

Securing Cisco Network Devices

Volume 2

Version 2.0

Student Guide

Text Part Number: 97-2360-01

**Corporate Headquarters**

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

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

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

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

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

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus • 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



© 2006 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0601R)

DISCLAIMER WARRANTY: THIS CONTENT IS BEING PROVIDED "AS IS." CISCO MAKES AND YOU RECEIVE NO WARRANTIES IN CONNECTION WITH THE CONTENT PROVIDED HEREUNDER, EXPRESS, IMPLIED, STATUTORY OR IN ANY OTHER PROVISION OF THIS CONTENT OR COMMUNICATION BETWEEN CISCO AND YOU. CISCO SPECIFICALLY DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE. This learning product may contain early release content, and while Cisco believes it to be accurate, it falls subject to the disclaimer above.

Table of Contents

Volume 2

<u>Securing LAN and WLAN Devices</u>	3-1
Overview	3-1
Module Objectives	3-1
<u>Applying Security Policies to Network Switches</u>	3-3
Overview	3-3
Objectives	3-3
Basic Switch Operation	3-4
Switches Are Targets	3-6
Securing Network Access to Layer 2 LAN Switches	3-8
Protecting Administrative Access to Switches	3-9
Protecting Access to the Management Port	3-12
Turning Off Unused Network Interfaces and Services	3-14
Summary	3-16
<u>Mitigating Layer 2 Attacks</u>	3-17
Overview	3-17
Objectives	3-17
Mitigating VLAN Hopping Attacks	3-18
Preventing STP Manipulation	3-21
Mitigating DHCP Server Spoofing with DHCP Snooping	3-24
Mitigating ARP Spoofing with DAI	3-26
Example: DAI Implementation	3-27
CAM Table Overflow Attacks	3-29
MAC Address Spoofing Attacks	3-34
Using Port Security to Prevent Attacks	3-35
Configuring Cisco Catalyst Switch Port Security	3-40
Layer 2 Best Practices	3-44
Summary	3-45
<u>Using Cisco Catalyst Switch Security Features</u>	3-47
Overview	3-47
Objectives	3-47
Security Features in Cisco Catalyst Switches	3-48
Identity-Based Networking Services	3-49
VLAN ACLs	3-51
Private VLANs	3-52
MAC Address Notification	3-53
Rate Limiting	3-54
SPAN for IPS	3-55
Management Encryption	3-56
Summary	3-57
<u>Securing WLANs</u>	3-59
Overview	3-59
Objectives	3-59
Introducing WLANs	3-60
Threats to WLANs	3-65
Evolution of 802.11 Security Features	3-66
Service Set Identifier	3-68
Wired Equivalent Privacy	3-69
Open Authentication	3-70
Shared Key Authentication	3-71
Passive or Weak Initialization Vector Attack	3-74
Active “Bit Flipping” or Replay Attack	3-75
Authentication Dictionary Attacks	3-75
Enhanced Methods for WLAN Threat Mitigation	3-76
WLAN IDS	3-82

Summary	3-83
Module Summary	3-84
References	3-85
Module Self-Check	3-86
Module Self-Check Answer Key	3-89
Cisco IOS Firewall Configuration	4-1
Overview	4-1
Module Objectives	4-1
Introducing Firewall Technologies	4-3
Overview	4-3
Objectives	4-3
Explaining a Firewall	4-4
Evolution of Firewall Technologies	4-5
Static Packet Filtering Firewalls	4-7
Circuit Level Firewalls	4-10
Application Layer or Proxy Firewalls	4-11
Dynamic or Stateful Packet Filtering Firewalls	4-18
Cut-Through Proxy Process	4-23
Implementing NAT on a Firewall	4-25
Application Inspection Firewall	4-31
Firewalls in a Layered Defense Strategy	4-37
Summary	4-38
Building Static Packet Filters with Cisco ACLs	4-39
Overview	4-39
Objectives	4-39
Access Control Lists	4-40
Cisco ACLs	4-41
Applying ACLs to Router Interfaces	4-47
Using ACLs to Filter Traffic	4-49
Filtering Router Service Traffic	4-52
Filtering Network Traffic to Mitigate Threats	4-56
Mitigating DDoS Attacks with ACLs	4-63
Combining Access Functions	4-69
Caveats	4-72
Summary	4-74
Configuring a Cisco IOS Firewall with the Cisco SDM Firewall Wizard	4-75
Overview	4-75
Objectives	4-75
Cisco SDM Firewall Wizard Tasks	4-76
Configuring a Basic Firewall	4-79
Configuring an Advanced Firewall	4-81
Configuring Firewall Inspection Rules	4-83
Application Security Policy Configuration	4-85
Default SDM Application Security Policies	4-85
Custom Application Security Policy Button	4-85
Delivering the Configuration to the Router	4-86
Editing Firewall Policies and ACLs	4-89
Summary	4-94
Defending Your Network with the Cisco Security Appliance Product Family	4-95
Overview	4-95
Objectives	4-95
Introducing the Cisco Security Appliance Product Family	4-96
Cisco IOS Firewall Features	4-98
When to Choose a Cisco IOS Firewall Solution	4-100
Introducing Cisco PIX 500 Series Security Appliances	4-102
Introducing Cisco ASA 5500 Series Adaptive Security Appliances	4-105

Developing an Effective Firewall Policy	4-108
Summary	4-111
Module Summary	4-112
References	4-113
Module Self-Check	4-114
Module Self-Check Answer Key	4-117
<u>Securing Networks with Cisco IOS IPS</u>	<u>5-1</u>
Overview	5-1
Module Objectives	5-1
<u>Introducing IDS and IPS</u>	<u>5-3</u>
Overview	5-3
Objectives	5-3
Introducing IDS and IPS	5-4
Types of IDS and IPS Sensors	5-10
Intrusion Prevention Technologies	5-12
Separate Monitoring Domains	5-14
Hierarchical Monitoring Structure	5-15
HIPS and Network IPS	5-17
Introducing Signatures	5-24
Examining SDFs and Signature Micro-Engines	5-29
Introducing Signature Alarms	5-39
<u>Configuring Cisco IOS IPS</u>	<u>5-45</u>
Overview	5-45
Objectives	5-45
Cisco IOS IPS Features	5-46
Configuring Cisco IOS IPS Using Cisco SDM	5-50
Using the Cisco SDM GUI for IPS	5-51
Configuring IPS Rules	5-53
Configuring IPS Signatures	5-55
Configuring Global Settings	5-59
Delivering the Configuration to the Router	5-60
Summary	5-61
<u>Defending Your Network with the Cisco IPS Product Family</u>	<u>5-63</u>
Overview	5-63
Objectives	5-63
Network IPS Solutions	5-64
HIPS Solutions	5-72
Positioning IPS Solutions	5-77
IPS Best Practices	5-79
Summary	5-82
Module Summary	5-83
References	5-84
Module Self-Check	5-86
Module Self-Check Answer Key	5-89
<u>Building IPsec VPNs</u>	<u>6-1</u>
Overview	6-1
Module Objectives	6-1
<u>Introducing IPsec VPNs</u>	<u>6-3</u>
Overview	6-3
Objectives	6-3
IPsec Overview	6-4
Internet Key Exchange	6-5
IKE: Other Functions	6-8
ESP and AH Protocols, Transport and Tunnel Modes	6-12
Message Authentication and Integrity Check	6-14

Symmetric vs. Asymmetric Encryption Algorithms	6-16
PKI Environment	6-21
Summary	6-27
Building a Site-to-Site IPsec VPN Operation	6-29
Overview	6-29
Objectives	6-29
Site-to-Site IPsec VPN Operations	6-30
Configuring IPsec	6-31
Site-to-Site IPsec Configuration—Phase 1	6-32
Site-to-Site IPsec Configuration—Phase 2	6-33
Site-to-Site IPsec Configuration—Apply VPN Configuration	6-34
Site-to-Site IPsec Configuration—Interface Access List	6-35
Summary	6-36
Configuring IPsec Site-to-Site VPNs Using Cisco SDM	6-37
Overview	6-37
Objectives	6-37
Introducing the Cisco SDM VPN Wizard Interface	6-38
Site-to-Site VPN Components	6-39
Launching the Site-to-Site VPN Wizard	6-41
Quick Setup	6-43
Step-by-Step Setup	6-44
Connection Settings	6-45
IKE Proposals	6-46
Transform Set	6-47
Defining What Traffic to Protect	6-48
Option 1: Single Source and Destination Subnet	6-48
Option 2: Using an ACL	6-49
Completing the Configuration	6-52
Testing the Tunnel Configuration and Operation	6-53
Monitoring Tunnel Operation	6-54
Advanced Monitoring	6-55
Troubleshooting	6-56
Summary	6-57
Building Remote-Access VPNs	6-59
Overview	6-59
Objectives	6-59
Cisco Easy VPN	6-60
Configuring Cisco Easy VPN Server	6-63
Managing Cisco Easy VPN Server Connections	6-76
Configuring Cisco Easy VPN Remote	6-81
Summary	6-83
Defending Your Network with the Cisco VPN Product Family	6-85
Overview	6-85
Objectives	6-85
Secure Connectivity—VPN Solutions	6-86
Secure Connectivity—Cisco VPN Product Family	6-91
Secure Connectivity—VPN Product Positioning	6-106
Cisco VPN Best Practices	6-109
Module Summary	6-116
References	6-117
Module Self-Check	6-118
Module Self-Check Answer Key	6-121

Securing LAN and WLAN Devices

Overview

LAN devices, including wireless devices operating at Layer 2, are open to attacks because of vulnerabilities that are inherent to Layer 2. This module describes Layer 2 vulnerabilities, how to account for vulnerabilities in a security policy, and how to enable the security features that are built into Cisco LAN and Cisco wireless LAN (WLAN) products.

The goal of this module is to configure LAN devices to control access, resist attacks, guard other network devices and systems, and protect the integrity and confidentiality of network traffic.

Module Objectives

Upon completing this module, you will be able to configure LAN devices to be secure. This ability includes being able to meet these objectives:

- Explain how to apply security policies to switches to mitigate Layer 2 attacks
- Explain how to mitigate attacks against network topologies and protocols
- Describe how to use the security features embedded in Cisco Catalyst switches to mitigate network threats
- Describe how to secure WLAN segments in your network

Lesson 1

Applying Security Policies to Network Switches

Overview

Anyone accessing a public network must be aware of hackers and their methods. Failure to understand what hackers do can leave you and your network exposed. While thieves and opportunists often attack an easy target versus a difficult or well-prepared target, some hackers intentionally go after very difficult targets, such as government offices or networking companies, solely for the prestige of doing so.

This lesson describes the steps needed to provide basic security to Cisco Catalyst switches in the network.


Objectives

Upon completing this lesson, you will be able to explain how to apply security policies to switches to mitigate Layer 2 attacks. This ability includes being able to meet these objectives:

- Explain how basic switch operation opens networks to attack at Layer 2
- Describe the vulnerabilities posed by unprotected network switches
- Describe the basic steps in securing network access to Layer 2 LAN switches
- Describe how to configure passwords to protect administrative access to switches
- Describe how to protect access to the management port on a switch
- Explain why unused network interfaces and services should be disabled

Basic Switch Operation


This topic explains how basic switch operation opens networks to attack at Layer 2.



Unlike hubs, switches are able to regulate the flow of data between their ports by creating “instant” networks that contain only the two end devices communicating with each other at that moment in time. When data frames are sent by end systems, their source and destination addresses are not changed throughout the switched domain. Switches maintain Content Addressable Memory (CAM) lookup tables to track the source addresses located on the switch ports. These lookup tables are populated by an address-learning process on the switch. If the destination address of a frame is not known, or if the frame received by the switch is destined for a broadcast or multicast address, the switch forwards the frame to all ports. Because of their ability to isolate traffic and create instant networks, switches can be used to divide a physical network into multiple logical segments, or VLANs, through the use of Layer 2 traffic segmenting.

Domino Effect

- If one layer is hacked, communications are compromised without the other layers being aware of the problem.
- Security is only as strong as your weakest link.
- When it comes to networking, Layer 2 can be a very weak link.



Layer 2 is the data link layer in the Open Systems Interconnection (OSI) model and is one of seven layers designed to work together but with autonomy. Layer 2 operates above the physical layer, but below the network and transport layers. Layer 2 independence enables interoperability and interconnectivity. However, from a security perspective, Layer 2 independence creates a challenge because a compromise at one layer is not always known by the other layers. If the initial attack comes in at Layer 2, the rest of the network can be compromised in an instant. Network security is only as strong as the weakest link—and that link may be the data link layer.

Switches Are Targets

This topic describes the vulnerabilities posed by unprotected network switches.

Switches Are Targets

Protection should include:

- **Constraining Telnet access**
- **SNMP read-only**
- **Turning off unneeded services**
- **Logging unauthorized access attempts**

VLANs are an added vulnerability:

- **Remove user ports from automatic trunking**
- **Use nonuser VLANs for trunk ports**
- **Set unused ports to a nonrouted VLAN**
- **Do not depend on VLAN separation**
- **Use private VLANs**

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0-3.5

Similar to the steps necessary for securing routers, these issues should be considered for securing switches:

- If an attacker can gain access to a Telnet prompt, the attacker can attempt unauthorized access. For general administrative functions, devices should have a direct local connection. Where remote access is necessary, the device should connect via a private encrypted tunnel over the production network. Such a tunnel should be preconfigured to communicate only across the specific ports required for management and reporting. The tunnel should also be locked down so that only appropriate hosts can initiate and terminate tunnels.
- Simple Network Management Protocol (SNMP) is a service used to perform network management functions. SNMP version 1 (SNMPv1) uses clear text community strings for access to information on the switch. If SNMP is necessary, consider providing read-only access to devices via SNMP and treat the SNMP community string with the same care with which you might treat a root password on a critical UNIX host. If this is not sufficient, configure the switch for SNMP version 3 (SNMPv3), which can use cryptographic hashes for authentication to protect the community string. Be aware that by introducing SNMP into your production network, you are introducing a potential vulnerability into your environment.
- The expression “less is more,” when applied to security, means that *fewer* services or ports running on a device make that device *more* secure. In keeping with this axiom, disable all unused ports on a switch. This setup prevents hackers from plugging into unused ports and communicating with the rest of the network.
- A consistent theme in tracking security compromises such as unauthorized access is logging unauthorized access attempts. You should make sure that you read the logs regularly.

VLANs, while providing segmentation, have vulnerabilities as well. Consider these factors when configuring and securing VLANs in your network:

- Ports without any need to trunk should have any trunk settings set to “off,” not “auto.” This setup prevents a host from becoming a trunk port and receiving all of the traffic that would normally reside on a trunk port.
- For ports that require trunking, you should always use a dedicated VLAN identifier and set the native VLAN to be different from any data (user) VLANs.
- When feasible for user ports, limit each port to associate a limited number of MAC addresses (perhaps two or three). This practice will mitigate MAC flooding and other attacks.
- VLANs do not provide security functions such as confidentiality and authentication. You could account for this need in a security policy that specifies filtering and stateful firewalling in addition to VLAN segmentation for a defense-in-depth approach to securing the access between two subnets.
- Procedures for carrying out change, control, and configuration analysis must be in place to ensure that a secure configuration results after changes are made. This process is especially valuable in cases where multiple organizational groups may control the same switch.
- Private VLANs (PVLANS) provide some added security. PVLANS work by limiting which ports within a VLAN can communicate with other ports in the same VLAN. Consider a standard public services segment with a web, FTP, and Domain Name System (DNS) server. If the DNS server is compromised, a hacker can pursue the other two servers, bypassing the firewall. When PVLANS are deployed, if one system is compromised, it cannot communicate with the other systems.

Securing Network Access to Layer 2 LAN Switches

This topic describes the basic steps in securing network access to Layer 2 LAN switches.

Securing Network Access at Layer 2

Follow these steps:

- **Protect administrative access to the switch.**
- **Protect the switch management port.**
- **Turn off unused network services.**
- **Lock down the ports.**
- **Use Cisco Catalyst switch security features.**

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0-3-6

The first steps in defending against Layer 2 attacks are to ensure that you configure every switch in the network with basic security in mind.

Using the Cisco Catalyst switch security features will be covered in the “Using Cisco Catalyst Switch Security Features” lesson.

Protecting Administrative Access to Switches

This topic describes how to configure passwords to protect administrative access to switches.

Protecting Administrative Access

- **Two access levels:**
 - **User level**—accessed via Telnet or SSH connections to a switch or via the console line on the switch
 - **Privileged level**—accessed after user level is established
- **The main vulnerability arises from poor password security.**

© 2006 Cisco Systems, Inc. All rights reserved.

SNB v2.0—3-7

By default, Cisco IOS switches have two levels of access: user (level 1) and privileged (level 15). The user level is typically accessed via Telnet or Secure Shell (SSH) connections to a switch or via the console line on the switch. The privileged level is typically accessed after the user level is established.


Each level is usually configured with a password. Here are some of the specific vulnerabilities associated with these passwords:

- By default, a Cisco switch shows the passwords in plain text for these settings in the configuration file: the enable password, the username password, the console line, and the vty. If an attacker collects the configuration file for the switch from the network using a network analyzer, these passwords can then be used to access this system.
- If the **enable secret** command is not used to set the enable password, or if the password on a Cisco switch is weak, an attacker may be able to obtain privileged-level access to retrieve or to change the configuration on the switch. Also, setting the same password for the enable secret passwords on multiple switches provides a single point of failure, because one compromised switch endangers other switches.
- Using the same password for both the **enable secret** command and other settings on a switch allows for a potential compromise because the password for certain settings (for example, Telnet) may be in plain text and can be collected on a network using a network analyzer. The attacker who can collect passwords going to a switch may be able to gain privileged-level access at a later time.

Password Encryption

Switch(config)#

`enable password password`



- Sets a local password to control access to various privilege levels

Switch(config)#

`enable secret [level level] {password | [encryption-type] encrypted-password}`

- Specifies an additional layer of security over the `enable password` command

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0-3.8

Using strong passwords is one of the first steps in defending switch configurations. Unfortunately, user passwords in Cisco IOS configuration files are encrypted using a scheme that is very weak by modern cryptographic standards. For that reason, the **enable password** command should no longer be used.

Use the **enable secret** command for better security. The only instance in which the **enable password** command might be tested is when the device is running in a boot mode that does not support the **enable secret** command.

Configure an **enable secret** password on each Cisco Catalyst switch.

Password Guidelines

- Use passwords at least 10 characters long
- Do not use real words
- Mix letters, numbers, and special characters
- Do not use a number for the first character of the password

Administrators should perform these tasks:

- Change passwords every 90 days
- Make sure that the enable secret password is unique for each switch
- Do not use enable secret passwords for anything else on the switch

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—3-9

Use these guidelines for creating a strong password:

- Passwords should be at least 10 characters long and not based on words.
- Include at least one character from each of the sets of letters, numbers, and special characters. Special characters include the following: ,./<>;':"[]\{}|~!@#\$%^&*()_+`-= .
- Do not use a number for the first character of the password.

The U.S. National Security Agency (NSA) recommends that administrators ensure that these policies are implemented:

- Change passwords at least once every 90 days
- Use a unique password for the enable secret password on each switch
- Use an enable secret password that is different from the passwords used for the other settings (for example, Telnet) on the same switch

Protecting Access to the Management Port

This topic describes how to protect access to the management port on a switch.

Protecting the Management Port

- **Assign a unique account for each administrator**
- **Use a strong and unique password on every switch**
- **Set a timeout**
- **Use a banner**
- **Use OOB management**

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—3-10

Every switch has a management port called the console line (line con 0) that provides direct administrative access to the switch. If the management port on the switch has settings that are too permissive, the switch is susceptible to attacks. The management port is a source of these potential vulnerabilities:

- A switch with a management port using a default user account allows an attacker to attempt to make connections using one or more of the well-known default user accounts (for example, administrator, root, and security). To mitigate this threat, set up a unique account for each administrator for access to the console line. Varying privilege levels, from 0 to 15, can be set on each administrator account. Privilege level 0 is the lowest level on Cisco switches and allows a very small set of commands.
- Bad passwords pose multiple vulnerabilities:
 - A missing or weak password allows an attacker to guess or crack the password and then retrieve or change the configuration on the switch.
 - Using the same password for the management port on multiple switches provides a single point of network failure. The attacker who compromises one switch can then compromise other switches.
 - Using the same password for the management port and other settings on a switch allows for potential compromise. If the password for certain settings (for example, Telnet) is in plain text, these passwords can be collected on a network using a network analyzer. The attacker who collects Telnet passwords from network traffic going to a switch may be able to access the switch management port at a later time.

- If the connections to a management port on a switch do not have a timeout period set or have a long timeout period, the connections are more available for an attacker hijack.
- A banner gives notice to anyone who connects to a switch that the network is for authorized use only and that any use of the network is monitored. Courts have dismissed cases against those who have attacked systems without banners. Not having a banner on a switch may lead to legal or liability problems.
- In terms of network design, use out-of-band (OOB) management. This approach separates management traffic from operational traffic, thus preserving operational bandwidth.

Turning Off Unused Network Interfaces and Services

This topic explains why unused network interfaces and services should be disabled.

Turning Off Unused Network Services

Enabled network services open vulnerabilities for these reasons:

- Many connections are unencrypted.
- Default user accounts allow unauthorized entry.
- Weak and shared passwords on services open doors for attackers.
- Extended timeouts allow hijacking.

“Less is more.”

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—3-11

Switches may be running network services such as SNMP and Cisco Discovery Protocol (CDP). Many services are typically not necessary for normal operation. Some services are enabled by default and others are sometimes left enabled even though they are no longer necessary. Leaving unused network services enabled increases the possibility of those services being maliciously exploited and susceptible to information gathering or to network attacks.

The figure provides several reasons for turning off or restricting access to these services. To improve network security, consider these factors:

- Connections to many of the services on a switch are not encrypted. Therefore, an attacker may be able to collect network traffic related to these services using a network analyzer. The traffic may contain usernames, passwords, or other configuration information related to the switch.
- Just like the management port, any other network service using a default user account allows an attacker to attempt to make connections using one or more of the well-known default user accounts.
- A network service set with no password or using a default password or a weak password presents vulnerability. Setting the same password for the network service on multiple switches provides a single point of failure. The attacker who compromises one switch can compromise other switches.
- Broad access that allows all systems or a large number of systems to connect to a network service on a switch makes the switch vulnerable to attack.
- Similar to the management port, all services should have a timeout to reduce hijack attempts.

Shutting Down Interfaces

```
Switch(config)# interface fastethernet 0/1  
Switch(config-if)# shutdown
```

- Shuts down a single interface

```
Switch(config)# interface range fastethernet 0/2 - 8  
Switch(config-if-range)# shutdown
```

- Shuts down a range of interfaces

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--3-12

The figure shows the use of the **shutdown** interface configuration command to disable an unused port or a group of unused ports.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- **Layer 2 vulnerabilities often escape notice, but network security is only as strong as its weakest link.**
- **Switches are targets because they can give attackers access to an entire network.**
- **Five basic steps can mitigate Layer 2 attacks.**
- **Use passwords to protect administrative access to switches.**
- **Protect the management port by assigning unique accounts and using strong passwords, timeouts, banners, and OOB management.**
- **Turn off unused network services and interfaces.**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—3-13

Mitigating Layer 2 Attacks

Overview

There are fundamental vulnerabilities inherent in Layer 2 topologies and protocols that expose all devices on a LAN to attacks. Attacks arising from these vulnerabilities must be mitigated.


Objectives

Upon completing this lesson, you will be able to explain how to mitigate attacks against network topologies and protocols. This ability includes being able to meet these objectives:

- Explain how to configure VLANs to mitigate VLAN hopping attacks
- Explain how to prevent STP manipulation
- Describe how to use the DHCP snooping feature to mitigate man-in-the-middle attacks
- Explain how to mitigate ARP spoofing with DAI
- Describe how an attacker can flood a switch by launching a CAM table overflow attack
- Describe how an attacker launches a MAC spoofing attack
- Describe port security as a key step in defending networks from Layer 2 attacks
- Describe how to configure port security on a Cisco Catalyst switch
- Explain how specific best practices mitigate attacks on specific areas of Layer 2 hardware and software components

Mitigating VLAN Hopping Attacks

This topic explains how to configure VLANs to mitigate VLAN hopping attacks.



VLAN architecture simplifies network maintenance and improves performance. However, VLAN operation opens the door to abuse. VLAN hopping allows traffic from one VLAN to be seen by another VLAN without first crossing a router. Under certain circumstances, attackers can sniff data and extract passwords and other sensitive information at will. The attack works by taking advantage of an incorrectly configured trunk port. By default, trunk ports have access to all VLANs and route traffic for multiple VLANs across the same physical link, generally between switches. The data moving across these links can be encapsulated with IEEE 802.1Q or an Inter-Switch Link (ISL).


In a basic VLAN hopping attack, the attacker takes advantage of the default automatic trunking configuration on most switches. By tricking a switch into thinking it is another switch with a need to trunk, an attacker can gain access to all the VLANs allowed on the trunk port. This attack requires a trunking-favorable setting such as “auto” to succeed. As a result, the attacker is a member of all the trunked VLANs on the switch and can “hop,” that is, send and receive traffic, on those VLANs.

A VLAN hopping attack can be launched in one of these two ways:

- **Spoofing Dynamic Trunking Protocol (DTP) messages from the attacking host to cause the switch to enter trunking mode:** From here, the attacker can send traffic tagged with the target VLAN, and the switch then delivers the packets to the destination.
- **Introducing a rogue switch and turning trunking on:** The attacker can then access all of the VLANs on the victim switch from the rogue switch.

The best way to prevent a basic VLAN hopping attack is to turn off trunking on all ports except the ones that specifically require trunking. On the required trunking ports, disable DTP (auto trunking) negotiations and manually enable trunking.

VLAN Hopping by Double Tagging



Note: This attack works only if the trunk has the same native VLAN as the attacker.

- The attacker sends double-encapsulated 802.1Q frames.
- The switch performs only one level of decapsulation.
- Only unidirectional traffic is passed.
- The attack works even if the trunk ports are set to "off".

Note: There is no way to execute these attacks unless the switch is misconfigured.

The double tagging (or double-encapsulated) VLAN hopping attack takes advantage of the way that hardware on most switches operates. Most switches perform only one level of IEEE 802.1Q decapsulation and allow an attacker, in specific situations, to embed a hidden 802.1Q tag inside the frame. This tag allows the frame to go to a VLAN that the outer 802.1Q tag did not specify. An important characteristic of the double-encapsulated VLAN hopping attack is that it works even if trunk ports are set to "off."

A double-tagging VLAN hopping attack follows these four steps:

- Step 1** The attacker sends a double-tagged 802.1Q frame to the switch. The outer header has the VLAN tag of the attacker, which is the same as the native VLAN of the trunk port. (For the purposes of this example, assume VLAN 10.) The inner tag is the victim VLAN, VLAN 20.
- Step 2** The frame arrives on the switch, which looks at the first 4-byte 802.1Q tag. The switch sees that the frame is destined for VLAN 10 and sends it out on all VLAN 10 ports (including the trunk), because there is no Content Addressable Memory (CAM) table entry. At this point, the second VLAN tag is still intact and has not been inspected by the first switch.
- Step 3** The frame arrives at the second switch but has no knowledge that it was supposed to be for VLAN 10. (Native VLAN traffic is not tagged by the sending switch as specified in the 802.1Q specification.)
- Step 4** The second switch looks at only the 802.1Q tag (the former inner tag that the attacker sent) and sees that the frame is destined for VLAN 20 (the victim VLAN). The second switch sends the packet on to the victim port or floods it, depending on whether there is an existing CAM table entry for the victim host.

The figure illustrates the attack. It is important to note that this attack is unidirectional and works only when the attacker and trunk port have the same native VLAN. Thwarting this type of attack is not as easy as stopping basic VLAN hopping attacks. The best approach is to ensure that the native VLAN of the trunk ports is different than the native VLAN of the user ports.

Mitigating VLAN Hopping Network Attacks

Example 1: If no trunking is required on an interface

```
Router(config-if)# switchport mode access
```

- **Disable trunking on the interface.**

Example 2: If trunking is required

```
Router(config-if)# switchport mode trunk  
Router(config-if)# switchport nonegotiate
```

- **Enable trunking but prevent DTP frames from being generated.**

Example 3: If trunking is required

```
Router(config-if)# switchport trunk native vlan vlan number
```

- **Set the native VLAN on the trunk to an unused VLAN.**

© 2006 Cisco Systems, Inc. All rights reserved.SND v2.0-3.5

To prevent a VLAN hopping attack that is using double 802.1Q encapsulation, the switch must look further into the packet to determine whether more than one VLAN tag is attached to a given frame. Unfortunately, the ASICs that are used by most switches are only hardware optimized to look for one tag and then to switch the frame. The issue of performance versus security requires administrators to balance their requirements carefully.

Mitigating VLAN hopping attacks using double 802.1Q encapsulation requires several modifications to the VLAN configuration. One of the more important elements is to use a dedicated native VLAN for all trunk ports. This attack is easy to stop if you follow the best practice that native VLANs for trunk ports should never be used anywhere else on the switch. Also, disable all unused switch ports and place them in an unused VLAN.


