivity to traditional networks in venues where connectivity was previously unavailable.

Wireless network connectivity is not limited to enterprise use. WLANs offer increased productivity before and after meetings, and outside the traditional office environment. Numerous wireless ISPs (WISPs) are appearing in airports, coffee shops, hotels, and conference and convention centers, enabling enterprise users to connect in public-access venues.

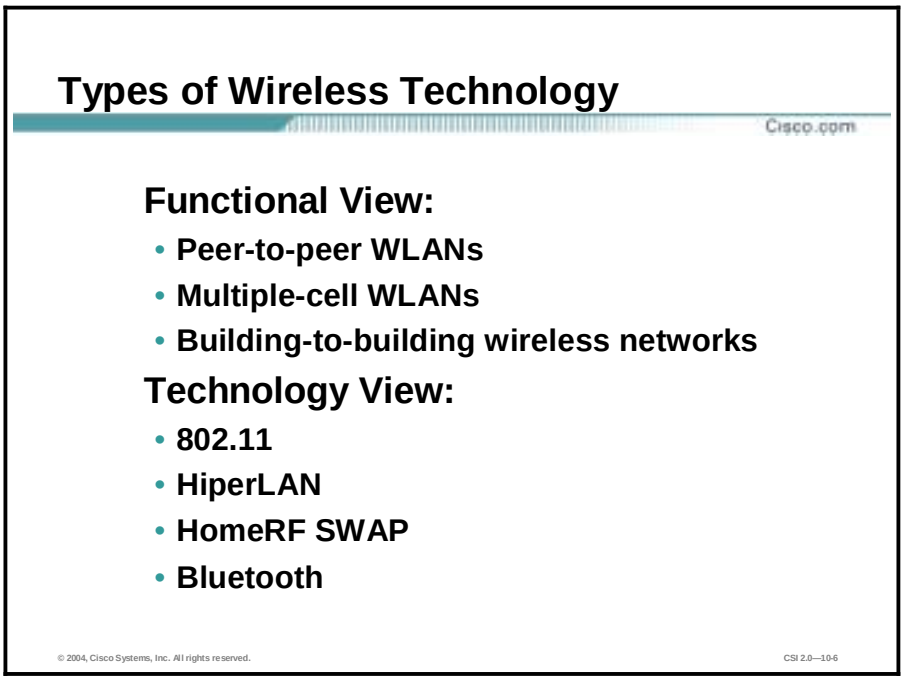

Wireless local-area networking has existed for many years, providing connectivity to wired infrastructures where mobility was a requirement to specific working environments. These early wireless networks were nonstandard implementations, with speeds ranging between 1 and 2 Mbps. Without any standards driving WLAN technologies, the early implementations of WLAN were relegated to vendor-specific implementations, with no provision for interoperability, thereby inhibiting the growth of standards-based WLAN technologies. Today, WLANs can be categorized based on their function and the technology they use.

From a functional viewpoint, WLANs can be categorized as follows:

- Peer-to-peer WLANs
- Multiple-cell WLANs
- Building-to-building wireless networks (point-to-point and point-to-multipoint)

From a technology viewpoint, WLANs can be categorized as follows:

- 802.11—The IEEE 802.11 committee developed a standard that would deliver a high-quality, high-performance product for different wireless data companies and some wired data LAN companies. IEEE 802.11-based WLANs present new challenges for network administrators and information security administrators alike. Unlike the relative simplicity of wired Ethernet deployments, 802.11-based WLANs broadcast radio frequency (RF) data for the client stations to hear. This presents new and complex security issues..
- HiperLAN—A European Telecommunications Standards Institute (ETSI) standard ratified in 1996. HiperLAN/1 standard operates in the 5-GHz radio band up to 24 Mbps. The ETSI has recently approved HiperLAN/2, which operates in the 5-GHz band at up to 54 Mbps using a connection-oriented protocol for sharing access among end-user devices.
- HomeRF SWAP—In 1988, The HomeRF SWAP Group published the Shared Wireless Access Protocol (SWAP) standard for wireless digital communication between PCs and consumer electronic devices within the home. SWAP supports voice and data over a

common wireless interface at 1 and 2 Mbps data rates using frequency-hopping and spread-spectrum techniques in the 2.4-GHz band.

- Bluetooth—A personal-area network (PAN) specified by the Bluetooth Special Interest Group for providing low-power and short-range wireless connectivity using frequency-hopping spread spectrum in the 2.4-GHz frequency environment.

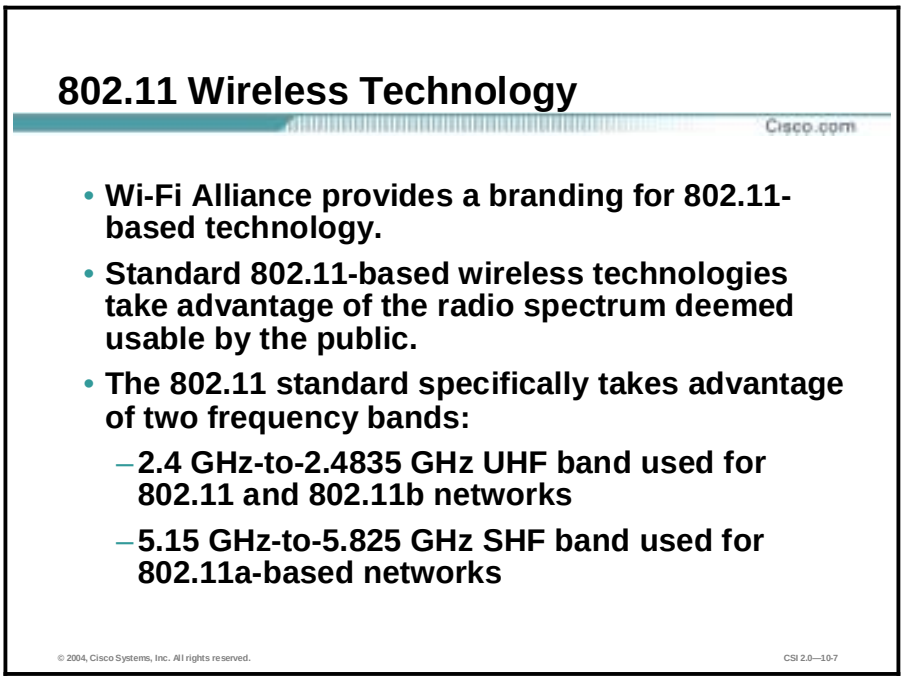

The IEEE maintains the 802.11-based standard, as well as other 802-based networking standards, such as 802.3 Ethernet. A nonprofit, vendor-neutral organization known as the Wireless Fidelity Alliance (Wi-Fi Alliance) provides a branding for 802.11-based technology known as Wi-Fi. A Wi-Fi compliant device must pass interoperability testing in the Wi-Fi laboratory. All vendor products that are Wi-Fi certified are guaranteed to work with all other Wi-Fi certified products—regardless of the vendor.

Standard 802.11-based wireless technologies take advantage of the radio spectrum deemed usable by the public. This spectrum is known as the Industrial, Scientific, and Medical (ISM) band. The 802.11 standard specifically takes advantage of two of the three frequency bands, the 2.4 GHz-to-2.4835 GHz UHF band used for 802.11b and 802.11g networks, and the 5.15 GHz-to-5.825 GHz SHF band used for 802.11a-based networks.

The spectrum is classed as unlicensed, meaning there is no one owner of the spectrum, and anyone can use it as long as that user's device complies with FCC regulations. Some of the areas the FCC governs include the maximum transmit power of the radios and the type of encoding and frequency modulations that can be used.

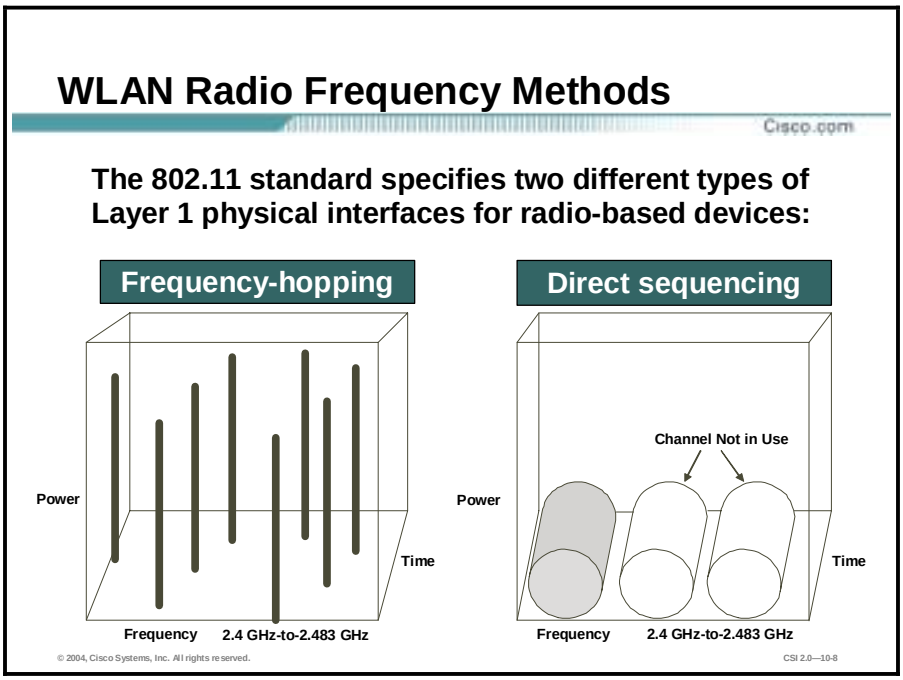

The 802.11 standard specifies two different types of Layer 1 physical interfaces for radio-based devices. One uses a frequency-hopping architecture, whereas the other uses a more straightforward single-frequency approach, known as direct sequencing.

- Frequency hopping—The frequency-hopping architecture makes use of the available frequency range by creating hopping patterns to transmit on one of 79 1-MHz-wide frequencies for no more than 0.4 seconds at a time. This setup allows for an interference-tolerant network. If any one channel stumbles across interference, it would be for only a small time slice because the frequency-hopping radio quickly hops through the band and retransmits data on another frequency. The major drawback to frequency hopping is that the maximum data rate achievable is 2 Mbps.
- Direct-sequencing—Direct sequencing provides 11 overlapping channels of 83 MHz within the 2.4-GHz spectrum. Within the 11 overlapping channels, there are three 22-MHz-wide non-overlapping channels. The large bandwidth along with advanced modulation based on complementary code keying (CCK) provided by direct sequencing is the primary reason why direct sequencing can support higher data rates than frequency hopping. Additionally, because the three channels do not overlap, three access points can be used simultaneously to provide an aggregate data rate of the combination of the three available channels. In 1999, the IEEE ratified the 802.11b standard, which provided newer, enhanced modulation types to allow direct-sequencing networks to achieve data rates as high as 11 Mbps, or 33 Mbps when the three non-overlapping channels are used together. Direct sequencing does have one disadvantage compared to frequency hopping: interference intolerance. Though both are affected by interference, throughput in a direct-sequencing network falls dramatically when interference is introduced.

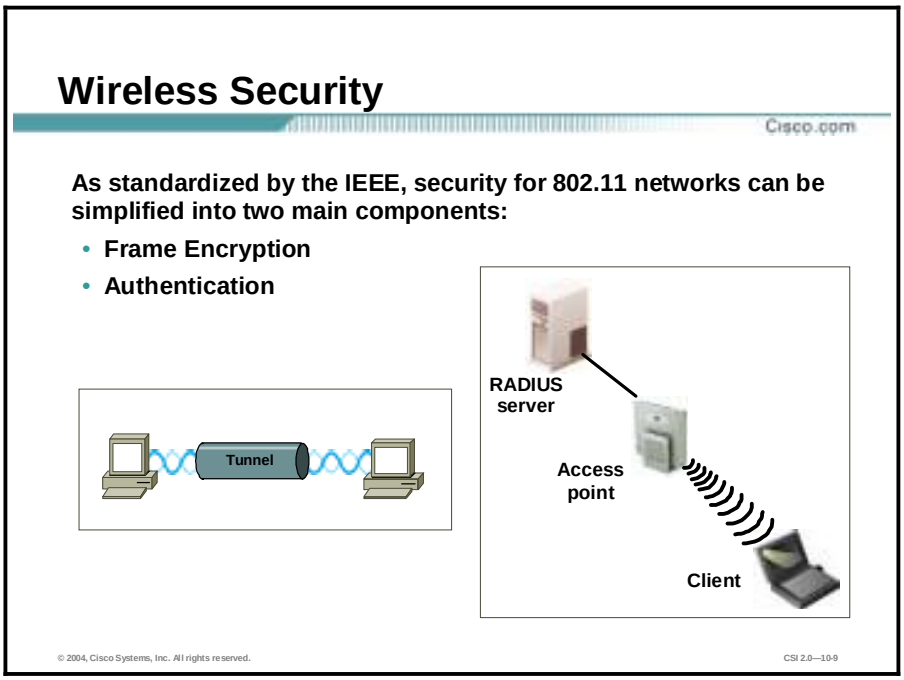

As standardized by the IEEE, security for 802.11 networks can be simplified into two main components: encryption and authentication. The implementation of these components has been proven and documented as insecure by the security community at large.

- Frame Encryption—Properly performed encryption allows for confidentiality. Encryption is the process of taking a message, referred to as clear text, and passing it through a mathematical algorithm to produce what is known as cipher text. Decryption is the reverse of the process. Encryption algorithms typically rely on a value, called a key, in order to encrypt and decrypt the data. Two major forms of encryption are used today:
  - Symmetric encryption (also known as shared-key encryption)
  - Asymmetric encryption (also know as public or private encryption)

Note Symmetric encryption is about 1000 times faster than asymmetric encryption, and is, therefore, used for bulk encryption of data. Generally with well-designed encryption algorithms, longer keys result in a higher degree of security because more brute force is required to try every possible key (known as the key space) in order to decrypt a message. The IEEE has specified that wired equivalent privacy (WEP) be the means to encrypt 802.11 data frames. WEP uses the RC4 stream cipher for encryption. The RC4 encryption algorithm is a symmetric stream cipher that supports a variable-length key.

- Authentication Mechanism—The IEEE specifies two authentication algorithms for 802.11-based networks. First is open authentication, which is a null authentication algorithm because any station requesting authentication is granted access. The second form of authentication is called shared-key authentication, which requires that both the requesting and granting stations be configured with matching WEP keys. The requesting stations send an authentication request to the granting station. The granting station sends a plaintext challenge frame to the requesting station. The requesting station WEP encrypts the challenge frame and sends it back to the granting station. The granting station attempts to

decrypt the frame, and if the resulting plaintext matches what the granting station originally sent, the requesting station has a valid key and is granted access.

Note	Shared-key authentication has a known flaw in its concept. Because the challenge packet is sent in the clear to the requesting station and the requesting station replies with the encrypted challenge packet, an attacker can derive the stream cipher by analyzing both the plaintext and the cipher text. This information can be used to build decryption dictionaries for that particular WEP key.

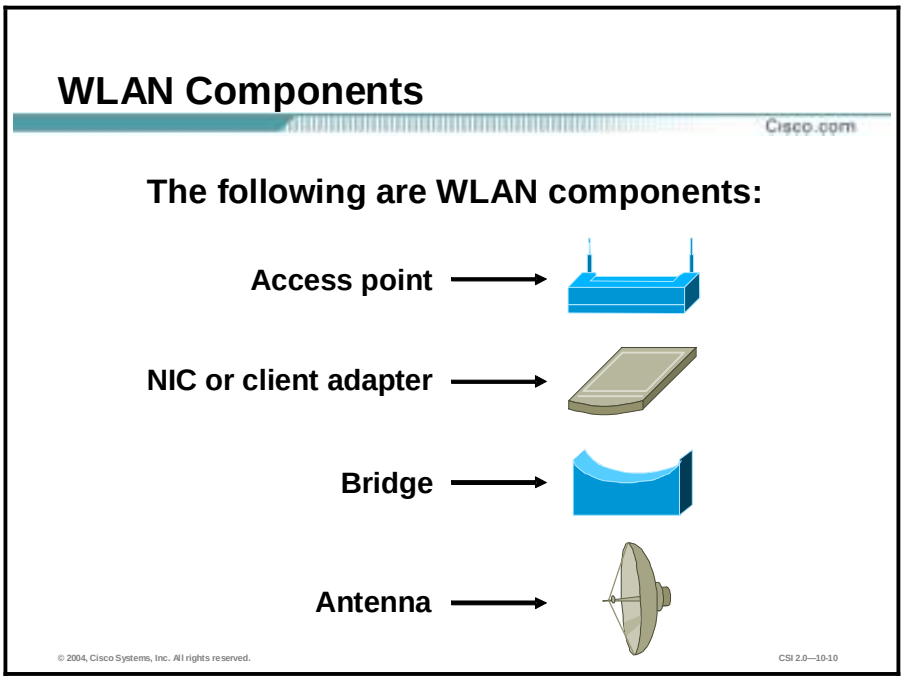

Components of a WLAN are access points, network interface cards (NICs) or client adapters, bridges, and antennas.

- Access point—An access point operates within a specific frequency spectrum and uses an 802.11 standard specified modulation technique. It also informs the wireless clients of its availability and authenticates and associates wireless clients to the wireless network. An access point also coordinates the wireless clients' use of wired resources.
- NIC or client adapter—A PC or workstation uses a wireless NIC to connect to the wireless network. The NIC scans the available frequency spectrum for connectivity and associates it to an access point or another wireless client. The NIC is coupled to the PC or workstation operating system using a software driver.
- Bridge—Wireless bridges are used to connect multiple LANs (both wired and wireless) at the Media Access Control (MAC) layer level. Used in building-to-building wireless connections, wireless bridges can cover longer distances than access points (IEEE 802.11 standard specifies one mile as the maximum coverage range for an access point).
- Antenna—An antenna radiates the modulated signal through the air so that wireless clients can receive it. Characteristics of an antenna are defined by propagation pattern (directional versus omni directional), gain, transmit power, and so on. Antennas are needed on both the access point and bridge and the clients.

SAFE Wireless LAN Caveats and Axioms

This topic discusses SAFE WLAN caveats and axioms.

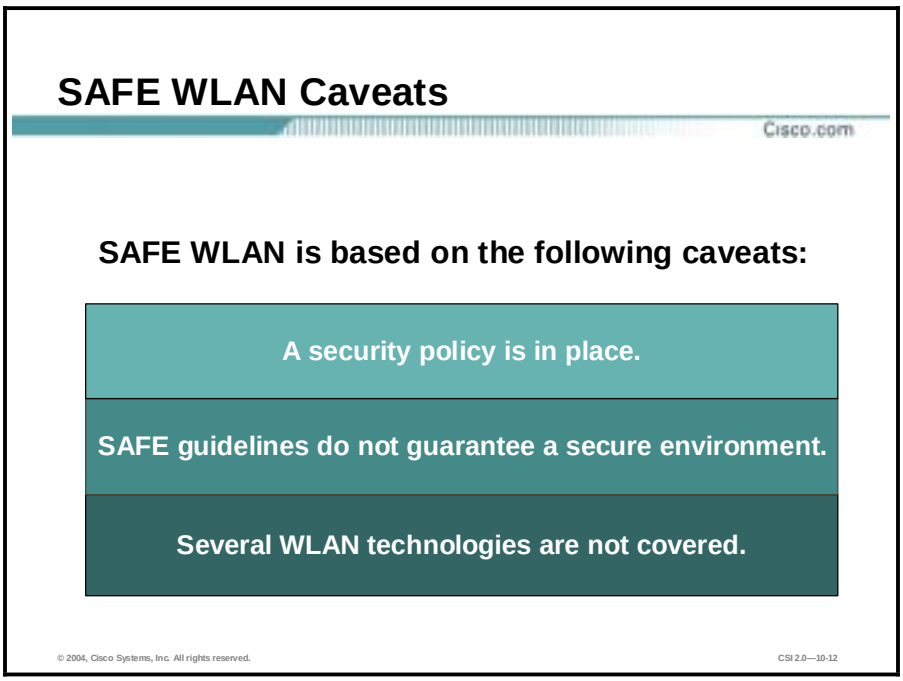

SAFE WLAN is based on the following caveats:

- A security policy is in place—Cisco does not recommend deploying WLANs without an associated security policy. Although network security fundamentals are mentioned in this document, they are not described in detail. Security within this document is always mentioned as it pertains to WLANs.
- SAFE WLAN guidelines do not guarantee a secure WLAN environment—Even though WLANs introduce security risks, many organizations choose to deploy WLANs because they bring user productivity gains and simplify deployment of small networks. Following the guidelines in this document does not guarantee a secure WLAN environment, nor does it guarantee that you will prevent all penetrations. By following the guidelines, you will mitigate WLAN security risks as much as possible.
- Several WLAN technologies are not covered—Though this lesson contains a large amount of detail on most aspects of wireless security, the discussion is not exhaustive. In particular, the document does not address wireless bridges, personal digital assistants (PDAs), or non-802.11-based WLAN technology. In addition, it does not provide specific best practices on general WLAN deployment and design issues that are not security related.

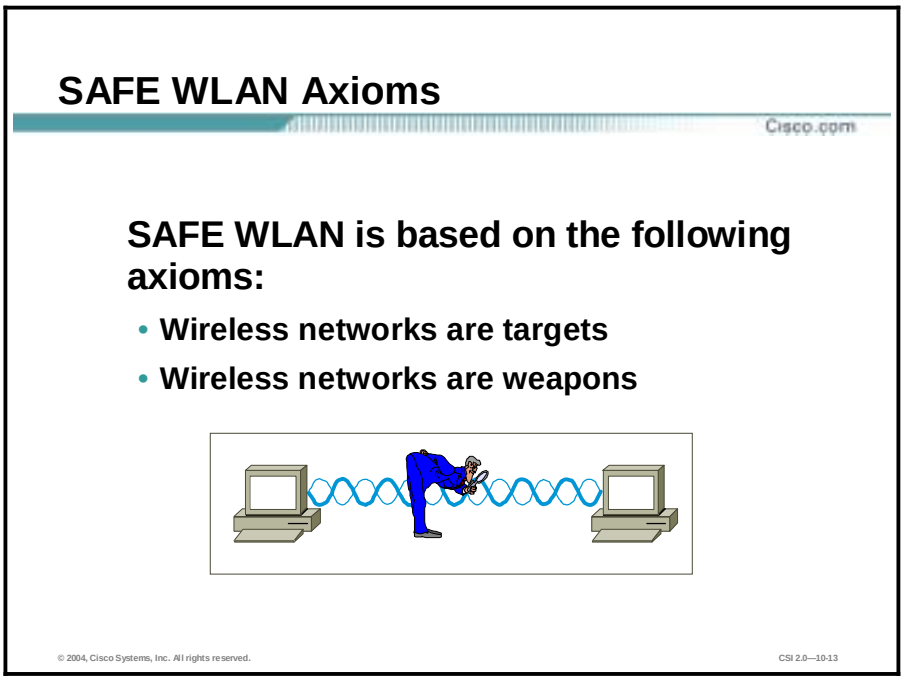

SAFE WLAN is based on the following axioms:

- Wireless networks are targets—Organizations today are deploying wireless technology at a rapid rate, often without considering all security aspects. Because WLAN devices ship with all security features disabled, increasing WLAN deployments have attracted the attention of the hacker community. Unlike a wired network, a WLAN sends data over the air and may be accessible outside the physical boundary of an organization. When WLAN data is not encrypted, the packets can be viewed by anyone within radio frequency range. Wireless networks are susceptible to the following weaknesses:
  - Interference and jamming
  - MAC authentication
  - Ad hoc versus infrastructure modes
  - Denial or degradation of service
- Wireless Networks Are Weapons—Rogue access points are major threats to wireless networks. A rogue access point is one that is accessible to an organization's employees but is not managed as a part of the trusted network. Most rogue access points are installed by employees for which IT is not providing WLAN access. A hacker, even one outside the physical boundaries of an organization's facilities, can gain access to the trusted network simply by associating with a rogue access point. Another type of rogue access point is one that masquerades as a trusted access point and tricks WLAN users into associating with it, thereby enabling a hacker to manipulate wireless frames as they cross the access point. The threat posed by rogue access points can be mitigated by preventing their deployment and detecting those rogue access points that are deployed. The following components are required in order to mitigate the threat of rogue access points:
  - Prevention
    - Corporate policy
    - Physical security

- Supported WLAN infrastructure
- 802.1x port-based security on edge switches

— Detection

- Using wireless analyzers or sniffers
- Using scripted tools on the wired infrastructure
- Physically observing WLAN access point placement and usage

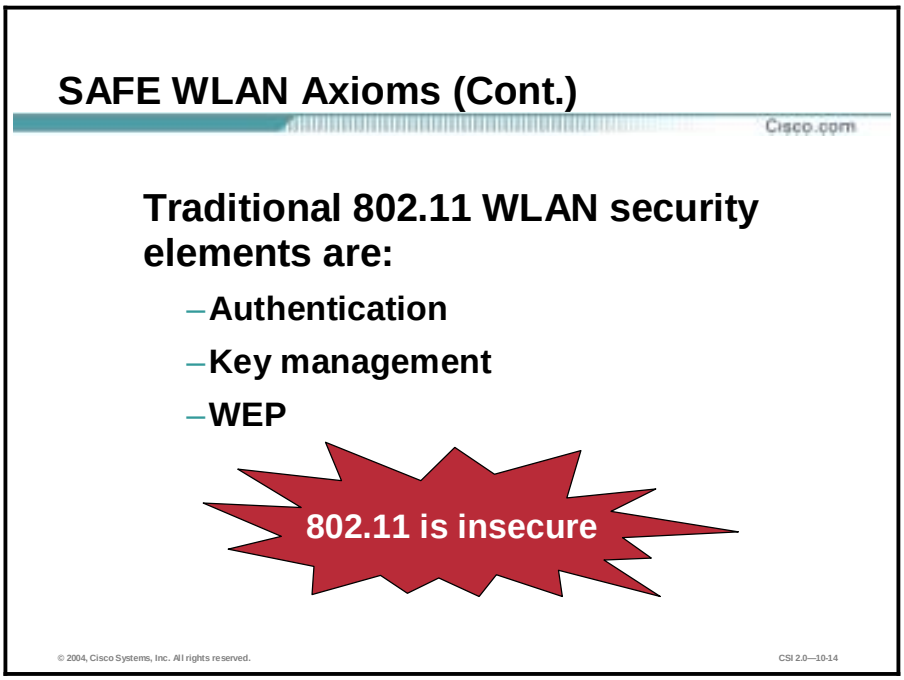

The 802.11 standard is insecure. The most widely deployed WLAN technologies today are 802.11b and 802.11a. Traditional 802.11 WLAN security includes the use of open or shared-key authentication and static WEP keys. This combination offers a rudimentary level of access control and privacy, but each element can be compromised. The following describes these elements and the challenges of their use in enterprise environments:

- Authentication—The 802.11 standard supports two means of client authentication: open and shared-key. Open authentication involves little more than supplying the correct service set ID (SSID).With open authentication, the use of WEP prevents the client from sending data to and receiving data from the access point unless the client has the correct WEP key. With shared-key authentication, the access point sends the client device a challenge text packet that the client must then encrypt with the correct WEP key and return to the access point. If the client has the wrong key or no key, authentication will fail and the client will not be allowed to associate with the access point. Shared-key authentication is not considered secure because a hacker who detects both the clear text challenge and the same challenge encrypted with a WEP key can decipher the WEP key.
- Key management—Another type of key that is often used but is not considered secure is a “static” WEP key. A static WEP key is a key composed of either 40 or 128 bits that is statically defined by the network administrator on the access point and on all clients that communicate with the access point. When static WEP keys are used, a network administrator must perform the time-consuming task of entering the same keys on every device in the WLAN. If a device that uses static WEP keys is lost or stolen, the possessor of the stolen device can access the WLAN. An administrator will not be able to detect that an unauthorized user has infiltrated the WLAN until the theft is reported. The administrator must then change the WEP key on every device that uses the same key used by the missing device. In a large enterprise WLAN with hundreds or even thousands of users, this can be a daunting task. Worse still, if a static WEP key is deciphered through a tool such as AirSnort, the administrator has no way of knowing that the key has been compromised by a hacker.