Set all user ports to nontrunking mode by explicitly turning off DTP on those ports that can be used to mitigate a VLAN hopping attack using switch spoofing.

To turn on control trunking for ports you have these options:

- If you do not intend to trunk across those links, use the **switchport mode access** interface configuration command to disable trunking.
- If you do intend to trunk across those links, take these actions:
 - Use the **switchport mode trunk** and **switchport nonegotiate** interface configuration commands to cause the interface to become a trunk but to not generate DTP frames.
 - Use the **switchport trunk native vlan** *vlan_number* command to set the native VLAN on the trunk to an unused VLAN. The default native VLAN is VLAN 1.

Preventing STP Manipulation

This topic explains how to prevent Spanning Tree Protocol (STP) manipulation.




Another attack against switches involves STP.

STP maintains a loop-free topology in a redundant Layer 2 infrastructure. Upon bootup, the switches begin a process of determining a loop-free topology. The switches identify one switch as a root bridge and block all other redundant data paths.

STP sends messages using bridge protocol data units (BPDUs) describing the configuration, topology change notification, and topology change acknowledgment.

STP Attack (Cont.)



By spoofing as the root bridge, the network attacker hopes to fool the other switches by acting as the root bridge in the topology. The attacker broadcasts STP configuration or topology change BPDUs in an attempt to force spanning-tree recalculations.

The BPDUs sent out by the attacker system announce that the attacking system has a lower bridge priority, which causes the attacker system to be elected as the root bridge. If successful, the attacker PC receives the user frames because each frame flows through the attacker PC posing as the root bridge.

The figure illustrates how a network attacker can use STP to change the topology of a network so that it appears that the attacker host is a root bridge. By transmitting spoofed STP BPDU packets, the attacker causes the switches to initiate STP recalculations that result in all traffic between the two switches flowing through the attacker PC.

Mitigating STP Attacks with bpduguard and guard root Commands

```
IOS(config)#spanning-tree portfast bpduguard
```

- Mitigates STP manipulation with bpduguard command

```
IOS(config-if)#spanning-tree guard root
```

- Mitigates STP manipulation with guard root command

© 2006 Cisco Systems, Inc. All rights reserved.SND v2.0—3-8

To mitigate STP root bridge manipulation, use the **spanning-tree guard root** interface configuration command.

The root guard feature provides a way to enforce the root bridge placement in the network. Root guard must be enabled on all ports where the root bridge should not appear. If the bridge receives superior STP BPDUs on a root guard-enabled port, this port is moved to a root-inconsistent STP state (effectively equal to listening state), and no traffic is forwarded across this port.


To prevent a rogue switch from connecting to a switch port, STP BPDU guard can be configured using the **spanning-tree portfast bpduguard default** global configuration command. BPDU guard allows network designers to keep the active network topology predictable. Globally enabled BPDU guard disables any PortFast port that receives a BPDU message. Because these PortFast ports are end-user ports, no BPDU messages should be sent to them. Although a BPDU guard may seem unnecessary because the administrator can set the bridge priority of a desired switch to zero, there is still no guarantee that it will be elected as the root bridge. There may still be a bridge with priority zero and a lower bridge ID. BPDU guard is best deployed toward user-facing ports to prevent rogue switch network extensions by an attacker.

BPDU guard and root guard are similar, but their impact is different. BPDU guard disables the port upon BPDU reception if PortFast is enabled on the port. This feature effectively denies devices behind the participation of such ports in STP. The port that is put into an error-disable state requires manual intervention to be re-enabled, or you must configure an error-disable timeout.

Root guard allows the device to participate in STP as long as the device does not try to become the root. If root guard blocks the port, subsequent recovery is automatic as soon as the offending device stops sending superior BPDUs.

Mitigating DHCP Server Spoofing with DHCP Snooping

This topic describes how to use the DHCP snooping feature to mitigate man-in-the-middle attacks.




One way that an attacker can gain access to network traffic is to spoof responses that would be sent by a valid DHCP server. The DHCP spoofing device replies to client DHCP requests. The legitimate server may also reply, but if the spoofing device is on the same segment as the client, its reply to the client may arrive first. The intruder DHCP reply offers an IP address and supporting information that designates the intruder as the default gateway or Domain Name System (DNS) server. In the case of a gateway, the clients will then forward packets to the attacking device, which will, in turn, send them to the desired destination. This is referred to as a man-in-the-middle attack, and it may go entirely undetected as the intruder intercepts the data flow through the network.

Another type of DHCP attack is the DHCP starvation attack in which the attacker PC continually requests IP addresses from a real DHCP server by changing their source MAC addresses. If successful, this kind of DHCP attack causes all of the leases on the real DHCP server to be allocated, thus preventing the real users (DHCP clients) from obtaining an IP address.

To prevent DHCP attacks, use the DHCP snooping and the port security feature on the Cisco Catalyst switches.

DHCP Snooping

- DHCP snooping allows the configuration of ports as trusted or untrusted.
 - Trusted ports can send DHCP requests and acknowledgements.
 - Untrusted ports can forward only DHCP requests.
- DHCP snooping enables the switch to build a DHCP binding table that maps a client MAC address, IP address, VLAN, and port ID.
- Use the `ip dhcp snooping` command.



DHCP snooping is a Cisco Catalyst feature that determines which switch ports can respond to DHCP requests. Ports are identified as trusted and untrusted. Trusted ports can source all DHCP messages; untrusted ports can source requests only. Trusted ports host a DHCP server or can be an uplink toward the DHCP server. If a rogue device on an untrusted port attempts to send a DHCP response packet into the network, the port is shut down. This feature can be coupled with DHCP option 82 in which switch information, such as the port ID of the DHCP request, can be inserted into the DHCP request packet.

Untrusted ports are those not explicitly configured as trusted. A DHCP binding table is built for untrusted ports. Each entry contains a client MAC address, IP address, lease time, binding type, VLAN number, and port ID recorded as clients make DHCP requests. The table is then used to filter subsequent DHCP traffic. From a DHCP snooping perspective, untrusted access ports should not send any DHCP server responses, such as DHCP OFFER, DHCP ACK, or DHCP NAK.

This brief series of steps illustrates how DHCP snooping is configured on Cisco IOS switches:

Step 1 Enable DHCP snooping:

```
Switch(config)#ip dhcp snooping
```

Step 2 Enable DHCP snooping for specific VLANs:

```
Switch(config)#ip dhcp snooping vlan number [number]
```

Step 3 Define ports as trusted or untrusted at the interface level by defining the trusted ports:


```
Switch(config-if)# ip dhcp snooping trust
```

Step 4 (Optional) Set a rate limit on the amount of DHCP messages allowed per second through untrusted ports to limit the rate at which an attacker can continually send bogus DHCP request to the DHCP server:

```
Switch(config-if)# ip dhcp snooping limit rate rate
```

Mitigating ARP Spoofing with DAI

This topic explains how to mitigate Address Resolution Protocol (ARP) spoofing with Dynamic ARP Inspection (DAI).



ARP spoofing attacks, or ARP cache poisoning, occurs when ARP allows a gratuitous reply from a host even if an ARP request is not received. After the attack, all traffic from the device under attack flows through the attacker computer and then to the router, switch, or host. An ARP spoofing attack can target hosts, switches, and routers connected to your Layer 2 network by poisoning the ARP caches of systems connected to the subnet and by intercepting traffic intended for other hosts on the subnet. The figure shows an example of ARP cache poisoning.


The figure shows that hosts A, B, and C are connected to the switch on interfaces A, B, and C, all of which are on the same subnet. Their IP and MAC addresses are also shown. In this example, host A uses IP address 10.1.1.2 and MAC address A.A.A.A. When host A needs to communicate to host B at the IP layer, host A broadcasts an ARP request for the MAC address associated with IP address 10.1.1.1.

When host B receives the ARP request, it populates its ARP cache with an ARP binding for a host with the IP address 10.1.1.2 and MAC address A.A.A.A. When host B responds, host A populates its ARP cache with a binding for a host with IP address 10.1.1.1 and MAC address B.B.B.B.

Host C can poison the ARP caches of host A and host B by broadcasting forged ARP responses with bindings for a host with an IP address of 10.1.1.2 (or 10.1.1.1) and a MAC address of C.C.C.C. Hosts with poisoned ARP caches use the MAC address C.C.C.C as the destination MAC address for traffic intended for 10.1.1.2 or 10.1.1.1. By poisoning the ARP caches of host A and host B, host C intercepts traffic intended for them. Because host C knows the true MAC addresses associated with the host A and host B IP addresses, host C can forward the intercepted traffic to those hosts by using the correct MAC address as the destination. Host C has inserted itself into the traffic stream from host A to host B, which is the topology of the classic man-in-the-middle attack.

Mitigating Man-in-the-Middle Attacks with DAI

MAC or IP Tracking Built on DHCP Snooping



DAI provides protection against attacks such as ARP poisoning using spoofing tools such as ettercap, dsniiff, and arpspoof.

The DAI feature of Cisco Catalyst switches stops ARP spoofing attacks. Like DHCP snooping, DAI uses the concept of trusted and untrusted ports to decide which ARP packets must be inspected. In this process, DAI intercepts all ARP packets and examines them for proper MAC-to-IP bindings using the DHCP binding table that was built when DHCP snooping was enabled. If an ARP packet arrives on a trusted port, no examination is made. If the packet arrives on an untrusted port, the ARP is examined and compared against the DHCP binding table, as explained in the “Dynamic ARP Inspection Commands” table.


Dynamic ARP Inspection Commands

Command	Description
Switch(config)# ip arp inspection vlan <i>vlan_id</i> [,<i>vlan_id</i>]	Enables DAI on a VLAN or range of VLANs
Switch(config-if)# ip arp inspection trust	Enables DAI on an interface and sets the interface as a trusted interface

Example: DAI Implementation

This example shows how to configure dynamic ARP inspection for hosts on VLAN 1, where client devices are located for switch 2. All client ports are untrusted by default. Only port 3/3 is trusted, because that is the only port where DHCP replies would be expected.

```
Switch S2(config)#ip arp inspection vlan 1
Switch S2(config)#interface fastethernet 3/3
Switch S2(config-if)#ip arp inspection trust
```




When a host sends an ARP request to resolve its own IP address, it is called gratuitous ARP (GARP). In a properly configured network, an ARP reply is not provided for a GARP request. However, if another host in the network appears to be configured with the same IP address as the source host, the source host gets an ARP reply. In this way, a host can determine whether another host is also configured with its IP address.

The figure shows a user with an IP address of 10.1.1.2 connected through a switch to a default gateway with an IP address of 10.1.1.1. An intruder residing on an untrusted port sends a GARP in an attempt to reset IP-to-MAC bindings so that all traffic from 10.1.1.2 to the 10.1.1.1 default gateway goes to the attacker. The attacker poisoned the ARP cache of 10.1.1.2, so 10.1.1.2 thinks the attacker MAC address is the MAC address of the 10.1.1.1 default gateway.

DAI examines the ARP packet and compares its information with the information in the switch DHCP binding table. In the DHCP binding table, by snooping the DHCP messages, the switch knows the proper MAC address to IP address bindings. Because in this case there will be no match for the 10.1.1.1 IP address to the attacker MAC address, the ARP packet is dropped, and the port is locked.

CAM Table Overflow Attacks


This topic describes how an attacker can flood a switch by launching a CAM table overflow attack.



The CAM table in a switch contains the MAC addresses available on a given physical port of a switch and the associated VLAN parameters for each. When a Layer 2 switch receives a frame, the switch looks in the CAM table for the destination MAC address. If an entry exists for the MAC address in the CAM table, the switch forwards the frame to the MAC address port designated in the CAM table. If the MAC address does not exist in the CAM table, the switch acts like a hub and forwards the frame out every port on the switch.

CAM table overflow attacks are sometimes referred to as MAC flooding attacks. To understand the mechanism of a CAM table overflow attack, recall the basic operation of a switch.


In the figure, host A (MAC A) sends traffic to host B (MAC B). The switch receives the frames and looks up the destination MAC address in its CAM table. If the switch cannot find the destination MAC in the CAM table, the switch then copies the frame and broadcasts it out every switch port.



Host B receives the frame and sends a reply to host A. The switch then learns that the MAC address for host B is located on port 2 and writes that information into the CAM table.


Host C also receives the frame from host A to host B, but because the destination MAC address of that frame is host B, host C drops that frame.

CAM Table Is Updated—Flooding Stops



Now, any frame sent by host A (or any other host) to host B is forwarded to port 2 of the switch and not broadcast out every port.

The key to understanding how CAM overflow attacks work is to know that CAM tables are limited in size. MAC flooding makes use of this limitation to bombard the switch with fake source MAC addresses until the switch CAM table is full. The switch then enters into what is known as a fail-open mode, starts acting as a hub, and broadcasts packets to all the machines on the network. As a result, the attacker can see all of the frames sent from a victim host to another host without a CAM table entry.




An attacker can use the normal operating characteristics of the switch to stop the switch from operating.

MAC flooding can be performed using macof, a utility that comes with the dsniff suite. Dsniff is a collection of tools for network auditing and penetration testing. A network intruder can use the macof tool to flood the switch with a large number of invalid source MAC addresses until the CAM table fills up. When the CAM table is full, the switch floods all ports with incoming traffic because it cannot find the port number for a particular MAC address in the CAM table. The switch, in essence, acts like a hub.

Dsniff (macof) can generate 155,000 MAC entries on a switch per minute. Depending on the switch, the maximum CAM table size will vary.

In the figure, the macof program is running on the host with MAC address C, in the bottom right of the screen. This tool floods a switch with packets containing randomly generated source and destination MAC and IP addresses. Over a short period of time, the CAM table in the switch fills up until it cannot accept new entries. When the CAM table fills up with these invalid source MAC addresses, the switch begins to forward all frames that it receives to every port.


The CAM Table Overflows—Switch Crumbles Under the Pressure



As long as macof is left running, the CAM table on the switch will remain full. When this happens, the switch begins to broadcast all received frames out every port so that frames sent from host A to host B are also broadcast out of port 3 on the switch.

MAC Address Spoofing Attacks

This topic describes how an attacker launches a MAC spoofing attack.



In a MAC spoofing attack, the network attacker uses a known MAC address to attempt to make the targeted switch forward frames destined for the remote host to the network attacker. By sending a single frame with the source Ethernet address of another host, the network attacker overwrites the CAM table entry so that the switch forwards packets destined for the host to the network attacker. From then on, the host being spoofed does not receive any traffic until it sends traffic to again reset the CAM table entry to point back to the original port.

This figure shows how MAC spoofing works, as discussed here:

- **Top-left illustration:** Under a normal operating environment, the switch learns that host A is on port 1, host B is on port 2, and host C is on port 3. The CAM table reflects this situation.
- **Top-right illustration:** In an attack, the network attacker causes host B to send a packet identifying itself using the IP address of host B but the MAC address of host A.
- **Bottom-left illustration:** The switch now moves the location of host A in its CAM table from port 1 to port 2. Traffic from host C destined to host A is now visible to host B and is therefore compromised.
- **Bottom-right illustration:** To correct this situation, host A must send out traffic on the switch port for the switch to relearn the port that is actually associated with the host A MAC address. However, until that happens, the door is open to intruders.

Using Port Security to Prevent Attacks

This topic describes port security as a key step in defending networks from Layer 2 attacks.

Using Port Security to Mitigate Attacks

Port security can mitigate attacks by these methods:

- **Blocking input to a port from unauthorized MAC addresses**
- **Filtering traffic to or from a specific host based on the host MAC address**

Port security mitigates these:

- **CAM table overflow attacks**
- **MAC address spoofing attacks**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—3-20

You can use port security to block input to an Ethernet, Fast Ethernet, or Gigabit Ethernet port when the MAC address of the station attempting to access the port is different from any of the MAC addresses specified for that port. Alternatively, you can use port security to filter traffic destined to or received from a specific host based on the host MAC address.

By limiting the number of valid MAC addresses allowed on a port, the port security feature is an effective mitigation against CAM table overflow and MAC address spoofing attacks.

Port Security Fundamentals

- This feature restricts input to an interface by limiting and identifying MAC addresses of end devices.
- Secure MAC addresses are included in an address table in one of these ways:
 - Use the switchport port-security mac-address *mac_address* interface configuration command to configure all secure MAC addresses
 - Allow the port to dynamically configure secure MAC addresses with the MAC addresses of connected devices
 - Configure some addresses and allow the rest to be configured dynamically
- Configure “restrict” or “shutdown” violation rules.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—3-21

A switch that does not provide port security allows an attacker to attach a system to an unused, enabled port and to perform information gathering or attacks. A switch can be configured to act like a hub, which means that every system connected to the switch can potentially view all network traffic passing through the switch to all systems connected to the switch. Thus, an attacker could collect traffic that contains usernames, passwords, or configuration information about the systems on the network.

Port security limits the number of valid MAC addresses allowed on a port. All switch ports or interfaces should be secured before the switch is deployed. In this way, you can set or remove the security features as required, instead of adding and strengthening features randomly or as the result of a security incident.

You can use the port security feature to restrict input to an interface by limiting and identifying MAC addresses of the end devices that are allowed to access the port. When you assign secure MAC addresses to a secure port, the port does not forward packets with source addresses outside the group of defined addresses. If you limit the number of secure MAC addresses to one and assign a single secure MAC address to that port, the workstation attached to that port is assured the full bandwidth of the port and only that workstation with that particular secure MAC address can successfully connect to that switch port.

If a port is configured as a secure port and the maximum number of secure MAC addresses is reached, a security violation occurs when the MAC address of a workstation attempting to access the port is different from any of the identified secure MAC addresses.

After you have set the maximum number of secure MAC addresses on a port, the secure addresses are included in an address table in one of these ways:

- You can configure all secure MAC addresses by using the **switchport port-security mac-address mac_address** interface configuration command when using a Cisco IOS Catalyst switch.
- You can allow the port to dynamically learn the secure MAC addresses with the MAC addresses of connected devices.
- You can configure a number of static secure MAC addresses and allow the rest to be dynamically learned.

You can configure the interface for one of these violation modes based on the action taken if a violation occurs:

- **Protect:** When the number of secure MAC addresses reaches the limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses or increase the number of maximum allowable addresses. You are not notified that a security violation has occurred.
- **Restrict:** When the number of secure MAC addresses reaches the limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses or increase the number of maximum allowable addresses. In this mode, you are notified that a security violation has occurred. Specifically, a Simple Network Management Protocol (SNMP) trap is sent, a syslog message is logged, and the violation counter changes incrementally.
- **Shutdown:** In this mode, a port security violation causes the interface to immediately become error-disabled, and turns off the port LED. The interface also sends an SNMP trap, logs a syslog message, and registers in the violation counter. When a secure port is in the error-disabled state, you can bring it out of this state by entering the **errdisable recovery cause psecure-violation** global configuration command, or you can manually re-enable it by entering the **shutdown** and **no shutdown** interface configuration commands. Shutdown is the default mode.

The “Security Violation Mode Actions” table provides a summary of these modes.

Security Violation Mode Actions

Violation Mode	Forwards Traffic	Sends SNMP Trap	Sends Syslog Message	Displays Error Message	Increases Violation Counter	Shuts Down Port
Protect	No	No	No	No	No	No
Restrict	No	Yes	Yes	No	Yes	No
Shutdown	No	Yes	Yes	No	Yes	Yes

Port Security Configuration

Secure MAC addresses are these types:

- Static secure MAC addresses
- Dynamic secure MAC addresses
- Sticky secure MAC addresses

Security violations occur in these situations:

- A station whose MAC address is not in the address table attempts to access the interface when the table is full.
- An address is being used on two secure interfaces in the same VLAN.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—3-22

Ports can be configured with these types of secure MAC addresses:

- **Static secure MAC addresses:** These addresses are manually configured by using the **switchport port-security mac-address *mac-address*** interface configuration command, stored in the address table, and added to the switch running configuration.
- **Dynamic secure MAC addresses:** These addresses are dynamically configured, stored only in the address table, and removed when the switch restarts.
- **Sticky secure MAC addresses:** These addresses are dynamically configured, stored in the address table, and added to the running configuration. The sticky secure MAC addresses do not automatically become part of the configuration file, which is the startup configuration used each time the switch restarts. If you save the sticky secure MAC addresses in the startup configuration file, when the switch restarts, the interface does not need to relearn these addresses. If you do not save the configuration, they are lost.

A security violation occurs in these situations:

- The maximum number of secure MAC addresses has been added to the address table and a station whose MAC address is not in the address table attempts to access the interface.
- An address that is learned or configured on one secure interface is seen on another secure interface in the same VLAN.

Port Security Defaults

Feature	Default Setting
Port security	Disabled on a port
Maximum number of secure MAC addresses	1
Violation mode	Shutdown (The port shuts down when the maximum number of secure MAC addresses is exceeded, and an SNMP trap notification is sent.)

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--3-23

The figure shows the default port security values on a Cisco Catalyst switch. To take full advantage of the port security feature, you must change these values.

Configuring Cisco Catalyst Switch Port Security

This topic describes how to configure port security on a Cisco Catalyst switch.

Configuring Port Security on a Cisco Catalyst Switch

1. Enter global configuration mode.
2. Enter interface configuration mode for the port that you want to secure.
3. Enable basic port security on the interface.
4. Set the maximum number of MAC addresses allowed on this interface.
5. Set the interface security violation mode. The default is shutdown. For mode, select one of these keywords:
 - shutdown
 - restrict
 - protect
6. Return to privileged EXEC mode.
7. Verify the entry.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—3-24

The figure lists the tasks required to configure port security on a Cisco Catalyst switch. The “Enabling Port Security with Cisco IOS Software Commands” table provides a description of the steps and commands required, including additional optional steps.

Enabling Port Security with Cisco IOS Software Commands

Step	Command	Description
1	<code>configure terminal</code>	This command opens the global configuration mode.
2	<code>Switch(config)# interface interface_id</code>	This command enables interface configuration mode. In this mode, the physical interface is configured (for example, gigabitethernet 3/1).
3	<code>Switch(config-if)# switchport mode access</code>	This command sets the interface mode as access. Port security is configured on the access port only.
4	<code>Switch(config-if)# switchport port-security</code>	This command enables port security on the interface.
5 (Optional)	<code>Switch(config-if)# switchport port-security maximum value</code>	This command sets the maximum number of secure MAC addresses for the interface. The range is 1 to 132 for a Cisco Catalyst 2950 Series Switch; 1 to 3072 for a Cisco Catalyst 4500 Series Switch. The default is 1.
6 (Optional)	<code>Switch(config-if)# switchport port-security violation {protect restrict shutdown}</code>	This command sets the violation mode. The protect option is platform-dependent or version-dependent.

Step	Command	Description
7	Switch(config-if)# switchport port-security limit rate invalid-source-mac	This command sets the rate limit for bad packets.
8 (Optional)	Switch(config-if)# switchport port-security mac-address mac_address	This command enters a secure MAC address for the interface. Use this command to enter the secure MAC addresses. If you configure fewer secure MAC addresses than the maximum, the remaining MAC addresses are dynamically learned.
9 (Optional)	Switch(config-if)# switchport port-security mac-address sticky	This command enables sticky learning on the interface.
10	Switch(config-if)# end	This command returns the console to privileged EXEC mode.
11	Switch# show port-security Switch# show port-security address interface interface_id Switch# show port-security address	This command verifies your entries.

Port Security Configuration Script

Use these configuration parameters:

- Enable port security on Fast Ethernet port 1
- Set the maximum number of secure addresses to 50
- Set violation mode to default
- No static secure MAC addresses needed
- Enable sticky learning

```
Switch# configure terminal
Switch(config)# interface fastethernet0/1
Switch(config-if)# switchport mode access
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security maximum 50
Switch(config-if)# switchport port-security mac-address sticky
Switch(config-if)# switchport port-security aging time 20
Switch(config-if)# end
```

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—3-25

MAC addresses are gathered dynamically with some switches supporting static entries and sticky entries. Static entries are manually entered for each port (for example, **switchport port-security mac-address mac-address**) and saved in the running configuration. Sticky entries are similar to static entries except that they are dynamically learned. Existing dynamic entries are converted to sticky entries when the **switchport port-security mac-address sticky** command is issued for a port. These former dynamic entries are entered into the running configuration using the command **switchport port-security mac-address sticky mac-address**. If the running configuration is then saved to the startup configuration, these MAC addresses do not need to be relearned on restart. Also, the maximum number of MAC addresses (for example, the command **switchport port-security maximum value**) for the port can be set.

This figure shows how to enable port security on Fast Ethernet port 0/1 and to set the maximum number of secure addresses to 50. The violation mode is the default, no static secure MAC addresses are configured, and sticky learning is enabled.

Verify the Configuration

```
Switch# show port-security interface fastethernet0/1
Port Security: Enabled
Port status: SecureUp
Violation mode: Shutdown
Maximum MAC Addresses :50
Total MAC Addresses: 11
Configured MAC Addresses: 0
Sticky MAC Addresses :11
Aging time: 20 mins
Aging type: Inactivity
SecureStatic address aging: Enabled
Security Violation count: 0
```

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--3-26

The figure shows the output of the verification step.

Layer 2 Best Practices

This topic describes how specific best practices mitigate attacks on specific areas of Layer 2 hardware and software components.

Layer 2 Best Practices

- **Restrict management access to the switch so that parties on nontrusted networks cannot exploit management interfaces and protocols such as SNMP.**
- **Avoid using clear text management protocols on a hostile network.**
- **Turn off unused and unneeded network services.**
- **Use port security mechanisms to limit the number of allowed MAC addresses to provide protection against a MAC flooding attack.**
- **Use a dedicated native VLAN ID for all trunk ports.**
- **Shut down unused ports in the VLAN.**
- **Prevent denial-of-service attacks and other exploits by locking down the Spanning Tree Protocol and other dynamic protocols.**
- **Avoid using VLAN 1, where possible, for trunk and user ports.**
- **Use DHCP snooping and DAI to mitigate man-in-the-middle attacks.**

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—3-27

The figure summarizes Layer 2 security best practices. These suggestions mitigate attacks on specific areas of Layer 2 hardware and software components.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- **Disabling auto trunking mitigates VLAN hopping attacks.**
- **The guard root command and the bpduguard command mitigate STP attacks.**
- **DAI can protect against man-in-the-middle attacks.**
- **To prevent DHCP attacks, use the DHCP snooping and the port security feature on the Cisco Catalyst switches.**
- **Mitigate CAM table overflow attacks with Cisco IOS software commands.**
- **Configuring port security can prevent MAC address spoofing attacks.**
- **Limiting the number of valid MAC addresses allowed on a port provides many benefits.**
- **Configure port security with Cisco IOS software commands.**
- **Following best practices mitigates Layer 2 attacks.**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0-3-28

Using Cisco Catalyst Switch Security Features

Overview

While firewalls and virtual private networks (VPNs) provide WAN security, Cisco Catalyst switches provide LAN security. However, the issue is not a simple one, because security threats come from inside and outside the immediately controllable network infrastructure. As the security needs for networks increase, so does the need for flexible access to the network for remote users and customers. Security features embedded in the Cisco Catalyst switch greatly reduce the chances of network attacks. This lesson describes some of the security features embedded in Cisco Catalyst switches.

Objectives

Upon completing this lesson, you will be able to describe how to use the security features embedded in Cisco Catalyst switches to mitigate network threats. This ability includes being able to meet these objectives:

- Describe the embedded security features of Cisco Catalyst switches
- Describe the function and benefits of the IBNS feature embedded in Cisco Catalyst switches
- Describe the function and benefit of the VACL feature embedded in Cisco Catalyst switches
- Describe the function and benefit of the PVLAN feature embedded in Cisco Catalyst switches
- Describe how MAC address notification can be used to enhance network security
- Describe the function and benefits of the rate-limiting feature embedded in Cisco Catalyst switches
- Describe the function and benefit of the SPAN feature embedded in Cisco Catalyst switches
- Describe the function and benefit of the management traffic encryption features embedded in Cisco Catalyst switches

Security Features in Cisco Catalyst Switches

This topic describes the embedded security features of Cisco Catalyst switches.

Switching Infrastructure and Security

Switching devices provide infrastructure protection through support for these:

- IBNS
- VACLs
- VLANs
- MAC address notification
- Rate limiting, also known as traffic policing
- SPANs
- Secure management protocols:
 - SSHv2
 - SNMPv3



© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0-3.3

LAN security is important. Research by the U.S. Federal Bureau of Investigation (FBI) and the Computer Security Institute (CSI) indicates that up to 60 percent of attacks are initiated on LANs as opposed to WANs. Clearly, a balanced focus on the LAN portion of any security plan is required to provide an added layer of protection. The Cisco Catalyst switch portfolio supports secure connectivity, perimeter security, intrusion protection, identity services, and security management as key elements in the Cisco Self-Defending Network architecture.

The security features introduced in this lesson are as follows:


- Cisco Identity-Based Networking Services (IBNS)
- VLAN access control lists (VACLs)
- Private VLANs (PVLANS)
- MAC address notification
- Rate limiting (also known as traffic policing)
- SPANs
- Secure management protocols: Secure Shell version 2 (SSHv2) and Simple Network Management Protocol version 3 (SNMPv3)

Identity-Based Networking Services

This topic describes the function and benefits of the IBNS feature embedded in Cisco Catalyst switches.

Identity-Based Networking Services

- **IBNS does the following:**
 - Using the 802.1x protocol with Cisco enhancements, the network grants privileges based on user login information, regardless of the user location or device.
- **The benefits of IBNS are as follows:**
 - Allows different people to use the same PC and have different capabilities
 - Ensures that users get only their designated privileges, no matter how they are logged into the network
 - Reports unauthorized access
- **Otherwise, there is no way to control who gets on the network and where they can go.**




© 2006 Cisco Systems, Inc. All rights reserved. SMD v2.0--3-4

Using 802.1x with Cisco enhancements allows you to limit access to network resources based on user login identity. User privileges remain the same, no matter how or where someone logs into the network. IBNS is recommended for organizations that have mobile users logging in using various devices from different ports.

Identity-Based Networking Services (Cont.)

IBNS functions as follows:

- Each user trying to enter the network must receive authorization based on a personal username and password.



© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0-3-5

The figure shows the topology and process for IBNS. The IBNS process is described in these four steps:


- Step 1** Each user logging in to the network must type in a username and password. Although the switch does not permit the person to log in to the network yet, it does pass the username and password to an authentication server (Cisco Secure Access Control Server [ACS]).
- Step 2** The Cisco Secure ACS looks up the username and password to determine its validity. The server also makes a note of which port and MAC address the person is using to log in.
- Step 3** If the username and password are correct, the authentication server sends a message to the switch to allow the person to proceed with the login process.
- Step 4** If the username and password are not correct, the server sends a message to the switch to block that port. After the port has been blocked, it cannot be opened until a correct username and password have been received.

The communications from the client to the switch use Extensible Authentication Protocol over LAN (EAPOL), and the communications from the switch to the authentication, authorization, and accounting (AAA) server use RADIUS.

To learn more about this feature and its configuration, refer to the Securing Networks with Cisco Routers and Switches (SNRS) course or the Cisco Catalyst switch configuration documentation.

VLAN ACLs

This topic describes the function and benefit of the VACL feature embedded in Cisco Catalyst switches.




VACLs, also known as VLAN maps, can be used to filter VLAN traffic. Unlike regular Cisco IOS standard or extended ACLs that are configured on router interfaces, VACLs are applied to any VLAN. When you configure a VACL and apply it to a VLAN, all packets entering the VLAN are checked against this VACL.

The list of commands here shows how to define and apply a VLAN access map. In this example, IP traffic matching ACL 100 is forwarded, and all other IP packets are dropped because of the default drop action. The map is applied to VLANs 12 to 16.

```
switch(config)# vlan access-map test 10
switch(config-access-map)# match ip address 100
switch(config-access-map)# action forward
switch(config-access-map)# exit
switch(config)# vlan filter test vlan-list 12-16
```

Private VLANs

This topic describes the function and benefit of the PVLAN feature embedded in Cisco Catalyst switches.




PVLANS work by limiting which ports within a VLAN can communicate with other ports in the same VLAN. Typically, PVLANS are deployed so that the hosts on a given segment can communicate only with their default gateway and not with the other hosts on the network. For example, if a web server is compromised by the Blaster worm, the server is not able to initiate infection attempts to other devices in the same VLAN even though they exist in the same network segment.

This access control, carried out by assigning hosts to either an isolated port or a community port, is an effective way to mitigate the effects of a single compromised host. Isolated ports can communicate only with promiscuous ports (typically the router). Community ports can communicate with the promiscuous port and other ports in the same community.

To learn more about this feature and its configuration, refer to the *Building Cisco Multilayer Switched Networks* (BCMSN) course or the Cisco Catalyst switch configuration documentation.

MAC Address Notification

This topic describes how MAC address notification can be used to enhance network security.



Network managers need a way to monitor who is using the network and where they are.


MAC address notification allows the network administrator to monitor the MAC addresses that are learned by the switch and the MAC addresses that are aged out and removed from the Content Addressable Memory (CAM) in the switch.

The MAC address notification feature sends SNMP traps to the network management station (NMS) whenever a new MAC address is added to or an old address is deleted from the forwarding tables. MAC notifications are generated only for dynamic and secure MAC addresses.

Use the **mac address-table notification** global configuration command to enable the MAC address notification feature on the switch.

Rate Limiting

This topic describes the function and benefits of the rate-limiting feature embedded in Cisco Catalyst switches.



Rate limiting (also referred to as traffic policing) involves creating a policing agent that specifies the upper bandwidth limit for traffic. Packets that exceed the limits are considered to be out-of-profile or nonconforming. Each policing agent decides on a packet-by-packet basis whether the packet is in profile or out-of-profile and specifies the actions taken on the packet. Actions include dropping the packet, modifying (marking down) the assigned differentiated services code point of the packet, and allowing the packet to pass through.

Rate limiting is similar to putting an upper speed limit on a car. Rate limiting ensures that no user can flood the network with too much traffic. Rate limiting also allows important applications and users to maintain a minimum network priority, which is useful when voice, video, and data are all deployed on a single network.


Rate limiting enables you to assign a bandwidth restriction to a category of traffic, such as Internet Control Message Protocol (ICMP), UDP, or to specific connection types, as a way to limit the damage from a denial of service (DoS) or a distributed denial of service (DDoS) attack while you are still working out a solution.

Traffic policing is configured using Modular QoS (quality of service) CLI (command-line interface) (MQC).

To learn more about this feature and its configuration, refer to the *Implementing Cisco Quality of Service (QoS)* course or the Cisco Catalyst switch configuration documentation.

SPAN for IPS

This topic describes the function and benefit of the Switched Port Analyzer (SPAN) feature embedded in Cisco Catalyst switches.



An IDS has the ability to detect misuse, abuse, and unauthorized access to networked resources. SPAN can be used to mirror traffic to another port where a probe or an IDS sensor is connected. When an IDS sensor detects an intruder, the sensor can send out a TCP reset that tears down the intruder connection within the network, immediately removing the intruder from the network.


You can analyze network traffic passing through ports or VLANs by using SPAN or Remote SPAN (RSPAN) to send a copy of the traffic to another port on the switch or on another switch that has been connected to a network analyzer or other monitoring or security device. SPAN copies (or mirrors) traffic received or sent (or both) on source ports or source VLANs to a destination port for analysis. SPAN does not affect the switching of network traffic on the source ports or VLANs. You must dedicate the destination port for SPAN use. Except for traffic that is required for the SPAN or RSPAN session, destination ports do not receive or forward traffic.

The example here shows how to set up SPAN session 1 for monitoring source port traffic to a destination port. First, any existing SPAN configuration for session 1 is deleted, and then bidirectional traffic is mirrored from source Gigabit Ethernet port 0/1 to destination Gigabit Ethernet port 0/2, retaining the encapsulation method.

```
Switch(config)# no monitor session 1
Switch(config)# monitor session 1 source interface gigabitethernet0/1
Switch(config)# monitor session 1 destination interface
gigabitethernet0/2 encapsulation replicate
Switch(config)# end
```

Management Encryption

This topic describes the function and benefit of the management encryption feature embedded in Cisco Catalyst switches.



Password and management traffic encryption is important if there are sophisticated users who also have malicious or mischievous intent using the network.

Catalyst switches support the use of SSHv2 to provide secure remote vty connections. Because SSHv2 sends no traffic in clear text, network administrators can conduct remote access securely.

Catalyst switches support SNMPv3, which is an interoperable standards-based protocol for network management. SNMPv3 provides secure access to devices by a combination of authenticating and encrypting packets over the network. The security features provided in SNMPv3 are as follows:

- **Message integrity:** Ensures that a packet has not been tampered with in transit
- **Authentication:** Determines that the message is from a valid source
- **Encryption:** Scrambles the contents of a packet to prevent it from being seen by an unauthorized source

Refer to the Cisco Catalyst switch configuration documentations for more information on SSH and SNMP configurations.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- The Cisco Catalyst switch portfolio supports secure connectivity, perimeter security, intrusion protection, identity services, and security management as key elements in the Cisco Self-Defending Network architecture
- The Cisco Catalyst IBNS feature provides user authentication using EAPOL and RADIUS.
- VACLs are used to filter VLAN traffic.
- PVLANs work by limiting which ports within a VLAN can communicate with other ports in the same VLAN
- MAC address notification enables you to track users on a network by storing the MAC address activity on the switch. Whenever the switch learns or removes a MAC address, an SNMP notification can be generated and sent to the NMS.
- Rate limiting (traffic policing) involves creating a traffic policing agent that specifies the upper bandwidth limit for the traffic.
- SPAN is used to mirror traffic to another port where a probe or an IDS sensor is connected.
- Management encryption features, such as SSHv2 and SNMPv3, prevent hackers from stealing usernames and passwords and device configuration information.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—3-12

Securing WLANs

Overview

Wireless LANs (WLANs) have created a new level of productivity and freedom both within and outside organizations. Both back-office and front-office applications rely on wireless connectivity. While productivity has increased, new challenges to security have arisen. Because wireless signals propagate beyond the physical boundaries of the organization, the traditional view that the inside of the organization is secure is now invalid. Signals from unsecured WLANs that extend outside the corporate network can be found and used by unauthorized personnel and malicious hackers. Although the wireless medium has specific, unique characteristics, essential WLAN security measures are not very different from those required to build strong wired security, and information technology (IT) administrators can maintain corporate privacy with the proper WLAN security measures employed.

The Cisco Self-Defending Network strategy protects against the new threats to security posed by wireless technologies by improving the ability of the network to automatically identify, prevent, and adapt to security threats. This lesson describes how to secure WLAN segments in your network.


Objectives

Upon completing this lesson, you will be able to describe how to secure WLAN segments in your network. This ability includes being able to meet these objectives:

- Describe basic 802.11 architecture and components
- Describe the threats to WLAN segments
- Describe the evolution of the security features of the 802.11 protocol
- Describe the function of the SSID in WLAN security
- Describe the WEP protocol, including its purpose, evolution, and the weaknesses that have limited its effectiveness as a standalone WLAN security protocol
- Describe how to mitigate risks to WLANs by applying countermeasures to address specific threats and vulnerabilities
- Describe how intrusion detection is done for WLANs

Introducing WLANs

This topic describes basic 802.11 architecture and components.



Wired LANs require that users locate in one place and stay there. WLANs are an extension to the wired LAN network. WLANs can be an overlay to or a substitute for traditional wired LAN networks.

WLANs provide mobile users with these features:

- Free movement around a facility
- Real-time access to the wired LAN at wired Ethernet speeds
- Access to all the resources of wired LANs


Comparing WLANs with LANs

Similarities:

- **A WLAN is an 802 LAN:**
 - Data over air instead of data over wire
 - Looks like a wired network to the user
- **The same protocols run over both LANs and WLANs:**
 - Simple Network Management Protocol
 - IPsec

Differences:

- **Use of radio frequency introduces country-specific regulations**
- **Clients are mobile**
- **Radio-frequency physical layer introduces privacy and connectivity issues**



The basic similarities between the WLAN and the LAN are as follows:

- A WLAN is an 802 LAN. WLAN technology and the WLAN industry date back to the mid-1980s when the U.S. Federal Communications Commission (FCC) first made the radio frequency (RF) spectrum available to industry. Early WLAN technologies were expensive, provided low data rates, were prone to radio interference, and were designed mostly for proprietary RF technologies. To overcome these WLAN technology limitations, the IEEE initiated the 802.11 project in 1990. The 802.11 international interoperability standard was approved in 1997. The 802.11a and the 802.11b wireless networking communication standards were ratified in 1999, and, finally, the 802.11g standard was approved in June 2003.
- The same protocols run over both LANs and WLANs.

The relevant differences between the WLAN and the LAN are as follows:

- Regulations for the use of RF spectrum sometimes vary from country to country.
- Client mobility introduces authentication issues as clients move between access points in a wireless network.
- The RF physical layer introduces privacy and connectivity issues such as coverage holes, multipath issues, interference, and noise.

WLAN Characteristics


Characteristic	Description
Physical layer	DSSS OFDM Infrared
Frequency band	2.4 GHz (ISM band) and 5 GHz
Data rates	802.11b—1 Mbps, 2 Mbps, 5.5 Mbps, 11 Mbps (DSSS) 802.11a—54 Mbps (OFDM) 802.11g—54 Mbps (OFDM)
Operating range	Up to 150 feet indoors and 1500 feet outdoors
Positive aspects	High data throughput without wires
Negative aspects	Throughput decreasing with distance and load; poor security in native mode

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—3-5

The IEEE 802.11a standard operates in the licensed 5-GHz band using orthogonal frequency division multiplexing (OFDM) technology. The 802.11b standard operates in the unlicensed 2.4-GHz to 2.5-GHz Industrial, Scientific, and Medical (ISM) frequency band using a direct sequence spread spectrum technology (DSSS). The ISM band is popular for wireless communications because it is available worldwide. The 802.11b WLAN technology permits transmission speeds of up to 11 Mbps; 802.11g is an extension to 802.11b and will broaden 802.11b data rates to 54 Mbps within the 2.4-GHz band using OFDM technology.

Typical WLAN Components and Topologies



There are four components in a WLAN:

- Clients
- Access points
- Controller
- Network infrastructure

The basic service area (BSA) is the area of RF coverage provided by an access point. The BSA is also referred to as a microcell, or alternatively as just a cell. In the figure, the BSA is called a wireless "cell."

An access point can be added to extend the BSA or to simply add wireless devices and extend the range of an existing wired system. As the name "access point" indicates, this unit is the point at which wireless clients can access the network.

The access point connects to a controller, which attaches to the Ethernet backbone and communicates with all the wireless devices in the cell area. The access point is the master for the cell and controls traffic flow to and from the network. The remote devices do not communicate directly with each other; they communicate with the access point.

If a single cell does not provide enough coverage, any number of cells can be added to extend the range. This is known as an extended service area (ESA). It is recommended that the ESA cells have 10 to 15 percent overlap to allow remote users to roam without losing RF connections. Bordering cells should be set to different nonoverlapping channels for best performance.

With the 15 percent overlap between cells, a shared service set identifier (SSID), and nonoverlap channels you create roaming capability. For example, as a client moves from the cell defined by channel 1 into the cell defined by channel 6, the client will switch from the access point on channel 1 to the access point on channel 6.

Cisco Unified Wireless Network

 Cisco Self-Defending Network	Unified Advanced Services Unified built-in support of leading-edge applications—not an afterthought; Cisco Wireless Location Appliance, Cisco WCS, SDN, NAC, Wi-Fi phones, and RF firewalls
	World-Class Network Management World-class NMS that visualizes and helps secure your air space; WCS
	Network Unification Seamless network infrastructure across a range of platforms; Cisco 2000 and 4400 Wireless LAN Controllers; future Cisco Catalyst 6500 Series WiSM, ISR, and 3750 integration
	Mobility Platform APs dynamically configured and managed through LWAPP. Cisco Aironet Access Points: 1500, 1300, 1240AG, 1230AG, 1130AG, and 1000. Bridges: 1400 and 1300.
	Client Devices Secure clients that work out of the box. These include Cisco Aironet clients and third-party devices that comply with the Cisco Compatible Extensions program.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0-3.7


In the figure, the components of WLANs are shown grouped into the roles that they each play in a WLAN. There are client devices, access points (shown in the figure as mobility devices), and controllers (shown in the figure as network unification devices). Finally, there are management solutions that help you administer your WLAN and advanced services that can be deployed in the WLAN.

Threats to WLANs

This topic describes the threats to WLAN segments.

Threats to WLANs

- **WLAN equipment is widely available and inexpensive**
- **802.11 standard designed for ease of use and deployment**
- **Availability of sniffers**
- **Statistics on WLAN securities**
- **Media hype about hotspots, WLAN hacking, war driving**
- **Nonoptimal implementation of encryption in standard WEP**
- **Authentication vulnerability**
- **Loss of mobile devices such as PDAs**



The illustration shows a person walking with a laptop, representing 'war driving'. They are finding WLAN hotspots (indicated by signal waves) near a 'UNIVERSITY' building and residential houses. Two speech bubbles from the person say 'I found another one.' The diagram is set against a background of a road and trees.

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—3-8

With the cost of IEEE 802.11b systems decreasing, it is inevitable that hackers will have many more unsecured WLANs to choose from.

802.11b sniffers enable network engineers (and hackers) to passively capture data packets to examine to correct system problems.

“War driving” is a phrase that has been used to describe someone who is using a cellular scanning device looking for cell phone numbers to exploit. Nowadays, war driving also refers to someone driving around with a laptop and an 802.11b client card looking for an 802.11b system to exploit.


Incidents have been reported of people using numerous open-source applications to collect and exploit vulnerabilities in the IEEE 802.11 standard security mechanism Wired Equivalent Privacy (WEP).

With basic WEP encryption enabled (or with no encryption enabled), it is possible to collect data and obtain sensitive network information such as user login information, account numbers, and personnel records.

Personal digital assistants (PDAs) and other mobile devices also pose specific threats to networks. Organizations must learn to treat these devices as full-fledged clients and create and apply the same sort of security policies that they have devised for desktop and laptop computers. Organizations tend to adopt PDAs informally, and that lack of formality can have disastrous consequences. To minimize the security threats from loss or theft of the device, organizations must mandate the use of power-on passwords and put in place remote destruct policies that erase data on the device remotely.

Evolution of 802.11 Security Features

This topic describes the evolution of the security features of the 802.11 protocol.



WLANs, which were at one time openly accessible, now have an array of security options available that can make them very secure. The language used to describe this array of security options can be confusing. For example, what is the difference between WEP and Wi-Fi Protected Access (WPA)? Or perhaps your company has a WLAN and you connect through Cisco Lightweight Extensible Authentication Protocol (LEAP); why do you still have to have an SSID if the network is using LEAP? The quickest way to understand all of your options is to review the emergence of WLAN security terminology during the evolution of WLAN security methods in the industry.

The figure shows four phases of WLAN security evolution. The first was the open access phase in which there was no real security except for the SSID. With this rudimentary mechanism, an access point would look for the SSID on a client machine. If it matched, the client was given access. However, it was not always true that a failure to match was denied access, because access points could—and still can—be configured to broadcast their SSID to clients that request it. Even if an access point was not configured to broadcast its SSID, it was still vulnerable, because the communication between a client and access point was in clear text and an eavesdropper could easily acquire a valid SSID.

Open access was followed by an initial phase in which the urgency for WLAN security resulted in a patchwork solution with WEP. In the now-famous paper *Security of the WEP Algorithm* (written in January 2001 by Nikita Borisov, Ian Goldberg, and David Wagner of the University of California, Berkeley and available at <http://www.isaac.cs.berkeley.edu/isaac/wep-faq.html>), the authors showed that the implementation of the RC4 algorithm in the WEP protocol was faulty and consequently easy to crack. Several attacks against WLANs that relied on WEP for security were described in the Berkeley paper, and a new WLAN security solution was required.

WEP was followed by WPA. Significantly, this approach to WLAN security was the first attempt to integrate 802.1x into the WLAN security suite and, along with WEP and SSID, provided an improved multitiered security approach. When WPA was released, the members of IEEE who were working on the standard recognized that WLAN security needed 802.1x, but that there was more work required on the standard than just adding it in. Consequently, WPA was released as an interim solution until the full 802.11i standard would be ready. Here are some of the features of WPA:

- 802.1x authentication framework (802.11 TGi baseline)
- Mutual authentication using dynamic, per-user, per-session WEP key
- Automatic, frequent reauthentication

In 2004, the 802.11i standard for WLAN security was released. The new standard incorporated all the WLAN security elements that had been used until then, further standardizing the use of 802.1x authentication and changing the encryption mechanism to the Advanced Encryption Standard (AES) from the Data Encryption Standard (DES). The Wi-Fi Alliance calls its implementation of the full 802.11i standard Wi-Fi Protected Access 2 (WPA2).

Service Set Identifier

This topic describes the function of the SSID in WLAN security.

Open Access Phase—SSID

SSID

- String of 32 ASCII characters
- If access point broadcasts SSID under 802.11, any client with a null string will associate to any access point regardless of SSID setting on access point
- Should not be considered a security feature

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—3-10

The SSID is a configurable parameter that is checked as part of the association process and should match on both the wireless client and the access point. The SSID serves to logically segment the users and access points that form part of a wireless subsystem.

Under 802.11 specifications, an access point may advertise, or broadcast, its SSID. During the association process, any 802.11 wireless clients with a null string (no value entered into the SSID field) may request that the access point broadcast its SSID. If the access point is configured to do so, it sends the SSID to the client. The client then uses this SSID to associate with the access point. For these reasons, the SSID should not be considered a security feature.


Wired Equivalent Privacy

This topic describes the WEP protocol, including its purpose, its evolution, and the weaknesses that have limited its effectiveness as a standalone WLAN security protocol.

Initial Phase—WEP

WEP

- The basic IEEE 802.11 security standard
- Uses 40-bit keys
- 128-bit keys optional
- Optional part of the association process
- Uses the RC4 stream cipher from RSA Security for encryption




© 2006 Cisco Systems, Inc. All rights reserved.SND v2.0—3-11

The 802.11 standard defines a type of security: WEP using 40-bit keys. WEP is based on a stream cipher called RC4. The RC4 method allows encryption up to 128 bits; however, IEEE 802.11 has chosen to use 40-bit keys.

WEP requires a wireless client and an access point to compare static 40-bit keys during the authentication process. If the client WEP key does not match the key of the access point, the client is not allowed to associate and cannot connect to the network.

The 802.11 standard provides two schemes for defining the WEP keys to be used on a WLAN. With the first scheme, as many as four default keys are shared by all stations (clients and access points) in a wireless subsystem. When a client obtains the default keys, that client can communicate securely with all other stations in the subsystem. (Cisco Systems uses this method.) Be aware that when default keys are widely distributed, they are more likely to be compromised. In the second scheme, each client establishes a “key mapping” relationship with another station. This method is more secure because fewer stations have the keys. However, distributing such unicast keys becomes more difficult as the number of stations increases.

802.11 Open Authentication



Two types of WEP encryption are defined: open and shared key. This section looks at both of these types and the process that the client undergoes during the authentication process.


Open Authentication

The open authentication method allows authorization and associations with or without a WEP key. If the client does not use a WEP key, the client undergoes the normal association process with the access point. The user is then granted access to the network.

If a WEP key is used, both the client and the access point must have matching WEP keys. If the client uses a WEP key that is different from the WEP key of the access point, data traffic cannot be passed because the data is encrypted. Keep in mind that the header is not encrypted; only the payload (or data) is encrypted.

Using open authentication, the client goes through the normal association process, regardless of whether the client is using a WEP key. After the client is associated and data transmission begins, a client using a WEP key encrypts the data. If the WEP key on the access point does not match, the access point is unable to decrypt the data, so it is impossible to send the data via the WLAN.

802.11 Shared Key Authentication



Shared Key Authentication


The figure shows the wireless client using shared key authentication to attempt to associate with an access point. Steps 1 through 3 are the same as those for open authentication. These additional four steps are required for shared key authentication:

- Step 1** The client sends an authentication request to access point A.
- Step 2** Access point A sends an authentication response. The authentication response from the access point to the client is sent containing challenge text. This packet is unencrypted.
- Step 3** The client then uses the text from the authentication response to form another authentication packet, which will be encrypted using one of the client WEP keys, and sends this as a response to the access point.
- Step 4** Access point A will then compare the encrypted challenge text against the access point copy of the encrypted challenge text. If the encrypted text is the same, the access point allows the client on the WLAN.

Shared key authentication is considered less secure than open authentication because of the challenge text packet. Because this packet is sent unencrypted and then returned as an encrypted packet, it may be possible to capture both packets and determine the stream cipher.

Basic 802.11 Security Issues

- **Protects against outside threats**
 - Checks for devices that do not possess key
 - Hardware theft may be an issue
- **One-way authentication**
 - Checks client key only
- **No way to dynamically generate keys**
- **No integration with existing network authentication methods**
- **Rogue access points**
 - May render either client or network vulnerable
 - Important to two-way (mutually) authenticate user and network



© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—3-14

Basic 802.11 WEP security is designed to guard against the threat to network security from unauthorized 802.11 devices outside the LAN. Any device with a valid WEP key is considered a legitimate and authorized user.


If the WEP key is obtained, either through hardware loss or theft or through a wireless security exploit, the network and wireless users are rendered vulnerable and keys must be changed. Note that persistent WEP keys may be assigned to a client adapter (keys stored in nonvolatile memory on the card itself) via most WLAN client utilities.

A commonly deployed protection against this vulnerability is MAC address authorization or filtering. However, because MAC addresses may be relatively easily spoofed, the WEP keys must be changed to ensure security. There is no way to remotely administer WEP keys, so this task could be very burdensome depending on the number of wireless devices to be managed.

A rogue access point is an access point that has been placed on a WLAN and that might be used to interfere with normal network operations (for example, denial of service [DoS] attacks). If this rogue access point is programmed with the correct WEP key, client data may be captured. This access point may also be configured to provide unauthorized users with information about the network, such as MAC addresses of clients (both wireless and wired), the ability to capture and spoof data packets, and, at worst, access to servers and files.

Basic 802.11 Security Issues (Cont.)

- **Device-based authentication**
 - User or user credential-based authentication more desirable
 - No simple integration with existing database to authenticate users
- **No method for WLAN account auditing**



Authentication can also be device-based. With this method, identification is based on the MAC address, not the username. Keys are typically stored in the flash memory of the card. As you have already seen, a stolen card could circumvent this authentication method. A more effective method is for authentication to be dependent on usernames and passwords, which are client-independent and which users may already have.

There is also no way to integrate existing network authentication methods with basic 802.11 WEP security, such as Lightweight Directory Access Protocol (LDAP) or RADIUS.

Basic 802.11 WEP security provides only one-way authentication. The client is authenticated with the access point (the WEP key is checked), but not vice versa. The client has no way of knowing whether the access point is a legitimate part of the WLAN or a rogue device that uses the same WEP key.

Even if authentication were based on username and password, you would still want to be able to audit and account for usage to warn against unusual activities, such as the following:

- Users who do not log in for long periods of time
- Users who transfer too much data and stay on too long
- Multiple simultaneous logins
- Logins from the “wrong” account

What is needed is the ability to administer and monitor wireless clients just as you would administer and monitor wired clients.

Exploits of 802.11 Security Vulnerabilities

Several attacks exploit vulnerabilities in 802.11 security:

- **Weak initialization vector attack**
- **Active “bit flipping” attack to inject traffic or to decrypt traffic**
- **Authentication dictionary attacks**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—3-16

These attacks are described in *Security of the WEP Algorithm*:

- Passive attacks to decrypt traffic based on statistical analysis
- Active attacks to inject new traffic from unauthorized mobile stations, based on known plain text
- Active attacks to decrypt traffic, based on tricking the access point
- Dictionary-building attacks, which, after analysis of traffic for about a day, allow real-time automated decryption of all traffic

Passive or Weak Initialization Vector Attack

An initialization vector (IV) is a 24-bit field that changes with each packet. Its purpose is to ensure that the same plain text data frame never generates the same WEP-encrypted data frame.

The IV is transmitted as plain text, and a user “sniffing” the WLAN can see the IV. Using the same IV over and over with the same WEP key, a hacker could capture the frames and derive information about the data in the frame and data about the network.

Static WEP keys have proven to be highly vulnerable to this type of attack. Therefore, WLANs should not use static WEP and should instead use the more advanced security features implementing 802.1x.

Active “Bit Flipping” or Replay Attack

The wireless attacker flips (changes or forges) arbitrary bits in an encrypted message and correctly adjusts the cyclic redundancy check (CRC) (a value or checksum of the total number of bits sent) to appear as a valid message. Here are the eight steps in this type of attack:

- Step 1** An attacker intercepts WEP encrypted packets.
- Step 2** The attacker flips bits in a packet and recalculates the checksum with a valid cyclic redundancy check (CRC) value.
- Step 3** The attacker transmits to the access point a bit-flipped frame with a known IV.
- Step 4** The access point receives the modified frame and accepts the frame based on a valid CRC value.
- Step 5** The access point forwards the frame to the Layer 3 router. The data in the frame is rejected because of a data error, and the router sends a retransmit response.
- Step 6** The access point encrypts the response and sends it to the attacker.
- Step 7** The attacker uses this response to derive the key or stream cipher data as a stream of bits rather than divided into blocks.
- Step 8** The attacker gains access.

Authentication Dictionary Attacks

Most password-based authentication algorithms are susceptible to online (active) and offline (passive) dictionary attacks. During a dictionary attack, an attacker tries to guess a password and gain network access by using every “word” in a dictionary of common passwords or possible combinations of passwords. A dictionary attack relies on the fact that a password is often a common word, name, or concatenation of words or names with a minor modification such as a trailing digit or two. Longer passwords with a variety of characters (such as “4yosc10cP!”) offer the greatest protection against dictionary attacks.

During an online dictionary attack, an attacker tries to actively gain network access by trying possible combinations of passwords for a specific user. Online dictionary attacks can be prevented using lockout mechanisms available on RADIUS servers to lock the user out after a certain number of invalid login attempts. Online attacks also provide some evidence that a breach or compromise is being attempted, allowing you to take corrective measures.