- Wired Equivalent Privacy (WEP)—The 802.11 standards define WEP as a simple mechanism to protect over-the-air transmission between WLAN access points and NICs. Working at the data link layer, WEP requires that all communicating parties share the same secret key. To avoid conflicting with U.S. export controls that were in effect at the time the standard was developed, 40-bit encryption keys were required by IEEE 802.11b, though many vendors now support the optional 128-bit standard. WEP can be easily cracked in both 40- and 128-bit variants by using off-the-shelf tools readily available on the Internet. On a busy network, 128-bit static WEP keys can be obtained in as little as 15 minutes, according to current estimates. Although traditional WLAN security that relies on open or shared keys and static WEP keys is better than no security at all, it is not sufficient for the enterprise organization. Only very small businesses, or those that do not entrust mission-critical data to their WLAN networks, can rely on these WLAN security types. All other enterprises and organizations must invest in a robust, enterprise-class WLAN security solution.

Wireless LAN Security Extensions

This topic describes the security extensions and protocols that help make WLANs more secure.

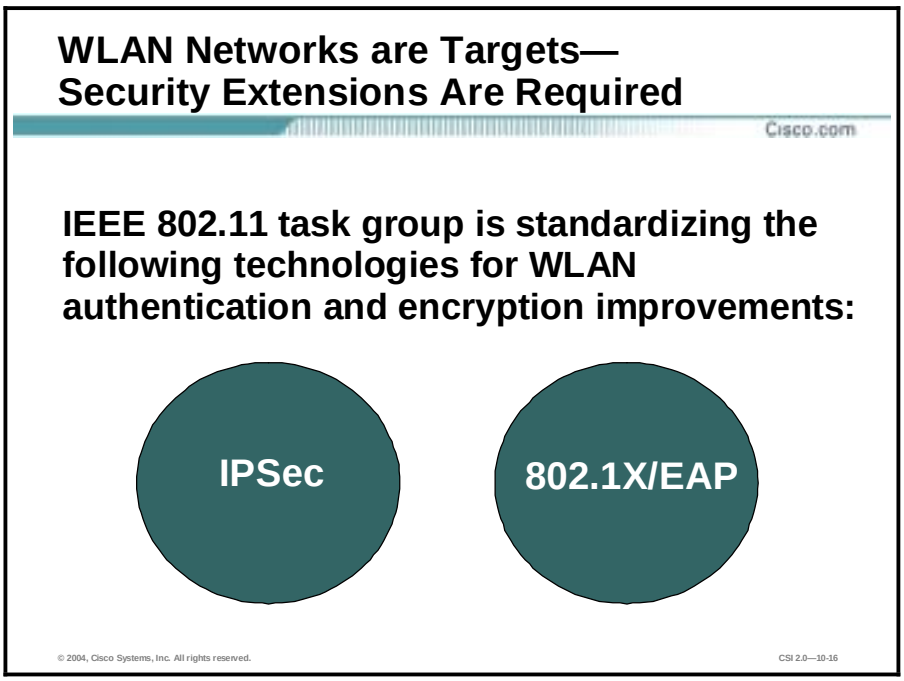

Cisco recommends deploying elements of three technologies as an alternative to WEP as specified by IEEE 802.11. These technologies include a network-layer encryption approach based on IP security (IPSec), a mutual authentication-based key distribution method using 802.1x, and some proprietary improvements to WEP recently implemented by Cisco, described as follows:

- IPSec—IPSec is a framework of open standards for ensuring secure private communications over IP networks. IPSec virtual private networks (VPNs) use the services defined within IPSec to ensure confidentiality, integrity, and authenticity of data communications across public networks such as the Internet. IPSec also has a practical application to secure WLANs by overlaying IPSec on top of clear text 802.11 wireless traffic. When deploying IPSec in a WLAN environment, an IPSec client is placed on every PC connected to the wireless network, and the user is required to establish an IPSec tunnel to route any traffic to the wired network. Filters are put in place to prevent any wireless traffic from reaching any destination other than the VPN gateway and Dynamic Host Configuration Protocol (DHCP) or Domain Name System (DNS) server. IPSec provides for confidentiality of IP traffic, as well as authentication and anti-replay capabilities. Confidentiality is achieved through encryption using a variant of the Data Encryption Standard (DES), called Triple-Data Encryption Standard (3DES), or the new Advanced Encryption Standard (AES). Though IPSec is used primarily for data confidentiality and device authentication, extensions to the standard allow for user authentication and authorization to occur as part of the IPSec process.
- 802.1x/Extensible Authentication Protocol (EAP)—An alternative WLAN security approach focusing on developing a framework for providing centralized authentication and dynamic key distribution. This approach is based on the IEEE 802.11i end-to-end framework using 802.1x and EAP to provide this enhanced functionality. Cisco has

incorporated 802.1x and EAP into its WLAN security solution: the Cisco Wireless Security Suite. The three main elements of an 802.1x and EAP approach are as follows:

- Mutual authentication between client and authentication RADIUS server
- Encryption keys dynamically derived after authentication
- Centralized policy control, where session time-out triggers re-authentication and new encryption key generation

When these features are implemented, a wireless client that associates with an access point cannot gain access to the network until the user performs a network logon. After association, the client and the network (access point or RADIUS server) exchange EAP messages to perform mutual authentication, with the client verifying the RADIUS server credentials, and vice versa. An EAP supplicant is used on the client machine to obtain the user credentials (user ID and password, user ID and one-time password [OTP], or digital certificate). Upon successful client and server mutual authentication, the RADIUS server and client then derive a client-specific WEP key to be used by the client for the current logon session. User passwords and session keys are never transmitted in the clear, over the wireless link.

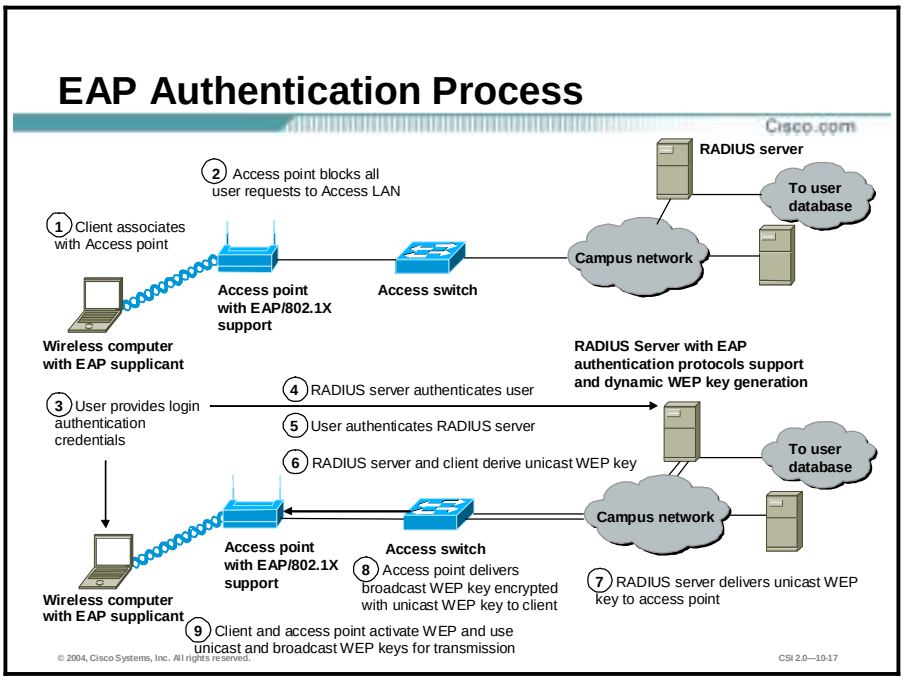

The sequence of events in the EAP authentication process is as follows:

Step 1 A wireless client associates with an access point.

Step 2 The access point blocks all attempts by the client to gain access to network resources until the client logs on to the network.

Step 3 The user on the client supplies network login credentials (user ID and password, user ID and OTP, or user ID and digital certificate) via an EAP supplicant.

Step 4 Using 802.1x and EAP, the wireless client and a RADIUS server on the wired LAN perform a mutual authentication through the access point in two phases. In the first phase of EAP authentication, the RADIUS server verifies the client credentials, or vice versa.

Step 5 In the second phase, mutual authentication is completed by the client verifying the RADIUS server credential, or vice versa.

Step 6 When mutual authentication is successfully completed, the RADIUS server and the client determine a WEP key that is distinct to the client. The client loads this key and prepares to use it for the logon session.

Step 7 The RADIUS server sends the WEP key, called a session key, over the wired LAN to the access point.

Step 8 The access point encrypts its broadcast key with the session key and sends the encrypted key to the client, which uses the session key to decrypt it.

Step 9 The client and access point activate WEP and use the session and broadcast WEP keys for all communications during the remainder of the session or until a time-out is reached and new WEP keys are generated.

Both the session key and broadcast key are changed at regular intervals. The RADIUS server at the end of EAP authentication specifies session key time-out to the access point and the broadcast key rotation time can be configured on the access point.

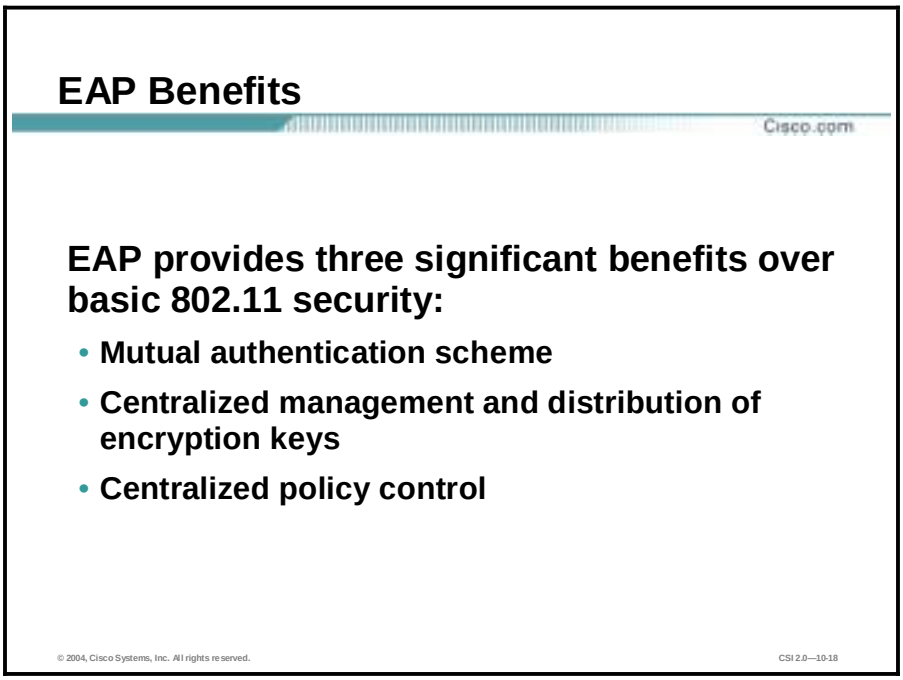

EAP provides three significant benefits over basic 802.11 security:

- Mutual authentication scheme—This scheme effectively eliminates man-in-the-middle attacks introduced by rogue access points and RADIUS servers.
- Centralized management and distribution of encryption keys—Even if the WEP implementation of RC4 had no flaws, there would still be the administrative difficulty of distributing static keys to all the access points and clients in the network. Each time a wireless device was lost, the network would need to be re-keyed to prevent the lost device from gaining unauthorized access.
- Ability to define centralized policy control—EAP allows for the ability to define centralized policy control when session time-out triggers re-authentication and new key derivation.

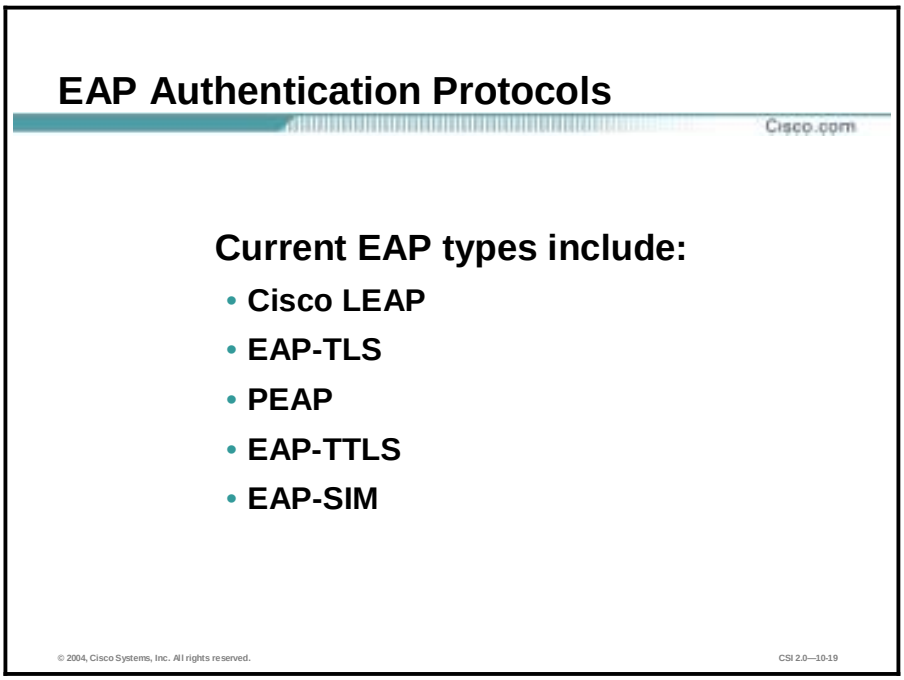

Numerous EAP types are available today for user authentication over wired and wireless networks. Current EAP types include:

- EAP-Cisco Wireless (Cisco LEAP)
- EAP-Transport Layer Security (EAP-TLS)
- Protected EAP (PEAP)
- EAP-Tunneled TLS (EAP-TTLS)
- EAP-Subscriber Identity Module (EAP-SIM)

In the Cisco SAFE wireless architecture, Cisco Light EAP (LEAP), EAP-TLS, and PEAP were tested and documented as viable mutual authentication EAP protocols for WLAN deployments.

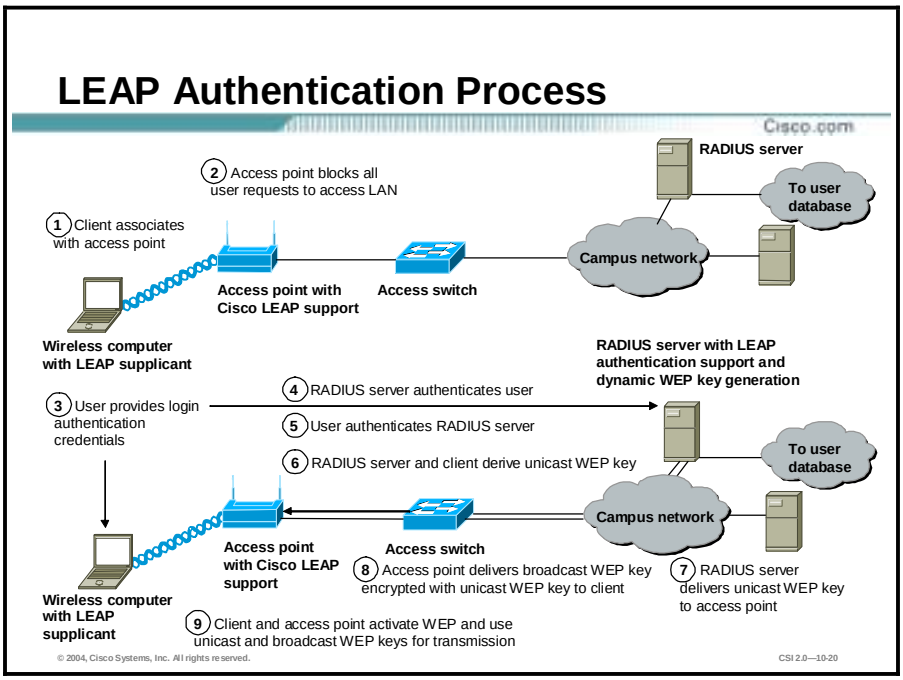

Cisco LEAP is the widely deployed EAP type in use today in WLANs. LEAP supports all three of the 802.1x and EAP elements mentioned previously. With LEAP, mutual authentication relies on a shared secret, the user's logon password, which is known by the client and the network.

As shown in the figure, the RADIUS server sends an authentication challenge to the client. The client uses a one-way hash of the user-supplied password to fashion a response to the challenge and sends that response to the RADIUS server. Using information from its user database, the RADIUS server creates its own response and compares that to the response from the client. When the RADIUS server authenticates the client, the process repeats in reverse, enabling the client to authenticate the RADIUS server. When this is complete, an EAP-Success message is sent to the client and both the client and the RADIUS server derive the dynamic WEP key.

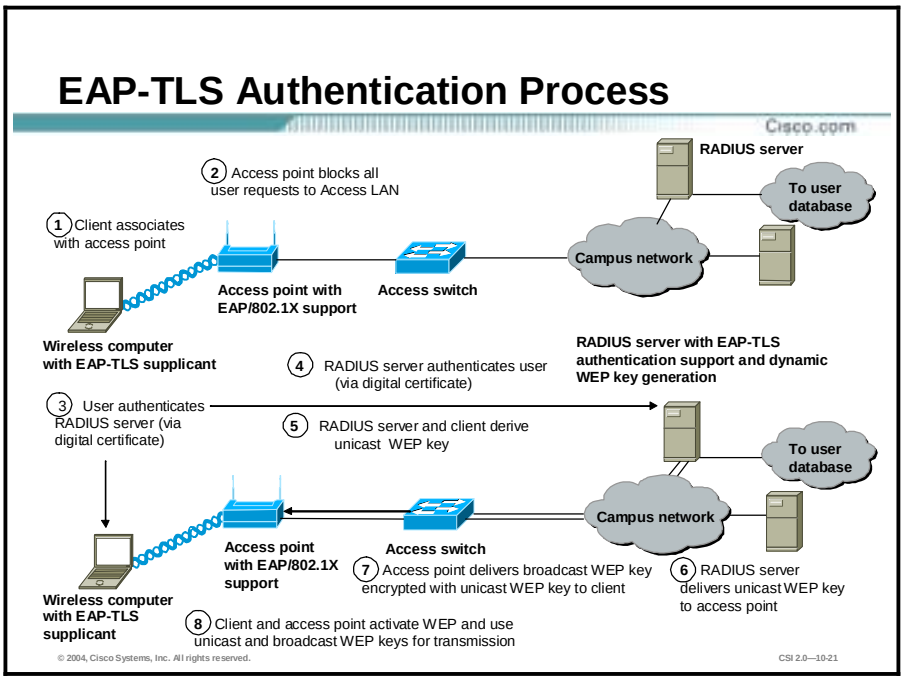

EAP-TLS is an Internet Engineering Task Force (IETF) standard (RFC 2716) that is based on the TLS protocol (RFC 2246). EAP-TLS uses digital certificates for both user and server authentication and supports the three key elements of 802.1x/EAP mentioned previously.

As shown in the figure, the RADIUS server sends its certificate to the client in phase 1 of the authentication sequence (server-side TLS). The client validates the RADIUS server certificate by verifying the issuer of the certificate, a certificate authority server entity, and the contents of the digital certificate. When this is complete, the client sends its certificate to the RADIUS server in phase 2 of the authentication sequence (client-side TLS). The RADIUS server validates the client's certificate by verifying the issuer of the certificate (certificate authority server entity) and the contents of the digital certificate. When this is complete, an EAP-Success message is sent to the client and both the client and the RADIUS server derive the dynamic WEP key.

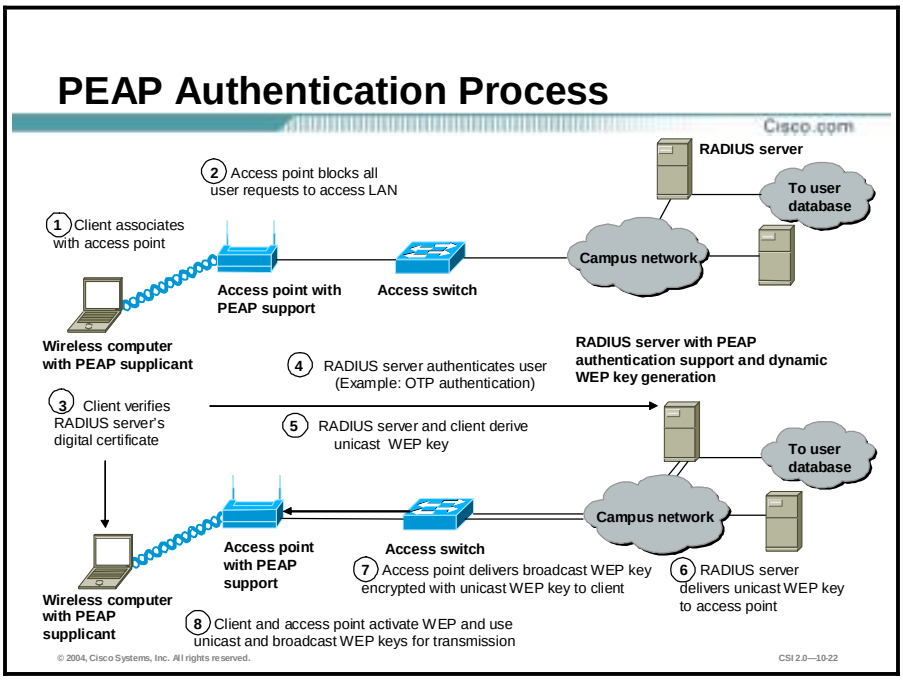

PEAP is an IETF draft RFC authored by Cisco Systems, Microsoft, and RSA Security. PEAP uses a digital certificate for server authentication. For user authentication, PEAP supports various EAP-encapsulated methods within a protected TLS tunnel. PEAP supports the three main elements of 802.1x/EAP, as mentioned previously.

As shown in the figure, phase 1 of the authentication sequence is the same as that for EAP-TLS (server-side TLS). At the end of phase 1, an encrypted TLS tunnel is created between the user and the RADIUS server for transporting EAP authentication messages. In phase 2, the RADIUS server authenticates the client through the encrypted TLS tunnel via another EAP type. As an example, a user can be authenticated using an OTP in the Extensible Authentication Protocol-Generic Token Card (EAP-GTC) subtype (as defined by the PEAP draft). In this case, the RADIUS server will relay the OTP credentials (user ID and OTP) to an OTP server to validate the user login. When this is complete, an EAP Success message is sent to the client and both the client and the RADIUS server derive the dynamic WEP key.