An offline dictionary attack is carried out in two phases to uncover a password. In the first phase, the attacker captures the challenge and response messages between the user and the network. In the second phase, the attacker looks for a password match by computing a list of possible challenge response messages (using a precomputed dictionary, usually with the aid of a password-cracking program) and comparing these messages against the captured challenge and response message. The attacker uses known authentication protocol vulnerabilities to reduce the size of the user password dictionary. Enforcing a strong password policy and periodically requiring that users change their passwords can significantly reduce the potential for a successful offline attack using these tools. Unlike online attacks, offline attacks are not easily detected.


Enhanced Methods for WLAN Threat Mitigation

This topic describes how to mitigate risks to WLANs by applying countermeasures to address specific threats and vulnerabilities.

Enhanced 802.11 Security

Authentication

- Prove that you belong to the network



Encryption

- Provide encryption keys after authentication



802.11 security is enhanced by adding methods for user authentication and data stream encryption.

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—3-17

Enhanced 802.11 security incorporates authentication and encryption to improve on standard or basic 802.11 security.

802.11 authentication uses the IEEE 802.1x standard that permits policy assignment to users as a result of the authentication transaction. Basing the authentication transaction on user credentials instead of machine credentials reduces the risk of a security compromise from lost or stolen equipment. 802.1x authentication also permits flexible credentials such as password, one-time token, public key infrastructure (PKI) certificate, or device ID to be used for client authentication. Using 802.1x for wireless client authentication also has the advantage that dynamic encryption keys may be distributed to each user each time that they authenticate to the network.

Encryption for 802.11 is enhanced with multiple mechanisms to aid in protecting the system from malicious exploits against the WEP key and protecting investment in the system by facilitating encryption improvements in existing hardware.

Interim Phase—WPA

History of WPA:

- WPA introduced in late 2003
- Prestandard implementation of IEEE 802.11i WLAN security
- Addresses currently known security problems with WEP
- Allows software upgrade on already deployed 802.11 equipment to improve security

Components of WPA:

- Authenticated key management using 802.1x: EAP authentication, and preshared key authentication
- Unicast and broadcast key management
- Standardized TKIP per-packet keying and Message Integrity Check protocol
- Initialization vector space expansion: 48-bit initialization vectors
- Migration mode—coexistence of WPA and non-WPA devices (optional implementation that is not required for WPA certification)

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—3-18

WPA is a standard that was developed in 2003 by the Wi-Fi Alliance, formerly known as the Wireless Ethernet Compatibility Alliance.

WPA provides a standard for authentication and encryption of WLANs that is intended to solve security problems that were known to exist through 2003. These problems include the well-publicized AirSnort and man-in-the-middle WLAN attacks.

The WPA standard is a step toward the 802.11i standard, WPA2. WPA uses many of the same components as WPA2, with the exception of AES.

WPA has these elements:

- In the mechanism for authenticated key management, the user is first authenticated, and then a “master key” is derived at the server and client. This master key is used to generate the actual keys used in encrypting the user session. The master key is not directly used.
- These key validation mechanisms are in place for both unicast and broadcast keys.
- WPA uses Temporal Key Integrity Protocol (TKIP), which includes per-packet keying and a message integrity check (MIC) for WPA.
- IV is expanded to 48 bits to prevent “collisions” or reuse of the same vector, which can be used in exploits to attempt to derive an encryption key. IV collisions are one of the primary mechanisms used by tools such as AirSnort.
- An optional mode allows migration (coexistence) of users from standard WEP encryption to WPA encryption.

Present Phase—WPA2

- **WPA2 is the full Wi-Fi Alliance implementation of the 802.11i standard.**
- **802.11i uses the AES block cipher that replaces the DES.**
- **WPA2 standardizes use of 802.1X for authentication.**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—3-19

AES is the next-generation encryption function approved by the U.S. National Institute of Standards and Technology (NIST). NIST solicited the cryptography community for new encryption algorithms. The algorithms had to be fully disclosed and available royalty-free. NIST judged candidates on cryptographic strength and practical implementation. The finalist, and the method that was finally adopted for AES, is known as the Rijndael algorithm.

The Cisco version of AES is called AES-Counter Mode with Cipher Block Chaining Message Authentication Code Protocol (AES-CCMP) and it uses IVs to augment the key stream. The IV increases by one after encrypting each block. This technique provides a unique key stream for each block. AES-CCMP also uses a message authentication check to verify packet integrity using frame length, destination and source addresses, and data-in input values.

One of the biggest benefits of 802.11x 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.

802.1x for WLANs

802.1x for 802.11

- 802.11i specifies use of 802.1x for client authentication
- Based on EAP framework
- Improved authentication credentials
 - Superior to device-based (such as MAC address) authentication
- Session-based encryption keys
- Centralized user administration

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--3-20

The IEEE has developed a supplement to the 802.1D standard that defines the changes that must be made to the operation of a MAC layer bridge to provide port-based network access control capability. This supplement is the 802.1x standard.

WLAN 802.1x uses these services, applications, or features:

- RADIUS and Extensible Authentication Protocol (EAP) for encapsulation of EAP packets within RADIUS
- Identification based on network access identifier
- Support for roaming access in public spaces
- RADIUS support for centralized authentication, authorization, and accounting (AAA)
- Dynamic rather than static WEP keys that no require user intervention-based management
- Compatibility with existing roaming technologies, enabling use in hotels and public places

These are the common 802.1x authentication protocols:

- **EAP-LEAP:** Cisco EAP-Lightweight Extensible Authentication Protocol
- **EAP-FAST:** EAP-Flexible Authentication via Secure Tunneling (Cisco protocol submitted as IEEE draft)
- **EAP-TLS:** EAP-Transport Layer Security
- **EAP-PEAP:** EAP-Protected Extensible Authentication Protocol

802.1x EAP Deployment Comparison

	LEAP	EAP-FAST	PEAP	EAP-TLS
Multiple operating system support	Yes	Limited	Limited	Limited
Single login using Microsoft Windows login	Yes	Yes	No*	Yes
Dynamic WEP key and mutual authentication	Yes	Yes	Yes	Yes
Static password support	Yes	Yes	Yes	No
One-time password support	No	No	Yes	No
Capability to tie login with non-Microsoft user databases (LDAP, Novell Directory Services, and so on)	No	Yes (LDAP)	Yes	Yes
Layer 3 roaming support	Yes	Yes	Yes	Yes
Works with WPA	Yes	Yes	Yes	Yes**

* Microsoft PEAP (EAP-Microsoft Challenge Handshake Authentication Protocol Version 2) supports single sign-on.

** WPA testing is done with EAP-TLS, but all EAP types can be used with WPA.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—3-21

There are a number of 802.1x EAP mechanisms available, and more are being developed.

Here are the three primary considerations for deploying an authentication mechanism:

- The back-end database that is used to authenticate users
- The desired user interface for clients to access the WLAN (for example, token card, password prompt, or single login)
- The EAP supplicants that are supported on the client devices

802.1x Advantages for WLANs

- **Mutual authentication**
 - The server is authenticated by the client, and the client is authenticated by the server.
- **Encryption keys derived dynamically**
- **Ability to refresh encryption keys**
 - RADIUS session timeout is used to give a fixed “validity” window for a user WLAN session key.
- **Centralized user and key management**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--3-22

A major advantage of EAP and the 802.1x standards is that they are designed to leverage existing standards. With support for EAP, WLANs can now offer the following:

- Support for RFC 2284, with password authentication. Users are authenticated based on usernames and passwords that are typically already stored in an active directory on the network. This directory is then connected to a certificate server, such as a RADIUS server or the Cisco Secure Access Control Server (ACS).
- One-time passwords, which take a plain text password and encrypt it. Thus, plain text passwords never have to be typed on a nonsecure connection. (Telnet and FTP use no encryption and, therefore, are not considered secure protocols.)


EAP support is designed to allow additional authentication methods to be deployed with no changes to the access point or client network interface card (NIC). Nothing beyond the latest versions of firmware and drivers is required for the Cisco Aironet equipment to take advantage of the benefits offered by EAP.

Note Some of the newer authentication protocols do require client software to perform authentication. This software is commonly referred to as a supplicant.

Dynamic keying, and the ability to centrally manage the user database, is a major advantage of 802.1x and EAP.

WLAN IDS

This topic describes how intrusion detection is performed for WLANs.



Someone who enjoys a wireless network at home might try to create the same freedom at work by plugging their own, inexpensive access point into a network jack at work without asking permission. These are known as rogue access points, and the majority of these are installed by employees, not malicious intruders. Even company-sanctioned access points, when configured improperly, can be security risks. Cisco WLAN platforms, supported by Cisco IOS software, provide intrusion detection as described in the figure.

WLAN intrusion detection systems (IDSs) have these features:

- **Excess Management Frame Detection:** This feature provides scanner access points the ability to detect that WLAN management and control frames exceed a configurable threshold.
- **Authentication Attack Detection:** This feature requires Cisco access points to detect and report on excessive attempted or failed authentication attempts (authentication failure detection and excess EAP over LAN [EAPOL] authentication).

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- WLANs are IEEE 802 LANs that allow mobile users to access a network.
- WLANs provide “ports” to users outside the confines of the organization buildings and, with growing availability of wireless equipment, can expose networks to attack.
- The IEEE 802.11 standard included specifications for security with WEP.
- The SSID is a configurable parameter that is checked as part of the association process and should match on both the wireless client and the access point. Under 802.11 specifications, an access point may advertise its SSID, so the SSID should not be considered a security feature.
- The implementation of the RC4 algorithm in WEP was inadequate and exposed WLANs to a variety of attacks that are based on exploiting the WEP initialization vector.
- An interim standard called WPA, which incorporated 802.1x authentication and improved encryption, was released prior to IEEE ratifying the 802.11i standard in June 2004. 802.1x has been implemented as WPA2.
- WLAN IDS look for excess management frames and excess authentication attempts to detect possible intrusions by attackers.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0-3-24

Module Summary

This topic summarizes the key points that were discussed in this module.

Module Summary

- To secure network access at Layer 2, follow these steps:
 - Protect administrative access to the switch.
 - Protect the switch management port.
 - Turn off unused network services.
 - Lock down the ports.
 - Use Cisco Catalyst switch security features.
- VLAN hopping, STP manipulation, ARP spoofing, CAM table overflow, and MAC spoofing are the Layer 2 attacks used to compromise LANs. Port security used along with security best practices will mitigate against these attacks.
- Cisco Catalyst switch security features greatly reduce the chances of network attack.
- WLANs are secured by applying 802.11i mechanisms, which include 802.1x authentication.

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0--3.1

The security of a network environment depends in part on the security of Layer 2 devices in the LAN, including wired and wireless devices. You must ensure that these LAN devices are accounted for in your organization security policy, and then you must configure the built-in features in these devices to secure your LAN.

References

For additional information, refer to these resources:

- Cisco Systems, Inc. *Configuring Layer 2 Ethernet Interfaces*.
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/12_1_11/config/layer2.htm.
- Cisco Systems, Inc. *Cisco Aironet Response to Press—Flaws in 802.11 Security*.
http://www.cisco.com/en/US/products/hw/wireless/ps430/prod_bulletin09186a0080088832.html.
- TechRepublic. Address These Seven Areas in Your Wireless LAN Security Policy.
http://techrepublic.com.com/5100-1009_11-5930876.html?tag=nl.e099.tp10-121205.
- TechRepublic. Strengthen Security by Implementing Network Address Translation.
http://techrepublic.com.com/5100-1009_11-5590551.html.
- Borza, A., D. Duesterhaus, C. Grabczynski, et al. National Security Agency. *Cisco IOS Switch Security Configuration Guide*. http://www.nsa.gov/snac/os/switch-guide-version1_01.pdf.
- Borisov, N., I. Goldberg, and D. Wagner. *Security of the WEP Algorithm*.
<http://www.isaac.cs.berkeley.edu/isaac/wep-faq.html>.

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) Match each of the mitigation techniques with the type of attack that it will mitigate by putting the letter of the technique in the space provided beside each type of attack.

(Source: Mitigating Layer 2 Attacks)

A) **use guard root** and **bpdu guard**

B) turn off DTP

_____ 1. VLAN hopping

_____ 2. STP manipulation

- Q2) Explain how VLAN configuration can mitigate VLAN hopping attacks. (Source: Mitigating Layer 2 Attacks)

- Q3) What is the effect of using the **guard root** and **bpdu-guard** enhancement commands? (Source: Mitigating Layer 2 Attacks)

- Q4) Match each of the mitigation techniques with the type of attack that it will mitigate by putting the letter of the technique in the space provided beside each type of attack.

(Source: Mitigating Layer 2 Attacks)

A) DAI

B) Port security

_____ 1. CAM table overflow

_____ 2. MAC address spoofing

Q5) Explain the role of the CAM table in switch security. (Source: Mitigating Layer 2 Attacks)

Q6) What ability does the **port security** command provide? (Source: Mitigating Layer 2 Attacks)

Q7) Indicate which Cisco Catalyst switch security feature must be employed to mitigate each of the security issues by putting the letter of the feature in the space provided beside each security issue. (Source: Using Cisco Catalyst Switch Security Features)

- A) IBNS
- B) PVLAN
- C) rate limiting
- D) port security feature
- E) management encryption
- F) SPAN

- _____ 1. Nosy users within the same VLAN can view neighbor traffic.
- _____ 2. Any unauthorized user with physical access can log in to the network.
- _____ 3. There is no way to control who gets on the network and where they can go.
- _____ 4. There can be a deliberate or accidental slowdown or freezing of the network because of a large of amount of abnormal traffic.
- _____ 5. This feature allows an IDS sensor to monitor traffic on a switch port.
- _____ 6. Snoopers can break into switches and bring down the network.

Q8) In what kind of a setting is password and management traffic encryption more important? (Source: Using Cisco Catalyst Switch Security Features)

Q9) What are the three types of secure MAC addresses that can be configured on a Cisco Catalyst switch port? (Source: Using Cisco Catalyst Switch Security Features)

Q10) Why should the SSID not be considered a security feature? (Source: Securing WLANs)

Q11) Why are security measures needed beyond 802.11 WEP security? (Source: Securing Wireless LANs)

Module Self-Check Answer Key

- Q1) A-2, B-1, C-3
- Q2) Mitigating VLAN hopping attacks requires several modifications to the VLAN configuration. One of the more important elements is to use a dedicated native VLAN for all trunk ports. Also, disable all unused switch ports and place them in an unused VLAN. Set all user ports to nontrunking mode by explicitly turning off DTP on those ports.
- Q3) The **guard root** and the **bpduguard** enhancement commands enforce the placement of the root bridge in the network and enforce the STP domain borders.
- Q4) A-2, B-1
- Q5) Switches maintain CAM lookup tables to track the source addresses located on the switch ports. These lookup tables are populated by an address-learning process on the switch. The CAM table in a switch contains the MAC addresses available on a given physical port of a switch. When a Layer 2 switch receives a frame, the switch looks in the CAM table for the destination MAC address. If an entry exists for the MAC address in the CAM table, the switch forwards the frame to the port designated in the CAM table for that MAC address. If the MAC address does not exist in the CAM table, the switch forwards the frame out every port on the switch, effectively acting like a hub. If a response is seen, the switch updates the CAM table.
- Q6) The **port security** command provides the ability to specify the MAC address of the system connected to a particular switch port. The command also provides the ability to specify an action to take if a port security violation occurs.
- Q7) A-2 and A-3, B-1, C-4, D-2, E-6, F-5
- Q8) A setting where there are sophisticated users on the network with a tendency toward mischievous intent. Universities are most at risk for this type of interruption.
- Q9) static secure, dynamic secure, and sticky secure MAC addresses
- Q10) An access point may broadcast its SSID so that any 802.11 wireless clients with no value entered into the SSID field may be automatically updated with an SSID by an access point.
- Q11) If the WEP key is obtained, either through hardware loss or theft or through a wireless security exploit, the network and wireless users are rendered vulnerable and keys must be changed. Authentication can also be device-based. With device-based authentication, identification is based on MAC address, not username. Keys are typically stored in the flash memory of the card. A stolen card could circumvent this authentication method. A more effective method is for authentication to be dependent on usernames and passwords that are client-independent and that users may already have.

Cisco IOS Firewall Configuration

Overview

Implementing network-wide security can be daunting depending on the size and business of the company. Organizations must balance the cost in staff and equipment to implement a network security policy against the costs and possibility of network security breaches. The Cisco IOS Firewall meets the needs of many organizations that choose not to use a firewall appliance due to financial constraints or technical complexity. The Cisco IOS Firewall provides a fully featured firewall implemented on Cisco routers using Cisco IOS software.

In this module, you will learn the basic configuration skills to implement a Cisco IOS Firewall on a Cisco router.

Module Objectives

Upon completing this module, you will be able to configure a Cisco IOS Firewall to perform basic security operations on a network. This ability includes being able to meet these objectives:

- Describe firewall technologies embedded in Cisco routers and security appliances
- Build static packet filters with Cisco ACLs
- Configure a firewall on your network using the Cisco SDM wizard
- Explain the best practices for deploying the hardware and software components of the Cisco security appliances.

Lesson 1

Introducing Firewall Technologies

Overview

A firewall protects network devices from intentional hostile intrusion that could threaten information assurance (that is, availability, confidentiality, and integrity) or lead to a denial of service (DoS) attack. A firewall may protect a hardware device or a software program running on a secure host computer. In either case, a firewall must have at least two network interfaces, one for the network it protects and one for the network to which it is exposed. A network firewall sits at the junction or gateway between two networks that are usually a private network and a public network such as the Internet. This lesson introduces the firewall technologies that Cisco uses in routers and security appliances.


Objectives

Upon completing this lesson, you will be able to describe firewall technologies embedded in Cisco routers and security appliances. This ability includes being able to meet these objectives:

- Describe the role of firewalls in securing networks
- Describe four generations of firewall technologies developed between 1983 and 1993 that form the foundation for the current Cisco firewall technology
- Describe how a static packet filter allows or blocks data packets as they pass through a network interface
- Describe the operation of a circuit level firewall
- Describe how application layer or proxy firewalls control or monitor inbound and outbound traffic
- Explain how dynamic or stateful inspection packet filtering improves network security and performance
- Describe the cut-through proxy process used by the Cisco IOS Firewall
- Describe how to implement NAT on a firewall device
- Describe application inspection firewalls, also called deep inspection firewalls
- Explain the role of firewalls in a layered defense strategy

Explaining a Firewall

This topic describes the role of firewalls in securing networks.



The term “firewall” is a metaphor. By segmenting a network into different physical subnetworks, firewalls can limit the damage that could spread from one subnet to another—just like fire doors and firewalls used by firefighters limit the spread of fire. In network security terms, a firewall is a software or hardware barrier between an internal (trusted) network and an external (untrusted) network. In this sense, a firewall is a set of related programs that enforces an access control policy between two or more networks.

In principle, a firewall is a pair of mechanisms that perform these two separate functions:


- One mechanism blocks traffic.
- The second mechanism permits traffic.

Specific firewall designs or concepts balance these two functions by either placing greater emphasis on blocking traffic or on permitting traffic based on your specifications. Firewalls implement an access control policy that is defined before implementing the selected firewall solution. Once deployed, the firewall enforces access to and from the firewall. The larger the network behind the firewall is, the more important the design.

A firewall can also manage public access to private network resources such as host applications. It can log all attempts to enter the private network and trigger alarms when hostile or unauthorized entry is attempted. Firewalls can filter packets based on their source, destination address, and port number. Firewalls can also filter specific types of network traffic based on protocol and forward or reject traffic depending on the protocol used (HTTP, FTP, or Telnet). Firewalls can also filter traffic by packet attribute or state.

Evolution of Firewall Technologies

This topic describes four generations of firewall technologies developed between 1983 and 1995 that form the foundation for the current Cisco firewall technology.



Before firewalls had the advanced capabilities of the Cisco PIX Security Appliance and Cisco IOS Firewall, all firewalls inspected network traffic using one of four architectural models defined by the information that they examine to make security-relevant decisions. The initial four firewall technologies are as follows:

- **Static packet filtering firewalls:** A packet filter firewall is first-generation firewall technology that analyzes network traffic at the transport protocol layer. Each IP network packet is examined to see if it matches one of a set of rules defining which data flows are allowed. The rules identify whether communication is allowed based on information contained within the network and transport layer headers and the direction in which the packet is headed (internal to external network or vice versa).
- **Circuit level firewalls:** A circuit level firewall is second-generation firewall technology that validates the fact that a packet is either a connection request or a data packet belonging to a connection, or virtual circuit, between two peer transport layers.
- **Application layer firewalls:** An application layer firewall is third-generation firewall technology that evaluates network packets for valid data at the application layer before allowing a connection. This firewall examines the data in all network packets at the application layer and maintains complete connection state and sequencing information. In addition, an application layer firewall can validate other security items that only appear within the application layer data, such as user passwords and service requests.

- **Dynamic packet filtering firewalls:** A dynamic packet filter firewall is fourth-generation firewall technology. Dynamic packet filters, or stateful firewalls, keep track of the actual communication process by using a state table. A stateful firewall operates at Layers 3, 4 and 5.


Each technology has advantages and disadvantages, and each one has a role to play depending on the needs of the security policy.

Cisco acquired the original Private Internet Exchange (PIX) technology in 1995 and continues to develop PIX capabilities. Cisco PIX appliances are network layer firewalls with stateful inspection. By design, these firewalls allow internal connections out (outbound traffic) and only allow inbound traffic that is a response to a valid request or is allowed by an access control list (ACL). You can configure Cisco PIX technology to perform many functions including Network Address Translation (NAT) and Port Address Translation (PAT).

Using the features of the Cisco IOS Firewall embedded in Cisco IOS software allows you to turn your router into an effective, robust firewall that shares many of the capabilities of the Cisco PIX Security Appliance.

Static Packet Filtering Firewalls


This topic describes how a static packet filter allows or blocks data packets as they pass through a network interface.



Packet filtering firewalls work at the network level of the Open Systems Interconnection (OSI) model, or the IP layer of TCP/IP. Packet filtering firewalls are usually part of a router firewall. As a Layer 3 device, packet filtering uses rules and ACLs to determine whether to permit or deny traffic based on source and destination IP addresses, source and destination port numbers, and packet type. These firewalls also use rules to reject any packet from the outside that claims to come from an address inside the network. Recall that each service relies on specific ports. Because this firewall filters according to static packet header information, packet filtering firewalls are sometimes called static filters. Packet filtering firewalls are usually part of a router firewall.

By restricting certain ports, you can restrict the services that rely on certain ports. For example, blocking port 25 on a specific workstation prevents an infected workstation from broadcasting e-mail viruses across the Internet.

Packet filtering firewalls are like packet filtering routers but with some differences in implementation. Packet filters are very scaleable, application-independent, and have high performance standards; however, they do not offer the complete range of security solutions required in modern networks.



Any device that uses ACLs can perform packet filtering. Cisco IOS router configurations commonly use ACLs, not only as packet filtering firewalls, but also to select specified types of traffic to be analyzed, forwarded, or influenced in some way.

The figure shows a simple packet filtering example using a Cisco router.

In most network topologies, you need to protect the Ethernet interface connecting to the internal (inside) network, while the serial interface that connects to the Internet (outside) is unprotected. In this example, the internal addresses that the firewall must protect are in the 12.23.34.x range (on the Ethernet interface). The subnet mask is 255.255.255.0, making the IP address of the Ethernet 0 interface 12.23.34.1 255.255.255.0.

The particular network security policy shown in the slide (ACL 101) allows all users from the inside to access Internet services on the outside. Therefore, all outgoing connections are accepted. The router only checks packets coming from the Internet (security policy ACL 102). In this case, the ACL allows Domain Name System (DNS), mail, FTP services, and the return of traffic initiated from the inside. ACL 102 denies access to all other services.

Advantages and Disadvantages of Packet Filters

Advantages:

- Based on simple permit or deny rule sets
- Low impact on network performance
- Easy to implement
- Supported by most routers
- Afford an initial degree of security at a low network layer
- Perform 90 percent of what higher-end firewalls do, but at a much lower cost

Disadvantages:

- Susceptible to IP spoofing
- Do not filter fragmented packets well
- Complex ACLs difficult to implement and maintain correctly
- Cannot filter certain services

© 2006 Cisco Systems, Inc. All rights reserved.SND v2.0—4-7

Packet filter firewalls (or packet filters) use a simple policy table lookup based on {source-ip, destination-ip, source-port, destination-port, SYN-seen yes/no} permit or deny rule sets. The firewalls are extremely fast because they do little computation. The rules are extremely easy to implement because they require little security expertise. Router manufacturers easily embed packet filtering logic in silicon and, consequently, packet filtering is a feature of most routers. Packet filtering firewalls are relatively inexpensive. Even if other firewalls are used, implementing packet filtering at the router level affords an initial degree of security at a low network layer.


Packet filters do not represent a complete firewall solution. However, considering that filtering on Layer 3 traffic represents 90 percent of what firewalls do, packet filters are a key element of a complete firewall solution.

Here are the disadvantages to packet filters:

- Packet filtering is susceptible to IP spoofing. Hackers send arbitrary packets that fit ACL criteria and pass through the filter.
- Packet filters do not filter fragmented packets well. Because fragmented IP packets carry the TCP header in the first fragment and packet filters filter on TCP header information, all fragments after the first fragment are passed unconditionally. Decisions to use packet filters assume that the filter of the first fragment accurately enforces the policy.
- Complex ACLs are difficult to implement and maintain correctly.
- Packet filters cannot dynamically filter certain services. For example, sessions that use dynamic port negotiations without opening access to a whole range of ports are difficult to filter.

Circuit Level Firewalls

This topic describes the operation of a circuit level firewall.



A circuit level firewall, also called a circuit level gateway, is second-generation firewall technology that validates that a packet is either a connection request or a data packet belonging to a connection or virtual circuit between two peer transport layers. In addition to allowing or disallowing packets, the circuit level firewall also determines whether the connection between both ends is valid according to configurable rules.

To validate a session, a circuit level firewall examines each connection setup to ensure that the connection follows a legitimate TCP handshake. In addition, the firewall will not forward data packets until the handshake is complete. Any information passed to a remote computer through a circuit level firewall appears to have originated from the gateway. This is useful for hiding information about protected networks.


A circuit level firewall maintains a table (including complete session state and sequencing information) of valid connections and allows network packets containing data to pass through when network packet information matches an entry in the virtual circuit table. Once the firewall terminates a connection, it removes the table entry for that connection, and the virtual circuit between the two peer transport layers is closed.

Circuit level filtering has an advantage over packet filtering because it can make up for the shortcomings of the UDP protocol in which the source address is never validated as a function of the protocol. This makes IP spoofing much more difficult.

However, because circuit level firewalls work at the transport layer, they require substantial modification to the programming that normally provides transport functions; for example, Windows Socket Interface (WinSock).

Application Layer or Proxy Firewalls

This topic describes how application layer or proxy firewalls control or monitor inbound and outbound traffic.



Application layer firewalls, also called proxy firewalls or application gateways, provide a higher level of security than circuit level firewalls because they allow the greatest level of control. Application level proxy servers work up on Layer 1 to Layer 7 of the OSI model.

Most application layer firewalls include specialized application software and proxy servers. A proxy is an application that does work on behalf of something else. Proxy services are special-purpose programs that manage traffic through a firewall for a specific service, such as HTTP or FTP. Proxy services are specific to the protocol that they are designed to forward, and they can provide increased access control and careful detailed checks for valid data and generate audit records about the traffic that they transfer.

Proxy firewalls act as intermediaries between networks to determine whether to allow the communication to proceed. There is no direct connection between an outside user and internal network resources. The proxy is the only device with a visible IP address on the Internet. The client connects to the proxy server and submits an application layer request. The application layer request includes the true destination and the data request itself. The proxy server analyzes the request and may filter or change the packet contents. The server makes a copy of each incoming packet, changes the source address, and sends the packet to the destination address. The destination server replies to the proxy server, and the proxy server passes the response back to the client.

Proxy servers control or monitor outbound traffic by protecting private network servers inside the network. These servers require users to communicate with a secure system. Users gain access to the network by going through a proxy that establishes the session state, user authentication, and authorized policy. This means that users connect to services through application programs (proxies) running on the gateway that connects to the outside, unprotected zone.

Application Layer Proxy Firewall

- An application layer firewall operates on OSI Layers 3, 4, 5, and 7.
- Advantages of application layer proxy firewalls:
 - This firewall authenticates individuals, not devices.
 - Hackers have a harder time with spoofing and implementing DoS attacks.
 - This firewall can monitor and filter application data.
 - This firewall can provide detailed logging.

Layer 7	Application
Layer 6	Presentation
Layer 5	Session
Layer 4	Transport
Layer 3	Network
Layer 2	Data Link
Layer 1	Physical

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0-4.10

Application layer firewalls filter information at Layers 3, 4, 5, and 7 of the OSI reference model. Because application layer firewalls process information at the application layer, they do most firewall control and filtering in the software. Locating the firewall at the application layer provides much more control over traffic than packet filtering, stateful, or application inspection firewalls do.


Sometimes, application layer firewalls support only a limited number of applications, or even just one application. Some of the more common applications that an application layer firewall might support include e-mail, web services, DNS, Telnet, FTP, USENET news, Lightweight Directory Access Protocol (LDAP), and finger.