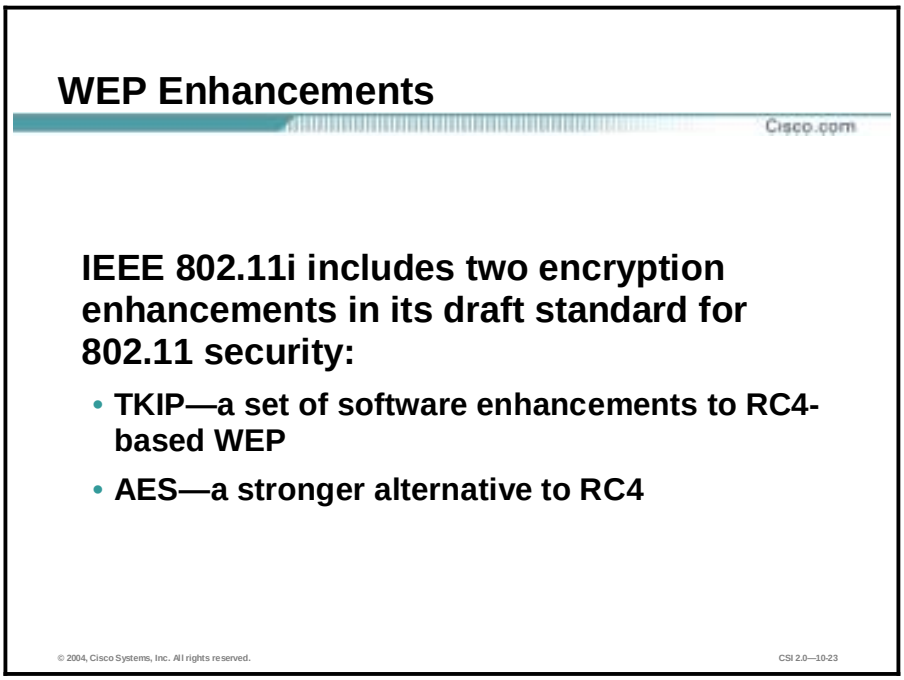

Cisco introduced support for Temporal Key Integrity Protocol (TKIP) as a component of the Cisco Wireless Security Suite with 802.11i standard. Cisco TKIP and the Wi-Fi Protected Access (WPA) TKIP include per-packet keying (PPK) and Message Integrity Check (MIC). WPA TKIP introduces a third element: extension of the initialization vector from 24 bits to 48 bits. IEEE 802.11i includes the following two encryption enhancements in its draft standard for 802.11 security:

- TKIP—A set of software enhancements to RC4-based WEP. Cisco TKIP improvements include the following:
  - Per-packet keying—Because the most popular attack against WEP relies on exploiting multiple weak initialization vectors in a stream of encrypted traffic using the same key, using different keys per packet is a potential way to mitigate the threat. The initialization vector and WEP key are hashed to produce a unique packet key (called a temporal key), which is then combined with the initialization vector and run through a mathematical function called XOR with the plain text. This scenario prevents the weak initialization vectors from being used to derive the base WEP key, because the weak initialization vectors allow you to derive only the per-packet WEP key. In order to prevent attacks due to initialization-vector collisions, the base key should be changed before the initialization vectors repeat. Because initialization vectors on a busy network can repeat in a matter of hours, mechanisms like EAP authentication protocols should be used to perform the re-key operation. Similar to a unicast key, the WLAN broadcast key is susceptible to attacks due to initialization vector collisions. Cisco access points support broadcast key rotation to mitigate this vulnerability. The access point dynamically calculates the broadcast WEP key and the new broadcast WEP key is delivered to clients using EAP over LAN (EAPOL)-Key messages. Thus, broadcast WEP key rotation can be enabled only with EAP protocols such as LEAP, EAP-TLS, and PEAP that support dynamic derivation of encryption keys.
  - MIC—Another concern with WEP is its vulnerability to replay attacks. The MIC protects WEP frames from tampering. The MIC is based on a seed value, destination

MAC, source MAC, and payload (that is, any changes to these will affect the MIC value). The MIC is included in the WEP-encrypted payload and uses a hashing algorithm to derive the resulting value. This is an improvement of the cyclic redundancy check (CRC)-32 checksum function as performed by standards-based WEP. With CRC-32, it is possible to compute the bit difference of two CRCs based on the bit difference of the messages over which they are taken. In other words, flipping bit *n* in the message results in a deterministic set of bits in the CRC that must be flipped to produce a correct checksum on the modified message. Because flipping bits carries through after an RC4 decryption, this allows the attacker to flip arbitrary bits in an encrypted message and correctly adjust the checksum so that the resulting message appears valid.

- AES—Advanced Encryption Standard is a stronger alternative to WEP/RC4. AES is a privacy transform for IPSec and Internet Key Exchange (IKE) and has been developed to replace the DES. AES is designed to be more secure than DES. AES offers a larger key size, while ensuring that the only known approach to decrypt a message is for an intruder to try every possible key. AES has a variable key length—the algorithm can specify a 128-bit key (the default), a 192-bit key, or a 256-bit key. Although AES is being developed to replace DES, the National Institute of Standards and Technology (NIST) anticipates that 3DES will remain an approved algorithm for the near future. This feature, which adds support for AES encryption to IPSec, introduces a new level of security strength and speed that was not present in the VPN marketplace.

Note The National Institute of Standards and Technology (NIST) has created AES, which is a new Federal Information Processing Standard (FIPS) publication that describes an encryption method.

Note Organizations should choose to deploy either IPSec or 802.1x/EAP with TKIP or Cisco TKIP, but generally not both. Organizations should use IPSec when they have the utmost concern for the sensitivity of the transported data, but remember that this solution is more complex to deploy and manage than 802.1x/EAP with TKIP. The 802.1x/EAP with TKIP should be used when an organization wants reasonable assurance of confidentiality and a transparent user security experience. The basic WEP enhancements can be used anywhere WEP is implemented. For the vast majority of networks, the security provided by 802.1x/EAP with TKIP is sufficient.

Cisco Wireless LAN Product Portfolio

This topic describes the Cisco WLAN product line.

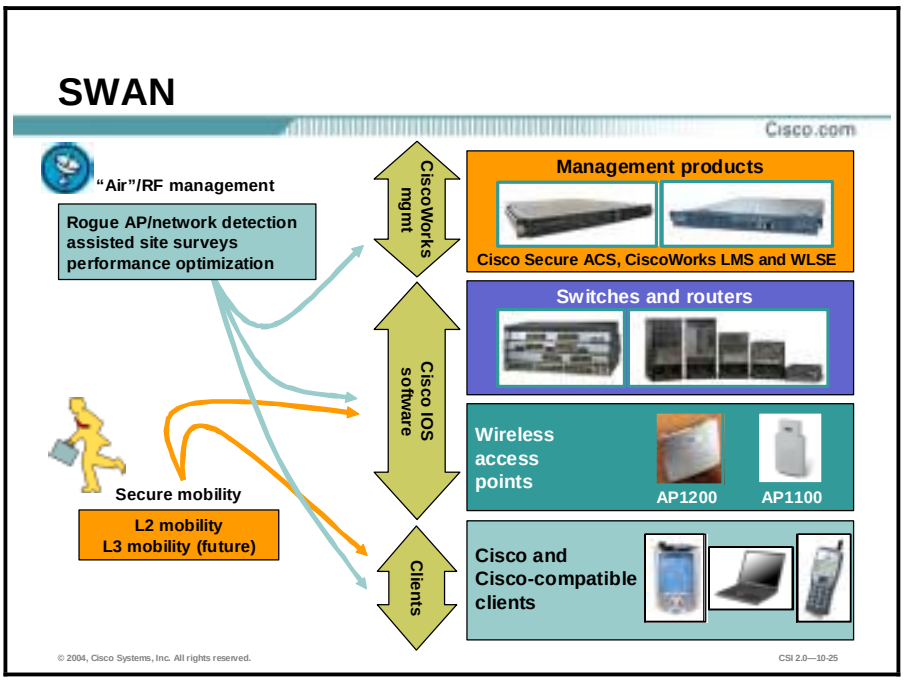

The Cisco Structured Wireless-Aware Network (SWAN) combines the Cisco switch and router infrastructure with the Cisco wireless network, making one integrated "wireless-aware" network. This secure, integrated wired and wireless network extends Cisco's proven LAN infrastructure capabilities to the WLAN, providing the security, management, ease-of-deployment, scalability, and reliability that our enterprise customers depend on for their core business applications.

Cisco SWAN infrastructure enhancements will be integrated in Cisco Aironet 1100 and 1200 Series access points, Cisco Catalyst 3750, 4500, and 6500 Series switches, and Cisco 2600XM and 3700 Series routers. Other components of the solution include CiscoWorks Wireless LAN Solution Engine (WLSE) for management and monitoring, Cisco IOS software, Cisco Secure Access Control Server for centralized authentication, and Cisco and Cisco-compatible client adapters for radio frequency (RF) monitoring and measurement.

Key SWAN Components are:

- Cisco Aironet Series WLAN access points
- Cisco Aironet Series WLAN client adapters
- Cisco-compatible client adapters
- Cisco Wireless Security Suite
- Cisco IOS Software
- Cisco Secure Access Control Server (ACS)
- CiscoWorks Wireless LAN Solution Engine version 2.0 and 2.5

The Cisco Aironet product family is available in a variety of forms to fit almost any application, and provides a complete solution to customers who require the mobility and flexibility of a WLAN to complement or replace a wired LAN. The products seamlessly integrate into wired Ethernet networks and fully comply with the IEEE standards. Cisco wireless solutions include the following:

- 802.11a and 802.11b access points
- 802.11a and 802.11b client adapters
- 802.11a and 802.11b wireless bridges
- 802.11b workgroup bridge

Note	The Cisco Aironet 802.11g Access Point Upgrade Kit allows customers to take advantage of the new 802.11g technology, with its higher 54-Mbps throughput and its 802.11b backward capability, while continuing to use their existing Cisco Aironet WLAN infrastructures. Because 802.11g and 802.11b operate in the same 2.4-GHz unlicensed band, migrating to 802.11g is an affordable choice for customers with existing Cisco Aironet 1100 Series and 1200 Series Access Points.

Cisco Aironet 1200 Series Access Point—Dual-Band

Cisco.com

Features include:

- **Compliant with IEEE 802.11a, 802.11b, and 802.11g standards**
- **Field-upgradable radio and software**
- **Rugged design and plenum rated**
- **Multiple mounting options: wall, ceiling, and desktop**
- **Secure with cable lock or padlock**
- **Inline and local power**
- **High performance**

CSI 2.0—10-27

Cisco Aironet 1200 Series access points have the following features:

- Secure, manageable, and reliable wireless connectivity—With simultaneous support for both 2.4 GHz and 5 GHz radios, the Cisco Aironet 1200 Series preserves existing IEEE 802.11b features while providing a migration path to the faster IEEE 802.11a standard and future IEEE 802.11g products.
- Modular design—Supports single and dual-band configurations, plus the field upgradeability to change these configurations as requirements change and technologies evolve.
- 802.11a support—The radio supports data rates of up to 54 Mbps and eight non-overlapping channels that offer high performance as well as maximum capacity and scalability.
- 802.11b and 802.11g support—The radio provides data rates up to 11 Mbps and three non-overlapping channels to support widely deployed 802.11b clients. The mini-PCI form factor of the 802.11b radio allows for upgrade to higher-speed 2.4 GHz technologies such as the IEEE 802.11g standard.
- Return on Investment—The Cisco Aironet 1200 Series platform offers investment protection through field-upgradeable Card Bus and mini-PCI radios. The Card Bus-based 802.11a modules can easily be fitted into installed Cisco Aironet 1200 Series access points.
- Cisco Aironet 1200 Series access point is UL 2403 plenum rated—This means the access point does not give out toxic smoke when burned. Thus, the access point can be installed above dropped ceilings.

Cisco Aironet 1100 Series Access Point—Cost-Effective, Enterprise-Class

Cisco.com

Features include:

- **Available in an IEEE 802.11g version or IEEE 802.11b version that is field upgradeable to 802.11g**
  - **Cisco IOS operating system**
  - **Runs latest Cisco IOS Release 12.2(11)JA code**
  - **WPA, Fast Secure Roaming, etc.**
- **Variety of mounting configurations**
- **Integrated antenna**
- **Plenum-rated plastic housing**
- **Inline Power-over-Ethernet**

CSI 2.0—10-28

Cisco Aironet 1100 Series access points offer following features:

- Affordable, intelligent, and upgradeable—The 2.4 GHz WLAN solution delivers enterprise-class security and manageability. Equipped for the IEEE 802.11b standard, a field-upgradeable design ensures a smooth migration to the IEEE 802.11g standard. The compact size, integrated antenna, and innovative bracket design of this Cisco IOS-based access point allow for quick, easy installation in a variety of orientations.
- Enterprise features—The 1100 Series Cisco IOS operating system includes features like VLAN, QoS, and proxy mobile IP that extend end-to-end intelligent networking to the enterprise WLAN.
- Security features—The Cisco Aironet 1100 Series offers the Cisco Wireless Security Suite of security solutions.
- Easy to use—The Cisco Aironet 1100 Series uses integrated antennas for installation in any orientation. The flexible mounting system allows the customer to install this in nearly any location. Cisco command-line interface (CLI) and the redesigned web GUI allow network managers to quickly set up access points.
- Return on Investment—The Cisco Aironet 1100 Series is engineered with extra system capacity, including memory, storage, and processing power. Companies can also upgrade the Cisco Aironet 1100 Series hardware to future radio technologies, such as those based on higher speed 802.11g specifications, unleashing increased performance and further investment value.

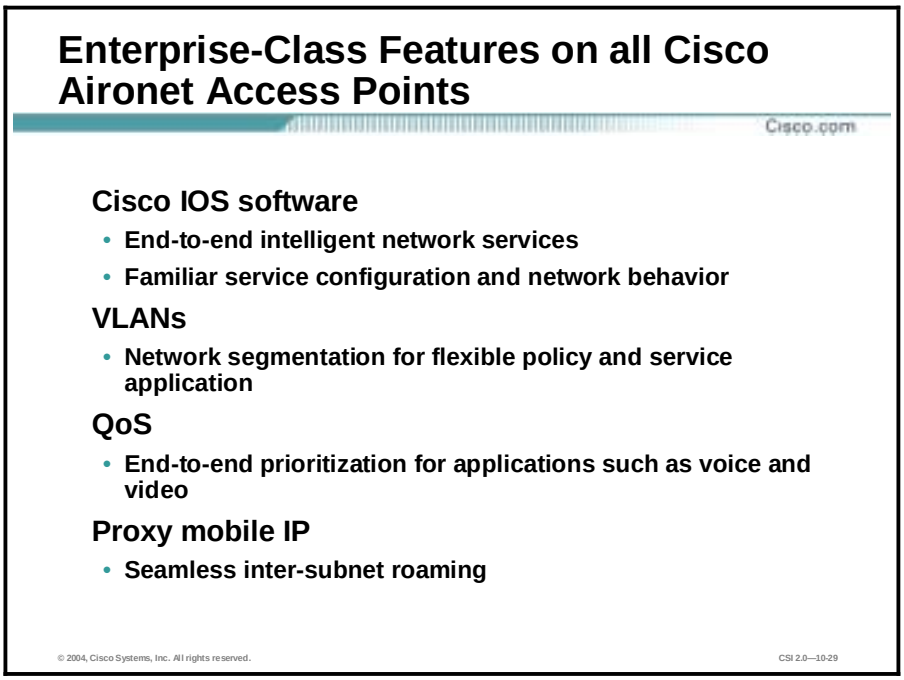

The enterprise-class features on all Cisco Aironet access points include the following:

- Cisco IOS software—Cisco Aironet products leverage the same Cisco IOS software that powers Cisco switches and routers, enabling customers to extend common services, management tools, and interfaces across their wired and wireless networks. Another key component of the Cisco Enterprise WLAN solution is the use of intelligent network services, including QoS, VLANs, and proxy mobile IP. With these features, the WLAN becomes an extension of the wired network because it offers much of the same functionality and management tools. These features can also enable new applications, such as video and voice, over wireless.
- VLANs—The Cisco Aironet family is capable of managing up to 16 VLANs per access point. This allows customers to vary WLAN policies and services, such as security and QoS, to accommodate different types of users and applications. For example, enterprise customers can use different wireless VLANs to separate employee traffic from guest traffic. VLANs can also be deployed in support of certain applications, such as time-sensitive voice traffic, to ensure peak performance.
- QoS—With support for IEEE 802.1p QoS, the Cisco Enterprise WLAN solution provides true end-to-end traffic prioritization. This feature is an essential building block for real-time applications and integrates seamlessly with the existing QoS features found in Cisco routers and switches. 802.1p QoS allows time-sensitive traffic, such as voice and video, to be prioritized over less critical packets, such as e-mail data, for an improved user experience and optimal bandwidth utilization. The Cisco Aironet family already supports voice prioritization schemes for 802.11b wireless phones, and for maximum investment protection, field upgrades will offer the ability to migrate to emerging QoS standards such as 802.11e.
- Proxy Mobile IP—With proxy mobile IP, users can roam from one wireless access point to another while maintaining seamless network connectivity. This is achieved by transporting individual IP addresses as users move from one subnet to the next. As a result, IT managers can partition the WLAN into distinct, easily managed segments without affecting user

mobility. This more closely mirrors the architecture of the wired network, yet the boundaries are transparent from the user's perspective.

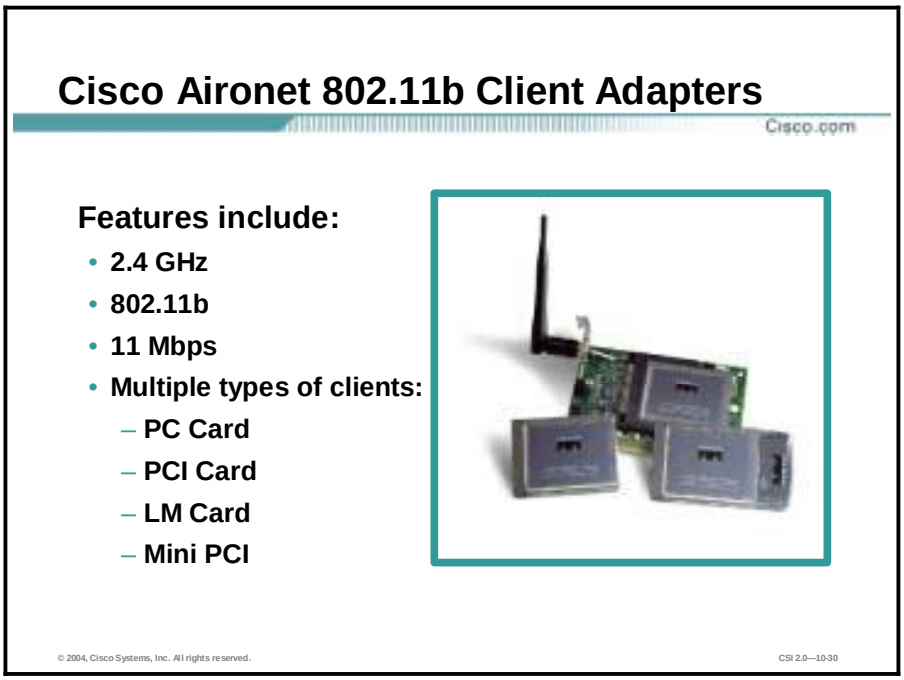

Cisco Aironet 350 Series client adapters complement Cisco Aironet 1200 Series and 1100 Series access points. Available in Personal Computer Memory Card International Association (PCMCIA) and PCI form factors, these 802.11b-compliant client adapters allow 11 Mbps of throughput. They quickly connect desktop and mobile computing devices to the WLAN.

Cisco Aironet offers a variety of clients for the 802.11b WLAN, including the following:

- PC Card—Standard PCMCIA product with attached end-cap antenna. Supports both 40 and 128-bit encryption.
- LM Card—PCMCIA card with MMCX connectors. This card allows the attachment of any of the Cisco antennas using the MMCX RP-TNC adapter cables. An end-cap antenna is offered for use with the LM card, giving it the same functionality as the PC card. The LM card is shipped without an antenna, and supported with both 40 and 128-bit encryption.
- PCI card—Standard PCI card typically used for desktop clients. The PCI card has the standard Cisco Aironet RP-TNC connector and can be used with all of the Cisco external antennas. The PCI card is shipped without an antenna, but is supported with both 40 and 128-bit encryption.

Note	All Cisco Aironet client devices support IEEE 802.1x security including LEAP and EAP-TLS, for mutual authentication and dynamic per-user, per-session WEP keys. All Cisco Aironet client devices are capable of load balancing when used with Cisco Aironet access points.

Cisco Aironet 802.11a Client Adapter

Cisco.com

Features include:

- **5 GHz/802.11a**
  - **54 Mbps**
- **Rate shifting channels:**
  - **6, 9, 12, 18, 24, 36, 48, or 54**
- **5 dBi patch antenna**
- **CardBus interface**
- **Transmit power settings:**
  - **20 mW, 10 mW, and 5 mW**

CSI 2.0—10-31

Cisco Aironet 5 GHz 54 Mbps Wireless LAN Client Adapters complement the Cisco Aironet 1200 Series access point, providing a solution that combines mobility with the security and manageability required by businesses. The features include the following:

- The 802.11a-compliant CardBus adapter operates in the UNII-1 and UNII-2 bands to allow 54 Mbps throughput.
- Rate shifting channels—The Cisco innovative radio and antenna design delivers industry-leading 802.11a enterprise performance. It provides for maximum capacity and scalability across the enterprise through eight nonoverlapping channels in the UNII-1 and UNII-2 bands.
- Integrated 5 dBi gain patch antenna—Optimizes range. The 68 dBm receive sensitivity at 54 Mbps provides high-data-rate range performance. Advanced signal processing in the Cisco Aironet 5 GHz Wireless LAN Client Adapter helps manage the multi-path propagation often found in office environments, and intelligent filtering addresses ambient noise and interference that can decrease network performance.
- CardBus Interface—IEEE 802.11a-compliant adapter operates in the UNII-1 and UNII-2 bands. The CardBus 32-bit interface handles 400-6000 Mbps.
- Transmit power settings—Various transmit power settings on the Cisco Aironet client adapter enable you to select range capabilities.

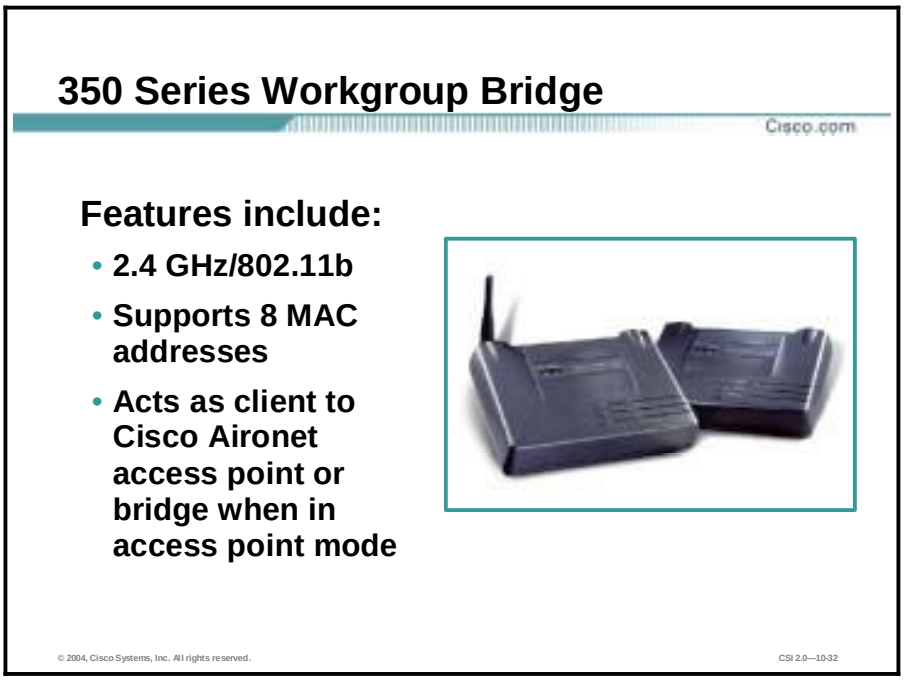

Cisco Aironet 350 Wireless Workgroup Bridges meet the needs of remote workgroups, satellite offices, and mobile users by bringing the freedom and flexibility of wireless connectivity to any Ethernet-enabled device. Features include:

- 2.4 Ghz and 802.11b support
- Supports 8 MAC addresses—The workgroup bridge quickly connects up to eight Ethernet-enabled devices, such as laptops, personal computers, and printers, to a WLAN, providing the link from these devices to any Cisco Aironet access point or wireless bridge.
- Access point mode—The Cisco Aironet Workgroup Bridge product connects to the Ethernet port of a device that does not have a PCI or PCMCIA slot available or to a hub with up to eight wired devices attached. It provides a wireless connection into an access point for up to eight MAC addresses, and onto the LAN backbone. It cannot be used in a peer-to-peer mode connection, and must communicate to a Cisco Aironet access point or to a Cisco Aironet bridge in access point mode.