Here are some of the advantages of application layer firewalls:

- **Application layer firewalls authenticate individuals, not devices:** These firewalls typically allow you to authenticate connection requests before allowing traffic to an internal or external resource. This process enables you to authenticate the user requesting the connection instead of authenticating the device.
- **It is harder for hackers to spoof and implement DoS attacks:** An application layer firewall enables you to prevent most spoofing attacks, and DoS attacks are limited to the application firewall itself. The application firewall can detect DoS attacks, reducing the burden on your internal resources.
- **Application layer firewalls can monitor and filter application data:** You can monitor all data on a connection, so you can detect application attacks such as malformed URLs, buffer overflow attempts, unauthorized access, and more. You can even control what commands or functions you allow an individual to perform based on the authentication and authorization information.


- **Application layer firewalls can provide detailed logging:** Using application layer firewalls, you can generate very detailed logs and monitor the actual data that the individual is sending across a connection. This can be extremely useful if a hacker finds a new type of attack, because you can monitor what the hacker does and how the machine does it and then address the attack. Besides using logging for security purposes, you can use it for management purposes by keeping track of who is accessing what resources, how much bandwidth is used, and how often a user accesses the resources.

Application Level Proxy Firewall



The figure shows a simple device acting as an application level proxy server. Application level proxy servers run at the application level of the network protocol stack for each different type of service (for example FTP or HTTP). An application level proxy firewall controls how internal users access the outside world (the Internet) and how Internet users access the internal network. In some cases, the proxy server blocks all outside connections and only allows internal users to access the Internet. The only packets allowed back through the proxy server are those that return responses to requests from inside the firewall. In other cases, the firewall allows both inbound and outbound traffic under strictly controlled conditions. This setup is like a virtual gap that exists in the firewall between the inside and outside networks. The proxy servers bridge this gap by working as agents for internal or external users.

Proxy Server Communication Process



- The proxy server requests connections between a client on the inside of the firewall and the Internet.
- Client requests are filtered on the basis of Layer 5 and Layer 7 information.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-12

The topology in the figure represents a typical proxy server deployment.

An application layer firewall usually has two network interfaces; one interface is used for the client connections, and a second interface is used for accessing the website from the Internet. Application proxies separate the trusted and untrusted networks either physically or logically.

The example in the figure shows a client inside the network requesting access to a website. The client browser uses a proxy server for all HTTP requests. As shown, the browser connects to the proxy server to make requests. Client-side DNS queries and client-side routing to the Internet are not needed when using a proxy server. The client only has to reach the proxy server to make the request.

When the proxy server receives the request from a client, it performs user authentication according to the rules applied to it and uses its Internet connection to access the requested website. It only forwards Layer 3 and Layer 4 packets that match the firewall rules. On the return route, the proxy server only forwards Layer 5 and Layer 7 messages and content that the server allows (that is, traffic that is not seen as malicious) according to the firewall rules.

In spite of how application layer firewalls work, this firewall provides only one service; it provides the highest level of filtering for a specific protocol. A disadvantage of the proxy server is that it slows network performance, because the server has to evaluate a significant amount of information embedded in many packets.

Limitations and Uses of Application Layer Firewalls

Limitations of application layer firewalls:

- Process packets in software
- Support a small number of applications
- Sometimes require special client software
- Are memory and disk space (logging) intensive

Uses for application layer firewalls:

- Use only for key applications where performance can be sacrificed for security

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0-4.13

The main limitation of application layer firewalls is that they are very process intensive. An application layer firewall requires many CPU cycles and a lot of memory to process every packet that needs inspection, which sometimes creates throughput problems. In addition, the detailed logging can create disk space problems. To address these issues, you can use one of these two solutions:

- Use a Context Transfer Protocol (CTP)
- Have the application layer firewall monitor only key applications


Using a CTP enables you to perform authentication and authorization only; you cannot monitor data on the connection. With the second solution, you limit the application layer firewall to processing only certain application types (such as e-mail, Telnet, FTP, or web services), and then, perhaps, processing only connections to specific internal resources. The problem with this approach is that you are not monitoring all applications and connections, and this creates a security weakness.

Application layer firewalls typically do not support all applications, but are generally limited to one or a small number of connection types, such as e-mail, Telnet, FTP, or web services. Therefore, you cannot monitor data on all connections.

Finally, application layer firewalls sometimes require you to install vendor-specific software on the client, which the firewall uses to handle the authentication process and any possible connection redirection. This limitation can create scalability and management problems if you need to support thousands of clients.

Dynamic or Stateful Packet Filtering Firewalls

This topic explains how dynamic or stateful inspection packet filtering provides improved network security and performance.



Stateful packet filters, or stateful firewalls, are the most versatile and therefore the most common firewall technologies in use. Stateful filtering provides dynamic packet filtering capabilities to firewalls. Stateful inspection is firewall architecture that works at the network layer. Unlike static packet filtering, which examines a packet based on the information in its header, stateful inspection tracks each connection traversing all interfaces of the firewall and confirms that they are valid. Stateful packet filtering maintains a state table. The state table is part of the internal structure of the firewall and tracks all sessions and inspects all packets passing through the firewall. If packets have the expected properties predicted by the state table, the firewall allows them to pass. The state table changes dynamically according to traffic flow.

Stateful firewalls keep track of the actual communication process by using a state table. Stateful firewalls operate at Layers 3, 4, and 5. From a transport layer perspective, the firewall examines information in the headers of Layer 3 packets and Layer 4 segments. For example, the firewall looks at the TCP header for SYN, RST, ACK, FIN, and other control codes to determine the state of the connection. In this scenario, the session layer is responsible for establishing and tearing down the connection.


When an outside service is accessed, the stateful packet filter firewall “remembers” certain details of the request by saving the state of the request in the state table. Each time a TCP or UDP connection is established for inbound or outbound connections, the firewall logs the information in a stateful session flow table. When the outside system responds to your request, the firewall server compares the received packets with the saved state to allow or deny network access.

The stateful session flow table contains the source and destination addresses, port numbers, TCP sequencing information, and additional flags for each TCP or UDP connection associated with that particular session. This information creates a connection object used by the firewall to compare all inbound and outbound packets against session flows in the stateful session flow table. The firewall permits data only if an appropriate connection exists to validate the passage of that data.

More advanced stateful firewalls include the ability to parse FTP port commands and update the state table to allow FTP to work transparently through the firewall. TCP sequence number interpretation and DNS query and response matching ensure that the firewall only allows packets to return in response to queries that originate from inside the network. These features reduce the threat of TCP RST flood attacks and DNS cache poisoning.

Stateful Filtering

Stateful Firewall



© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-15

There is a potential disadvantage of using stateful filtering that you must consider. While stateful inspection provides speed and transparency, packets inside the network must make their way to the outside network. This can possibly expose internal IP addresses to potential hackers. Most firewalls incorporate stateful inspection, NAT, and proxy servers for added security.

To overcome this disadvantage, stateful firewalls keep track of the state of a connection and whether the connection is in an initiation, data transfer, or termination state. This information is useful when you want to deny the initiation of connections from external devices but allow your users to establish connections to these devices and permit the responses to come back through the stateful firewall.

The example in the figure shows a successfully established HTTP TCP session that leads to a dynamic ACL rule entry on the outside interface and permits response packets from the web server to the client.

Limitations and Uses of Stateful Firewalls

Use stateful firewalls:

- As a primary means of defense
- As an intelligent first line of defense
- As a means of strengthening packet filtering
- To improve routing performance
- As a defense against spoofing and DoS attacks

Limitations:

- Stateful firewalls cannot prevent application layer attacks.
- Not all protocols contain state information (for example, UDP and ICMP).
- Some applications open multiple connections, some of which use dynamic port numbers for the additional connections .
- Stateful firewalls do not support user authentication of connections.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0-4.18

Use stateful packet filtering firewalls in these applications:


- **As a primary means of defense:** In most situations, a stateful firewall is used as a primary means of defense by filtering unwanted, unnecessary, or undesirable traffic.
- **As an intelligent first line of defense:** Networks use routing devices supporting a stateful function as a primary line of defense or as an additional security boost on perimeter routers.
- **As a means of strengthening packet filtering:** Stateful filtering provides more stringent control over security than does packet filtering without adding too much cost.
- **To improve routing performance:** Stateful packet filtering devices perform better than packet filters or proxy servers. Stateful firewalls do not require a large range of port numbers to allow returning traffic back into the network. The state table determines whether a packet is returning traffic. If it is not returning traffic, the filtering table filters the traffic.
- **As a defense against spoofing and DoS attacks:** Stateful packet filtering works on packets and connections. In particular, stateful firewalls track the state of the connection in the state table listing every connection or connectionless transaction. By determining whether packets belong to an existing connection or are from an unauthorized source, stateful firewalls only allow traffic from connections listed in the table. Once the firewall removes a connection from the state table, the firewall will not allow any more traffic from that device. In addition, the stateful firewall can log more information than a packet filtering firewall can, including when a connection was set up, how long it was up, and when it was torn down. This logging makes connections harder to spoof.

Stateful firewalls have these limitations:

- **Stateful firewalls cannot prevent application layer attacks:** For example, your network might allow traffic to port 80 to a web server. Your stateful firewall will examine the destination address in the Layer 3 packet and the destination port number in the segment. If there is a match, the stateful firewall allows the incoming and outgoing traffic. One problem with this approach is that the stateful firewall does not examine the actual contents of the HTTP connection.
- **Not all protocols are stateful:** UDP and Internet Control Message Protocol (ICMP) are not stateful and cannot be monitored by stateful firewalls. For example, UDP has no defined process for how to set up, maintain, and tear down a connection. Routers define UDP connections on an application-by-application basis.
- **Some applications open multiple connections:** With FTP, if the client is inside the network and the server is outside the network, both stateful and packet filtering firewalls have problems dealing with the data connection that the FTP server establishes to the client. You would have to open a whole range of ports to allow this second connection.
- **Stateful firewalls do not support user authentication:** Stateful firewall technology itself does not support user authentication.

Cut-Through Proxy Process

This topic describes the cut-through proxy communication process used by the Cisco IOS Firewall.



Cisco's firewall technology performs dramatically better than competing firewalls. A proprietary process called cut-through proxy is the fastest way for a firewall to authenticate a user. Using the cut-through proxy feature of the Cisco PIX Security Appliance or Cisco IOS Firewall helps alleviate performance issues inherent in proxy server design. Firewalls using a cut-through proxy gain dramatic performance advantages over proxy servers. Cut-through proxy is a method of transparently verifying the identity of users at the security appliance and permitting or denying access to any TCP- or UDP-based application. After the packet and user “pass” the firewall security policy at the application layer, the user is “cut-through” the firewall, and all traffic flows directly and quickly between the server and the client while maintaining session state information at Layer 3 and Layer 4. This method eliminates the price and performance changes that packet filtering firewalls impose in similar configurations.

The example in the figure shows the process that a cut-through proxy uses. In this example, Richard tries to access the internal web server (200.1.1.1). In Step 1, the cut-through proxy intercepts the connection request and authenticates Richard. After authentication (in Step 2), the cut-through proxy adds the authenticated connection and any other authorized connections to the filtering rules table. From here, the filtering rules at Layer 3 and Layer 4 handle any traffic from Richard to the web server (Step 3). This filtering provides a significant boost in throughput. However, the downside is that the cut-through proxy does not examine application layer data and, therefore, it cannot detect application layer attacks.

Typically, the cut-through proxy supports FTP, Telnet, HTTP, and HTTP Secure (HTTPS) for handling the initial authentication by having the user set up authentication to the cut-through proxy itself. Optionally, some cut-through proxies intercept certain connections and respond with authentication prompts. After authentication, the cut-through proxy allows the connection initiation request to the internal resource.

Implementing NAT on a Firewall

This topic describes how to implement NAT on a firewall device.

Implementing NAT on a Firewall

NAT operates at OSI Layer 3 and Layer 4.

Advantages of implementing NAT on a firewall:

- NAT hides your network addressing design.
- NAT controls the traffic entering and leaving your network.
- NAT allows for the use of private addressing.

Layer 7	Application
Layer 6	Presentation
Layer 5	Session
Layer 4	Transport
Layer 3	Network
Layer 2	Data Link
Layer 1	Physical


© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0--4-18

NAT enables a LAN to use one set of IP addresses for internal traffic and a second set of addresses for external traffic. By translating the internal IP addresses, NAT actually hides your network addressing design from the outside. Firewalls often have NAT functionality, and the hosts protected behind a firewall commonly use private address space.

A firewall using NAT can provide these advantages:

- NAT hides your network addressing design.
- NAT controls the traffic entering and leaving your network. NAT can serve as a “choke point” if your internal devices are using private addresses.
- NAT allows for the use of private addressing, providing millions of IP addresses at your disposal including 1 Class A, 16 Class B, and 256 Class C network numbers.

Network Address Translation



NAT translates the source address of a device inside a network to a public source address (“SA” in the figure).


© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-19

NAT allows a host on your private network that does not have a valid registered IP address to communicate with other hosts through the Internet. There are three types of NAT to consider:

- **Static NAT:** In static NAT, a private IP address is mapped to a public IP address, where the public address is always the same IP address (that is, a static address). A static address allows an internal host, such as a web server, to have an unregistered (private) IP address and still be reachable over the Internet via a fixed outside public address.
- **Dynamic NAT:** Like static NAT, in dynamic NAT the NAT router creates one-to-one mapping between an inside local and outside global address and changes the IP addresses in packets as they flow through the firewall. However, the mapping of an inside local address to an outside global address happens dynamically. Dynamic NAT sets up a pool of possible inside global addresses and defines criteria for the set of outside local IP addresses whose traffic NAT will translate. Typically, the NAT router keeps a table of registered IP addresses. When a private IP address requests access to the Internet, the router chooses an available IP address from the table. Dynamic NAT helps to secure a network as it masks the internal configuration of a private network and makes it difficult for someone outside the network to monitor individual use patterns. Dynamic NAT allows a private network to use inside private IP addresses that are invalid on the Internet but useful as internal addresses.

Port Address Translation




PAT extends NAT from “1 to 1” to “many to 1” by associating the source port with each flow.

- **PAT:** PAT is a type of network address translation. During PAT, each computer on a LAN is translated to the same global IP address but with a different port number assignment; for example, in the figure, 10.1.1.1:1024 is mapped to 217.2.2.1:2048.

Configuring NAT with Cisco SDM

1. In Cisco SDM, choose Configure.
2. Choose the NAT wizard on the task bar.
3. Choose Basic NAT or Advanced NAT.
4. Click the Launch the Selected Task button.



You can use the Cisco Router and Security Device Manager (SDM) NAT wizard to guide you in creating a NAT rule.

Choose the Basic NAT wizard if you want to connect your network to the Internet (or the outside) and your network has hosts but no servers. If your network is made up only of PCs that require access to the Internet, choose **Basic NAT** and click the **Launch the Selected Task** button.

Choose the Advanced NAT wizard if you want to connect your network to the Internet (or the outside), and your network has hosts and servers, and the servers must be accessible to outside hosts (hosts on the Internet). If your network has e-mail servers, web servers, or other types of servers and you want them to accept connections from the Internet, choose **Advanced NAT** and click the **Launch the Selected Task** button.

Note If you do not want your servers to accept connections from the Internet, you can use the Basic NAT wizard.

Limitations and Uses of NAT

Uses:

- When you have a private IP addressing scheme in your internal network
- When you need to separate two or more networks

Limitations:

- Delay is introduced because of packet manipulations.
- Some applications do not work with address translation.
- Using multiple layers of NAT is complicated.
- Tracing and troubleshooting become more difficult.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0-4.22

Address translation differs from other types of firewall techniques, such as application layer firewalls, stateful firewalls, and packet filtering firewalls. These latter firewalls filter traffic based on filtering rules you define that can be very specific. Address translation implemented in a firewall only translates the address (and sometimes port) information.

Use NAT effectively in these scenarios:

- **On networks using private IP addresses:** If you are using private IP addresses, you need to use some type of address translation device to allow traffic into and out of your network. An address translation firewall provides extra security when performing this process.
- **For controlling traffic between networks:** Use NAT when you need to separate two or more networks and control traffic between them. For example, one of the networks might be using private addresses, or the two networks might be using an overlapping address space.

Like other solutions, NAT implemented by a firewall has limitations. Here are some examples:

- **Packet manipulation introduces delay:** Because address translation must change the IP address in the IP header, the firewall will recompute the header checksum. If the firewall uses PAT on TCP or UDP segment headers, it will recompute the checksums. This process is very time-intensive and adds delay to your data stream.

- **Some applications do not work with address translation:** Not all protocols function correctly or at all when the router performs address translation on the packet and segment contents. Some network applications assume that the IP address and port assigned to the client will always be globally routable and can be used on the Internet directly. In many cases, the addresses are private IP addresses from Internet Engineering Task Force (IETF) reserved address ranges. The application will include this private IP address or port in the payload of packets sent to the server. The server may use this embedded address as the address to contact the client. If the server attempts to reply using the embedded IP private address and port instead of the mapped address and port supplied by the NAT, the server drops the packet. This action occurs because the embedded IP address is nonroutable. If the network application could discover the presence of a NAT and retrieve the external IP address and external port mapping that it needs, the application could embed the right information in the packet.
- **Using multiple layers of NATs becomes difficult:** If a client is behind a NAT that is behind another NAT, new problems that are beyond the scope of this course appear.
- **Tracing and troubleshooting become more difficult:** Tracing and troubleshooting becomes difficult when using NAT. By hiding the addressing scheme, the network is more secure. However, hackers can use this situation to their advantage by hiding their source addressing information. Thus when a hacker attacks, it becomes much more difficult to track down the culprit. Also, troubleshooting is not an easy process when dealing with address translation, because you must know both the IP address assigned to a device on its network interface coprocessor and its translated address when trying to find the source and destination of a connection.

Application Inspection Firewall

This topic describes application inspection firewalls, also called deep inspection firewalls.

Application Inspection Firewall

An application inspection firewall operates on OSI Layers 3, 4, 5, and 7. Application inspection firewalls are essentially stateful firewalls with intrusion detection system capabilities.

Application inspection firewalls:

- Are aware of the Layer 5 state of a connection
- Check the conformity of application commands on Layer 5
- Are able to check and affect Layer 7 (for example, Java applet or peer-to-peer filtering)
- Prevent more kinds of attacks than stateful firewalls

Layer 7	Application
Layer 6	Presentation
Layer 5	Session
Layer 4	Transport
Layer 3	Network
Layer 2	Data Link
Layer 1	Physical

© 2006 Cisco Systems, Inc. All rights reserved. SNO v2.0—4-23

Application inspection firewalls ensure the security of applications and services. Some applications require special handling by the firewall application inspection function. Applications that require special application inspection functions are those that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports.

The application inspection function works with NAT to help identify the location of embedded addressing information. This arrangement allows NAT to translate embedded addresses and to update any checksum or other fields that are affected by the translation.

The application inspection function also monitors sessions to determine the port numbers for secondary channels. Many protocols open secondary TCP or UDP ports. The initial session on a well-known port is used to negotiate dynamically assigned port numbers. The application inspection function monitors these sessions, identifies the dynamic port assignments, and permits data exchange on these ports for the duration of the specific session. For example, the FTP client opens a control channel between its port 2008 and the FTP server port 21. When data is to be exchanged, the FTP client alerts the FTP server through the control channel that it expects the data to be delivered back from FTP server port 20 to its port 2010. If FTP inspection is not enabled, the return data from FTP server port 20 to FTP client port 2010 is blocked by the stateful firewall. With FTP inspection enabled, however, the stateful firewall inspects the FTP control channel to recognize that the data channel will be established to the new FTP client port 2010 and temporarily creates an opening for the data channel traffic for the life of the session.

An application inspection firewall behaves in different ways according to each layer.

- **Transport layer mechanism:** From a transport layer perspective, the application inspection firewall acts like a stateful firewall by examining information in the headers of Layer 3 packets and Layer 4 segments. For example, the application inspection firewall looks at the TCP header for SYN, RST, ACK, FIN, and other control codes to determine the state of the connection.
- **Session layer mechanism:** From a session layer perspective, the application inspection firewall checks the conformity of commands within a known protocol. For example, when the application inspection firewall checks the Simple Mail Transfer Protocol (SMTP), only acceptable message types on Layer 5 are allowed (that is, DATA, HELO, MAIL, NOOP, QUIT, RCPT, RSET). In addition, the application inspection firewall checks whether the command attributes that are used (for example, length of a message type) conform to the internal rules. These rules often trust the RFC of a specific protocol.
- **Application layer mechanism:** From an application layer perspective, the application inspection firewall protocol is rarely supported. Sometimes, application layer firewalls provide protocol support for HTTP, and the application inspection firewall can determine whether the content is really an HTML website or a tunneled application, such as Kazaa Media Desktop or eDonkey. In the case of a tunneled application, the application inspection firewall would block the content or terminate the connection. Future development will provide more application inspection support for more protocols on an application inspection firewall.


Here are some of the advantages of an application inspection firewall:

- Application inspection firewalls are aware of the state of Layer 4 and Layer 5 connections. For example, the application inspection firewall knows that a Layer 5 SMTP **mail-from** command always follows a HELO command.
- Application inspection firewalls check the conformity of application commands on Layer 5.
- Application inspection firewalls are able to check and affect Layer 7, as explained before.
- Application inspection firewalls can prevent more kinds of attacks than stateful firewalls can. For example, application inspection firewalls stop an attacker from trying to set up a virtual private network (VPN) tunnel (triggered from inside the network) through your application firewall by way of tunneled HTTP requests.

Application Inspection Firewall Operation

Inspection engines:

- protocol support through firewalls
- conformity of commands through checks



	Inspect Outgoing Traffic	Inspect Incoming Traffic			
1	Session Initiation Protocol	Source	Destination	Protocol	Action
	To: INVITE sip:cch@62.3.3.3 SIP/2.0	62.3.3.3	75.1.1.1	TCP 5060	Permit
	From: <sip:bill@75.1.1.1>;tag=4c101d	Source	Destination	Protocol	Action
2	Media Port: 33005	62.3.3.3	75.1.1.1	UDP 33005	Permit
	HTTP	Filtered Java Applet			
	GET / HTTP/1.1\r\n	<applet code="fbun.class" width=550 height=300			
	Host: www.magazin.com\r\n	align="left">			
		</applet>			

The examples in the figure show inspection engines, a subset of the application inspection firewall. One inspection engine is responsible for checking a specific protocol.

The first example shows how a client establishes a pre-Fast Serial Interface Processor (pre-FSIP) session to the pre-FSIP server followed by a voice call controlled by pre-FSIP. The application inspection firewall dynamically inspects and allows response traffic from the pre-FSIP server. Layer 5 traffic is also being inspected. The pre-FSIP inspection engine recognizes a pre-FSIP call setup by understanding the pre-FSIP protocol INVITE message on Layer 5. The inspection engine dynamically reads the used media port for the Real-Time Transport Protocol (RTP) data stream and dynamically allows that specific traffic to pass the firewall.

The second example shows a user opening a website on a web server. The web server responds by providing access to a website. The HTTP inspection engine on the application inspection firewall recognizes on Layer 7 that the site contains a Java applet and filters the applet because of filtering rules.

Limitations and Uses of Application Inspection Firewalls

Limitations:

- Able to prevent simple application layer attacks
- Usually do not support user authentication of connections
- Size of state table

Uses:

- As a secondary means of defense
- Where more stringent controls over security than packet filtering are needed

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-25

Here are some of the Limitations of an application inspection firewall:


- Because of the complexity of filtering on Layer 7 content, there are only a few inspection engines that support Layer 7 filtering (for example, HTTP).
- Application inspection firewall technology by itself does not support user authentication.
- An application inspection firewall might be busy building and maintaining the state table, which puts an extra burden on the processing capacity of the firewall. The more connections that your application inspection firewall must monitor, the more horsepower your application inspection firewall needs to maintain the table, thus increasing operation cost.

Application inspection firewalls are used effectively in these scenarios:

- **As a secondary means of defense:** In most situations, an application inspection firewall is used as a secondary means of defense by filtering unwanted, malicious, or undesirable traffic.
- **When you need more stringent controls over security than stateful filtering provides:** Application inspection firewalls are more stringent than stateful firewalls but do not add significant cost to your implementation. Application inspection firewalls also provide more control than stateful filtering firewalls do, still at a minimal increase in cost.

Content Filtering Using Websense

1. End user sends HTTP request.
2. Cisco IOS firewall forwards the request to the web server and sends a look up request of the requested URL to the Websense server.
3. The Websense server compares the URL to its database. It returns a permit or deny status via a look up response to the Cisco IOS Firewall.
4. If permitted, the user receives an http request. If denied, the user is directed to an internal web server on the Websense server.



The firewall Websense URL filtering feature enables your Cisco IOS Firewall to interact with the Websense URL filtering software. This allows you to prevent users from accessing specified websites based on your security policy. The Cisco IOS Firewall works with the Websense server to know whether a particular URL should be allowed or denied.

Websense URL filtering follows these steps:

- Step 1** The end user browses a page on the web server, and the browser sends an HTTP request.
- Step 2** After the Cisco IOS Firewall receives this request, it forwards the request to the web server while simultaneously extracting the URL and sending a lookup request to the Websense server.
- Step 3** After the Websense server receives the lookup request, it checks its database to see whether it should permit or deny the URL; it returns a permit or deny status via a lookup response to the Cisco IOS Firewall.
- Step 4** After the Cisco IOS Firewall receives this lookup response, if the lookup response permits the URL, it sends the HTTP response to the end user. If the lookup response denies the URL, the Websense server redirects the user to its own internal web server, which displays a message that describes the category under which the URL is blocked; thereafter, the connection is reset at both ends.


The Websense server must be part of the protected network, and requests from the Cisco IOS Firewall will not travel over any unprotected network to reach the Websense server.

These filtering methods are supported:

- Global filtering, which is applied to all users, groups, and IP addresses
- User- or group-based filtering, which is applied to a specific user or group
- Keyword-based filtering, which is applied on the basis of specific keywords (For example, a user can configure a policy for which all URLs with the keyword “dog” will be denied.)
- Category-based filtering, which is applied on the basis of specific categories
- Customized filtering, which allows the user to apply a policy for customized URLs

Firewalls in a Layered Defense Strategy

This topic describes the role of firewalls in a layered defense strategy.



In a layered defense scenario, firewalls provide perimeter security of the entire network and of internal network segments in the core. For example, system administrators can use a firewall to separate the human resources or financial networks of an organization from other networks or network segments within the organization.

A layered defense uses different types of firewalls combined in layers to add depth to the information defense of an organization. For example, traffic that comes in from the untrusted network first encounters a packet filter on the outer router. The traffic goes to the screened host firewall or bastion host system that applies more rules to the traffic and discards suspect packets. The traffic now goes to an interior screening router. Only after this routing does the traffic move to the internal destination host. This type of demilitarized zone (DMZ) setup is called a screened subnet configuration.

The common misconception is that a layered firewall topology is all that you need to declare your internal network to be safe. This myth is probably encouraged by the booming firewall business; however, you need to consider these factors when building a complete defense in depth:

- A significant number of intrusions come from hosts within the network. For example, firewalls often do little to protect against viruses downloaded through e-mail.
- Firewalls do not protect against rogue modem installations. In addition, and most importantly, a firewall is no substitute for informed administrators and users.
- Firewalls do not replace backup and disaster recovery mechanisms resulting from attack or hardware failure. An in-depth defense also includes offsite storage and redundant hardware topologies.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- A firewall is a pair of mechanisms: one mechanism blocks traffic, and the second mechanism permits traffic
- There are four firewall technologies: packet filtering, proxy server, dynamic packet filtering and application inspection, each with strengths and weaknesses.
- Static packet filters provide an effective firewall capability in a layered defense architecture by examining source, destination, port, and service details.
- A circuit level firewall validates that a packet is either a connection request or a data packet belonging to a connection or virtual circuit between two peer transport layers.
- Proxy firewalls and servers provide additional security by inspecting the contents of packets.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-28

Summary (Cont.)

- Stateful firewalls are more efficient than static filters and proxies.
- The Cisco IOS Firewall cut-through proxy feature helps alleviate performance issues inherent in proxy server design.
- NAT hides internal IP addresses from users outside the network.
- Application inspection firewalls ensure the security of applications and services for applications that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports.
- Firewalls are part of a layered defense strategy and are deployed throughout the network.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-29

Building Static Packet Filters with Cisco ACLs

Overview

Cisco provides basic traffic filtering capabilities with access control lists (ACLs). You can configure ACLs for all routed network protocols to filter packets as the packets pass through a router or security appliance. There are many reasons to configure ACLs. For example, you can use ACLs to restrict the contents of routing updates or to provide traffic flow control. One of the most important reasons to configure ACLs is to provide security for your network; this is the reason focused on in this lesson.

This lesson outlines the types of ACLs that are available and provides guidelines that help create ACLs to provide network security.