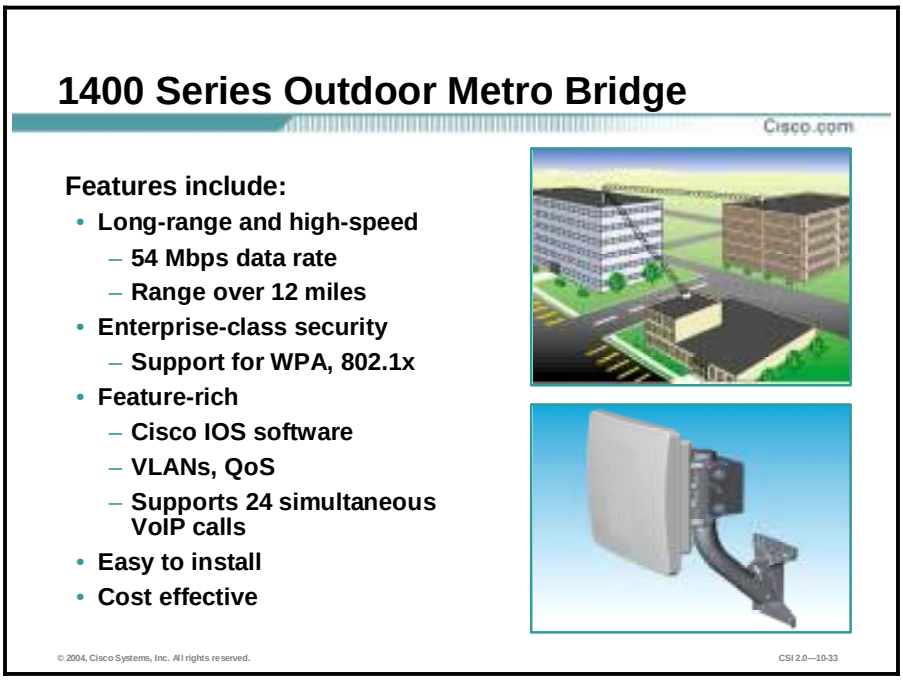

The Cisco Aironet® 1400 Series Wireless Bridge creates a new benchmark for wireless bridging by providing a high-performance and feature-rich solution for connecting multiple LANs in a metro area. Designed to be a cost-effective alternative to leased lines, it is engineered specifically for harsh outdoor environments. Key features include the following:

- Provides 54 Mbps data rate throughput and has a range of over 12 miles
- Support for both point-to-point or point-to-multipoint configurations
- Enhanced security mechanisms based on 802.11 standards
- Supports Cisco IOS software, VLANs, QoS, and supports 24 simultaneous VoIP calls
- Rugged enclosure optimized for harsh outdoor environments with extended operating temperature range
- Integrated or optional external antennas for flexibility in deployment
- Designed specifically for ease of installation and operation

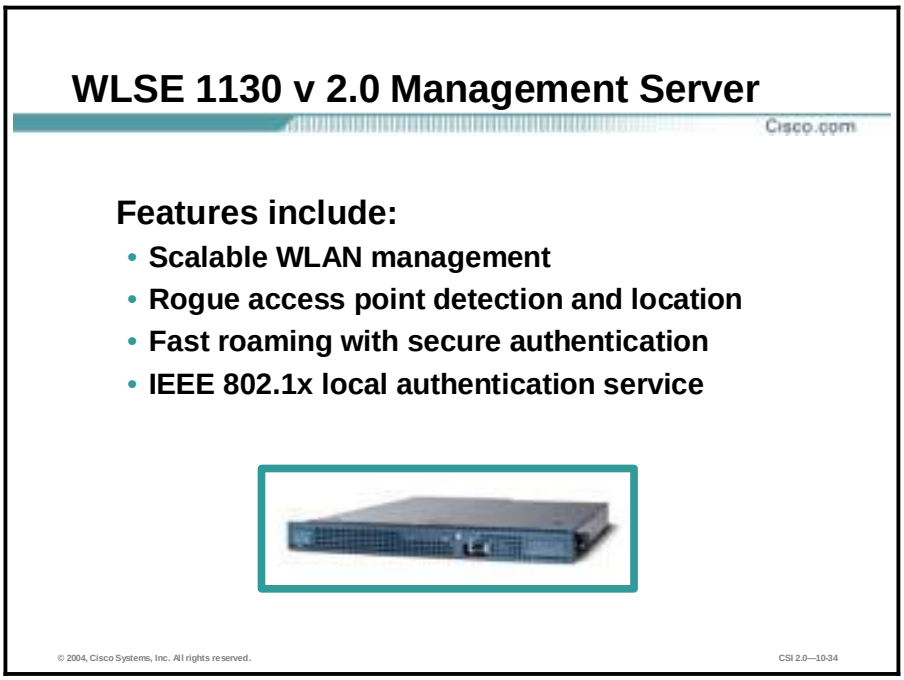

The Wireless LAN Solution Engine (WLSE) provides intuitive and comprehensive reports for troubleshooting and capacity planning. It assists in pinpointing problems with utilizations and client associations to help maximize network uptime. It also allows client tracking, WLAN performance reports, and fault monitoring to make network troubleshooting easier. Management and operations features include the following:

- Scalable WLAN management—Managing up to 2500 access points is as easy as managing a few access points.
- Rogue access point detection and location—The Cisco Structured Wireless-Aware Network detects, isolates, and mitigates rogue access points. These rogue access points create an unsecured WLAN connection that puts the entire wired network at risk.
- Fast secure roaming—Fast secure roaming allows authenticated client devices to roam securely from one access point to another without any perceptible delay during re-association. Fast secure roaming supports latency-sensitive applications such as wireless VoIP, enterprise resource planning (ERP), or Citrix-based solutions, Without Dropping Service (WDS) during roaming. WDS provides fast, secure handoff services to access points for under 150ms roaming within a subnet. Cisco fast secure roaming requires Cisco or Cisco-compatible client devices that support the Cisco Centralized Key Management (CCKM) protocol.
- Fast secure roaming across subnets (Layer 3 mobility)—For deployments that require mobility across subnet boundaries, Cisco is developing an easy-to-configure, scalable solution that builds on the solid foundation of the Cisco Structured Wireless-Aware Network. This solution will be supported across Cisco devices, including Cisco Aironet access points, and specific Cisco Catalyst LAN switches and Cisco routers.
- Centralized management—Facilitates mass conversion of Cisco Aironet 1200 Series access point VxWorks configuration files into Cisco IOS software configuration files using an expanded version of the Cisco Aironet Conversion Tool for Cisco IOS software.
- IEEE 802.1x local authentication service—With IEEE 802.1x local authentication service, Cisco Aironet access points are configured to act as a local RADIUS server to authenticate

wireless clients when the authentication, authorization, and accounting (AAA) server is not available. This provides authentication services for remote or branch office WLANs without RADIUS server and backup authentication services during WAN link or server failure to provide access to local resources like file servers or printers. IEEE 802.1x local authentication can support the authentication of up to 50 accounts for a given deployment in the local Cisco LEAP authentication database on the access point. One account is equal to one user name and password. The configuration of the IEEE 802.1x local authentication database can be centrally managed with the CiscoWorks WLSE 2.x management appliance.

Note The access point with the IEEE 802.1x local authentication service does not need to be dedicated to the IEEE 802.1x local authentication service. This access point can function as a regular access point in addition to providing IEEE 802.1x local authentication.

Note Cisco Aironet 350 or 1200 Series Access Points are non-IOS (VxWorks) access points and their configuration to Cisco IOS operation requires conversion. The conversion tool is a special utility that is required by administrators to create a Cisco IOS configuration using the configuration of an existing non-IOS 350 or 1200 Series Access Point.

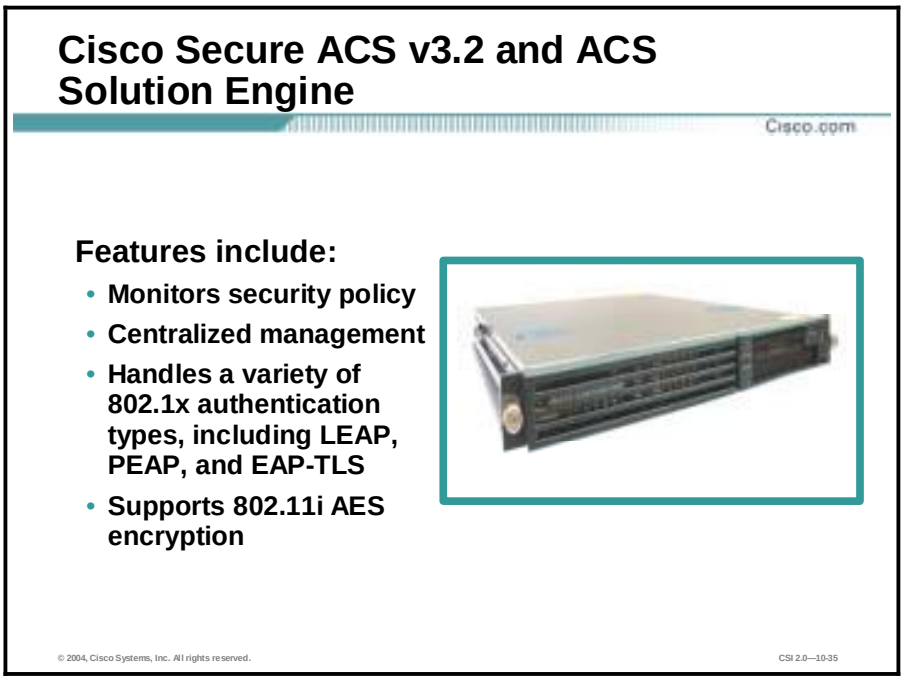

A wide selection of RADIUS servers can be used for scalable, centralized user management. When Cisco Aironet access points are used in conjunction with Cisco Catalyst intelligent switches, many of the security features of the wired LAN can be extended to the WLAN, further protecting organizations against internal and external threats.

In addition to providing IEEE 802.1x local authentication service, the Cisco Structured Wireless-Aware Network provides extensive security management features that use the Cisco Wireless Security Suite, including the following:

- Monitors security policy—Monitors security policies for predefined Cisco Wireless Security Suite parameters across all access points. Alerts are generated for violations in areas such as service set IDs (SSIDs), broadcasts, 802.1x EAP settings, and WEP. Alerts can be delivered via e-mail, Syslog, or Simple Network Management Protocol (SNMP) trap notifications.
- Centralized management of security settings—Parameters such as 802.1x EAP, WEP, and WPA are ensured through centralized WLAN management of all local and remote access point settings. Notification of 802.1x EAP RADIUS or AAA server management thresholds are managed via e-mail, Syslog, and SNMP trap notifications.
- Monitoring of the 802.1x EAP RADIUS or AAA servers—The RADIUS or AAA server providing support for Cisco LEAP and PEAP is monitored and the availability of Cisco Secure ACS and Committed Access Rate (CAR) EAP servers is verified.
- IEEE 802.11i AES encryption support—Future support for IEEE 802.11i AES encryption is planned.

Cisco Compatible Extensions Program wireless networking solutions allow customers to do the following:

- Implement cost-effective, scalable, and secure wireless applications
- Improve service productivity and accuracy with an integrated wireless architecture that supports voice or data information available at the point of delivery
- Manage enterprise WLAN deployments and operations to ensure network reliability, security, and availability

Wireless LAN Design Approach

This topic describes the SAFE WLAN design approach.

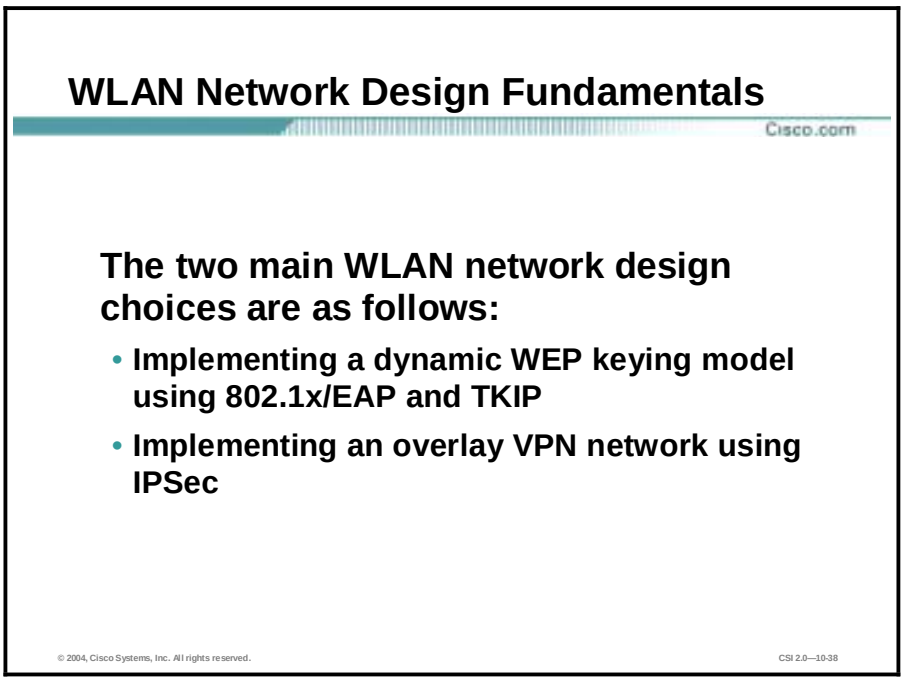

The size and security concerns of the specific WLAN design dictate the mitigation techniques that are applied to a WLAN design. Therefore, the network designer is offered a choice of the mitigation technology to implement, considering the advantages and disadvantages of the technologies specific to the SAFE design. The mitigation technologies are consistent across all the SAFE designs, so a review of the networking elements of each of the two main technology choices is presented in the following. The two main design choices are as follows:

- Implementing a dynamic WEP keying model using 802.1x/EAP and TKIP
- Implementing an overlay VPN network using IPSec

Standard WLAN Design

This topic describes the WAN design.

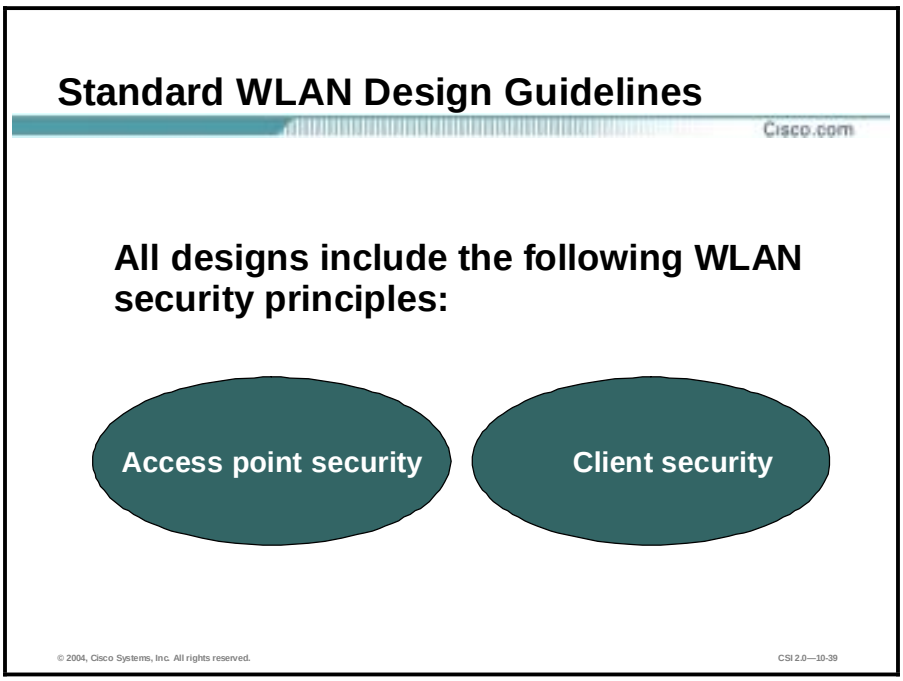

This topic outlines the generic elements of WLAN designs because so many of them are common throughout the SAFE designs. The basic concepts can be included with specific variances and alternatives discussed in the specific SAFE design. In the standard WLAN designs, it is assumed that all WLAN devices are connected to a unique IP subnet to enable end-user mobility throughout various designs. An assumption is made in the designs that most services available to the wired network are also available to the wireless network addition. All designs include the following WLAN security principles:

- Access point security recommendations:
  - Enable centralized user authentication (RADIUS, Terminal Access Controller Access Control System Plus [TACACS+]) for the management interface.
  - Choose strong community strings for SNMP and change them often.
  - Consider using SNMP Read Only if your management infrastructure allows it.
  - Disable any insecure and nonessential management protocol provided by the manufacturer.
  - Use secure management protocols, such as Secure Shell (SSH).
  - Limit management traffic to a dedicated wired subnet.
  - Isolate management traffic from user traffic and encrypt all management traffic where possible.
  - Enable wireless frame encryption where available.
  - Physically secure the access point.

- Client security recommendations:
  - Disable ad hoc mode.
  - Enable wireless frame encryption where available.

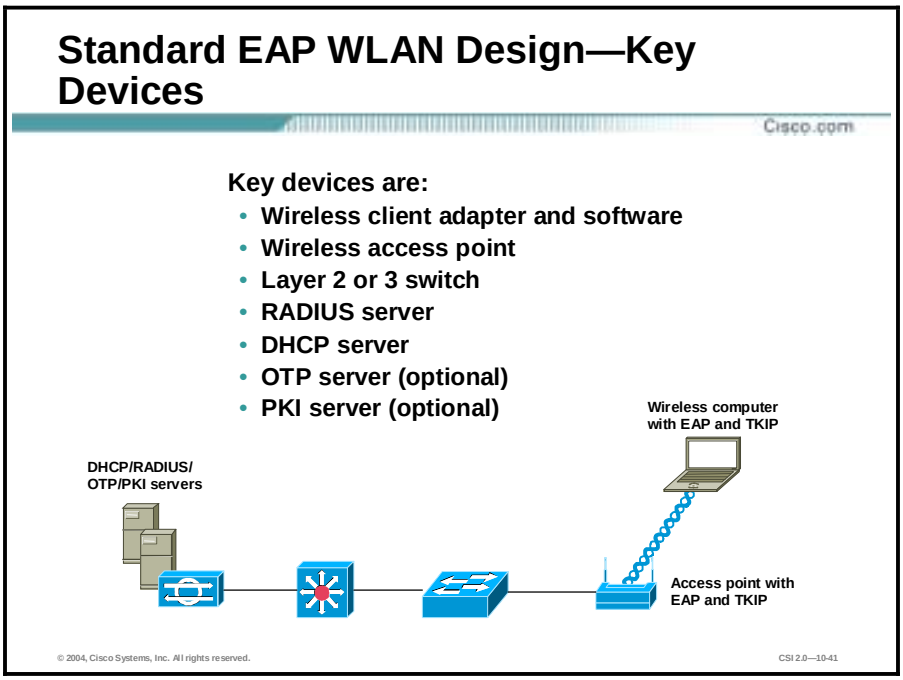

This design details a generic method for using EAP with TKIP as a security mechanism to access the production corporate network. The following are the key EAP devices:

- Wireless client adapter and software—A software solution that provides the hardware and software necessary for wireless communications to the access point. It provides mutual authentication to the access point via an EAP mutual authentication type. An EAP supplicant is required on the client machine to support the appropriate EAP authentication type.
- Wireless access point—Mutually authenticates wireless clients via EAP. It can support multiple Layer 2 VLANs for user differentiation.
- Layer 2 or 3 switch—Provides Ethernet connectivity and Layer 3 or 4 filtering between the WLAN access point and the corporate network.
- RADIUS server—Delivers user-based authentication for wireless clients and access point authentication to the wireless clients. Additionally, the RADIUS server can be used to specify VLAN access control parameters for users and user groups.
- DHCP server—Delivers IP configuration information for wireless LEAP clients.
- OTP server (optional)—Authorizes OTP information relayed from the RADIUS server (for PEAP clients only).
- Public-key infrastructure (PKI) server (optional)—Provides X.509v3 digital certificate for user and server identification.

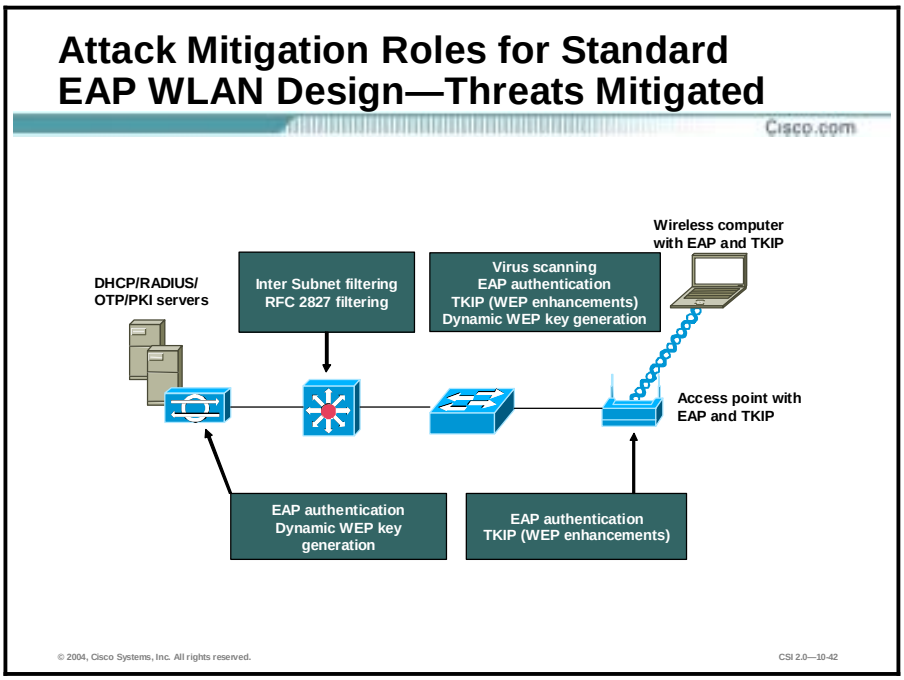

The threats mitigated in the standard EAP WLAN design are as follows:

- Wireless packet sniffers—Wireless packet sniffers can take advantage of any of the known WEP attacks to derive the encryption key. These threats are mitigated by WEP enhancements (specifically per-packet keying as part of TKIP) and key rotation using EAP.
- Unauthenticated access—Only authenticated users are able to access the wireless and wired network. Optional access control on the Layer 3 switch limits wired network access.
- Man-in-the-middle—The mutual authentication nature of several EAP authentication types combined with the MIC can prevent hackers from inserting themselves in the path of wireless communications.
- IP spoofing—Hackers cannot perform IP spoofing without first authenticating to the WLAN. Optional RFC 2827 filtering on the Layer 3 switch restricts any spoofing to the local subnet range.
- Address Resolution Protocol (ARP) spoofing—Hackers cannot perform ARP spoofing without first authenticating to the WLAN. After authenticating, ARP spoofing attacks can be launched in the same manner as in a wired environment to intercept data of other users.
- Network topology discovery—Hackers cannot perform network discovery if they are unable to authenticate. The attacker can note that a WLAN network exists by looking for or observing the access point SSID, but cannot access the network. When authenticated via EAP, standard topology discovery can occur in the same way that is possible in the wired network.

Threats Not Mitigated are:

- Password attack—Several EAP types take into consideration that an attacker can passively monitor the 802.1x/EAP exchanges between the client and the access point, and they mitigate this risk via various methods. PEAP mitigates this by establishing a TLS tunnel from the client to the server before asking for user authentication credentials. Also, because EAP-PEAP takes advantage of other EAP types for client-to-server authentication, the

designer may choose to implement a strong authentication method such as OTP. EAP-TLS mitigates this via public key cryptography.