Objectives

Upon completing this lesson, you will be able to build static packet filters with Cisco ACLs. This ability includes being able to meet these objectives:

- Explain how ACLs are used to control access in networks
- Identify the types and formats of ACLs that the Cisco IOS Firewall uses to restrict access and filter packets
- Describe how to apply ACLs to router interfaces
- Explain how to filter traffic with ACLs to block services that hackers use to gather information about your network
- Explain how to use ACLs to filter IP traffic destined for Telnet, SNMP, and RIP
- Explain how to implement ACLs to mitigate a range of threats
- Explain how to configure router ACLs to help reduce the effects of various DDoS attacks
- Describe how to combine many ACL functions into two or three larger ACLs
- Explain some of the caveats that you need to consider when creating ACLs

Access Control Lists

This topic explains how to use ACLs to control access to networks.



ACLs provide packet filtering for routers and firewalls to protect internal networks from the outside world. However, as presented in the “Introducing Firewall Technologies” lesson, ACLs filter network traffic in both directions by controlling whether to forward or block packets at the router interfaces based on the criteria that you specified within the ACLs. ACL criteria could be the source address of the traffic, the destination address of the traffic, the upper-layer protocol, or other information. Be aware, however, that sophisticated users (hackers) can sometimes successfully evade or fool basic ACLs because no authentication is required.

ACLs provide a basic level of security for accessing your network. If you do not configure ACLs on your router, all packets passing through the router could get to all parts of your network. You can 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. An ACL on the router allows one host to access a part of your network and prevents another host from accessing the same area. The ACL shown in the figure allows host A to access the human resources network and prevents host B from accessing the human resources network.

To provide the security benefits of ACLs, you should, at a minimum, configure ACLs at the perimeter 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 network edge routers, you should configure ACLs for each network protocol configured on the router interfaces.

Cisco ACLs

This topic describes the types and formats of ACLs that Cisco IOS Firewall uses to restrict access and filter packets.

Standard and Extended ACLs

Cisco routers support two basic types of ACLs:

- **Standard ACL: Filters IP packets based on the source address only**

```
access-list 10 permit 192.168.3.0 0.0.0.255
```

- **Extended ACL: Filters IP packets based on several attributes, such as:**
 - Source and destination IP addresses
 - Source and destination TCP and UDP ports
 - Protocol type (IP, ICMP, UDP, TCP, or protocol number)

```
access-list 101 permit tcp 63.36.9.0 0.0.0.255 any  
eq 80
```

© 2005 Cisco Systems, Inc. All rights reserved. SMD v2.0-4.4

The Cisco ACL is probably the most commonly used object in Cisco IOS software. This ACL is not only used for packet filtering (a type of firewall), but also for selecting types of traffic to be analyzed, forwarded, or influenced in some way.

The ACL is a group of statements. Each statement defines a pattern in an IP packet. As each packet comes through an interface with an associated ACL, the router scans the list from top to bottom in the exact order in which it appears for a pattern that matches the incoming packet. A permit or deny rule associated with the pattern determines whether the packet is allowed into the network or denied entry to the network.

Cisco routers use ACLs as packet filters to decide which packets can access a router service or cross an interface. Packets allowed across an interface are permitted packets. Packets not allowed across an interface are denied packets.

An ACL enforces one or more corporate security policies. For example, a corporate security policy may allow only packets using source addresses from within the trusted network to access the Internet. Once this policy is written, you can develop an ACL that includes certain statements that, when applied to a router interface, can implement this policy.

Cisco router security depends strongly on well-written ACLs to restrict access to router network services and to filter packets as they traverse the router.

Cisco routers support standard IP ACLs and extended IP ACLs. The figure shows an example of these two types:

- **Standard ACL:** A standard ACL only allows you to permit or deny traffic from specific IP addresses. The destination of the packet and the ports involved do not matter. The example in the figure allows traffic from all addresses in the range 192.168.3.0 to 192.168.3.255.
- **Extended ACLs:** An extended ACL is a series of statements created in global mode. This list can filter IP packets based on several attributes; for example, protocol type, source and IP address, destination IP address, source TCP or UDP ports, destination TCP or UDP ports, and optional protocol type information for finer granularity of control. The example shown in the figure configures ACL 101 to permit traffic originating from any address on the 63.36.9.0/24 network to any destination host port 80 (HTTP).

Identifying ACLs

Cisco routers can identify ACLs using two methods:

- **ACL:** The number of the ACL determines which protocol it is filtering:
 - (1 to 99) and (1300 to 1999): Standard IP ACL
 - (100 to 199) and (2000 to 2699): Extended IP ACL
- **ACL name (Cisco IOS Releases 11.2 and later):** You provide the name of the ACL:
 - Names contain alphanumeric characters.
 - Names cannot contain spaces or punctuation and must begin with an alphabetic character.
 - You can add or delete entries within the ACL.

© 2005 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-5

Prior to Cisco IOS Release 11.2, you had to assign a number to each ACL as you created it. Since then, either a number or a name can identify Cisco ACLs and the protocols that they filter.


Using numbered ACLs is an effective method on smaller networks with more homogeneously defined traffic. Because each ACL type is limited to an assigned range of numbers, it is easy to determine the type of ACL that you are using. There can be up to 99 standard IP ACLs ranging in number from 1 to 99. The extended IP ACL number ranges are 100 to 199, and 2000 to 2699. The “ACL Number and Type” table lists the number range and the type of associated ACL.

ACL Number and Type

ACL Number	Type
1–99	IP standard ACL
100–199	IP extended ACL
200–299	Protocol type-code ACL
300–399	DECnet (developed by Digital Equipment Corporation) ACL
400–499	Xerox Network Systems (XNS) standard ACL
500–599	XNS extended ACL
600–699	AppleTalk ACL
700–799	48-bit MAC address ACL
800–899	Internetwork Packet Exchange (IPX) standard ACL
900–999	IPX extended ACL
1000–1099	IPX Service Advertisement Protocol (SAP) ACL
1100–1199	Extended 48-bit MAC address ACL
1200–1299	IPX summary address ACL
1300–1999	IP standard ACL (expanded range)
2000–2699	IP extended ACL (expanded range)

Beginning with Cisco IOS Release 11.2, you can identify ACLs with an alphanumeric string (a name) rather than a number. Software releases before Cisco IOS Release 11.2 do not recognize named ACLs. Named ACLs allow you to configure more ACLs in a router than if you were to use numbered ACLs alone. If you identify your ACL with a name rather than a number, the mode and command syntax are slightly different. Currently, only packet and route filters can use a named list.

Enable Turbo ACLs on Cisco 7200, 7500, and 12000 Series Routers



```
Router(config)#
```

```
access-list compiled
```

```
Router#
```

```
show access-list compiled
```

```
R2(config)# access-list compiled
```

```
R2(config)# exit
```

```
R2# show access-list compiled
```

© 2005 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-6

Routers search ACLs sequentially to find a matching rule. Because of increasing needs and requirements for security filtering and packet classification, ACLs can expand to the point that searching the ACL adds a significant amount of time and memory when the router is forwarding packets. Additionally, the time taken by the router to search the list is not always consistent, which adds a variable latency to the packet forwarding. A high CPU load is necessary for searching an ACL with several entries.

The Turbo ACL feature, supported by Cisco 7200 Series, Cisco 7500 Series, and Cisco 12000 Series Routers, processes ACLs into lookup tables. Turbo ACLs use the packet header to access these tables in a small, fixed number of lookups, independent of the existing number of ACL entries. The benefits of the Turbo ACL feature are as follows:

- The CPU load is lower for ACLs larger than three entries when matching the packet to the predetermined packet matching. In fact, the Turbo ACL feature fixes the CPU load, regardless of the size of the ACL. This allows the use of larger ACLs without incurring additional CPU overhead penalties. The larger the ACL, the greater the benefit.
- The time taken to match the packet is fixed so that latency of the packets is smaller (significantly in the case of large ACLs) and, more importantly, the time taken to match is consistent, which allows better network stability and more accurate transit times.

If your router supports Turbo ACLs, you should use the **access-list compiled** command in global configuration mode, as shown in the figure, whenever you develop ACLs with more than three statements.

The **access-list compiled** command has no keywords or arguments.

To view the status of your Turbo ACLs, use the **show access-list compiled** command in privileged EXEC mode, as shown in the figure.

The **show access-list compiled** command has no keywords or arguments.

Guidelines for Developing ACLs

Guideline 1: Base your ACL on the security policy.

Guideline 2: Write it out on paper.

- Write out what you need the ACL to accomplish.
- Think about the potential problems that the ACL might cause.

Guideline 3: Set up a development system.

- This allows you to copy and paste statements easily.
- It also allows you to develop a library of ACLs.
- Store the files as ASCII text files.

Guideline 4: Apply ACLs to a router and test.

- If at all possible, run your ACLs in a test environment before placing them into production.

© 2005 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-7

Before you start to develop any ACLs, consider these basic rules:

- **Guideline 1—Base your ACLs on your security policy:** Unless you anchor the ACL in a comprehensive security policy, you cannot be certain that it will effectively control access in the way that access needs to be controlled.
- **Guideline 2—Write it out:** Never sit down at a router and start to develop an ACL without first spending some time in design. The best ACL developers suggest that you write out a list of things that you want the ACL to accomplish. Starting with something as simple as, “This ACL must block all Simple Network Management Protocol (SNMP) access to the router except for the SNMP host at 16.1.1.15.”
- **Guideline 3—Set up a development system:** Whether you use your laptop PC or a dedicated server, you need a place to develop and store your ACLs. Word processors or text editors of any kind are suitable, as long as you can save the files in ASCII text format. Build a library of your most commonly used ACLs and use them as sources for new files. ACLs can be pasted into the router running configuration (requiring console or Telnet access) or can be stored in a router configuration file. The system that you choose should support TFTP to make it easy to transfer any resulting configuration files to the router.


Note

Hackers love to gain access to router configuration development systems or TFTP servers that store ACLs. A hacker can discover a lot about your network from looking at these easily read text files. For this reason, it is imperative that the system where you choose to develop and store your router files be a secure system.

- **Guideline 4—Test:** If possible, test your ACLs in a secure environment before placing them into production. Testing is a common-sense approach to any router configuration changes. Most enterprises maintain their own network test beds. While testing may appear to be an unnecessary cost, over time it can save time and money.

Applying ACLs to Router Interfaces

This topic describes how to apply ACLs to router interfaces.



You must apply packet filtering ACLs to a router interface for the ACL to take effect. You apply the ACLs on an interface based on the direction of the data flow as shown in the figure. You can apply the list to incoming packets (an “in ACL”) or to outgoing packets (an “out ACL”).

- **Inbound (in):** The packet filtering ACL applies to packets received on the router interface.
- **Outbound (out):** The packet filtering ACL applies to packets transmitted out of the router interface. For out ACLs, you need to set up the filter only on the one outgoing interface rather than on the individual incoming interfaces. This improves performance because only the network you are protecting will force a lookup on the ACL.

Applying ACLs to Interfaces

```
Router(config-if)#
```

```
ip access-group {access-list-number | access-  
list-name} {in | out}
```

```
Tulsa(config)# interface e0/1
```

```
Tulsa(config-if)# ip access-group 2 in
```

```
Tulsa(config-if)# exit
```

```
Tulsa(config)# interface e0/2
```

```
Tulsa(config-if)# ip access-group mailblock out
```

```
Tulsa(config-if)# end
```

© 2005 Cisco Systems, Inc. All rights reserved.

SND v2.0-4-9

Before applying a packet filtering ACL to a router interface, make sure you know in which direction it will filter.

Apply ACLs to router interfaces using the **ip access-group** command in interface configuration mode, as shown in the figure.


The syntax for the **ip access-group** command is as follows:

ip access-group {access-list-number | access-list-name} {**in** | **out**}

Command Element	Description
<i>access-list-number</i>	This is the number of the IP standard numbered or IP extended numbered ACL. This number is a decimal number from 1 to 199 or from 1300 to 2699.
<i>access-list-name</i>	This is the name of the IP standard named or IP extended named ACL as specified by the ip access-list command.
in	This command element filters on inbound (flowing toward router interface) packets.
out	This command element filters on outbound (flowing away from router interface) packets.

Using ACLs to Filter Traffic

This topic explains how to filter traffic with ACLs to block services that hackers use to gather information about your network.



To review, always apply these general rules when deciding how to handle router services, ports, and protocols:

- **Disable unused services, ports, or protocols:** In the case where no one, including the router itself, needs to use an enabled service, port, or protocol, disable that service, port, or protocol.
- **Limit access to services, ports, or protocols:** In the case where a limited number of users or systems require access to an enabled router service, port, or protocol, limit access to that service, port, or protocol using ACLs.

ACLs are important because they act as traffic filters between the corporate (trusted) network and the Internet (untrusted network). Using ACLs, the router enforces corporate security policies by rejecting protocols and restricting port usage.

The “Blocked Services” table contains a list of common router services used to gather information about your network and lead to an attack. Unless your network configuration specifically requires one of these services, do not allow them to traverse the router. Use ACLs to block these services inbound to the protected network and outbound to the Internet.

Blocked Services

Service	Port	Transport
tcpmux	1	TCP and UDP
echo	7	TCP and UDP
discard	9	TCP and UDP
systat	11	TCP
daytime	13	TCP and UDP
netstat	15	TCP
chargen	19	TCP and UDP
time	37	TCP and UDP
whois	43	TCP
Bootstrap Protocol (BOOTP)	67	UDP
TFTP	69	UDP
DC-OK		
SUPDUP	93	TCP
SunRPC	111	TCP and UDP
loc-srv	135	TCP and UDP
NetBIOS Name Service (NBNS)	137	TCP and UDP
NetBIOS datagram service (NetBIOS-DGM)	138	TCP and UDP
NetBIOS session service (NetBIOS-SSN)	139	TCP and UDP
X-Display Manager Client Protocol (XDMCP)	177	UDP
NetBIOS	445	TCP
rexec	512	TCP
line printer remote (LPR)	515	TCP
talk	517	UDP
ntalk	518	UDP
UNIX-to-UNIX Copy Program (UUCP)	540	TCP
Microsoft UPnP SSDP	1900, 5000	TCP and UDP
Network File System (NFS)	2049	UDP
Xwindow System	6000-6063	TCP
Internet Relay Chat (IRC)	6667	TCP
NetBus	12345	TCP
NetBus	12346	TCP
Back Orifice	31337	TCP and UDP

The “Deny Services” table contains a list of common services that reside either on the corporate protected network or on the router itself. Use ACLs to deny these services to untrusted clients.

Deny Services


Service	Port	Transport
finger	79	TCP
SNMP	161	TCP and UDP
SNMP trap	162	TCP and UDP
rlogin	513	TCP
who	513	UDP
remote shell protocol (rsh), remote copy protocol (rcp), rdist, rdump	514	TCP
syslog	514	UDP
new-who	550	TCP and UDP

Here are two ways to control access to router services:

- **Disable the service itself:** Once a router service is disabled, no one can use that service. Disabling a service is safer, and more reliable, than attempting to block all access to the service using an ACL.
- **Restrict access to the service using ACLs:** If your situation requires limited access to a service, build and test appropriate ACLs that you can apply to the service.

Filtering Router Service Traffic


This topic explains how to use ACLs to filter IP traffic destined for Telnet, SNMP, and Routing Information Protocol (RIP).



This figure shows the network topology referenced in the remainder of this lesson.

For the sake of clarity, the next lesson topics depict ACLs as individual ACLs. Generally, you would not build a succession on small ACLs, as this lesson does. Most likely, you would build at least one ACL for the outside router interface, one for the inside router interface, and one or more ACLs for general router use. Do not attempt to combine the small examples shown here into these larger lists, because the statements tend to contradict one another. A sample router configuration at the end of this lesson details how to combine these functions into one logical ACL.

vty Filtering




```
R2(config)# access-list 90 permit host 16.2.1.3 log
R2(config)# access-list 90 deny any log
R2(config)# line vty 0 4
R2(config-line)# login authentication vty-sysadmin
R2(config-line)# transport input ssh
R2(config-line)# access-class 90 in
R2(config-line)# end
```

Systems administrators use Secure Shell (SSH) to remotely access the router console for configuration and maintenance. You should restrict which hosts have access to the vty lines of the router by using an ACL statement, as shown in the figure.

In this example, IP standard ACL 90 allows only host 16.2.1.3 to access router R2 using SSH (port 22). The ACL denies SSH access to R2 by other hosts. This ACL also logs all successful and unsuccessful attempts to access R2 using SSH.

SNMP Service Filtering



```
R2(config)# access-list 80 permit host 16.2.1.3
R2(config)# snmp-server community snmp-host1 ro 80
```

© 2005 Cisco Systems, Inc. All rights reserved.


SND v2.0—4-13

Because of the inherent lack of authentication in SNMP version 1 (SNMPv1), you should only use this version of SNMP on protected, internal networks. You should limit access to a router SNMP agent using an ACL statement, as shown in the figure.

In the example, only the SNMP server with an IP address of 16.2.1.3 may access the router R2 SNMP agent. The **snmp-server** command specifies that the SNMP server must use a community string of **snmp-host1**.

Note The latest Cisco IOS software versions support SNMP version 3 (SNMPv3), which offers more secure SNMP operations. You should implement SNMPv3 rather than older SNMP versions whenever possible.

RIPv2 Route Filtering



```
R1(config)# access-list 12 deny 16.2.2.0 0.0.0.255
R1(config)# access-list 12 permit any
R1(config)# router rip
R1(config-router)# distribute-list 12 out
R1(config-router)# version 2
R1(config-router)# no auto-summary
R1(config-router)# end
```


Cisco routers share routing table update information to provide directions on where to route traffic. Use ACLs to limit which routes a router accepts (takes in) or advertises (sends out) to its counterparts.

The example in the figure shows a standard IP ACL applied to RIP. In this example, access-list 12 is used to prevent R1 from advertising any routes of the 16.2.2.0 demilitarized zone (DMZ) network out of interface e0/0.

Filtering Network Traffic to Mitigate Threats

This topic explains how to implement ACLs to mitigate a range of threats.

IP Address Spoof Mitigation—Inbound



```
R2(config)# access-list 150 deny ip 16.2.1.0 0.0.0.255 any log
R2(config)# access-list 150 deny ip 127.0.0.0 0.255.255.255 any log
R2(config)# access-list 150 deny ip 0.0.0.0 0.255.255.255 any log
R2(config)# access-list 150 deny ip 10.0.0.0 0.255.255.255 any log
R2(config)# access-list 150 deny ip 172.16.0.0 0.15.255.255 any log
R2(config)# access-list 150 deny ip 192.168.0.0 0.0.255.255 any log
R2(config)# access-list 150 deny ip 224.0.0.0 15.255.255.255 any log
R2(config)# access-list 150 deny ip host 255.255.255.255 any log
R2(config)# access-list 150 permit ip any 16.2.1.0 0.0.0.255
R2(config)# interface e0/0
R2(config-if)# ip access-group 150 in
R2(config-if)# exit
```

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—4-15

You can use ACLs to mitigate many threats, including those listed here:

- IP address spoofing—inbound
- IP address spoofing—outbound
- Denial of service (DoS) TCP SYN attacks—blocking external attacks
- DoS TCP SYN attacks—using TCP intercept
- DoS smurf attacks
- Filtering Internet Control Message Protocol (ICMP) messages—inbound
- Filtering ICMP messages—outbound
- Filtering traceroute

As a rule, do not allow any IP packets containing the source address of any internal hosts or networks inbound to a private network. The figure shows ACL 150 for router R2. In this example, the ACL denies all packets containing these IP addresses in their source field:

- Any addresses from the internal 16.2.1.0 network
- Any local host addresses (127.0.0.0/8)
- Any reserved private addresses (RFC 1918, *Address Allocation for Private Internets*)
- Any addresses in the IP multicast address range (224.0.0.0/4)

Apply this ACL inbound to the external interface (e0/0) of router R2.

IP Address Spoof Mitigation—Outbound

Corporate LAN 16.1.0/16



```
R2(config)# access-list 105 permit ip 16.2.1.0 0.0.0.255 any
R2(config)# access-list 105 deny ip any any log
R2(config)# interface e0/1
R2(config-if)# ip access-group 105 in
R2(config-if)# end
```

Be a good citizen and prevent your network from being spoofed.

© 2005 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-18

As a rule, you should not allow any outbound IP packets with a source address other than a valid IP address of the internal network.


The example in the figure shows ACL 105 for router R2. This ACL permits only those packets that contain source addresses from the 16.2.1.0/24 network and denies all others.

This ACL is applied inbound to the inside interface (e0/1) of router R2.

Note	Cisco routers running Cisco IOS Release 12.0 and later may use IP Unicast Reverse Path Forwarding (RPF) verification as an alternative IP address spoof mitigation mechanism.

DoS TCP SYN Attack Mitigation—Blocking External Access

Corporate LAN 16.1.0/16



```
R2(config)# access-list 109 permit tcp any 16.2.1.0 0.0.0.255
              established
R2(config)# access-list 109 deny ip any any log
R2(config)# interface e0/0
R2(config-if)# ip access-group 109 in
R2(config-if)# end
```

© 2005 Cisco Systems, Inc. All rights reserved.


SND v2.0—4-17

TCP SYN attacks involve sending large numbers of TCP SYN packets from a spoofed source into the internal network, which results in the flooding of the TCP connection queues of the receiving nodes.

The ACL in the figure prevents inbound packets, with the SYN flag set, from entering the router. However, the ACL does allow TCP responses from the outside network for TCP connections that originated on the inside network. The **established** command option is used for TCP protocol only. It indicates return traffic from an established connection. For example, a match occurs if the TCP datagram has the ACK control bits set.

DoS Smurf Attack Mitigation

Corporate LAN 16.1.0/16



```
R2(config)# access-list 111 deny ip any host 16.2.1.255 log
R2(config)# access-list 111 deny ip any host 16.2.1.0 log
R2(config)# access-list 111 permit ip any any
R2(config)# interface e0/0
R2(config-if)# ip access-group 111 in
R2(config-if)# end
```

© 2005 Cisco Systems, Inc. All rights reserved.


SND v2.0-4.18

Smurf attacks consist of large numbers of ICMP packets sent to a router subnet broadcast address using a spoofed source IP address from that same subnet. If you configure some routers to forward these broadcasts to other routers in the protected network, you can cause performance degradation. The ACL shown in the figure prevents this forwarding process and halts the smurf attack. It blocks all IP packets originating from any host destined for the subnet broadcast addresses specified (16.2.1.255 and 16.2.1.0).

Note Cisco IOS Releases 12.0 and later now have the “no ip directed-broadcast” command enabled by default, which prevents this type of ICMP attack. Therefore, you may not need to build an ACL as shown here.

Filtering ICMP Messages—Inbound

Corporate LAN 16.1.0/16



```
R2(config)# access-list 112 deny icmp any any echo log
R2(config)# access-list 112 deny icmp any any redirect log
R2(config)# access-list 112 deny icmp any any mask-request log
R2(config)# access-list 112 permit icmp any 16.2.1.0 0.0.0.255
R2(config)# interface e0/0
R2(config-if)# ip access-group 112 in
R2(config-if)# end
```

© 2005 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-19


There are several types of ICMP message types that hackers use to attack a network. Unfortunately, there are many legitimate ICMP messages as well. Various management applications use ICMP messages. Network management uses ICMP messages automatically generated by the router.

Hackers use ICMP echo packets to discover subnets and hosts on the protected network and to generate DoS floods. Hackers use ICMP redirect messages to alter host routing tables. Both ICMP echo and redirect messages should be blocked inbound by the router.

The ACL statement shown in the figure blocks all ICMP echo and redirect messages. As an added safety measure, this ACL also blocks ICMP mask request messages. This ACL allows all other ICMP messages inbound to the 16.2.1.0/24 network.

Filtering ICMP Messages—Outbound

Corporate LAN 16.1.0/16



```
R2(config)# access-list 114 permit icmp 16.2.1.0 0.0.0.255 any echo
R2(config)# access-list 114 permit icmp 16.2.1.0 0.0.0.255 any
parameter-problem
R2(config)# access-list 114 permit icmp 16.2.1.0 0.0.0.255 any
packet-too-big
R2(config)# access-list 114 permit icmp 16.2.1.0 0.0.0.255 any
source-quench
R2(config)# access-list 114 deny icmp any any log
R2(config)# interface e0/1
R2(config-if)# ip access-group 114 in
R2(config-if)# end
```

© 2005 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-20

These ICMP messages are required for proper network operation and you should allow them outbound:


- **Echo:** Allows users to ping external hosts
- **Parameter problem:** Informs the host of packet header problems
- **Packet too big:** Required for packet maximum transmission unit (MTU) discovery
- **Source quench:** Throttles down traffic when necessary

As a rule, you should block all other ICMP message types outbound.

The ACL shown in the figure permits all of the required ICMP messages outbound to the e0/1 interface while denying all others.

Filtering UDP Traceroute Messages

Corporate LAN 16.1.0/16



```
R2(config)# access-list 120 deny udp any any range 33400 34400 log
R2(config)# access-list 120 permit ip any any
R2(config)# interface e0/0
R2(config-if)# ip access-group 120 in
R2(config-if)# end
R2(config)# access-list 121 permit udp 16.2.1.0 0.0.0.255 any range
33400 34400 log
R2(config)# interface e0/1
R2(config-if)# ip access-group 121 in
R2(config-if)# end
R2(config)# interface e0/0
R2(config-if)# ip access-group 121 out
R2(config-if)# end
```

© 2005 Cisco Systems, Inc. All rights reserved.


SND v2.0—4-21

The traceroute feature uses some of the ICMP message types to complete several tasks. Traceroute displays the IP addresses of the routers that a packet encounters along its path (at hops) from source to destination. In the ACL displayed, the router blocks traceroute packets coming into the less secure interface e0/0, but permits traceroute packets from the more secure to the less secure interface. Attackers can use ICMP responses to the UDP traceroute packets to discover subnets and hosts on the protected network.

As a rule, you should block all inbound UDP traceroute messages, as shown in the figure (UDP ports 33400 to 34400).

Mitigating DDoS Attacks with ACLs

This topic explains how to configure router ACLs to help reduce the effects of distributed denial of service (DDoS) attacks, including Trin00, Stacheldraht, Trinity v3, and SubSeven by blocking selected ports.



The figure shows how a DDoS attack occurs, as described here:

- Behind a client is a person who launches the attack.
- A handler is a compromised host running the attacker program. Each handler is capable of controlling multiple agents.
- An agent is a compromised host that is running the attacker program. Each agent is responsible for generating a stream of packets it directs toward the intended victim.

Routers can help reduce the number of DDoS attacks by using ACLs to filter known attack ports. Apply these ACL rules to inbound and outbound traffic between the protected network and the Internet.


A DDoS attack compromises several hundred to several thousand hosts. The hosts are usually Linux and Sun Microsystems computers. However, hackers can also port the attack tools to other platforms. The process of compromising a host and installing the tool is automated. A DDoS attack proceeds as follows:

- Step 1** The attacker initiates a scan phase in which a large number of hosts (approximately 100,000 or more) are probed for a known vulnerability.
- Step 2** The attacker compromises the vulnerable hosts to gain access.
- Step 3** The attacker installs the tool on each host.
- Step 4** The attacker uses the compromised hosts for further scanning and compromises.

Because an automated process is used, attackers can compromise and install the tool on a single host in under 5 seconds, and several thousand hosts in under an hour.

DDoS Attack Mitigation—Trin00

Corporate LAN 16.1.0/16



```
R2(config)# access-list 190 deny tcp any any eq 1524 log
R2(config)# access-list 190 deny tcp any any eq 27665 log
R2(config)# access-list 190 deny udp any any eq 27444 log
R2(config)# access-list 190 deny udp any any eq 31335 log
R2(config)# access-list 190 permit ip any any
R2(config)# interface e0/0
R2(config-if)# ip access-group 190 in
R2(config-if)# end
R2(config)# interface e0/1
R2(config-if)# ip access-group 190 in
R2(config-if)# end
```

© 2005 Cisco Systems, Inc. All rights reserved.

SND v2.0—4.23

Trin00 is a distributed SYN DoS attack. The attack method is a UDP flood. The Trin00 attack sets up communications between clients, handlers, and agents using these ports:

- TCP 1524
- TCP 27665
- UDP 27444
- UDP 31335

The mitigation tactic for the Trin00 attack, and for the other DDoS attacks considered in this topic, is to block both interfaces in the “in” direction. The goal is to prevent infected outside systems from sending messages to the reference network and to prevent any infected internal systems from sending messages out of the reference network to the vulnerable ports.


For example, in the figure, the command **access-list 190 deny tcp any any eq 27665 log** translates to “ACL number 190 will deny any TCP traffic going from any network to any network that has the port equivalent to 27665, and the denial will be logged.”

You need to specify these ports if you need specific information about the exact incoming and outgoing network. For example, if the IP address of the inside network is 10.0.1.0 and you want to block all traffic going from this inside network to the Internet, the command would be **access-list 190 deny tcp 10.0.1.0 0.0.0.255 any eq 27665 log**.

However, you must consider that blocking these ports may have an impact on regular network users, because some high port numbers that are blocked are used by legitimate network clients. You may wish to wait to block these port numbers until a particular threat presents itself.