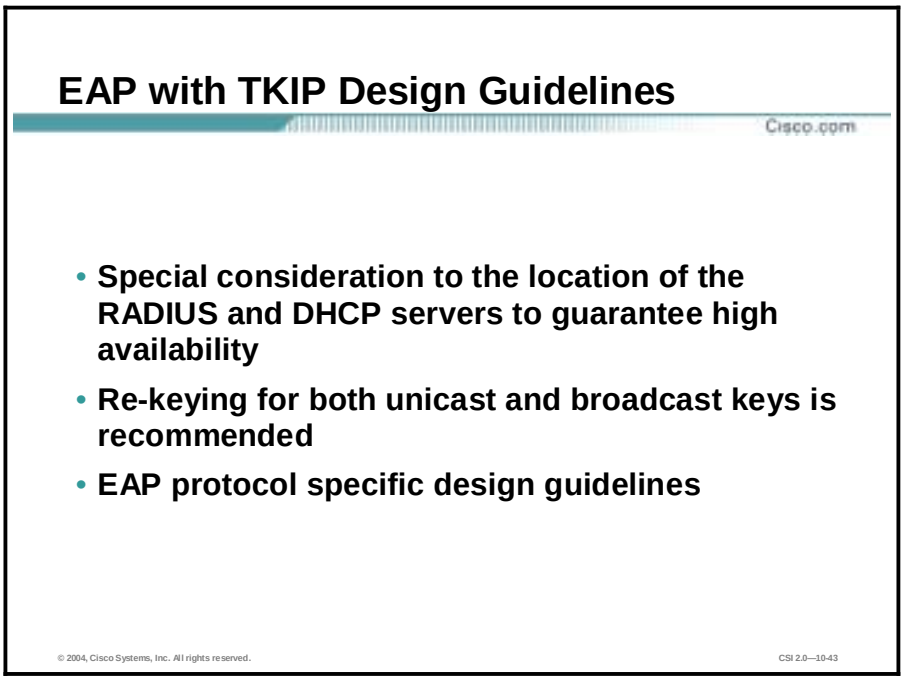

The design guidelines for EAP with TKIP are as follows:

- RADIUS and DHCP server considerations—In most cases, WLAN access points are connected to existing Layer 2 access switches. RADIUS and DHCP servers are located in the network services module of the corporate network. Security in the design is maintained by preventing network access to unauthenticated clients, including events of a RADIUS service failure. Because most of the mitigation against security risks relies on the RADIUS service, this behavior is required. Overall, management of the solution is hindered if DHCP services fail. The wireless clients and access points use EAP to authenticate the WLAN client devices and end users against the RADIUS servers. For scalability and manageability purposes, the WLAN client devices are configured to use the DHCP protocol for IP configuration. DHCP occurs after the device and end user are successfully authenticated via an EAP protocol. After successful DHCP configuration, the wireless end user is allowed access to the corporate network and filtering occurs, if configured. Network designers should give special consideration to the location of the RADIUS and DHCP servers used by EAP in order to guarantee high availability of the network services for WLAN users.
- In order to prevent attacks due to initialization vector collisions, re-keying for both unicast and broadcast keys is recommended. For EAP with Cisco TKIP, re-keying time of 4 hours and 40 minutes is recommended for both unicast and broadcast WEP keys.
- EAP protocol specific design guidelines follow:
  - For EAP-TLS, use a private PKI to issue digital certificates. This allows for integration of the PKI infrastructure with existing back-end user databases (for example, Microsoft Windows 2000 Active Directory [AD]) for certificate management.
  - For EAP-TLS and EAP-PEAP, configure wireless clients with the trusted certificate server's digital certificate and prevent the normal user from modifying these settings. Only the IT administrator should have the privilege levels to modify these settings on the EAP supplicant in a wireless client. If the trusted certificate

authority's certificate was not configured, man-in-the-middle attacks are possible via the use of identity spoofing.

- For EAP-LEAP and EAP-PEAP (when using static passwords), after a small number of incorrect login attempts, lock the account to prevent brute-force attacks from occurring on the user account. The number of attempts is specified on the RADIUS server, and you should age these passwords aggressively. The network designer can additionally mitigate this risk by requiring OTPs for client authentication.
- For EAP-TLS, configure the RADIUS server to check the certificate authority's certificate revocation list (CRL) for expired client certificates.

Optionally, network designers could consider implementation of unique wireless VLANs with the EAP design. Dynamic VLAN assignment can be implemented for EAP users using the RADIUS server and user-group settings. This has the advantage of segregating wireless users into user communities and enforcing policies for these groups at the distribution layer. With the use of VLANs on access points, the management traffic can also be isolated from user traffic with the implementation of management VLAN on the access points.

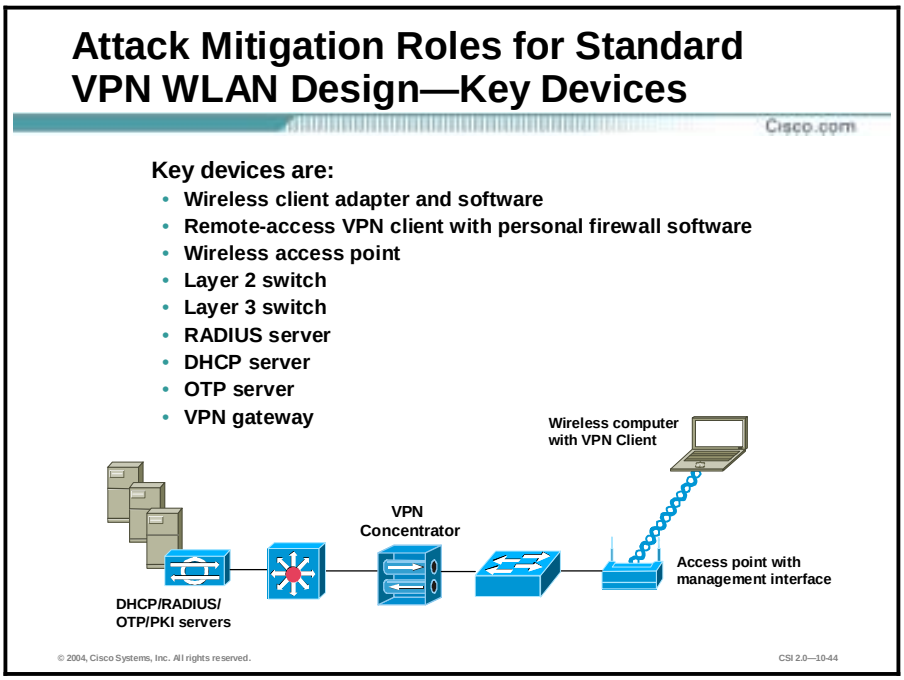

Key VPN devices in the standard VPN WLAN design are as follows:

- Wireless client adapter and software—A software solution that provides the hardware and software necessary for wireless communications to the access point.
- Remote-access VPN client with personal firewall software—A software client that provides end-to-end encrypted tunnels between individual PCs and the corporate wireless VPN gateways. Personal firewall software provides device-level protection for individual PCs.
- Wireless access point—Provides initial IP protocol filtering between the WLAN and corporate network.
- Layer 2 switch—Provides Ethernet connectivity between the WLAN access points and the corporate network. Additionally, recent models of access layer switches have the capability to implement a technology called VLAN Access Control List (VACL), which can provide an additional layer of IPSec filtering.
- Layer 3 switch—Routes and switches production network data from one module to another. It provides additional policy enforcement via protocol-level filtering for wireless traffic.
- RADIUS server—Authenticates wireless users terminating on the VPN gateway. Optionally, it talks to an OTP server.
- OTP server—Authorizes OTP information relayed from the RADIUS server.
- DHCP server—Delivers IP configuration information for wireless VPN clients before and after VPN establishment.
- VPN gateway—Authenticates individual remote users and terminates their IPSec tunnels and can also provide DHCP relay functionality for wireless clients.

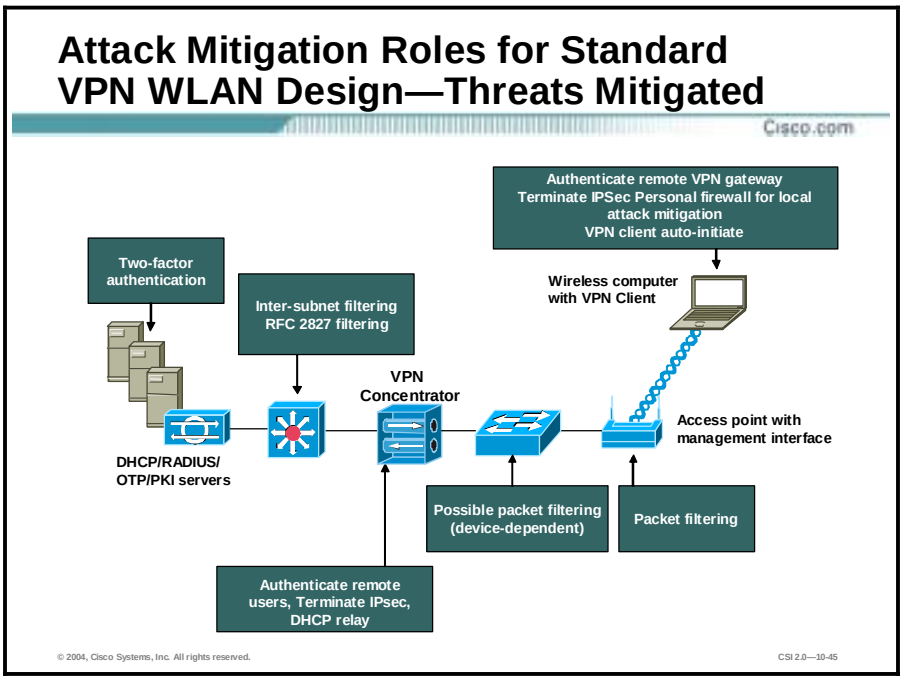

The threats mitigated in the standard VPN WLAN design are as follows:

- Wireless packet sniffers—These threats are mitigated by IPSec encryption of wireless client traffic. Also, new features in VPN client software allow the designer to specify that the VPN tunnel is automatically initiated when the correct WLAN IP address is assigned to the client. This eliminates user interaction to bring up the IPSec tunnel and also protects the client PC from broadcasting traffic onto the wireless media that could be used for inference-based attacks.
- MITM—These threats are mitigated by IPSec encryption and authentication of wireless client traffic.
- Unauthorized access—The only known protocols for initial IP configuration (DHCP) and VPN access (DNS, IKE, and Encapsulating Security Payload [ESP]) are allowed from the WLAN to the corporate network through filtering at the access point and access layer switch. Authorization policies can be optionally enforced on the VPN gateway for individual user groups.
- IP spoofing—Hackers can spoof traffic on the WLAN, but only valid, authenticated IPSec packets will ever reach the production wired network.
- ARP spoofing—ARP spoofing attacks can be launched. However, data is encrypted to the VPN gateway, so hackers will be unable to read the data.
- Password attacks—These threats are mitigated through good password policies and auditing, and, optionally, OTP.
- Network topology discovery—Only IKE, ESP, DNS, and DHCP are allowed from this segment into the corporate network. Internet Control Message Protocol (ICMP) should be allowed only to the outside interface of the VPN Concentrator for troubleshooting purposes.

Threats Not Mitigated are:

- MAC or IP spoofing from unauthenticated users—ARP spoofing and IP spoofing are still effective on the WLAN subnet until the wireless client uses IPSec to secure the connection.

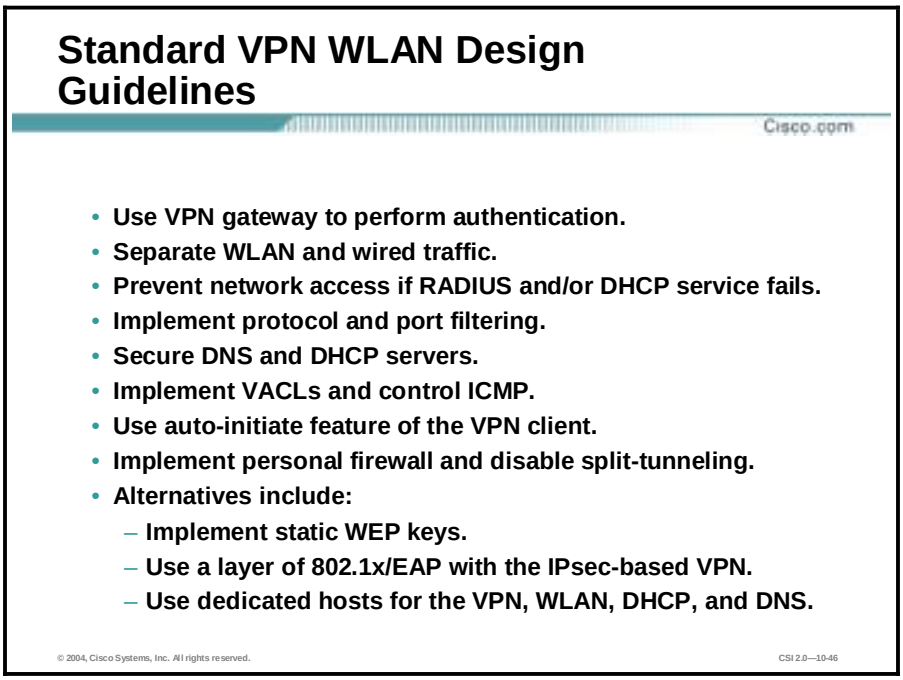

In the standard VPN WLAN design, WLAN access points connect to Layer 2 switches in the building module on a dedicated wired VLAN and forward IPSec traffic from the WLAN client. When the wireless client is communicating with the campus network, but before the IPSec tunnel is established, the client traffic is not considered secure. All the noted WLAN security issues are still present until the wireless client can secure communications with an IPSec VPN. Therefore, the following mitigation techniques are recommended:

- Use the VPN gateway to perform authentication—After the initial Layer 3 configuration, the VPN tunnel authenticates to the VPN gateway. The VPN gateway can use digital certificates or pre-shared keys for wireless device authentication. If the VPN gateway uses pre-shared keys for authentication, then OTPs are recommended to authenticate users to it. Without OTP, the VPN gateways are open to brute-force login attempts by hackers who have obtained the shared IPSec key used by the VPN gateway.
- Separate WLAN and wired traffic—The WLAN traffic should be kept separate from normal wired traffic until it is decrypted by the VPN termination device. Because WEP is not enabled in this design, the wireless network itself is considered an untrusted network, suitable only as a transit network for IPSec traffic. In order to isolate this untrusted network, administrators should not mix the VLAN for the WLAN users with a wired network.
- Prevent network access if RADIUS or DHCP service fails—The VPN gateway takes advantage of RADIUS services, which in turn contact the OTP server for user authentication. The VPN gateway uses DHCP for IP address configuration in order for the WLAN client to communicate through the VPN tunnel. Security in the design is maintained by preventing network access if a VPN gateway or RADIUS service fails. Both services are required in order for the client to reach the wired network with production traffic.
- Implement protocol and port filtering—The access point should be configured with ether type, protocol, and port filters based on a company's wireless usage policy. SAFE WLAN recommends restrictive filters that allow only the necessary protocols required for establishing a secure tunnel to a VPN gateway. These protocols include: DHCP for initial client configuration, DNS for name resolution of the VPN gateways, and the VPN-specific

protocols: IKE (UDP port 500), ESP (IP protocol 50), and ICMP for troubleshooting purposes.

- Secure DNS and DHCP servers—Extra care should be taken to ensure that these systems are as secure as possible at the host level, because they are still open to direct attack on the application protocols themselves. This includes keeping them up-to-date with the latest operating system and application patches and running a host-based intrusion detection system (HIDS).
- Implement VACLs—Recent models of access layer switches have the capability to implement a VLAN ACL (VACL). Implementing VACLs for VPN-related protocols and specific IP addresses of the VPN Concentrators can provide an additional layer of filtering to guarantee that only IPSec traffic destined for the appropriate enterprise VPN Concentrators crosses the switch. The DNS traffic is optional, depending on whether the VPN client needs to be configured with a DNS name for the VPN gateway, or if only an IP address is suitable.
- Control ICMP—Allow ICMP only to the outside interface of the VPN Concentrator for troubleshooting purposes, and implement Path Maximum Transmission Unit (MTU) Discovery.
- Use the auto-initiate feature of the VPN client—This feature should be used in order to minimize the amount of time before the VPN tunnel is established.
- Implement a personal firewall—Implement a personal firewall on the wireless client to protect the client while it is connected to the un-trusted WLAN network, without the protection of IPSec. In general, the VPN gateway distinguishes between the trusted wired network and the untrusted WLAN. The wireless client establishes a VPN connection to the VPN gateway to start secure communication to the corporate network. In the process of doing so, the VPN gateway provides device and user authentication via the IPSec VPN.
- Disable split tunneling—All traffic must traverse the tunnel.
- Alternatives include the following:
  - Implement static WEP keys—Network designers may still consider enabling static WEP keys on all devices in an effort to add an additional deterrent against hackers. The management overhead of dealing with static key changes makes this alternative less than ideal for large WLAN deployments.
  - Use a layer of 802.1x/EAP with the IPSec-based VPN deployment to secure the WLAN environment. The primary drawback with this alternative is the necessity of managing two separate security infrastructures for WLAN deployments.
  - To further secure the DNS and DHCP services, use dedicated hosts for the VPN, WLAN, DHCP, and DNS deployment. This mitigates against two potential threats that could affect wired resources:
    - Denial of service (DoS) attacks against the DHCP and DNS services that could affect wired users
    - Network reconnaissance through the use of DNS queries or reverse lookups

Note As an alternative to dedicated DNS servers, you may consider hard-coding the IP address of the VPN gateway for the VPN clients. The drawback of this solution is if the IP address of the VPN gateway changes, every client will need to update his gateway entry.

Enterprise Wireless LAN Design

This topic describes the enterprise Wireless LAN design.

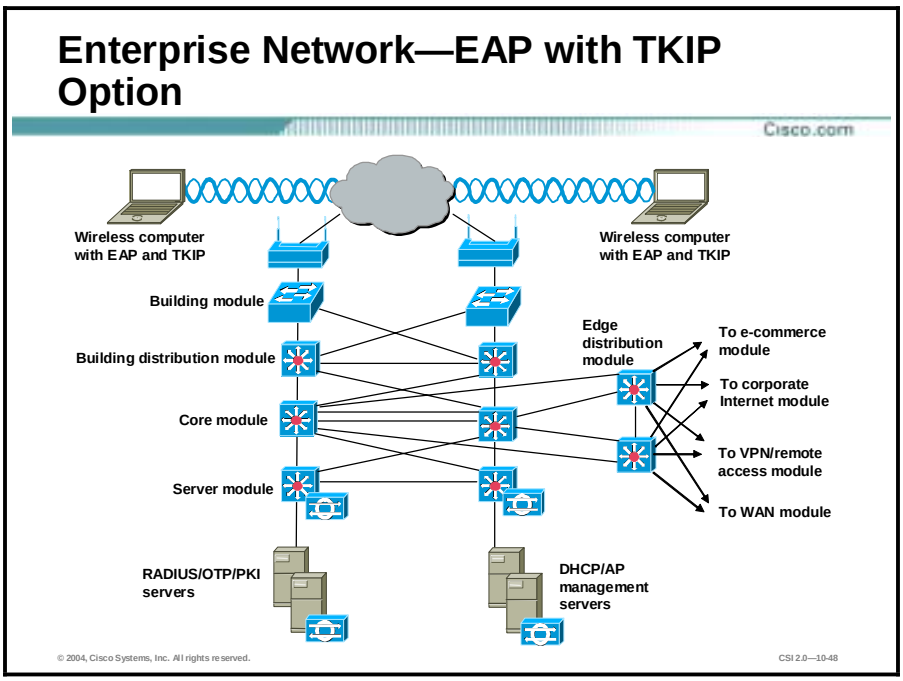

The large-enterprise WLAN network design overlays WLANs on top of the campus portion of the SAFE enterprise blueprint. All of the components for implementing the mitigation techniques are contained within the large-enterprise building, distribution, and server modules. These components are intended to allow WLAN access for enterprise end users within the enterprise campus.

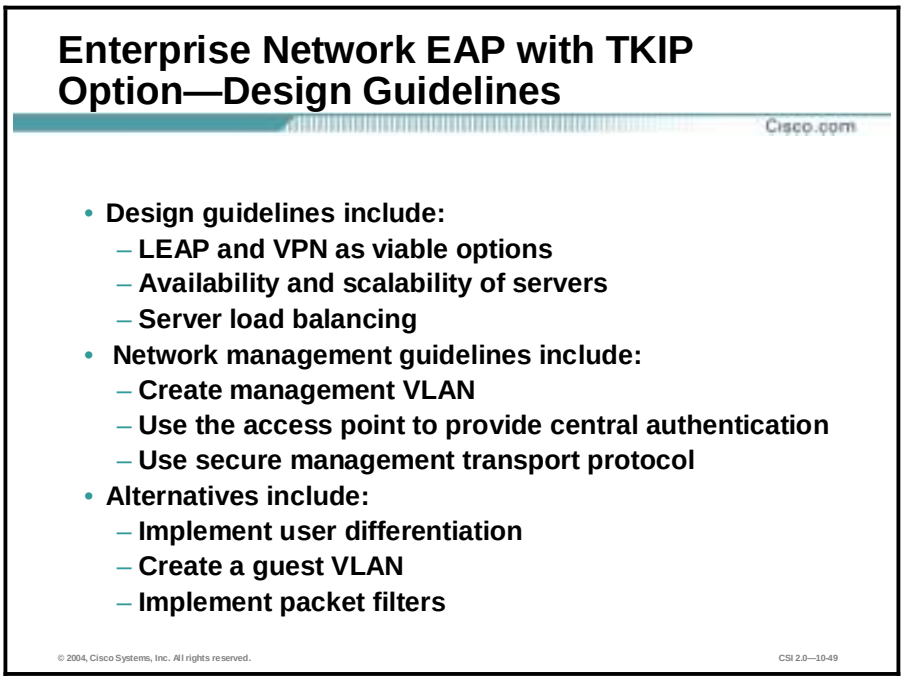

In the large-enterprise WLAN design, scalability and high availability are primary concerns when implementing the mitigation technologies. The following are specific design recommendations:

- LEAP and VPN—Both LEAP and VPN are considered viable security options for large-enterprise WLAN designs. You should weigh the business benefits of both technologies with the company security policy before selecting the technology that is best suited for your network.
- Availability and scalability of servers—The primary concern for EAP in a large WLAN design is the availability and scalability of the network configuration and authentication servers. The RADIUS, OTP, PKI, and DHCP servers are deployed in a redundant fashion on differing network subnets to ensure high availability and scalability.
- Server load balancing—Server load balancing can increase the scalability of the RADIUS servers by spreading the RADIUS authentication requests evenly across a farm of RADIUS servers.

Network management guidelines include the following:

- Create a management VLAN—In order to isolate management traffic from user traffic, use VLANs on the access points, creating a management VLAN for the access points and restricting access to the access points via the management subnet by implementing ACLs in the building distribution Layer 3 switch. The ACLs should be specific to the IP address and protocols that the centralized multidevice access point configuration tool requires.
- Use the access point to provide central authentication—Because access points support only one wired interface, all management was done in band, versus the out-of-band management as recommended by SAFE enterprise. This setup contains a security risk because some management traffic (SNMP, Trivial File Transfer Protocol [TFTP], HTTP) must be sent in the clear in order to manage each access point via a central management station. The access point should be configured to provide central authentication, authorization, and accounting

(AAA) of access point administrators via RADIUS or TACACS+, depending on the support of the deployed access points.

- Use a secure management transport such as SSH in order to manage the access points from the command line.

Alternatives include the following:

- Implement user differentiation—EAP design along with VLANs on the access points enables you to implement user differentiation (using wired and wireless VLANs). A separate VLAN can be created for traditional wireless devices that support only static WEP, and an appropriate security policy can be enforced for this VLAN as well.
- Enforce VLAN assignment for users and user groups using the RADIUS server.
- Create a guest VLAN—Create a guest VLAN to allow guests of the enterprise access to the corporate network for access to a limited set of resources, or a VPN to gain access to the guest's corporate network.
- Implement packet filters—Use packet filters on the access and Layer 3 building distribution switch (or wherever the guest VLAN is terminated) to allow only traffic that conforms to the enterprise guest security policy (that is, IPSec traffic only).

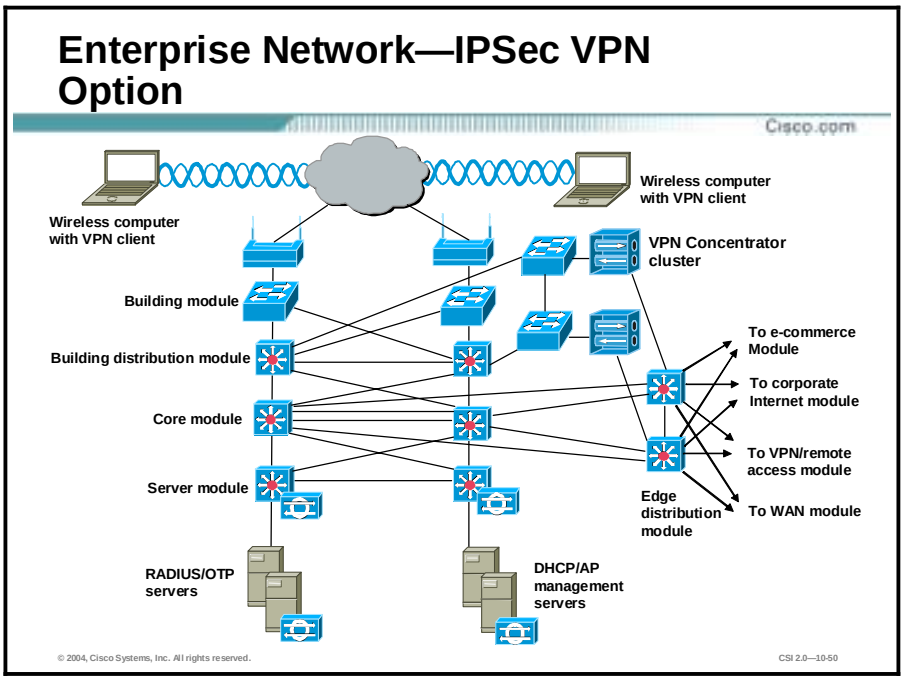

IPSec VPN access via the wireless network uses several modules from the SAFE enterprise architecture, as follows:

- Building module
- Building distribution module
- Edge distribution module
- Server module

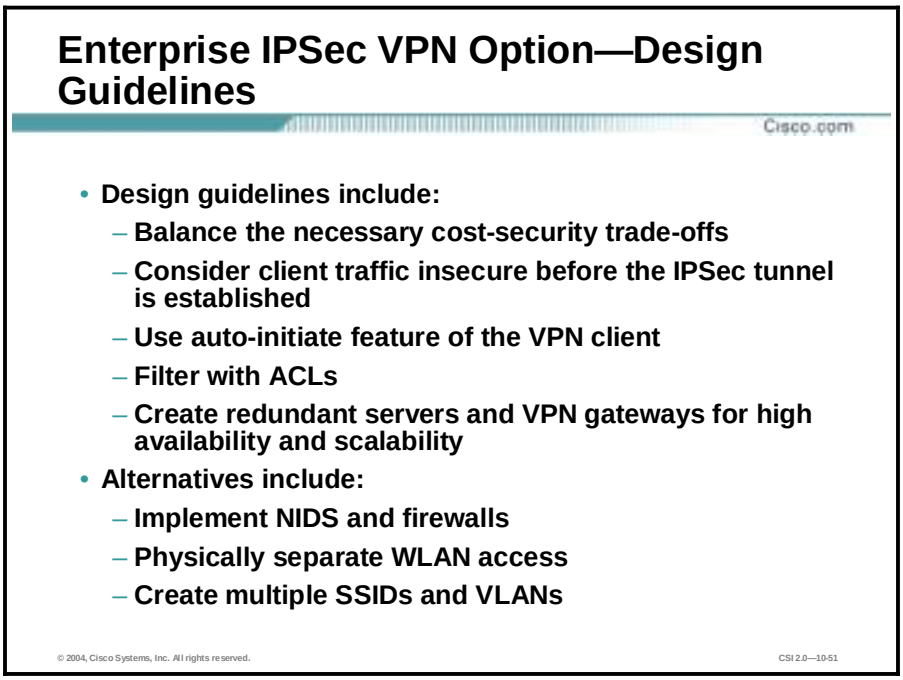

Design guidelines for the Enterprise IPSec VPN option include the following:

- Balance the necessary cost-security trade-offs—The primary objective in the large WLAN design involves balancing the mitigation of security risks with the creation of a scalable design that a business can afford to implement. In the context of a large WLAN environment, the guidelines described would be cost-prohibitive for most businesses because of the requirement for a separate Layer 2 switching infrastructure and cabling. Therefore, security trade-offs are made in order to make a VPN WLAN feasible in a large environment.
- Consider client traffic insecure before the IPSec tunnel is established—The WLAN clients associate with a wireless access point in the building module to establish connectivity to the campus network at Layer 2. The wireless clients then use DHCP and DNS services in the server module to establish connectivity to the campus at Layer 3. It should be noted that when the wireless client is communicating with the WLAN network, but before the IPSec tunnel is established, the client traffic is not considered secure. All the noted WLAN security issues are still present until the wireless client can secure communications with an IPSec VPN.
- Use the auto-initiate feature of the VPN client—This feature should be used in order to minimize the amount of time before the VPN tunnel is established.
- Filter with ACLs—In addition to the filters on the access point, the building distribution module Layer 3 switches are configured with ACLs to permit only protocols necessary for VPN connectivity and management.
- Create redundant VPN gateways for high availability and scalability—The wireless client establishes a VPN connection to the VPN gateways connecting the building distribution and edge distribution modules. The redundant VPN gateways are configured in a load-balancing configuration to provide high availability and scalability. These VPN gateways are a centralized resource shared by potentially multiple Layer 2 building modules.
- Server redundancy—The RADIUS, OTP, and DHCP servers used by the VPN gateways are deployed in a redundant fashion on different network subnets within the server module

to ensure high availability and scalability of their respective services to the VPN client tunnels.

Alternatives include the following:

- Implement NIDS and firewalls—An organization can further its security posture by deploying a network-based IDS (NIDS) and firewalls behind the VPN gateways before wireless user traffic hits the production wired network. This setup allows the network to audit, inspect, and filter user traffic that is being sent from the wireless clients to the enterprise network as defined by an organization's security policy. After providing the device and user authentication, the VPN gateway can optionally provide user authorization rights based on the group with which the wireless user is associated. All these security improvements are strongly recommended if the VPN user authentication policy chooses not to use OTP.
- Physically separate WLAN access—Consider the benefits of building a physically separate infrastructure for WLAN access. Physically separate Layer 2 and 3 segments on dedicated networking hardware are used to totally isolate the untrusted WLAN until traffic is decrypted at the VPN gateways and routed into the production wired network.
- Create multiple SSIDs and VLANs—The creation of multiple SSIDs and VLANs on the access point allows you to create a guest VLAN in order to allow guests of the enterprise access to the corporate network, to either gain access to a limited set of resources or VPN across the Internet to gain access to the guest's corporate network. In either of these cases, implement packet filters on the access and Layer 3 building distribution switch (or wherever the guest VLAN is terminated) to allow only traffic that conforms to the enterprise guest security policy. Similarly, a VLAN can be created for traditional wireless devices that support only static WEP, and an appropriate security policy can be enforced for this VLAN as well.

Medium Wireless LAN Design

This topic describes medium WLAN design.

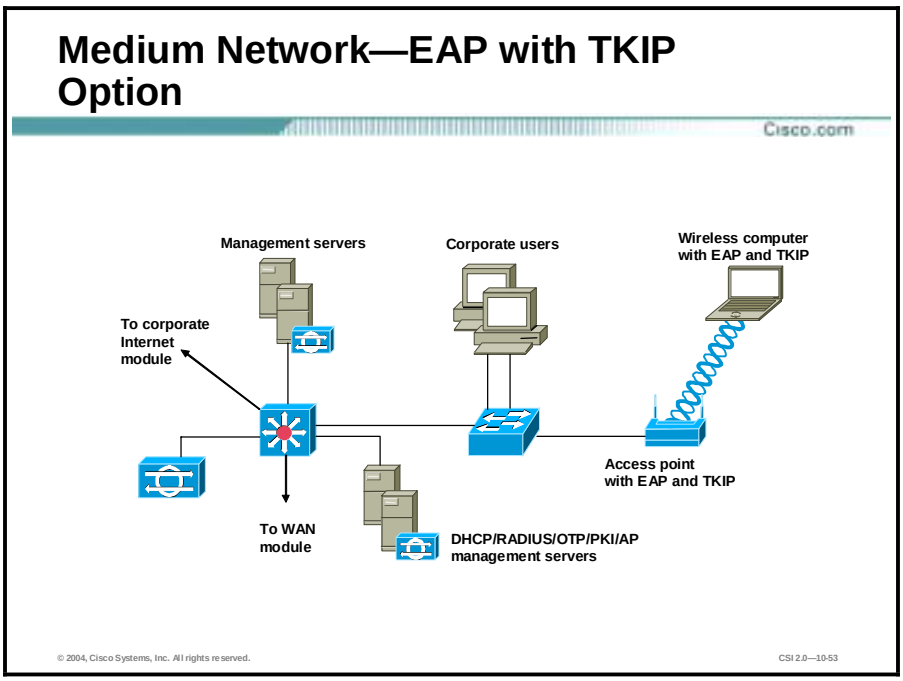

The medium network WLAN overlays wireless on top of the campus portion of the SAFE medium network design. All the components for implementing the mitigation techniques are contained within the medium-campus module. These components are intended to allow WLAN access for end users within the medium-network campus.

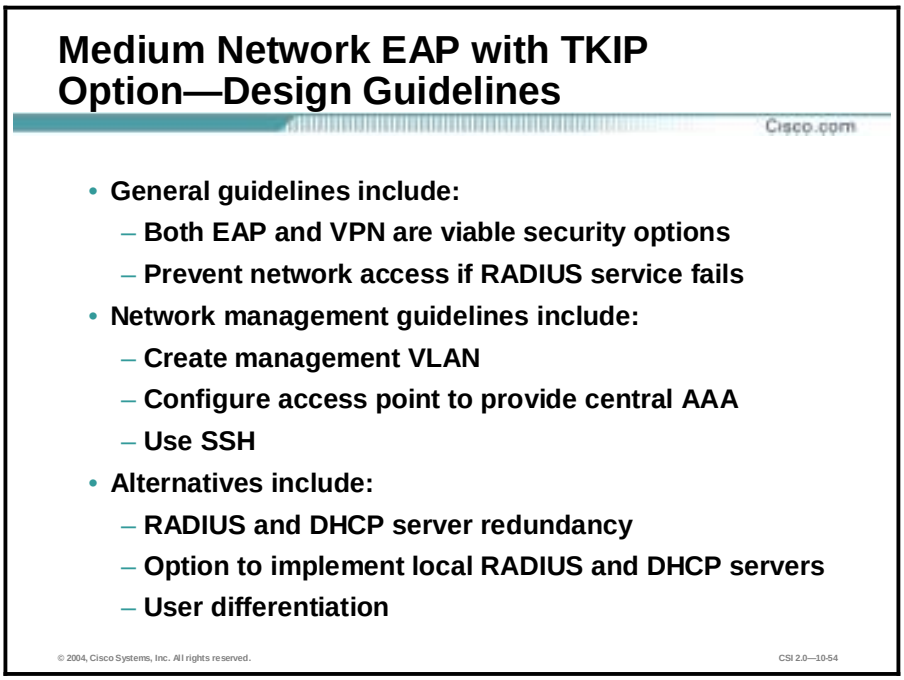

In the medium WLAN design, it is assumed that all WLAN devices are connected to a single IP subnet to enable end-user mobility throughout the medium WLAN design. An assumption is made in the designs that most services available to the medium wired network are also available to the medium WLAN design. Keeping with the design foundation for the SAFE medium network, the medium WLAN design does not offer high availability.

Design guidelines for the medium network EAP with TKIP option include the following:

- Both EAP and VPN are considered viable security options for a medium WLAN design—Key devices for both the EAP and VPN options are supported in the campus module of the SAFE medium network design. For both options, network designers should give special consideration to the location of the RADIUS and DHCP servers used by the EAP and VPN WLAN solutions. The location of the servers will depend on the type of office the medium-network WLAN represents: main business office, branch office, or corporate office. Options include the following:
  - If the medium network is the main business office, the DHCP and RADIUS servers will be located on the local network.
  - If the medium network is a branch office, the DHCP and RADIUS servers might reside at the corporate office, with connectivity via the WAN module or through a VPN in the corporate Internet module.
  - If the DHCP and RADIUS servers are located at the corporate office, wireless users will be denied access to the local network if the access point or VPN gateways cannot communicate with the RADIUS server for any reason, such as loss of WAN connectivity. Also, if the DHCP servers are unavailable to the medium network, the wireless clients will not be able to establish IP connectivity with the campus network.
- Prevent network access if RADIUS service fails—Because most of the mitigation against security risks relies on the RADIUS service, this behavior is required. Overall, management of the solution is hindered if DHCP services fail.

Network management guidelines include the following:

- Create management VLAN—In order to isolate management traffic from user traffic, it is recommended to use VLANs on the access points. Creating a management VLAN on the access point and restricting access to the access point via the management subnet by implementing ACLs in the building distribution Layer 3 switch is recommended. The ACLs should be specific to the IP address and protocols that the centralized multidevice access point configuration tool require. Note that because access points support only one wired interface, all management was done in band, versus the out-of-band management as recommended by SAFE enterprise. This setup contains a security risk because some management traffic (SNMP, TFTP) must be sent in the clear in order to manage each access point via a central management station.
- Configure access point to provide central AAA—AAA of access points can be provided via RADIUS or TACACS+, depending on the support of the deployed access points.
- Use a secure management transport protocol—Use a protocol such as SSH in order to manage the access points from the command line. EAP access in the medium WLAN design has wireless access points connected to the existing Layer 2 access switch in the medium-campus module. RADIUS and DHCP servers are also located in the campus module, but off a distinct Layer 3 subnet on the central-campus Layer 3 switch. The wireless EAP users will require DHCP and RADIUS authentication services to access the medium-campus network. If the medium network is a branch office, the DHCP and RADIUS servers may reside at the corporate office.

Alternatives include the following:

- RADIUS and DHCP server redundancy—In the case of a branch office where RADIUS and DHCP servers reside at the corporate office, RADIUS and DHCP server redundancy needs to be considered.
- Option to implement local RADIUS and DHCP servers—To provide WLAN access in the event of WAN link connectivity failure to the corporate network. If this alternative is chosen, the network designer needs to consider administration and maintenance of multiple RADIUS and DHCP servers.
- User differentiation—EAP design along with VLANs on access points enables the network designer to implement user differentiation (using wired and wireless VLANs) and enforce VLAN assignment for users and user groups using the RADIUS server. Furthermore, appropriate Layer 3 filters can be enforced at the access and distribution layer for each user group.

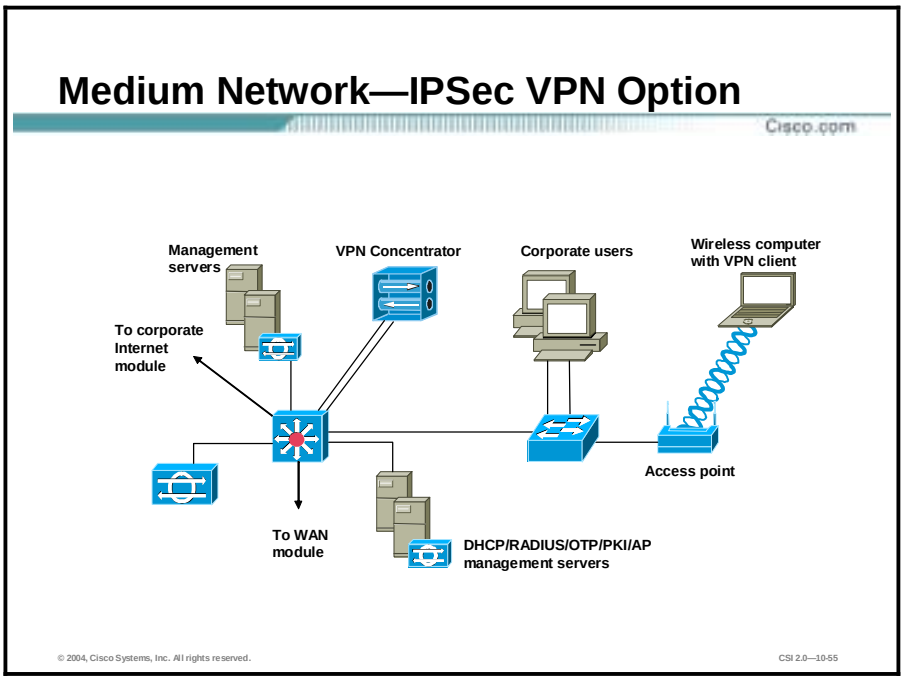

The IPSec VPN option in the medium network is very similar to the VPN option for the large WLAN design. The primary differences are in the physical connectivity of the VPN gateway that divides the wireless network from the wired network. The VPN gateway connects its interfaces to the campus-module Layer 3 switch using two different VLANs. This VLAN-based option was chosen because the alternative was not financially viable for businesses likely to deploy a midsize network.

The following alternatives provide for a more secure option using additional equipment:

- The VPN gateway connects its public interface to one VLAN that can connect to the wireless access points. As in the standard VPN WLAN design, it is recommended that the VPN Concentrator perform DHCP relay between the public side and the private side of the Concentrator. This allows the network designer to more effectively deploy and manage the DHCP services for the WLAN. The private interface of the VPN gateway connects to a VLAN with access to the wired network. The wireless access points connect to existing Layer 2 switches in the campus-module access layer on a dedicated VLAN and forward traffic from the WLAN to the VLAN with VPN connectivity. Like the large WLAN and general VPN WLAN design, the access and building distribution module Layer 3 switches are configured with ACLs to permit only protocols necessary for VPN connectivity and management.
- The wireless client establishes an IPSec connection to the wireless VPN gateway. In the process of doing so, the VPN gateway provides device and user authentication via the IPSec VPN. The VPN gateway can use digital certificates or pre-shared keys for wireless client device authentication. The VPN end user employs OTPs to authenticate to the VPN gateway. The VPN gateway uses RADIUS services, which in turn contact the OTP server for user authentication. The VPN gateway uses DHCP for IP addressing information in order for the WLAN client to communicate through the VPN tunnel.
- Deploy NIDS and firewall behind the VPN gateways before wireless user traffic hits the production wired network. This setup allows the network to audit, inspect, and filter user traffic that is being sent from the wireless clients to the medium network as defined by an

organization’s security policy. Both of these security improvements are strongly recommended if the VPN user authentication policy chooses not to use OTP.

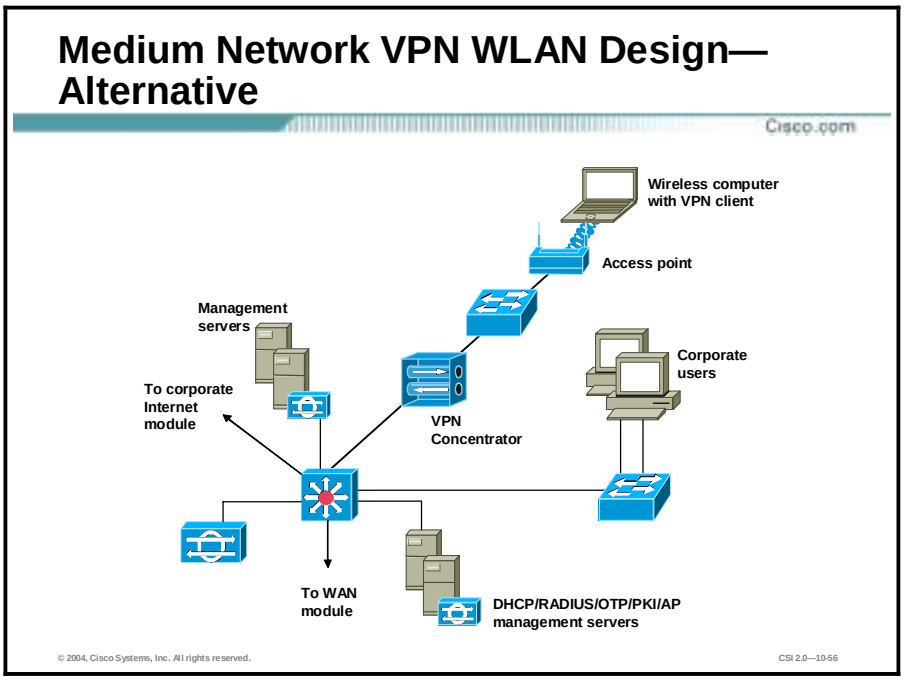

A network designer looking for more security than these designs provide should consider the benefits of a design similar to the standard VPN WLAN option. A design specific to the medium WLAN is depicted in the slide. The primary benefit is the clear delineation between the public and private interfaces of the VPN gateway. The primary detraction from this design is the potentially high cost of deploying additional Layer 2 switches just to connect the wireless access points. New features in VPN Concentrators allow the Concentrator to provide DHCP relay services to a DHCP server behind the VPN Concentrator. This deployment option does open the DHCP server to the security risks outlined in the standard VPN design, but is recommended rather than deploying a DHCP server outside the VPN Concentrator.

Small Wireless LAN Design

This topic describes small WLAN design.

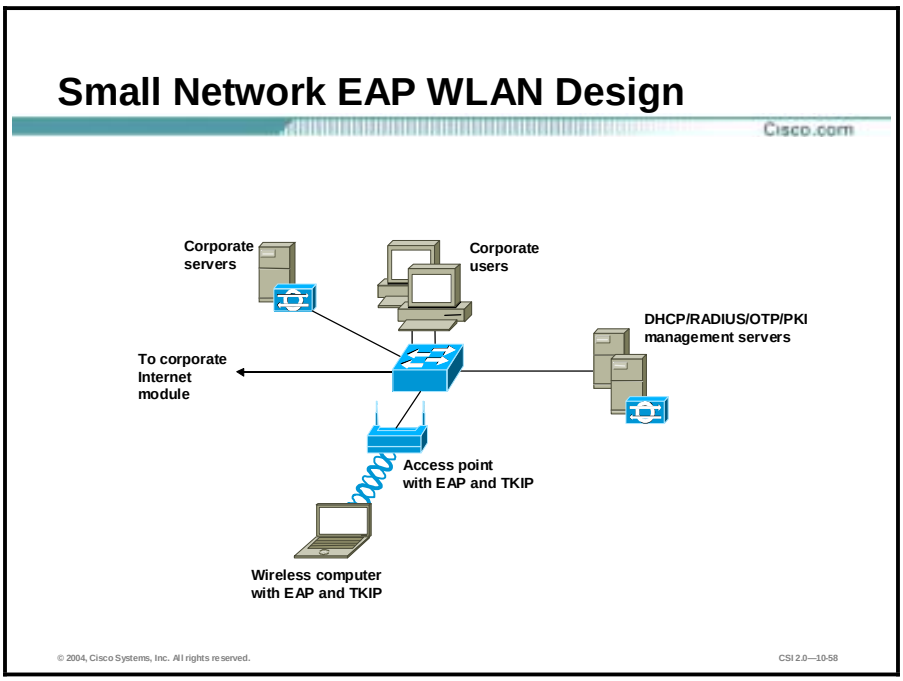

The small WLAN network design overlays WLAN on top of the SAFE small network design. The small WLAN design is contained within the campus module. This topic discusses one option, EAP with TKIP, for providing WLAN users connectivity to the wired campus. IPSec is not presented as an option because of the financial burden of implementing a dedicated WLAN VPN in a network of this size.

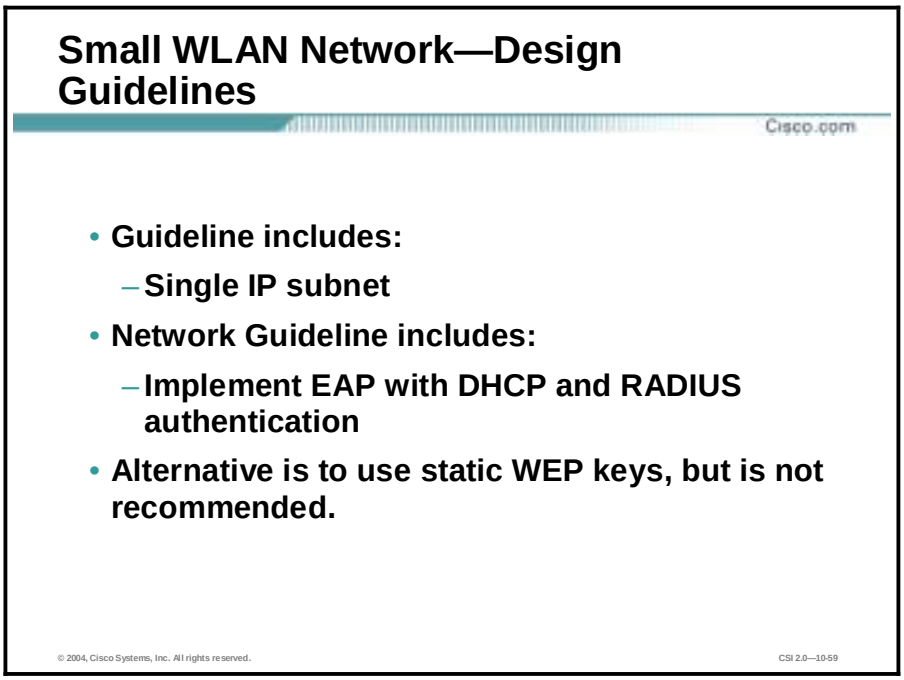

Design guidelines for the small WLAN network include a single IP subnet. Because the small-network design has a single Layer 2 switch for its campus connectivity, all devices are assumed to have a single IP subnet network to enable access point roaming.

Network management guidelines are to implement EAP with DHCP and RADIUS authentication. Network management traffic from the management hosts to the access points is unrestricted because of the lack of a Layer 3 device in the small campus. Some management traffic is sent in the clear to each access point, as is done for the rest of the SAFE small design. Cisco EAP access in the small WLAN design has wireless access points connected to the existing Layer 2 access switch in the small-campus module. The wireless EAP users require DHCP and RADIUS authentication services to access the small-campus network. Because of the single-site nature of small networks, the RADIUS and DHCP servers reside locally, connected to the Layer 2 switch in the campus module.

An alternative is to use static WEP. Although not recommended, if an organization is comfortable with managing the key distribution issues, static WEP can be used as an alternative to EAP.

Remote Wireless LAN Design

This topic describes remote WLAN design.