DDoS Attack Mitigation—Stacheldraht

Corporate LAN 16.1.0/16



```
R2(config)# access-list 190 deny tcp any any eq 16660 log
R2(config)# access-list 190 deny tcp any any eq 65000 log
R2(config)# access-list 190 permit ip any any
R2(config)# interface e0/0
R2(config-if)# ip access-group 190 in
R2(config-if)# end
R2(config)# interface e0/1
R2(config-if)# ip access-group 190 in
R2(config-if)# end
```

© 2005 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-24

Stacheldraht is a DDoS tool that appeared in 1999 and combines features of Trin00 and TFN2K. Stacheldraht also contains some advanced features, such as encrypted attacker-master communication and automated agent updates. The possible attacks are similar to those of TFN, namely, ICMP flood, SYN flood, UDP flood, and smurf attacks.

A Stacheldraht attack sets up communication between clients, handlers, and agents using these ports:

- TCP 16660
- TCP 65000
- ICMP echo
- ICMP echo-reply


Note The ports listed above are the default ports for this tool. Use these ports for orientation and example only, because the hacker can easily choose other port numbers.

This figure shows an example that mitigates a Stacheldraht DDoS attack by blocking traffic on these ports:

- TCP 16660
- TCP 65000

DDoS Attack Mitigation—Trinity v3

Corporate LAN 16.1.0/16



```
R2(config)# access-list 190 deny tcp any any eq 33270 log
R2(config)# access-list 190 deny tcp any any eq 6667 log
R2(config)# access-list 190 permit ip any any
R2(config)# interface e0/0
R2(config-if)# ip access-group 190 in
R2(config-if)# end
R2(config)# interface e0/1
R2(config-if)# ip access-group 190 in
R2(config-if)# end
```

© 2005 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-25


Trinity is capable of launching several types of flooding attacks on a victim site, including UDP, fragment, SYN, RST, ACK, and other floods. Trinity communicates from the handler or intruder to the agent using IRC or ICQ from AOL. Trinity appears to primarily use port 6667 and has a backdoor program that listens on TCP port 33270.

This figure shows an example that mitigates a Trinity v3 DDoS attack by blocking traffic on these ports:

- TCP 33270
- TCP 6667

DDoS Attack Mitigation—SubSeven

Corporate LAN 16.1.0/16



```
R2(config)# access-list 190 deny tcp any any range 6711 6712 log
R2(config)# access-list 190 deny tcp any any eq 6776 log
R2(config)# access-list 190 deny tcp any any eq 6669 log
R2(config)# access-list 190 deny tcp any any eq 2222 log
R2(config)# access-list 190 deny tcp any any eq 7000 log
R2(config)# access-list 190 permit ip any any
R2(config)# interface e0/0
R2(config-if)# ip access-group 190 in
R2(config-if)# end
R2(config)# interface e0/1
R2(config-if)# ip access-group 190 in
R2(config-if)# end
```

© 2005 Cisco Systems, Inc. All rights reserved.


SND v2.0—4-26

Depending on the version of attack, an attacker will try to exploit ports 2222, 6669, 6711, 6712, 6776, and 7000. The permit ACL entry that allows the desired traffic is not shown in this example. The figure shows an example that mitigates a SubSeven DDoS attack by blocking traffic these ports:

- TCP 2222
- TCP 6669
- TCP range 6711 to 6712
- TCP 6776
- TCP 7000

Combining Access Functions

This topic describes how to combine many ACL functions into two or three larger ACLs.



This is an example of a possible configuration for router R2 in the reference network. This partial configuration file contains several ACLs that contain most of the ACL features already explained in this lesson. View this partial configuration as an example of how to integrate multiple ACL policies into a few main router ACLs.

The partial configuration file that follows shows how to combine many ACL functions into two or three larger ACLs.

```
!
hostname R2
!
interface Ethernet0/0
 ip address 16.1.1.2 255.255.0.0
 ip access-group 126 in
!
interface Ethernet0/1
 ip address 16.2.1.1 255.255.255.0
 ip access-group 128 in
!
router ospf 44
 network 16.1.0.0 0.0.255.255 area 0
 network 16.2.1.0 0.0.0.255 area 1
```

```

!
! Access list 80 applies to SNMP hosts allowed to access this router
no access-list 80
access-list 80 permit host 16.2.1.2
access-list 80 permit host 16.2.1.3
!
!snmp-server community snmp-host1 ro 80

! Access list 126 applies to traffic flowing from external networks to
! the internal network or to the router itself
no access-list 126

! comment - this entry below prevents any IP packets containing the
! source address of any internal hosts or networks, inbound to the
! private network.
access-list 126 deny ip 16.2.1.0 0.0.0.255 any log

! comment - this set of entries below prevents any IP packets
! containing the invalid source address such as the local loopback
access-list 126 deny ip 127.0.0.0 0.255.255.255 any log
access-list 126 deny ip 0.0.0.0 0.255.255.255 any log
access-list 126 deny ip 10.0.0.0 0.255.255.255 any log
access-list 126 deny ip 172.16.0.0 0.15.255.255 any log
access-list 126 deny ip 192.168.0.0 0.0.255.255 any log
access-list 126 deny ip 224.0.0.0 15.255.255.255 any log
access-list 126 deny ip any host 16.2.1.255 log
access-list 126 deny ip any host 16.2.1.0 log
access-list 126 permit tcp any 16.2.1.0 0.0.0.255 established
access-list 126 deny icmp any any echo log
access-list 126 deny icmp any any redirect log
access-list 126 deny icmp any any mask-request log
access-list 126 permit icmp any 16.2.1.0 0.0.0.255
access-list 126 permit ospf 16.1.0.0 0.0.255.255 host 16.1.1.2
access-list 126 deny tcp any any range 6000 6063 log
access-list 126 deny tcp any any eq 6667 log
access-list 126 deny tcp any any range 12345 12346 log
access-list 126 deny tcp any any eq 31337 log
access-list 126 permit tcp any eq 20 16.2.1.0 0.0.0.255 gt 1023
access-list 126 deny udp any any eq 2049 log
access-list 126 deny udp any any eq 31337 log
access-list 126 deny udp any any range 33400 34400 log
access-list 126 permit udp any eq 53 16.2.1.0 0.0.0.255 gt 1023
access-list 126 deny tcp any range 0 65535 any range 0 65535 log
access-list 126 deny udp any range 0 65535 any range 0 65535 log

```

```

access-list 126 deny ip any any log
!
! Access list 128 applies to traffic flowing from the internal network
! to external networks or to the router itself
no access-list 128

access-list 128 permit icmp 16.2.1.0 0.0.0.255 any echo
access-list 128 permit icmp 16.2.1.0 0.0.0.255 any parameter-problem
access-list 128 permit icmp 16.2.1.0 0.0.0.255 any packet-too-big
access-list 128 permit icmp 16.2.1.0 0.0.0.255 any source-quench
access-list 128 deny tcp any any range 1 19 log
access-list 128 deny tcp any any eq 43 log
access-list 128 deny tcp any any eq 93 log
access-list 128 deny tcp any any range 135 139 log
access-list 128 deny tcp any any eq 445 log
access-list 128 deny tcp any any range 512 518 log
access-list 128 deny tcp any any eq 540 log
access-list 128 permit tcp 16.2.1.0 0.0.0.255 gt 1023 any lt 1024
access-list 128 permit udp 16.2.1.0 0.0.0.255 gt 1023 any eq 53
access-list 128 permit udp 16.2.1.0 0.0.0.255 any range 33400 34400
log
access-list 128 deny tcp any range 0 65535 any range 0 65535 log
access-list 128 deny udp any range 0 65535 any range 0 65535 log
access-list 128 deny ip any any log
!
snmp-server community snmp-host1 ro 80
!

```

Caveats

This topic explains some of the caveats that you need to consider when creating ACLs.

ACL Caveats	
Statement	Caveat
Implicit deny all	You may not see this statement, but it does exist.
Standard access list limitation	You may need to create extended ACLs to implement security policies.
Statement evaluation order	ACL statements are evaluated from the top down, so always consider the order of the statements.
Order of ACL statements	Place more specific ACL statements higher in the access list. Ensure that statements at the top of the ACL do not negate any statements found lower in the list.
Directional filtering	Always double-check the direction (inbound or outbound) of the data that your ACL is filtering.

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—4-28

Here are the caveats that you should consider when working with ACLs:

- **Implicit deny all:** All Cisco ACLs end with an implicit deny all statement. Although you may not actually see this statement in your ACLs, they do exist.
- **Standard ACL limitation:** Because standard ACLs are limited to packet filtering on source addresses only, you may need to create extended ACLs to implement your security policies.
- **Statement evaluation order:** ACL statements are evaluated in a sequential (top-down) order, starting with the first entry in the list. This process means that it is very important to consider the order in which you place statements in your ACLs.
- **Order of specific statements:** Certain ACL statements are more specific than others and, therefore, you need to place them higher in the ACL. For example, blocking all UDP traffic at the top of the list negates the blocking of SNMP packets lower in the list. Take care that statements at the top of the ACL do not negate any statements found lower in the list.
- **Directional filtering:** Cisco ACLs have a directional filter that determines whether they examine inbound packets (toward the interface) or outbound packets (away from the interface). Always double-check the direction of data that your ACL is filtering.

ACL Caveats (Cont.)

Statement	Caveat
Modifying numbered ACLs	Adding new statements may require that a new ACL be created (Cisco IOS Release 12.2 and earlier).
Special packets	If filtering router-generated packets is part of the security policy, they must be acted upon by inbound ACLs on adjacent routers or through other router filter mechanisms using ACLs.
Extended ACL placement	Always consider placing extended ACLs on routers as close as possible to the source being filtered.
Standard ACL placement	Always place standard ACLs as close to the destination as possible.

- **Modifying ACLs:** Always append new statements added to an existing ACL to the bottom of the ACL. Because of the inherent top-down statement evaluation order of ACLs, these new entries may render the ACL unusable. When a new statement does render the ACL unusable, you must create a new ACL with the correct statement ordering. Delete the old ACL and assign the new ACL to the router interface.
- **Special packets:** Router-generated packets, such as routing table updates, are not subject to outbound ACL statements on the source router. If your security policy requires filtering these types of packets, inbound ACLs on adjacent routers or other router filter mechanisms using ACLs must do the filtering task.
- **Extended ACL placement:** If you use extended ACLs on routers too far from the source that you need to filter, packets flowing to other routers and interfaces may be adversely affected. Always consider placing extended ACLs on routers as close as possible to the source that you are filtering.
- **Standard ACL placement:** Because standard ACLs filter packets based on the source address, placing these ACLs too close to the source can adversely affect packets destined to other destinations. Always place standard ACLs as close to the destination as possible.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- **ACLs control traffic in and out of routers and firewalls.**
- **Cisco routers and firewalls use standard, extended, enhanced, named, and numbered ACLs that you can refer to by name or number. You must follow the rules and guidelines for building ACLs.**
- **You must apply the ACLs to interfaces based on the direction of the data flow.**
- **You can use ACLs to block services that hackers use to gather information about your network .**
- **You can use ACLs to filter IP traffic destined for specific services such as Telnet, SNMP, and RIP.**
- **ACLs can be used to mitigate IP spoofing and DoS attacks coming from TCP SYN floods, ICMP floods and smurf attacks, and UDP traceroute reconnaissance.**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-30

Summary (Cont.)

- **ACLs can be used to mitigate DDoS attacks including Trin00, Stacheldraht, Trinity v3, and SubSeven.**
- **Many ACL functions can be combined into two or three larger ACLs.**
- **There are many caveats to be considered when creating ACLs.**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-31

Configuring a Cisco IOS Firewall with the Cisco SDM Firewall Wizard

Overview

A firewall is a set of rules used to protect the resources of your network. These rules filter the packets arriving at the router. If a packet does not meet the criteria specified in the rule, the router drops it. If it does meet the criteria, the router allows the packet to pass through the interface on which the rule is applied. The Cisco Router and Security Device Manager (SDM) Firewall Wizard helps you create a firewall for your LAN by answering prompts in a set of screens.


Objectives

Upon completing this lesson, you will be able to configure a firewall on your network using the Cisco SDM Firewall Wizard. This ability includes being able to meet these objectives:

- Describe the tasks that you can complete with the Cisco SDM Firewall Wizard
- Explain how to configure a basic firewall using the Cisco SDM Firewall Wizard
- Explain how to configure a firewall with a DMZ
- Explain how to configure firewall inspection rules
- Explain how to configure application security policy
- Explain how to deliver the configuration to the router
- Explain how to use the firewall and ACL policy editor to customize default ACL settings

Cisco SDM Firewall Wizard Tasks

The Cisco SDM Firewall Wizard allows you to configure a firewall to meet your needs. This topic describes the tasks you can complete with the Cisco SDM Firewall Wizard.




Choose the Cisco IOS Firewall that meets your network security needs—basic or advanced.

The figure shows the top part of the Cisco IOS Firewall configuration screen. You can choose either a basic firewall or an advanced firewall. Once you choose the firewall that you want, the Use Case Scenario figure appears for each selection.

The two firewalls differ in these ways:

- **Basic Firewall:** A basic firewall is a firewall that uses default rules. The Use Case Scenario figure shows a typical network configuration in which a basic firewall is used.
- **Advanced Firewall:** An advanced firewall is a firewall that includes the option of creating a demilitarized zone (DMZ) network and specifying inspection rules. The insert on the figure shows the Use Case Scenario figure for an advanced firewall.

SDM Firewall Wizard Help Screens



The lower part of the Create Firewall tab provides the “How do I ...” task list. The list includes these task choices:

- How Do I View Activity on My Firewall?
- How Do I Configure a Firewall on an Unsupported Interface?
- How Do I Configure a Firewall After I Have Configured a VPN?
- How Do I Permit Specific Traffic Through a DMZ Interface?
- How Do I Modify an Existing Firewall to Permit Traffic from a New Network or Host?
- How Do I Configure NAT on an Unsupported Interface?
- How Do I Configure NAT Passthrough for a Firewall?
- How Do I Permit Traffic Through a Firewall to My Easy VPN Concentrator?
- How Do I Associate a Rule with an Interface?
- How Do I Disassociate an Access Rule from an Interface?
- How Do I Delete a Rule That Is Associated with an Interface?
- How Do I Create an Access Rule for a Java List?
- How Do I View the IOS Commands I Am Sending to the Router?
- How Do I Permit Specific Traffic onto My Network if I Don't Have a DMZ Network?


Step-by-Step Help Screens



When you choose one of the “How do I . . .” tasks from the drop-down menu, a step-by-step help screen for that question appears, as shown in the figure. You can also view and print a PDF version of the help screen.

Configuring a Basic Firewall

This topic explains how to configure a basic firewall using the Cisco SDM Firewall Wizard.




Cisco SDM will protect the LAN with a default firewall when you choose the basic firewall option. The Cisco SDM Firewall Wizard will guide you through the steps of defining inside and outside interfaces.

- **Outside (untrusted) interface:** Choose the router interface connected to the Internet or to your WAN.
- **Allow secure SDM access from the outside interfaces:** Creating a firewall policy can block SDM access to the router from the outside interface. Choosing Allow Secure SDM Access from the Outside Interfaces permits secure HTTP access to the outside (untrusted) interface. Because it is a secure SDM connection to the firewall, you will not be able to browse the outside (untrusted) interface via HTTP after the firewall wizard completes the configuration.
 - You can specify the router outside interface or interfaces that you want to use for remote management access and the hosts from which administrators can log via HTTPS on to SDM to manage the router by performing these tasks:
 - **Choose the outside interface:** Choose the interfaces through which users are to launch the SDM.

- **Choose Host Address, Network Address, or Any:** If you want to allow only a single host to have SDM access to the firewall, choose **Host Address** and enter the IP address of the host you want. Choose **Network Address** and enter the address of a network and a subnet mask to allow hosts on that network SDM access to the firewall. The host or network must be accessible from the interfaces that you specify. Choose **Any** to exempt any host connected to the specified interfaces.
- **Choose inside (trusted) interfaces:** Choose the physical and logical interfaces connecting to the LAN. You can choose multiple interfaces.

Configuring an Advanced Firewall

This topic explains how to configure a firewall with a DMZ.



The advanced firewall configuration allows you to secure your private network by applying access and inspection rules to inside (trusted), outside (untrusted), and DMZ interfaces. A DMZ network is a buffer zone used to isolate traffic that comes from an untrusted network. If you have a DMZ network, choose the interface that connects to it when using an advanced firewall.

The Cisco SDM Firewall Wizard will guide you through the steps to complete these tasks:

- **Define inside and outside interfaces:** Check outside and inside to identify each interface as an outside or inside interface.
 - Outside interfaces connect to your WAN or to the Internet.
 - Inside interfaces connect to your LAN.
- **Configure DMZ interface:** A DMZ network is a buffer zone used to isolate traffic that comes from an untrusted network. If you have a DMZ network, choose the interface that connects to it.
- **DMZ Service:** This window allows you to view rule entries that specify which services available inside the DMZ you want to make available through the outside interfaces. You can specify which traffic service types to allow through the outside interfaces into the DMZ network.
 - **DMZ Service Configuration window:** This window shows these DMZ service entries configured on the router:
 - **Start IP Address:** Enter the first IP address in the range that specifies the hosts in the DMZ network.

- **End IP Address:** Enter the last IP address in the range that specifies the hosts in the DMZ network. If there is no value listed in this column, the Cisco SDM Firewall Wizard assumes that the IP address in the Start IP Address column is to be the only host in the DMZ network. The range can specify a maximum of 254 hosts.
- **Service Type:** Specify the type of service, either TCP or UDP.
- **Service:** Specify the name of the service, such as Telnet, FTP, or a protocol number.
 - **Configure a DMZ service entry:** Click **Add** and create the entry in the DMZ Service Configuration window.
 - **Edit a DMZ service entry:** Choose the service entry and click **Edit**. Then, edit the entry in the DMZ Service Configuration window.
- **DMZ Service Configuration:** Use this window to create or edit a DMZ service entry by filling in these fields:
 - **Host IP Address:** Enter the address range that will specify the hosts in the DMZ to which this entry applies. The firewall will allow traffic for the specified TCP or UDP service to reach these hosts.
 - **Start IP Address:** Enter the first IP address in the range; for example, 172.20.1.1. If Network Address Translation (NAT) is enabled, you must enter the NAT-translated address, known as the inside global address.
 - **End IP Address:** Enter the last IP address in the range; for example, 172.20.1.254. If NAT is enabled, you must enter the NAT-translated address.
 - **Service:** Click **TCP** or **UDP** for the service that you want.
 - **Service:** Enter the service name or number in this field. If you do not know the name or number, click the button and choose the service from the list that is displayed.

Configuring Firewall Inspection Rules

This topic explains how to configure firewall inspection rules.

Configuring Firewall Inspection Rules

Inspection rules allow returning traffic that would otherwise be blocked.

If you want to:	Do this:
Examine an existing inspection rule	Choose the rule name from the Inspection Rule Name list. The inspection rule entries appear in a separate dialog box.
Edit an existing inspection rule	- Choose the rule name from the Inspection Rule Name list and click Edit. - Then edit the rule in the Inspection Rule Information window.
Create a new inspection rule	Choose the rule name from the Inspection Rule Name list, click New, and create the rule in the Inspection Rule Information window.

© 2006 Cisco Systems, Inc. All rights reserved.SDM v2.0—4.8

Access rules in the firewall may deny return traffic on sessions started inside the firewall because of the type of service that the traffic uses. Outgoing traffic can leave the router, but if you do not explicitly permit return traffic of the same type, the firewall will block the traffic on its return to the LAN.

Inspection rules provide a means to allow return traffic onto the network. These rules cause the router to examine outgoing packets for specified types of traffic. The firewall compares traffic arriving at the outside interface against the traffic types in the inspection rule. The firewall allows traffic onto the network if the traffic is associated with a session started on the LAN and is of a type specified in the inspection rules. In this way, inspection rules create temporary holes in the firewall so that hosts on the LAN can receive return traffic.

When you view the Inspection Rules screen in SDM, the screen shows you the default inspection rule that SDM provides plus any user-configured inspection rules, and enables you to add or modify user-configured inspection rules.


An inspection rule is a named list of inspection rule entries. Each entry consists of a protocol specification, an alert switch, and an audit switch, defined in the list that follows. Choose the inspection rule whose entries you want to view.

- **Protocol:** This section defines the protocol that the inspection rule entry will inspect. For example, if the protocol is FTP, the rule inspects incoming FTP traffic if it is associated with a session started from inside the firewall.
- **Alert:** Turn this button to **On** if the router is to generate alerts when it encounters traffic of this type; turn it **Off** if no alert is required. If you enabled logging in the Router Properties Logging window, the router saves alerts in a syslog file.

- **Audit Trail:** Turn the audit trail **On** if the router is to generate an audit trail when it encounters traffic of the selected type. Turn the audit trail **Off** if no audit trail is required. The router saves audit trails in a syslog file.

Application Security Policy Configuration

This topic explains how to configure an application security policy.



Default SDM Application Security Policies

Cisco SDM provides preconfigured application security policies that you can use to protect the network. Use the slider bar to choose the security level that you want and to view a description of the security that it provides. The wizard summary screen displays the policy name, SDM_HIGH, SDM_MEDIUM, or SDM_LOW, and the configuration statements for that policy. The Preview Commands button allows you to view the Cisco IOS commands that make up the policy that you choose. You can also view the details of the policy by clicking first the Firewall and ACL icon then the Application Security tab and choosing the name of the policy.

Custom Application Security Policy Button


The Custom Application Security Policy button and the Policy Name field are visible if you are completing the Advanced Firewall wizard. Choose this option if you want to create your own application security policy. If the policy already exists, enter the name in the field or click the button on the right, choose **Select an Existing Policy**, and choose the policy that you want from the list. To create a policy, click the **Custom Appliance Security Policy** button, choose **Create a New Policy**, and create the policy in the dialog box that appears.

You must configure the router with the IP address of at least one Domain Name System (DNS) server for application security to work. Click **Enable DNS-Based Hostname-to-Address** translation and provide the IP address of the primary DNS server. If a secondary DNS server is available, enter the IP address in the **Secondary DNS Server** field.

The IP addresses that you enter will be visible in the DNS Properties window under Additional Tasks.

Delivering the Configuration to the Router

This topic explains how to deliver the configuration to the router.



- **Review and verify the configuration.**
- **Use the Back button to edit entries.**
- **Click Finish to complete the configuration.**


The Cisco SDM Firewall Wizard Summary screen summarizes the firewall information. The summary appears in three sections:

- **Inside (Trusted) Interface(s):** The Cisco SDM Firewall Wizard summary lists the logical and physical interfaces of the router that you designated as the inside interfaces in this wizard session, along with their IP addresses. Below the list, SDM describes which access and inspection rules were associated with these interfaces. Here are examples of inspection rules:
 - Apply access rule to the inbound direction to deny spoofing traffic.
 - Apply access rule to the inbound direction to deny traffic sourced from broadcast, local loopback address.
 - Apply access rule to the inbound direction to permit all other traffic.
 - Apply default inspection rule to the inbound direction of inside (trusted) interface. (This example applies to an advanced firewall.)
- **Outside (Untrusted) Interface(s):** The Cisco SDM Firewall Wizard summary lists the logical and physical interfaces of the router that you designated as outside interfaces in this wizard session, along with their IP addresses. Below the list, SDM describes which access and inspection rules were associated with these interfaces. Here are examples of inspection rules:
 - Apply default inspection rule to the outbound direction. (This example applies to a basic firewall.)
 - Turn on Unicast Reverse Path Forwarding (RPF) check.
 - Apply access rule to the inbound direction to permit IPsec tunnel traffic if necessary.

- Apply access rule to the inbound direction to deny spoofing traffic.
 - Apply access rule to the inbound direction to deny traffic sourced from broadcast, local loopback, and private address.
 - Apply access rule to the inbound direction to deny all other traffic.
- **DMZ Interface:** If you configured an advanced firewall, this area shows you the DMZ interface that you designated, along with its IP address. The Cisco SDM Firewall Wizard summary describes which access and inspection rules are associated with this interface. Here are examples of inspection rules:
- Apply Context-Based Access Control (CBAC) inspection rule to the outbound direction.
 - Apply access rule to the inbound direction to deny all traffic.

You can review the information in the summary screen and use the Back button to return to screens in the wizard to make changes.

Delivering the Commands to the Router




Commands are delivered to the running configuration and are not saved on exit.

The final step is to save the configuration to the router running configuration and leave the wizard, as described here:

- Click **Deliver**. SDM saves the configuration changes to the router running configuration. The changes will take effect immediately but will be lost if the router is turned off.
- Check the **Save Running Config. to Router's Startup Config** check box. This keeps the configuration from being lost when the router is turned off.
- If you checked the Preview Commands before Delivering to Router **check box** in the User Preferences window, the Deliver Configuration to Router window appears. In this window, you can view the command-line interface (CLI) commands that you are delivering to the router.

Editing Firewall Policies and ACLs

Even experienced administrators find configuring firewall policies to be a grueling and tedious task. The key advantage of Cisco SDM is its GUI for setting up firewall policies and associated access control lists (ACLs). You have already learned how to edit firewall policies. This topic explains how to use the ACL editor to customize default ACL settings.




The firewall and ACL policy editor is a powerful tool. The Edit Firewall Policy/ACL screen gives a high-level view of each policy based on direction of traffic flow.

In the screen shot on the left in the figure, the current configuration of interfaces (fe0/0 and fe0/1) protects the network from incoming traffic. The window at the bottom of the screen displays the ACL rules applied to that traffic flow. You can select and edit these rules as required.

By toggling the view, you can look at the policy applied to returning traffic as shown in the screen shot on the right.


Editing the Application Security Policy



To edit an application policy, click the **Firewall and ACL** icon and choose the **Application Protocol** tab from the lower window. This opens a full description of the policy as shown in the figure. In this example, the default policy is SDM_HIGH. You can change that to another default policy using the Policy Name drop-down menu.

The left side of the screen lists the applications that the policy applies to—e-mail, instant messaging, point-to-point, HTTP (headers and content), and other applications and protocols. The screen shot on the right side of the figure shows the default settings for instant messaging. You can change these settings using the drop-down menus. In some instances, the Cisco IOS image in your router may not support the requested feature.


Editing the Application Security Policy (Cont.)



Use the **Applications/Protocols** menu for applications and protocols that are not shown.

Choose **Applications/Protocols** for settings not found in the other windows.

Editing the Application Security Policy (Cont.)



Example: You can prevent internal defacing of a web page by choosing HTTP > Header Options to block put commands and send an alarm.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-15

As an alternative to accepting the SDM default settings, click the **Action** button (Add Delete Clone) and create your own custom policies. You simply need to clone the policy and save it under a new name.

The example shown in the figure refers to a policy that prevents defacing a web page by internal users. Choose **HTTP > Header Options** to make the changes. Click the **Apply Changes** button to complete this task.

Editing Firewall Global Settings

The screenshot shows a window titled "Global Timeouts and Thresholds" with the following sections:

- Global Timeout Values:**
 - TCP connection timeout value: 60 secs
 - TCP FIN-wait timeout value: 5 secs
 - TCP idle timeout value: 3600 secs
 - UDP idle timeout value: 30 secs
 - DNS timeout value: 5 secs
- SYN Flooding DoS Attack Thresholds:**

The threshold values define the limits for triggering deletion of existing half-open sessions. An unusual number of half-open sessions may indicate a SYN Flooding DoS attack.

 - One Minute Session Threshold:**
 - Low: 400 High: 500
 - Maximum incomplete session threshold:**
 - Low: 400 High: 500
 - TCP maximum incomplete sessions per host:** 50
 - Blocking Time:** 0 mins
- Other:**
 - ☒ Enable alert globally
 - ☐ Enable audit globally

Buttons at the bottom: OK, Cancel, Help.

Global settings should not be changed without care and attention.

This figure shows the Global Timeouts and Thresholds settings screen. You can change these settings as required.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- You can use Cisco SDM Firewall Wizard help screens to guide you through most configuration tasks.
- You need to define only inside and outside interfaces to create a basic firewall.
- An advanced firewall allows you to configure a DMZ.
- You may need to configure the firewall inspection rules to allow certain types of returning traffic.
- Cisco SDM allows you to customize default application security policies.
- Once policies have been selected, changes need to be delivered to the router, and you need to save the configuration.
- The firewall and ACL editor allows you to edit application policies and global settings.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-17

Defending Your Network with the Cisco Security Appliance Product Family

Overview

With the explosion of broadband access, organizations are increasingly demanding security. In addition, small businesses are increasingly tapping the power of the Internet for competitive advantage. This lesson describes the components of the Cisco security appliances product family so that you can protect your network from malicious attacks.


Objectives

Upon completing this lesson, you will be able to explain the best practices for deploying the hardware and software components of the Cisco security appliances family. This ability includes being able to meet these objectives:

- Describe the main components of the Cisco security appliance product family
- Describe the features of the Cisco IOS Firewall
- Explain when to choose a Cisco IOS Firewall over an appliance-based solution
- Describe the security features of the Cisco PIX 500 Series Security Appliances and the FWSM
- Describe the security features of the Cisco ASA 5500 Series Adaptive Security Appliances
- Explain how to develop an effective firewall policy based on firewall best practices

Introducing the Cisco Security Appliance Product Family

This topic describes the main components of the Cisco security appliance product family.