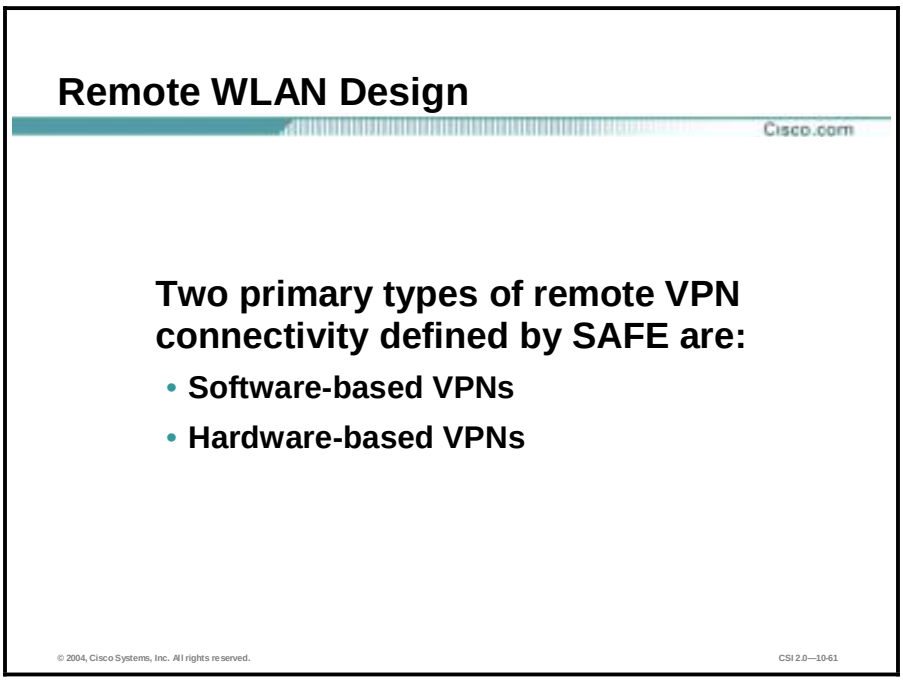

The remote WLAN design shows remote wireless solutions for the two primary types of remote VPN connectivity defined by SAFE:

- Software-based VPNs
- Hardware-based VPNs

This topic discusses these two options for providing WLAN users connectivity to a central office (small, medium, or enterprise) within the SAFE design.

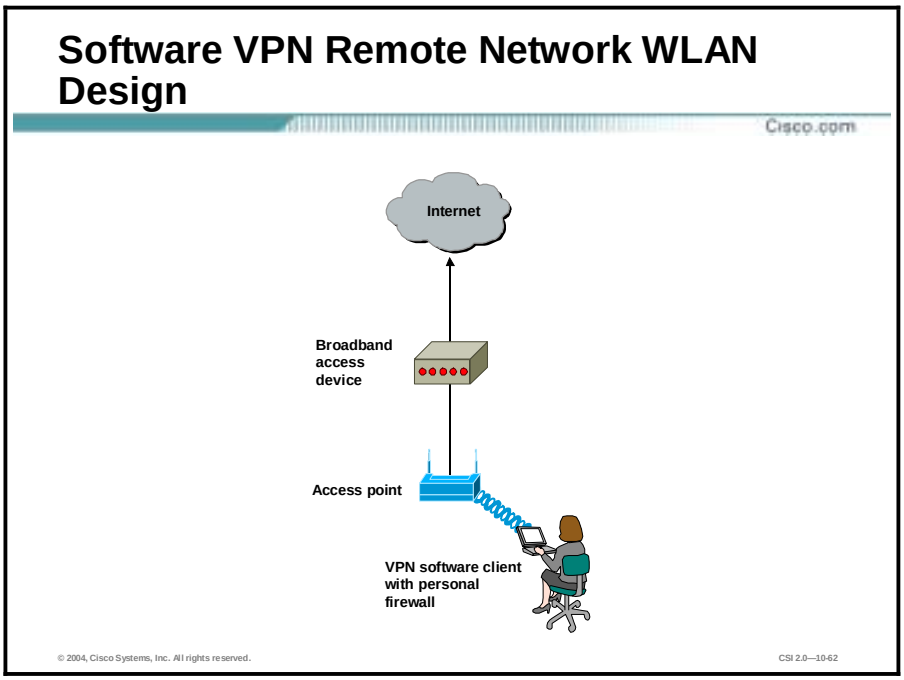

The IPSec VPN option in the remote network is recommended when the wireless user requires security from the wireless device to the corporate network. This is the most common configuration for remote workers who may not have IT-managed hardware resources at their remote location. It includes the following features:

- The design is suitable for part-time teleworkers.
- The access point can be set up with almost any configuration that allows connectivity to the broadband device, because the security is handled via the VPN client with personal firewall software.
- Network designers can consider putting filters on the access point to allow only IPSec, DHCP, and DNS traffic in order to mitigate attacks from the WLAN to the wired LAN.

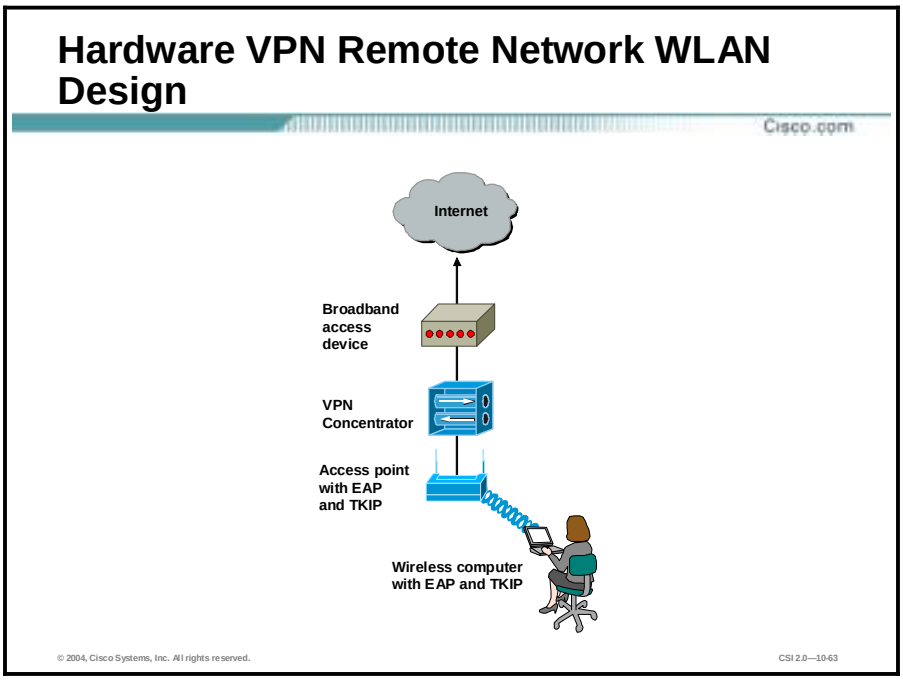

For configurations where an organization's IT department manages VPN and wireless gear at a user's remote location, using EAP from the PC to the access point and then IPSec from the hardware VPN device to the central office provides a robust security solution for a remote worker. It includes the following features:

- Full-time teleworkers are the individuals most likely to take advantage of this configuration.
- When the remote location is using hardware VPN and EAP for wireless, the design is nearly identical to that of the small WLAN design.
- A caveat regarding RADIUS access applies where wireless users are denied access to the local network if the access point cannot communicate with the RADIUS server for any reason, such as loss of IPSec VPN connectivity.
- This design also requires that the remote network have a unique IP range to facilitate IT management of the remote access point. If the hardware device uses Network Address Translation (NAT) for all traffic from the remote site to one IP address, the IT department cannot manage the access point.

SAFE WLAN Implementation

This topic describes basic security configurations on the Cisco Aironet access point, Wireless Client Adapter, and ACS.

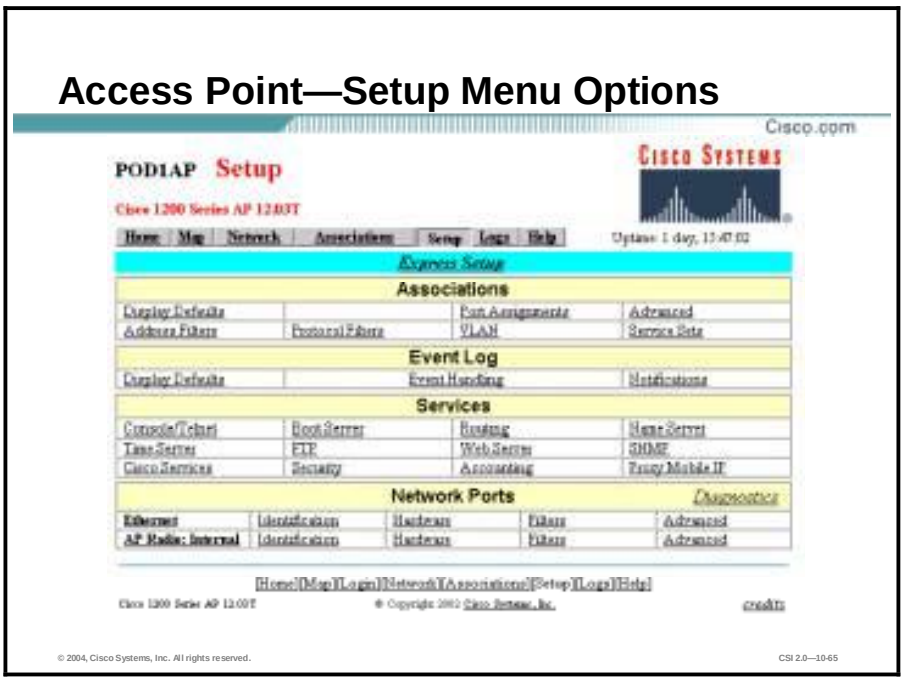

The setup screen displays menu options for Express Setup, Associations, Event Log, Services, and Network Ports.

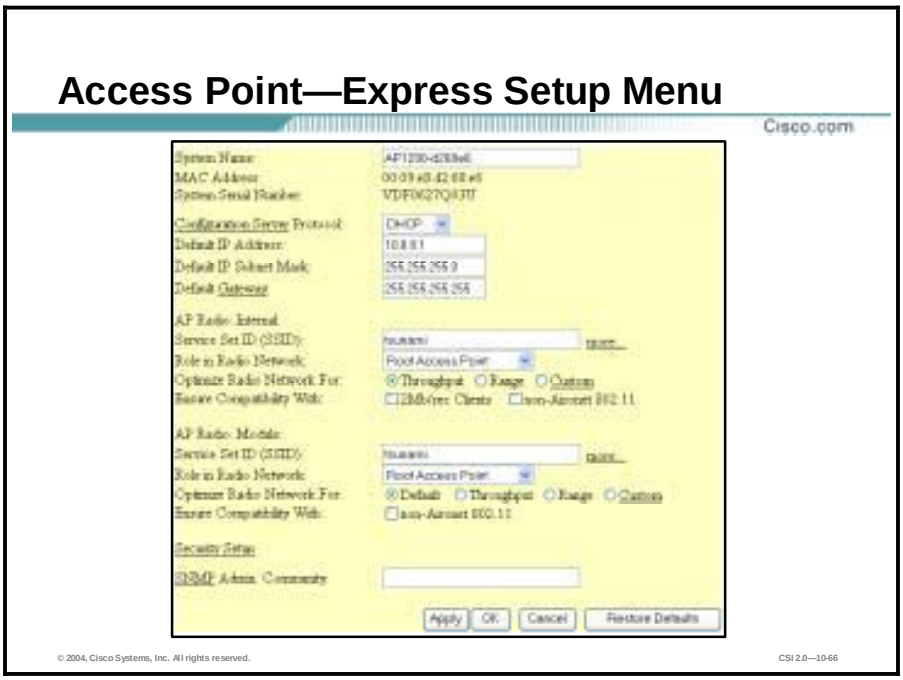

The first time the access point is turned on, this is the default web page for the access point. It will remain the default page until a configuration is successfully applied or you click **OK**. The Express Setup page contains the following settings:

- System Name
- MAC Address
- System Serial Number
- Configuration Server Protocol
- Default IP Address
- Default IP Subnet Mask
- Default Gateway
- Radio Service Set ID (SSID)
- Role in Radio Network
- Radio Network Optimization (Optimize Radio Network For)
- Radio Network Compatibility (Ensure Compatibility With)
- Security Setup Link
- SNMP Admin. Community

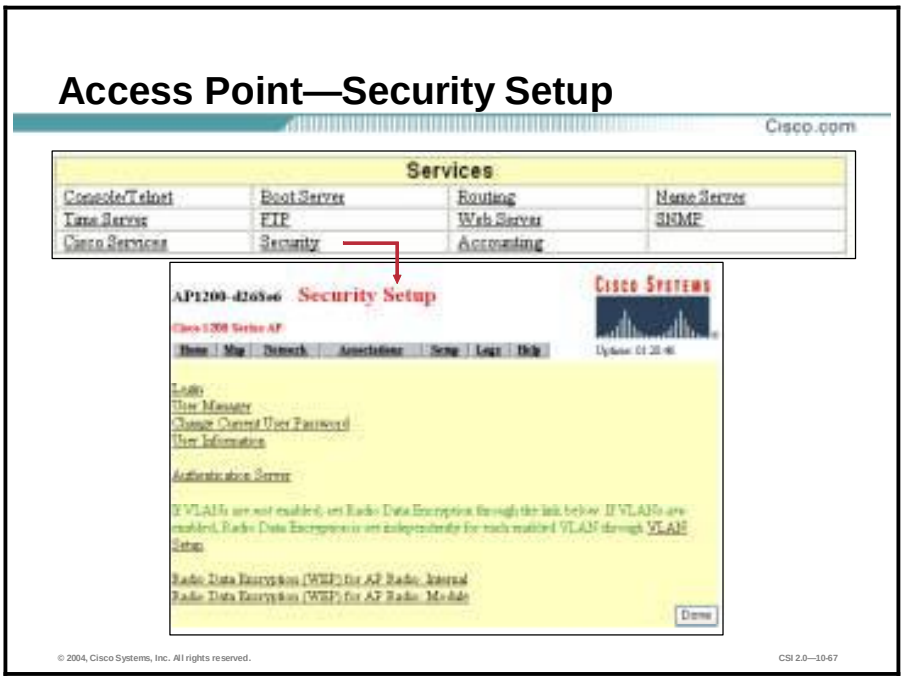

The Services menu has the option to configure access point security. The security features protect wireless communication between the access point and other wireless devices, control access to your network, and prevent unauthorized entry to the access point management system. In order to maximize security on the WLAN, a number of features need to be enabled and configured. These features are as follows:

- Login—Requires users to log in to the access point.
- User Manager—Options include enabling or disabling the user manager. User manager provides the access point with a level of security that is common to wired networking components. An administrator would not want just anyone to Telnet into their company's switch and make changes to the configuration. The same precautions should be taken with all access points.
- Change Current User Password—Allows user to change their password.
- User Information—This option will display all users and their capabilities. Clicking on a user's name will take you to the properties screen for that user.
- Authentication server—Before a wireless client device can communicate on your network through the access point, it must authenticate to the access point and to your network. The access point uses four authentication mechanisms or types and can use more than one at the same time:
  - EAP authentication
  - MAC address authentication
  - User authentication
  - MIP authentication
- VLAN setup—Provides VLAN summary status on an access point. VLAN Summary Status link has following options:
  - VLAN (802.1Q) Tagging

- 802.1Q Encapsulation Mode
- Maximum Number of Enabled VLAN IDs
- Native VLAN ID
- Single VLAN ID which allows Unencrypted packets
- Optionally allow Encrypted packets on the unencrypted VLAN
- VLAN ID
- VLAN Name
- Existing VLANs

- WEP—WEP is the first line of defense against intruders. Cisco recommends that you use full encryption on your wireless network. WEP encryption scrambles the communication between the access point and client devices to keep the communication private. Both the access point and client devices use the same WEP key to encrypt and decrypt radio signals.

Access Point—WEP Setup

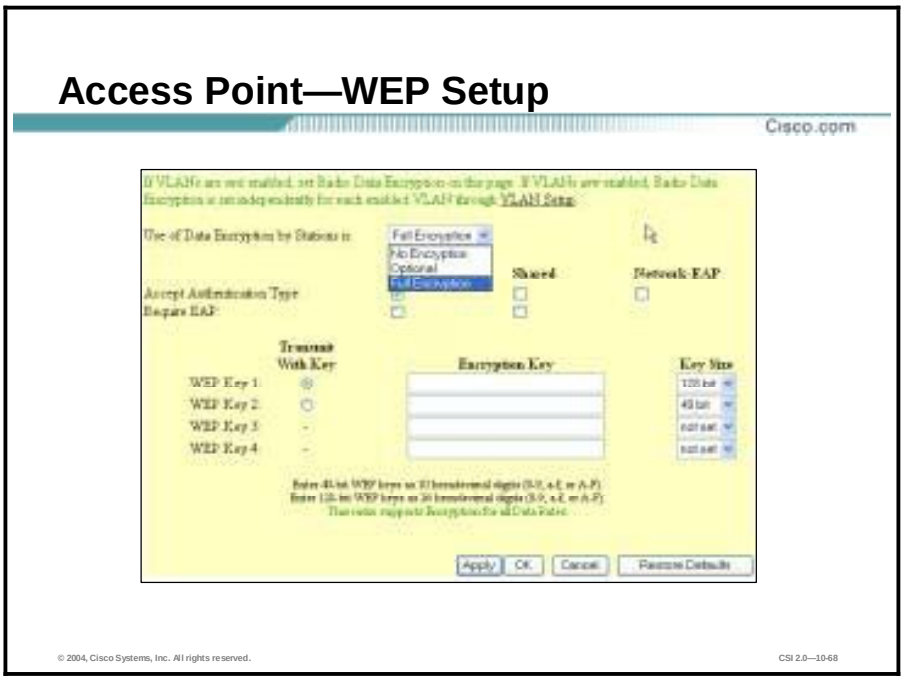

To configure WEP, go to the **Security Setup** screen and click on **Radio Data Encryption (WEP)** to launch the WEP configuration screen. Data encryption options are as follows:

- Use of Data Encryption by Stations options are:
  - No Encryption (default)—The access point communicates only with client devices that are not using WEP.
  - Optional—Client devices can communicate with the access point either with or without WEP.
  - Full Encryption—Client devices must use WEP when communicating with the access point. Clients not using WEP are not allowed to communicate.

Note If you use Network-EAP as the authentication method, a key must be set in the WEP Key 1 slot. This is the key that is used for multicast packets and is sent during the authentication process. The access point is not restricted to use of only 40-bit or 128-bit keys. Any combination of 40-bit and 128-bit keys may be used. To configure WEP, an encryption type must be chosen by checking the appropriate box.

- Accept Authentication Type options are:
  - Open (default)—Allows any device, regardless of its WEP settings, to authenticate and then attempt to communicate with the access point.
  - Shared Key—The access point sends a plain text, shared-key query to any device attempting to communicate with the access point. This query can leave the device open to a known-text attack from intruders, however, and is therefore not as secure as the Open setting.
  - Network-EAP—The access point uses EAP to interact with an EAP-compatible RADIUS server on your network to provide authentication for wireless client devices.

Note The standard 802.11 WEP can be used without using EAP or an authentication server, allowing for data encryption between the clients and the access point. Using 802.11 WEP does not encrypt all data on the network. Only the data sent between the client and the access point will be encrypted.

- Other WEP key options are:
  - — Transmit With Key—These buttons allow you to select the key this access point will use when transmitting data. Only one key can be selected at a time. All set keys can be used to receive data. The selected key must already be set before it can be specified as the Transmit key.
  - — Encryption Key—These fields allow you to enter the WEP keys. Type ten hexadecimal digits (any combination of 0-9, a-f, or A-F) for 40-bit WEP keys. Type 26 hexadecimal digits (any combination of 0-9, a-f, or A-F) for 128-bit WEP keys. To protect WEP key security, existing WEP keys do not appear in the entry fields. You can write over existing keys, but you cannot edit or delete them.
  - — Key Size—Use this setting to set the keys to either 40 or 128-bit WEP. If “not set” appears for this selection, the key has not been set.

Note Use the Restore Defaults button to remove all WEP Keys.

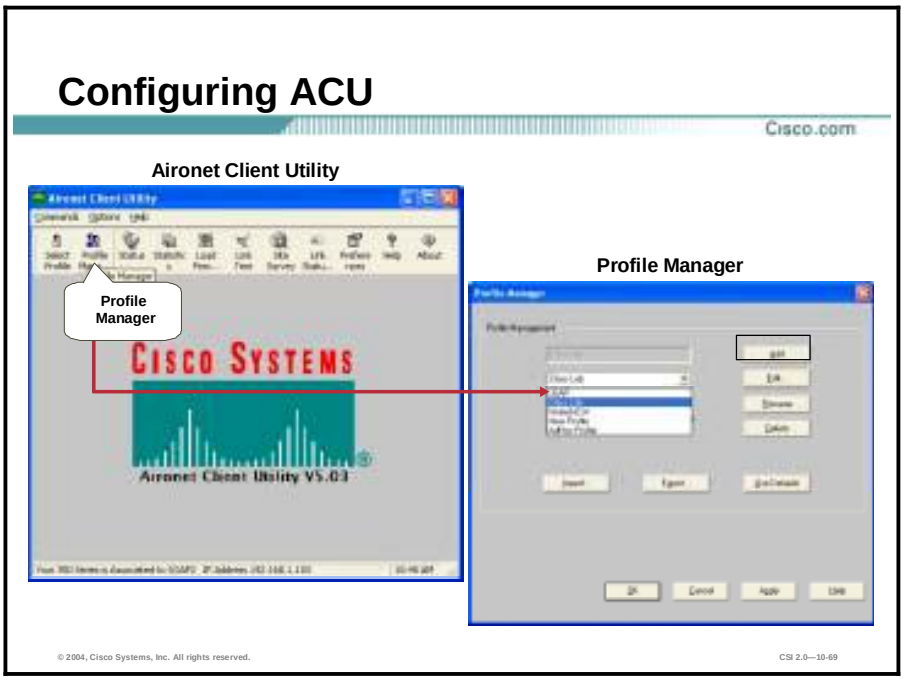

Wireless client is configured using the Aironet Client Utility (ACU). On the Aironet Client Utility screen, click **Profile Manager**. This will launch the Profile Manager screen. From the Profile Manager screen, choose the desired profile from the drop down box and click the **Edit** button.

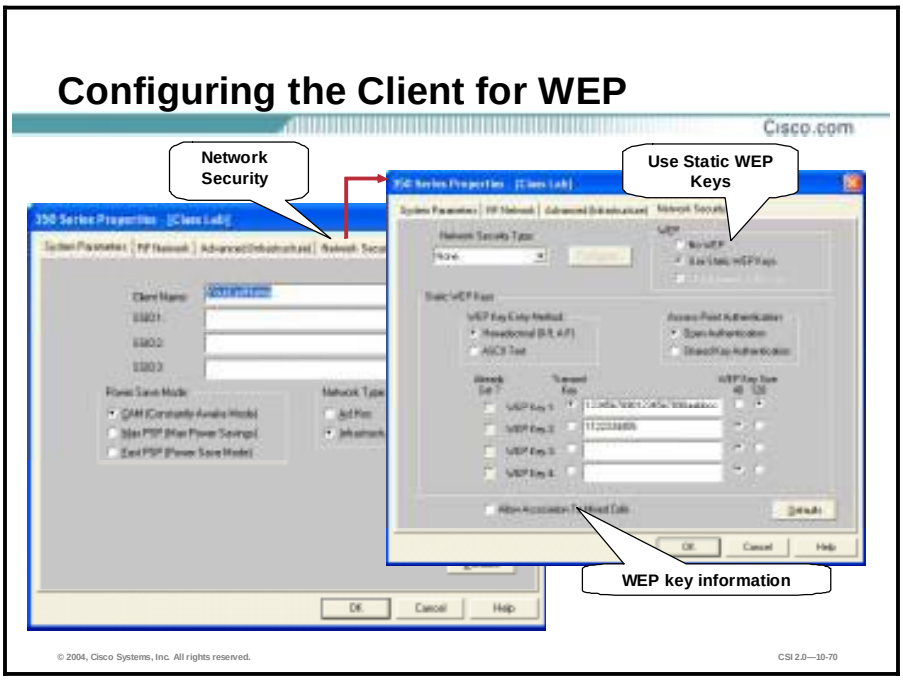

Static WEP keys are set on the client using the ACU. From the Aironet Client Utility screen, click **Profile Manager**. This will launch the Profile Manager screen. From the Profile Manager, choose the desired profile from the drop down box and click the **Edit** button.

Once the desired profile is brought up, click the **Network Security** tab. This will allow you to view and edit the security settings for this profile.

To set up static WEP keys, click the **Use Static WEP Keys** radio button under WEP. Once this button has been checked, the WEP keys can be entered. The WEP keys are entered here just as they are on the access point—26 hexadecimal characters (13 bytes) for 128-bit, 10 hexadecimal characters (5 bytes) for 40-bit. Choose the key size by checking the appropriate radio button next to the WEP key entry box. Once a WEP key is entered, it can be overwritten, but it cannot be edited or deleted. Up to four keys may be configured. Other options include the following:

- Access Point Authentication—Choose which type of authentication will be used, Open Authentication (default, more secure) or Shared Key Authentication (less secure).
- WEP Key Entry Method—You can choose to enter the WEP key as hexadecimal characters, or as ASCII text. The default is hexadecimal characters.
- Allow Association to Mixed Cells—If a client is to associate with an access point using Optional WEP (Open Authentication supporting both encrypted and non-encrypted clients), this box must be checked.

Note	If this box is not checked, the client will be able to communicate with clients configured for Full Encryption only.

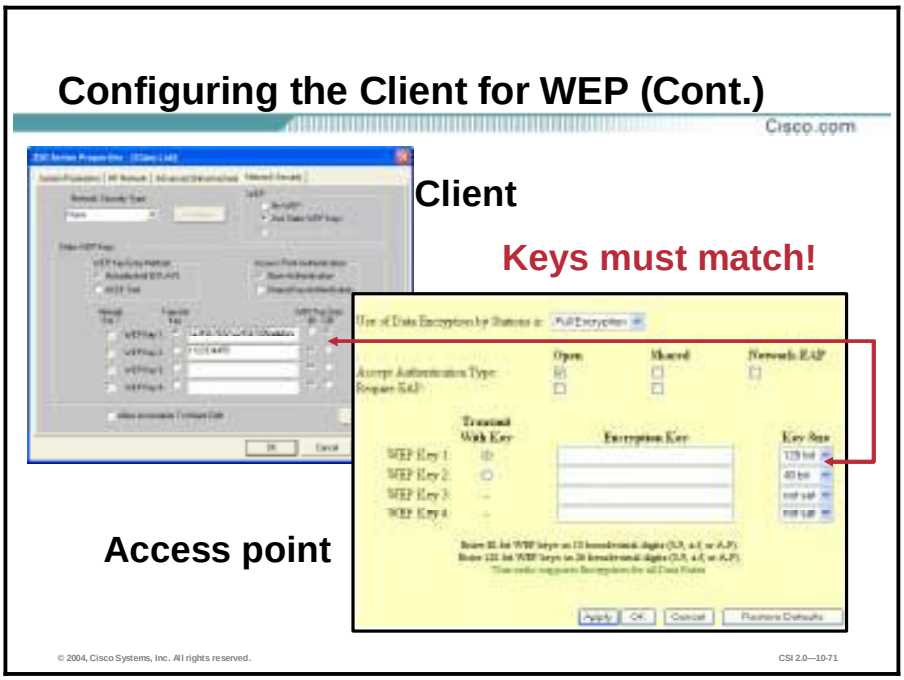

No matter which type of authentication is used, the WEP keys entered on the client and the access point must match. The key(s) themselves must match, and the order of the key(s) must match (i.e., 40 bit key entered as Key 1 on the client must match the 40 bit key entered as Key 1 on the access point).

The reason the order of the keys must match is because a transmit key will have to be chosen. When sending encrypted data, the client (or access point) will use the transmit key to encrypt the packet. The transmit key information is included in the packet's header. This lets the access point (or client) know which key to use to decrypt the packet.

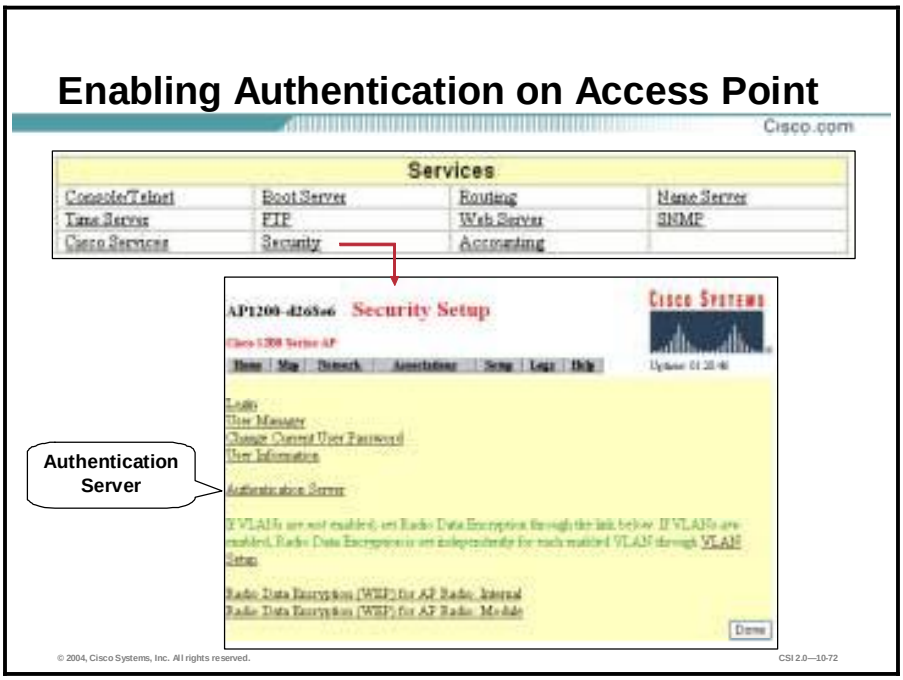

One of the main concerns with implementing WLAN technology into networks is that WLANs expose your ports to the outside of the facility, meaning that the wireless signal extends beyond the building. Without any form of security, an intruder could potentially use any 802.11-compliant card to access the network. Even when using WEP security, it might be possible for an intruder to capture network traffic outside the building and learn the system WEP keys. Because wireless traffic is broadcasted, and not directed to an individual, anyone with a wireless card could potentially get into the system.

Using the security features on the Cisco Aironet products allows for two-way verification. With RADIUS server support, the client verifies that the access point is an allowed access point while at the same time the access point verifies that the client is allowed. This means a secure channel and secure transmissions.

A user may associate to an access point but would not be granted access to network resources until the user performed a network logon. All attempts to gain access to the network resources will be blocked until the network logon is performed. And because all data is encrypted, a user trying to capture data outside the facility would not be able to use the data.

One of the biggest benefits of 802.1x is that it provides very strong authentication. Stealing or deriving a WEP key or spoofing a MAC address is no longer sufficient for gaining access to the WLAN.

In order to configure the access point for authentication using a RADIUS server, start at the Security Setup screen again. Click on the **Authentication Server** link to set the parameters for the server.

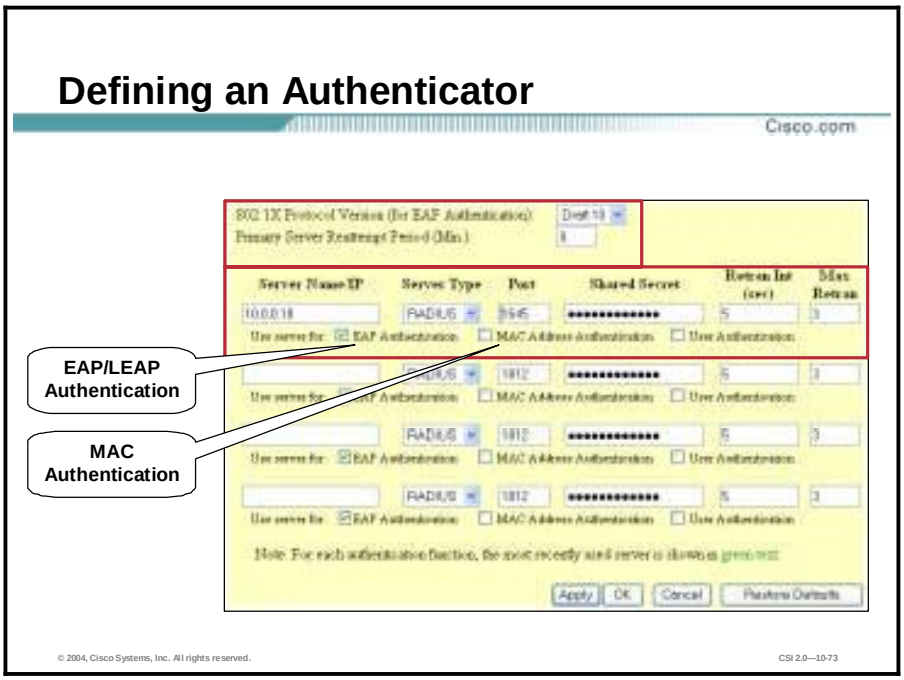

In order to use an authentication server, the access point must be configured to communicate with the server. From the Authentication Server screen, configure the following:

- 802.1X Protocol Version (for EAP Authentication)—This must match on both the APs and the clients.
- EAP Authentication—Check here if the access point will be supporting clients that must authenticate using EAP.
- MAC Address Authentication—Check here if the access point will be supporting clients that must authenticate by MAC address.
- Server Name/IP—Enter the name or IP address of the RADIUS server in the Server Name/IP entry field.
- Port—Enter the port number your RADIUS server uses for authentication. The default 1812 is the port setting for many RADIUS servers. 1645 is the port setting for Cisco's RADIUS server, the ACS. Check your server's product documentation to find the correct port setting.
- Shared Secret—Enter the shared secret used by your RADIUS server in the Shared Secret entry field. The shared secret on the access point must match the shared secret on the RADIUS server.
- Retran Int (sec)—Enter the number of seconds the access point should wait before retransmitting authentication challenges upon failure.
- Max Retran—Enter the maximum number of retransmits of authentication challenges upon failure.

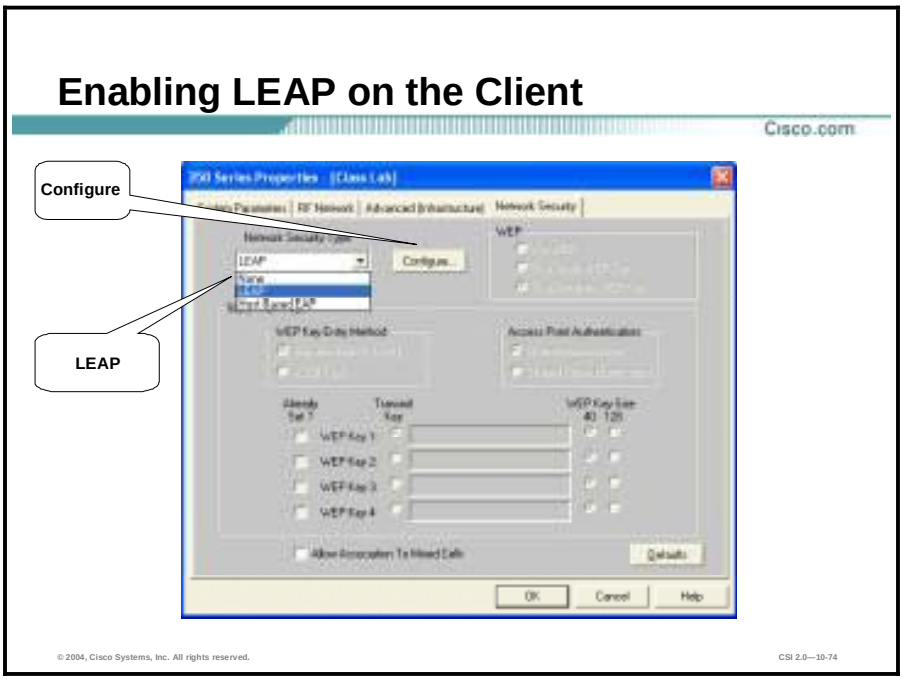

From the Profile Manager on ACU, choose the profile you wish to configure for Server Based Authentication, and click the **Edit** button. This will launch the Properties screens for the profile. Click the **Network Security** tab.

Network Security Type:

- None—Allows you to use Static WEP.
- LEAP—Leverages Cisco software and firmware to cause your network logon to trigger server-based authentication using your user name and password. Requires a LEAP-enabled RADIUS server running on the network. You can also choose the LEAP Settings.
- Host Based EAP—Allows you to use Windows XP's built-in EAP, such as EAP-TLS or EAP-MD5 or PEAP. Refer to Windows XP Help for information on how to set up EAP.

To configure LEAP, select **LEAP** from the list and click the **Configure** button.

Enabling LEAP on the Client (Cont.)

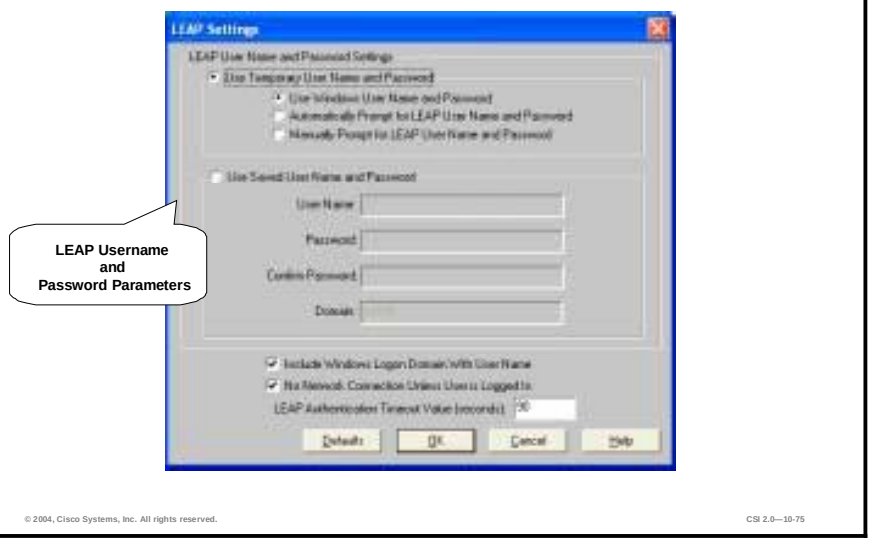

All Cisco WLAN products can be configured to take advantage of LEAP. The WLAN products are classified into two categories, Client and Authenticator.

- Clients—A Client card uses the ACU to set up and enable security. A Non-Root Bridge, Repeater Access Point, and Workgroup Bridge can be configured with a LEAP user name and password. The preset user name and password will respond automatically when they are challenged for LEAP credentials.
- Authenticators—A Root Access Point and Root Bridge can be configured to pass Client and Authentication Server credentials. Until the client's credentials are confirmed, the client is denied access to the LAN.

Cisco ACS—Main Screen

Cisco.com

CSI 2.0—10-76

To begin configuring Cisco ACS to work with Cisco Aironet access points as network access servers (NASs), open the Cisco ACS main screen.

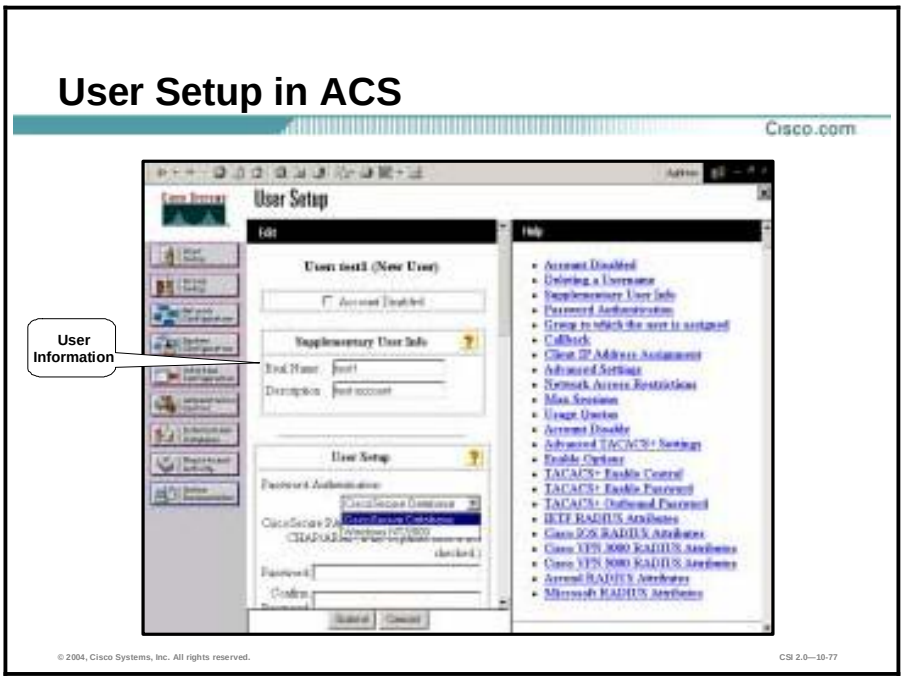

To configure accounts on Cisco ACS, click the **User Setup** button from the main screen. This will launch the User Setup screen. Enter information about the user.

Scroll down to the Password Authentication entry and choose **Cisco Secure Database** from the drop down menu. This indicates that the user account will be stored on Cisco ACS. Type the password in the **Password** box. Retype the password in the **Confirm Password** box.

Scroll down to the Group to which the user is assigned link. From the dropdown menu, choose which group the user will belong to. Unless otherwise specified, all users are assigned to the Default Group.

Scroll down to **Client IP Address Assignment**. Choose how the user will be assigned an IP address.

Scroll down to the **Account Disable** link. Choose how and if the account will be disabled. By default the account is never disabled. Using this feature allows you to set up temporary accounts with an expiration date.

When finished, click the **Submit** button. The user account is now set up and ready for use.

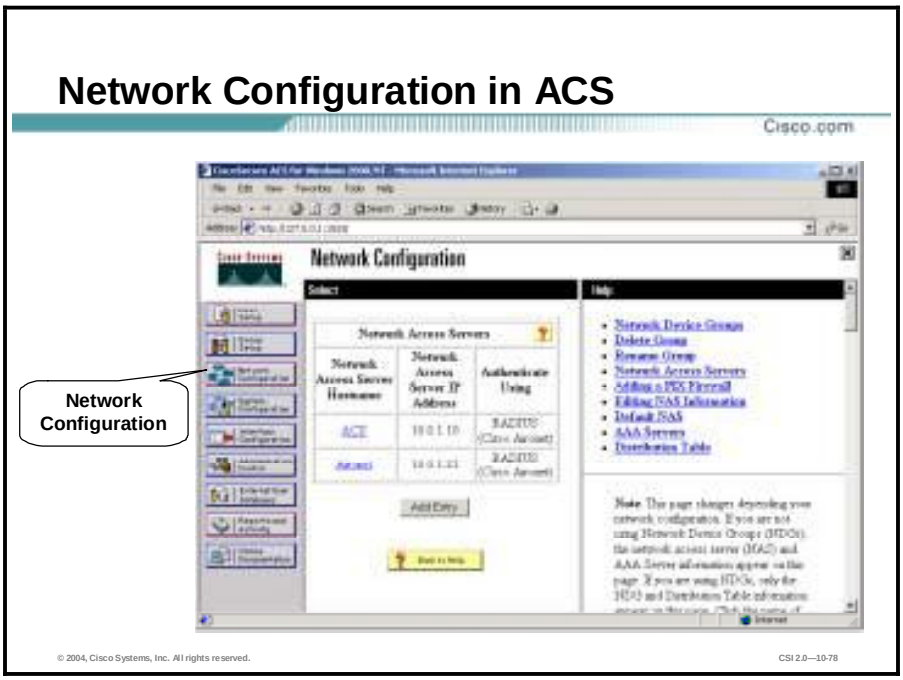

The Network Configuration screen lists all NASs that are currently configured. To add a NAS, click the **Add Entry** button. This will launch the Add Access Server screen.

Network Configuration in ACS (Cont.)

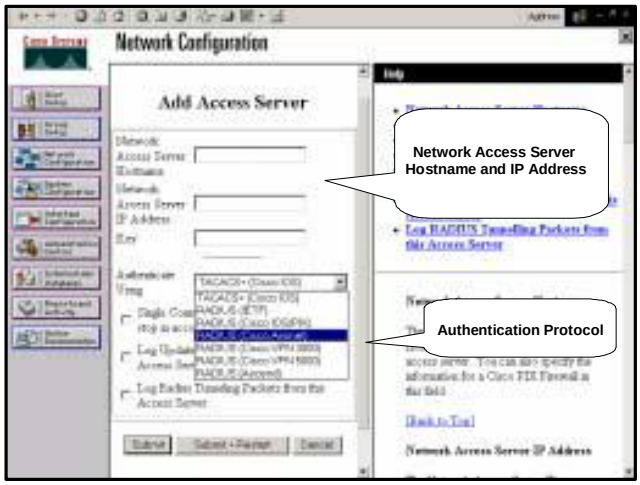

To configure Cisco ACS for use with the Cisco Aironet product as a NAS, perform the following:

On the ACS menu (left side of screen) click on the **Network Configuration** button then click **ADD Access Server**. This will bring up the Add Access Server screen. Each individual access point is considered a NAS. To configure a NAS, enter the following:

- Network Access Server Hostname—DNS name of the specified access point
- Network Access Server IP Address—IP address of the specified access point
- Key—Shared secret between the server and this individual access point
- Authenticate Using—Must select “RADIUS (Cisco Aironet)”

Each key can be different on a per-access point basis but must match the setting in the specified access point.

When finished configuring the NAS, click the **Submit+Restart** button. Cisco ACS is now ready to receive authentication requests via the NAS.

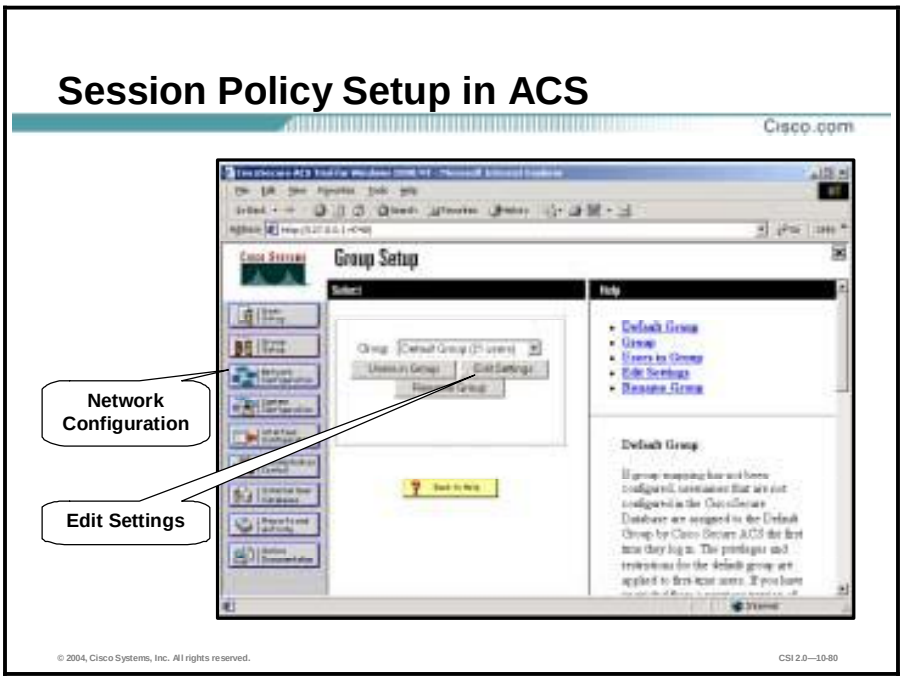

To adjust the timeout value for the client session (requiring re-authentication and new WEP key derivation as in LEAP), click the **Group Setup** button from the ACS main screen. This will launch the Group Setup screen.

From the Group Setup screen, choose a group (usually the default group) and click the **Edit Settings** button.

Session Policy Setup in ACS (Cont.)

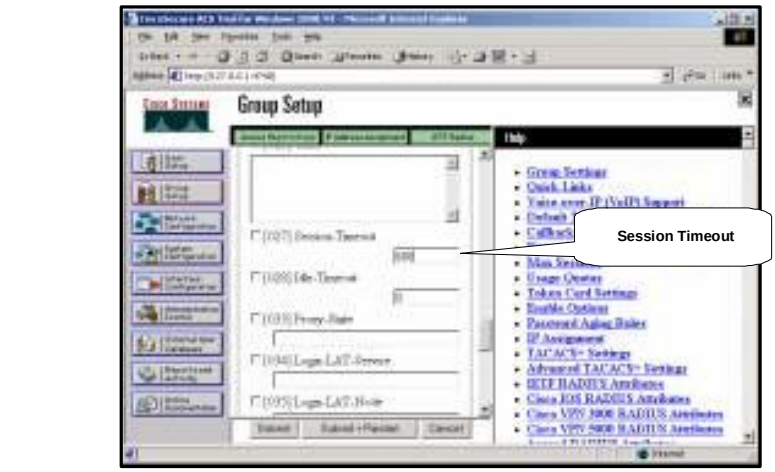

Scroll down through the setup menu and find the [027] Session Timeout entry. Enter the timeout value in seconds. The timeout value will depend upon the number of users typically attached to the access point, as well as the amount of traffic the clients will typically be sending. The larger the number of users, or the more traffic the users are sending, the smaller the value needs to be to insure that the WLAN is protected. By setting smaller values it is possible to prevent a hacker from being able to capture enough packets to hack a WEP key.

Recommended values:

Using WEP only with 802.1x, key rotation time: 15 minutes

Using WEP and TKIP with 802.1x key rotation time: 240 minutes

Note	These values apply to both session keys and broadcast key

When finished, click the **Submit+Restart** button.

Summary

Cisco.com

- IEEE 802.11 is the standard used by wireless technologies.
- Security for IEEE 802.11 networks can be simplified into two main components:
  - Encryption
  - Authentication
- There are four WLAN components.
- Described security extensions for SAFE WLAN.
- Listed two main WLAN network design choices:
  - Implementing a dynamic WEP keying model using 802.1x/EAP and TKIP
  - Implementing an overlay VPN network using IPSec

Summary (Cont.)

Cisco.com

- Discussed design considerations for small, medium, enterprise, and remote WLANs.
- The mitigation roles identified for each threat in SAFE white paper are integral to a successful WLAN implementation.
- The design process is often a series of trade-offs. Some of these trade-offs are made at the module level, while others are made at the component level.