The Cisco security appliances family includes these products:

- **Cisco IOS Firewall:** The Cisco IOS Firewall provides robust, integrated firewall and intrusion detection functionality for every perimeter of the network. The Cisco IOS Firewall is available for a wide range of Cisco IOS software-based routers and offers sophisticated security and policy enforcement for connections within an organization (intranet) and between partner networks (extranets) and for securing Internet connectivity for remote and branch offices.
- **Cisco PIX 500 Series Security Appliances:** From compact “plug-and-play” appliances for small and home offices to modular carrier-class gigabit appliances for enterprise and service provider environments, the Cisco PIX 500 Series Security Appliances provide robust, enterprise-class integrated network security services that create a strong multilayered defense for fast-changing network environments.
- **Cisco Catalyst 6500 Series Firewall Services Module (FWSM):** The FWSM includes cards designed for the chassis of Cisco Catalyst 6500 Series Switches and Cisco 7600 Series Routers. These cards provide firewall services along with a range of network services in one chassis. They use the latest versions of Cisco PIX Security Appliance Software.

- **Cisco ASA 5500 Series Adaptive Security Appliances:** The Cisco ASA 5500 Series is a modular platform that provides the next generation of security and virtual private network (VPN) services. The Cisco ASA 5500 Series Adaptive Security Appliances leverage technology developed for the Cisco PIX 500 Series Security Appliances, the Cisco Intrusion Prevention System (IPS) 4200 Series Sensors and Cisco VPN 3000 Series Concentrators. These technologies converge on the Cisco ASA 5500 Series Adaptive Security Appliances to deliver a platform that stops the broadest range of threats. The Cisco ASA 5500 Series Adaptive Security Appliances deliver application security, anti-X defense, network containment and control, and “clean” VPN connectivity across a network product portfolio.

Cisco IOS Firewall Features

The Cisco IOS Firewall combines the functions of packet inspection and proxy firewalls to provide an optimal security solution on one chassis. This topic describes the features of the Cisco IOS Firewall.

Cisco IOS Firewall Features

- **Stateful inspection firewall**
- **Application and protocol inspection and control**
- **Dynamic, per-user authentication and authorization**
- **Dynamic and static NAT and PAT**
- **Content filtering**
- **Remote management**
- **Administrative access control with AAA**
- **Multiple DMZ support**
- **Extensive multimedia support, including streaming video, streaming audio, and voice applications**
- **DoS protection**
- **Secure dynamic routing**
- **Firewall virtualization**

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—4-4

The Cisco IOS Firewall is a stateful inspection firewall option available for Cisco routers. It provides a single-box security and routing solution for protecting the WAN entry point into the network. Although the hub is a common location at which to position a firewall and inspect traffic for attacks, it is not the only location to consider when deploying security. Branch offices are also an important location in your network to both place a firewall and inspect traffic for attacks.

Firewall integration in Cisco IOS routers augments the inherent capabilities of a router, including multipoint interfaces, industry-standard routing protocols, a broad range of services, and an expanding group of other security features such as VPN and IPS features. The Cisco IOS Firewall interoperates with other Cisco IOS software technologies, including Network Address Translation (NAT), quality of service (QoS), IPsec and Secure Sockets Layer (SSL) VPN, to become a vital component of an end-to-end network security infrastructure.

The “Cisco IOS Firewall Features and Benefits” table explains the extent of the security that the Cisco IOS Firewall provides.

Cisco IOS Firewall Features and Benefits

Feature	Benefit
Stateful inspection firewall	This feature uses access control lists (ACLs) to enforce administrator-defined access control policies while performing deep packet inspection and tracking the state of all network communications.
Application and protocol inspection and control	This feature delivers enhanced application and protocol security by using specialized inspection engines capable of examining data streams at Layer 4 to Layer 7.
Dynamic, per-user authentication and authorization	This feature provides flexible user authentication and authorization via the high performance cut-through proxy mechanism and integration with Cisco Secure Access Control Sever (ACS) using RADIUS and TACACS+ protocols to allow for integration into numerous user databases, including Microsoft Active Directory, Microsoft Windows NT domains, Lightweight Directory Access Protocol (LDAP) directories, and one-time password (OTP) systems.
Dynamic and static NAT and Port Address Translation (PAT)	This feature provides extensive NAT application and protocol support and protects internal network addresses from the outside, providing an additional level of security.
Content filtering	This feature improves employee productivity through integration with leading third-party URL filtering solutions and supports URL filtering and blocks malicious Java applets.
Remote management	This feature offers remote management methods to configure, monitor, and troubleshoot.
Administrative access control with authentication, authorization, and accounting (AAA)	This feature provides granular control for administrative access based on the AAA services provided by the TACACS+ and RADIUS protocols. This allows administrators to enforce access policies to the level of what services and commands that they allow each administrative user or group.
Multiple demilitarized zone (DMZ) support	This feature supports additional physical or virtual network interfaces to provide protected access to multiple servers including web, e-mail, FTP, and Domain Name System (DNS) on a shared network.
Extensive multimedia support, including streaming video, streaming audio, and voice applications	This feature provides rich stateful inspection firewall services for a wide range of VoIP standards and other multimedia standards. Businesses can take advantage of the many benefits that converged data, voice, and video networks provide, such as improved productivity and competitive advantage.
Denial of service attack (DoS) protection	This feature provides several mechanisms to block and mitigate DoS attacks, such as TCP Intercept, TCP SYN cookies, DNS Guard, Flood Defender, Flood Guard, Mail Guard, and Unicast Reverse Path Forwarding (RPF).
Secure dynamic routing	Cisco IOS Firewall supports Message Digest 5 (MD5)-based and plain text routing authentication for Routing Information Protocol (RIP) and Open Shortest Path First (OSPF), preventing route spoofing and various routing-based DoS attacks.
Firewall virtualization	This feature enables partitioning of a single device into multiple virtual firewalls, or security contexts. Organizations can manage each of these virtual firewalls separately and can segregate business units or other functional areas on the same physical infrastructure. Similarly, service providers can leverage firewall virtualization to support and segregate multiple customers on a single physical device.

When to Choose a Cisco IOS Firewall Solution


This topic explains when to choose a Cisco IOS Firewall solution over an appliance-based solution.

When to Use a Cisco IOS Firewall

Choose the Cisco IOS Firewall when you need:

- A one-box solution with powerful security, QoS, multiprotocol routing, integrated WAN interfaces, and voice application support
- To leverage network infrastructure for security
- Extensive VPN support integrated with a firewall in a single device

Environment	Routers and Switches
Small and home office	Cisco 800 Series, 1700 Series, and 1800 Series Routers
Branch and extranet environments	Cisco 2600 and 3600 Series Multiservice Platforms and Cisco 2800 Series Integrated Services Routers, Cisco 3700 Series Multiservice Access Routers, and 3800 Series Integrated Services Routers
VPN and WAN aggregation points; high-throughput environments	Cisco 7200 Series Routers, 7301 Router, and 7400 Series Routers; Cisco Route Switch Modules; Cisco Catalyst 5000 and Catalyst 6000 Series Switches



© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—4.5

The Cisco IOS Firewall offers sophisticated security and policy enforcement for connections within an organization (intranet), between partner networks (extranets), and for securing Internet connectivity for remote and branch offices.

The Cisco IOS Firewall is available for the applications listed here and detailed in the “Cisco IOS Firewall Application” table:


- **Small or home offices:** Cisco 800, 1700, and 1800 Series Routers
- **Branch and extranet environments:** Cisco 2600 and 3600 Series Multiservice Platforms and Cisco 2800 Series Integrated Services Routers, Cisco 3700 Series Multiservice Access Routers, and Cisco 3800 Series Integrated Services Routers
- **VPN and WAN aggregation points or other high-throughput environments:** Cisco 7200 Series Routers, 7301 Router, and 7400 Series Routers; Cisco Route Switch Module (RSM); Cisco Catalyst 5000 and Catalyst 6000 Series Switches

Cisco IOS Firewall Application

Customer Requirement	Cisco IOS Firewall Benefit
One-box solution	The Cisco IOS Firewall provides a comprehensive, integrated security solution, including stateful packet filtering, intrusion detection and protection, per-user authentication and authorization, VPN capability, extensive QoS mechanisms, multiprotocol routing, voice application support, and integrated WAN interface support in one box.
Leverage existing network infrastructure	The Cisco IOS Firewall can be loaded on existing Cisco IOS routers.
Extensive VPN support integrated with firewall in a single device	Deploying the Cisco IOS Firewall with Cisco IOS encryption and QoS VPN features enables secure, low-cost transmissions over public networks. The Cisco IOS Firewall provides the most extensive VPN support, including but not limited to Dynamic Multipoint VPN (DMVPN), IPsec stateful failover, Cisco Easy VPN Remote, Cisco Easy VPN Server, site-to-site VPNs, Advanced Encryption Standard (AES), VPN acceleration cards, Voice and Video Enabled VPN (V3PN), and VPN QoS.

Introducing Cisco PIX 500 Series Security Appliances

This topic describes the security features of the Cisco PIX 500 Series Security Appliances and the FWSM.



Cisco PIX 500 Series Security Appliances scale to meet a range of requirements and network sizes and currently consists of these five models:

- The Cisco PIX 501 Security Appliance has an integrated 10/100BASE-T Ethernet port (100BASE-T option available in Cisco IOS Release 6.3) and an integrated four-port 10/100 switch.
- The Cisco PIX 506E Security Appliance has dual integrated 10/100BASE-T Ethernet ports (100BASE-T option available in Cisco IOS Release 6.3 for the Cisco PIX 506E Security Appliance only).
- The Cisco PIX 515E Security Appliance supports single-port or four-port 10/100 Ethernet cards.
- The Cisco PIX 525 Security Appliance supports single-port or four-port 10/100 Fast Ethernet and Gigabit Ethernet.
- The Cisco PIX 535 Security Appliance supports Fast Ethernet and Gigabit Ethernet.

Cisco PIX 500 Series Security Appliance Features

Features and uses are as follows:

- Typically used for site-to-site VPNs
- Restricts access to network resources
- Implemented at the physical perimeter between customer intranet and the intranet of the other company.
- Determines whether traffic crossing in either direction is authorized
- Contains limited intrusion detection system capability
- Provides a dedicated hardware appliance
- Has little or no impact on network performance




© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-7

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 security concerns such as protecting access to critical network resources. Network administrators need to know who is accessing what resources, so that they can establish clear perimeters to control that access. An effective security policy balances accessibility with protection. The Cisco PIX 500 Series Security Appliances enforce security policies at network perimeters. Often, network specialists think of a perimeter as the boundary between an internal network and the Internet, but you can establish a perimeter anywhere within a private network or between your network and a partner network. A solid perimeter security solution enables communications across the perimeter as defined by the security policy yet protects network resources from breaches or attacks. A perimeter security solution controls multiple network entry and exit points and increases user assurance by implementing multiple layers of security.

Cisco Catalyst 6500 Series Firewall Services Module

- **Runs in Cisco Catalyst 6500 Series Switches and Cisco 7600 Series Routers**
- **Designed for high-end enterprise and service providers**
- **Based on Cisco PIX Security appliance technology**
- **Provides feature parity with Cisco PIX Firewall Software Version 7.0**
- **Supports multiple performance and redundancy features**



Firewall Services Module for
Cisco Catalyst 6500 Series



Cisco Catalyst 6500 Series, Cisco
7600 Router Series

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0-4-8


The Cisco FWSM is a high-speed, integrated firewall module for Cisco Catalyst 6500 Series Switches and Cisco 7600 Series Routers. FWSM is based on Cisco PIX firewall technology and provides large enterprises and service providers with unmatched security, reliability, and performance. The “FWSM Key Features” table lists the key FWSM features.

FWSM Key Features

Feature	Description
Integrated module	The FWSM allows any port on the device to operate as a firewall port and integrates firewall security inside the network infrastructure.
Future proof	The FWSM can handle up to 5 Gbps of traffic, providing unsurpassed performance to meet future requirements without requiring a system overhaul. You can add up to three additional FWSMs to a chassis to meet growing demands.
Reliability	The FWSM uses Cisco PIX technology and uses the same time-tested Cisco PIX operating system, a secure, real-time operating system.
Lower cost of ownership	The FWSM offers the best price and performance of any firewall. Because Cisco bases the FWSM on the Cisco PIX security appliance and there are few boxes to manage, the cost of training and management is low.
Ease of use	The device manager (Cisco PIX Device Manager [PDM] or Cisco Adaptive Security Device Manager [ASDM]—FWSM revision dependent) is an intuitive GUI used to manage and configure the features within the FWSM. The FWSM is supported by the Cisco management framework and by Cisco Architecture for Voice, Video and Integrated Data (Cisco AVVID) partners for configuring and monitoring.
Efficiency and productivity gains	Virtualized FWSM delivers multiple firewalls on one physical hardware platform. Network administrators can configure, deploy, and manage these functions as if they were separate devices. Using virtualization to reduce the number of physical devices in a network significantly increases network efficiency and productivity.


Introducing Cisco ASA 5500 Series Adaptive Security Appliances

This topic describes the firewall features of Cisco ASA 5500 Series Adaptive Security Appliances.



The threat environment is constantly changing as new attacks (worms, viruses, and so on) emerge. Unlike most security appliances on the market, Cisco ASA 5500 Series Adaptive Security Appliances include future service extensibility. The ability to add completely new security services to the device as needed to address threats in the future is a key differentiator. Because Cisco ASA 5500 Series Adaptive Security Appliances do not require complete replacement when new services are required, these security appliances provide great investment protection.

The figure maps the old world of technology silos to the new world of Cisco Adaptive Threat Defense (ATD) solutions. The Cisco ASA 5500 Series Adaptive Security Appliances integrate market-proven technologies as shown. Many multifunction security devices are strong in one area, but weak in others. Customers tend to use only the stronger features, passing weaker functions to separate appliances or applications. Cisco ASA 5500 Series Adaptive Security Appliances are strong in all areas.



This figure shows more clearly the technical advantages of Cisco ASA 5500 Series Adaptive Security Appliances in terms of using proven best-of-breed solutions.

Cisco ASA 5500 Series Adaptive Security Appliances deliver complete consistency with Cisco PIX 500 Series Security Appliances allowing customers to leverage their extensive Cisco PIX knowledge while taking advantage of the many benefits that are unique to Cisco security appliances. In terms of leveraging user experience, consider these features:

- **Consistent firewall and IPsec VPN services:** The Cisco family of security appliances delivers advanced Cisco PIX Software Version 7.0 firewall-based appliances.
- **Consistent management:** The Cisco family of security appliances uses the same web-based management interface, Cisco ASDM, and command-line interface (CLI) as Cisco PIX 500 Series Security Appliances.
- **Consistent monitoring:** The Cisco family of security appliances has the same consistent monitoring capabilities as Cisco PIX 500 Series Security Appliances running Cisco PIX Software Version 7.0. These monitoring capabilities support syslog, Simple Network Management Protocol (SNMP), HTTP, HTTP Secure (HTTPS), and Cisco security appliance-specific extensions.

Migrating from Cisco PIX to Cisco Security Appliance

Key business and technology drivers:

- Lower total operating expenditures
- Lower capital expenditures
- High-performance worm, spyware, malware, and attack mitigation services
- Adaptive solution with converged best-of-breed security services
- Highly flexible and scalable VPN services
- Consistent user experience

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0-4.11

There are many reasons that support a migration from Cisco PIX-based security architectures to a Cisco security appliance-based architecture. The key business drivers supporting such an effort include these benefits:

- **Lower total operating expenditures:** Unified management and monitoring, added to a single platform, decreases complexity and simplifies deployments and ongoing support.
- **Lower capital expenditures:** The Cisco Technology Migration Program lowers capital expenditures. Leasing promotions further reduce costs and deliver new solutions as required.
- **High-performance worm, spyware, malware, and attack mitigation:** Advanced IPS and network antivirus services mitigate a wide range of threats.
- **Adaptive solution with converged best-of-breed security services:** Converged best-of-breed services combine market-proven firewall, IPS, IPsec, and SSL VPN services along with adaptive architecture to allow for future extensions in services. This combination of features protects businesses with a superior network security posture while providing strong investment protection.
- **Highly flexible and scalable VPN services:** Scalable VPN services give businesses ultimate VPN deployment flexibility by offering both IPsec and Cisco IOS WebVPN services, allowing businesses to tailor secure connectivity services based on their growing connectivity and scalability requirements.
- **Consistent user experience:** The Cisco security appliances leverage existing customer knowledge of Cisco PIX 500 Series Security Appliances for easy migration to Cisco ASA 5500 Series solutions.

Developing an Effective Firewall Policy

This topic explains how to develop an effective firewall policy based on firewall best practices.

Best Practices for Firewall Policy Development

- **Trust no one**
- **Base all filtering decisions on a sound firewall policy that balances security and business needs**
- **Deny physical access to firewall devices**
- **Only allow necessary protocols**
- **Use logs and alerts**
- **Segment security zones**
- **Do not use a firewall as a server**
- **Do not use a firewall as a workstation**
- **Set connection limits**
- **Restrict access to firewalls**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-12

Your firewall is an implementation of your policy, not the other way around. Your firewall policy comes first and details what traffic to filter and the nature of network connectivity needed before you start to set up your firewalls. Defending unplanned decisions after you have set up a firewall always complicates firewall administration.

As an example, suppose that a firewall configuration blocks Microsoft Remote Procedure Call (RPC)-based traffic from entering or leaving a protected subnet. Later, users in that subnet need RPC services to contact hosts on the outside. If there are no RPC filtering rules, it is difficult to deny access to these people, especially if a decision to deny access impairs productivity. Once an administrator makes an exception, more exceptions are likely to come into practice, and filtering rules become complex and unmanageable.

This example shows that filtering and connectivity policies must incorporate the security and business needs of your organization. By locating the firewall at the Internet gateway, as opposed to a subnet, you can give users the RPC access that they need without jeopardizing overall security.

There are a number of best practices lists available on the Internet that will guide you to develop a sound firewall policy. Here is a summary of a number of key points to add to the advice given thus far:

- **Trust no one:** Deny all traffic by default and only enable those services needed to conduct business. You need to consider the information that users need and provide them access only to that information. Use the least-privilege principle to give no more privilege than is necessary to perform a job. Keep your firewall configuration as simple as possible. Eliminate unneeded or redundant rules to ensure that your configuration supports your specific needs.

- **Deny physical access to firewall devices:** If a malicious user can obtain physical access to the firewall, anything can happen. Ensure that physical access to the firewall is controlled.
- **Only allow necessary protocols:** What protocols do you need to support to allow business operations to connect to other networks and subnetworks? You should allow only the necessary protocols.
- **Use logs and alerts:** Determine the level and type of logging needed and monitor logs on all firewalls on a regular basis. Use a secure remote syslog server to make log modification and manipulation difficult for malicious users.
- **Segment security zones:** Firewalls are useful for protecting internal systems from internal misuse in addition to their traditional role of protecting public servers from the dangers of being accessible from the Internet. You can use firewalls to create DMZs to limit access to defined security zones within your organization.
- **Do not use a firewall as a server:** Never include firewalls in server consolidation plans. Disable or uninstall any unnecessary services and software on the firewall that are not specifically required. Remove management tools from firewalls to prevent hackers from installing Trojan horse software or back doors. Run antivirus, content filtering, VPN, DHCP, and authentication software on other dedicated systems behind the firewall.
- **Do not use a firewall as a workstation:** Workstations use client applications (Microsoft Internet Explorer, Microsoft Outlook Express, FTP, and so on) that expose a firewall to viruses, worms, and other exploits. Only use robust management applications to help troubleshoot problems.
- **Set connection limits:** Cisco security appliance firewall capabilities can mitigate worm and other automated attacks if you enforce connection limits. You can change default connection limits in the global settings.
- **Restrict access to firewalls:** Restrict firewall accounts to administrator use. Do not allow network logins. Use strong passwords or use challenge-response and OTP cards. Use a unique user ID instead of “administrator” or “root.” Use a different user ID and password on every firewall device.

Best Practices for Firewall Policy Development (Cont.)

- **Combine firewall technologies**
- **Use firewalls as part of a comprehensive security solution**
- **Maintain your installation**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-13

- **Combine firewall technologies:** Do not rely on packet filtering alone. Use stateful inspection, protocol inspection, and application inspection, as applicable.
- **Use firewalls as part of a comprehensive security solution:** Do not depend entirely on firewalls—they are an adjunct to other security devices. Integrate firewalls with other technologies including these possibilities:
 - Network intrusion detection system (IDS) and IPS
 - Host IPS (HIPS)
 - Personal firewalls
 - Antivirus software
 - E-mail and web content filtering software
 - URL filtering software
 - Third-party authentication systems
- **Maintain your installation:** Keep software patches and updates current. Security is only as good as the latest security patch; therefore, system maintenance should be regular and timely. Firewalls are not install-and-forget devices. Patch the network operating system and application software with the latest code on a regular basis. However, make sure that you test these updates in a controlled, nonproduction environment whenever possible. As application requirements change, update firewall configurations to match those changes.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Cisco firewall products range from Cisco IOS Firewalls on routers, Cisco PIX 500 Series Security Appliances, FWSM for Cisco Catalyst 6500 Series Switches and Cisco 7600 Series Routers, and Cisco ASA 5500 Series Adaptive Security Appliances.
- The Cisco IOS Firewall provides stateful, application, and protocol inspection along with other key firewall features.
- The Cisco IOS Firewall provides security solutions to small and home offices, branch, extranet, and VPN and WAN aggregation points using Cisco routers, Cisco Catalyst switches, and Cisco Route Switch Modules.
- Cisco PIX 500 Series Security Appliances provide a range of requirements and network sizes. The FWSM has comparable features and can be installed in the Cisco Catalyst 6500 Series Switches or Cisco 7600 Series Routers.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-14

Summary (Cont.)

- Cisco ASA 5500 Series Adaptive Security Appliances delivers complete consistency with Cisco PIX 500 Series Security Appliances providing:
 - Firewall and IPsec VPN services
 - Web-based Cisco ASDM, and CLI management capabilities
 - Support for the same monitoring capabilities as the Cisco PIX 500 Series Security Appliances running Cisco PIX Software Version 7.0
- Using industry experience and best practices is the best way to develop an effective firewall policy.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—4-15

Module Summary

This topic summarizes the key points that were discussed in this module.

Module Summary

- Cisco firewall technology is built on generations of technology development and experience. Cisco firewall technology is a key component of every network security solution.
- Static packet filtering with Cisco ACLs provides a first line of defence against a wide range of security threats.
- Cisco SDM can be used to complete basic firewall configurations that can then be tailored to meet specific needs.
- Cisco provides a wide range of advanced security appliance products to meet all network needs.

© 2005 Cisco Systems, Inc. All rights reserved. SND v2.0-4.1

A firewall is a set of rules designed to protect network devices from intentional hostile intrusion that could threaten information assurance or lead to a denial of service (DoS) attack. Firewall rules are created to implement your security policies. Cisco Systems provides a range of firewall products that help you implement your security policies in a cost-effective way. Cisco extends your ability to maintain your firewall capabilities with the Cisco Router and Security Device Manager (SDM) Firewall Wizard, which helps you create a firewall for your LAN using a GUI.

References

For additional information, refer to these resources:

- Cisco Systems, Inc. Cisco ASA 5500 Series Adaptive Security Appliances: At-a-Glance.
http://www.cisco.com/application/pdf/en/us/guest/products/ps6120/c1031/cdccont_0900aecd80285492.pdf.
- Cisco Systems, Inc. Cisco Catalyst 6500 and Cisco 7600 Firewall Services Module At-a-Glance.
http://www.cisco.com/application/pdf/en/us/guest/products/ps4452/c1031/cdccont_0900aecd80356e40.pdf.
- Cisco Systems, Inc. Cisco IOS Firewall.
http://www.cisco.com/en/US/products/sw/secursw/ps1018/prod_white_papers_list.html
- Cisco Systems, Inc. *Cisco IOS Firewall Design Guide*.
http://www.cisco.com/en/US/products/sw/secursw/ps1018/products_implementation_design_guide09186a00800fd670.html.
- Cisco Systems, Inc. Cisco Pix Security Appliances: At-a-Glance.
http://www.cisco.com/application/pdf/en/us/guest/products/ps2030/c1031/ccmigration_09186a008007d065.pdf.
- Cisco Systems, Inc. Cisco Pix Security Appliance Threat Defense: Technology At-a-Glance.
http://www.cisco.com/application/pdf/en/us/guest/products/ps2030/c1031/cdccont_0900aecd800eb525.pdf.
- Cisco Systems, Inc. *Cisco PIX Firewall Configuration Guide, Version 6.0*.
http://www.cisco.com/en/US/partner/products/sw/secursw/ps2120/products_configuration_guide_book09186a0080143567.html.
- Cisco Systems, Inc. *PIX Device Manager 1.1 Online Help*.
http://www.cisco.com/application/pdf/en/us/guest/products/ps2032/c1626/ccmigration_09186a0080129fb0.pdf.
- Cisco Systems, Inc. *Evolution of the Firewall Industry*.
<http://www.cisco.com/univercd/cc/td/doc/product/iaabu/centri4/user/scf4ch3.htm>.
- The SANS Institute. The Twenty Most Critical Internet Security Vulnerabilities (Updated): The Experts Consensus. <http://www.sans.org/top20.htm>.

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) Which firewall technology uses a special piece of software designed to relay application layer requests and responses between endpoints? (Source: Introducing Firewall Technologies)
-
- Q2) Which firewall technology defines sets of rules and ACLs that determine which traffic is permitted or denied from being routed across a firewall by examining protocol header information up to the transport layer? (Source: Introducing Firewall Technologies)
-
- Q3) Which two statements describe a disadvantage of packet filtering technology? (Choose two.) (Source: Introducing Firewall Technologies)
- A) Packet filtering technology requires deep packet inspections up to the application layer.
 - B) Packet filtering requires complex ACLs, which can be difficult to implement and maintain correctly.
 - C) Packet filtering technology requires high CPU usages to support applications that negotiate dynamic ports.
 - D) Packet filtering technology requires high memory requirements to maintain the state table.
 - E) Packet filtering is susceptible to IP spoofing.
- Q4) On which protocol does a circuit level firewall depend? (Source: Introducing Firewall Technologies)
-
- Q5) Which four statements are true for stateful firewalls? (Choose four.) (Source: Introducing Firewall Technologies)
- A) They should not be used as a primary means of defense.
 - B) They can be used as an intelligent first line of defense.
 - C) They are not suitable as a means of strengthening packet filtering.
 - D) They can improve routing performance.
 - E) They are a defense against spoofing and DoS attacks.
 - F) They prevent application layer attacks.
 - G) All protocols contain state information, making their application almost universal.
 - H) They do not support user authentication of connections.

Q6) Which two ACL number ranges represent a standard ACL? (Choose two.) (Source: Building Static Packet Filters with Cisco ACLs)

- A) 1 to 99
- B) 100 to 199
- C) 1300 to 1999
- D) 2000 to 2699

Q7) Explain what the command statement **access-list 10 permit 192.168.3.0 0.0.0.255** does. (Source: Building Static Packet Filters with Cisco ACLs)

Q8) Explain what the command statement **access-list 101 permit tcp 63.36.9.0 0.0.0.255 any eq 80** does. (Source: Building Static Packet Filters with Cisco ACLs)

Q9) List the four types of enhanced ACLs. (Source: Building Static Packet Filters with Cisco ACLs)

Q10) What range of addresses does the sample ACL statement allow? (Source: Building Static Packet Filters with Cisco ACLs)

```
access-list 1 permit 172.30.16.0 0.0.15.255
```

- A) 172.30.16.0 through 172.30.16.255
- B) 172.30.16.0 through 172.30.29.255
- C) 172.30.16.0 through 172.30.255.255
- D) 172.30.16.0 through 172.30.31.255

Q11) Which ACL command statement will deny access to all users from network address 172.16.0.0.? (Source: Building Static Packet Filters with Cisco ACLs)

- A) **access-list 10 deny 172.16.0.0 0.0.255.255**
- B) **access-list 15 deny 172.16.0.0 0.0.255.255 any**
- C) **access-list 20 deny 172.16.0.0 255.255.0.0**
- D) **access-list 50 deny 172.16.0.0 0.0.0.0**

Q12) Describe what these ACL statements accomplish.(Source: Building Static Packet Filters with Cisco ACLs)

```
R2(config)# access-list 90 permit host 16.2.1.3 log
R2(config)# access-list 90 permit host 16.2.1.2 log
R2(config)# access-list 90 deny any log
R2(config)# line vty 0 4
R2(config-line)# access-class 90 in
R2(config-line)# end
```

-
- Q13) Explain how this ACL applied to an inside interface prevents IP spoofing. (Source: Building Static Packet Filters with Cisco ACLs)

```
R2(config)# access-list 105 permit ip 16.2.1.0 0.0.0.255 any
R2(config)# access-list 105 deny ip any any log
R2(config)# interface e0/1
R2(config-if)# ip access-group 105 in
R2(config-if)# end
```

-
- Q14) Explain how this ACL mitigates DDoS attacks using the Trin00 attack. (Source: Building Static Packet Filters with Cisco ACLs)

```
R2(config)# access-list 190 deny tcp any any eq 1524 log
R2(config)# access-list 190 deny tcp any any eq 27665 log
R2(config)# access-list 190 deny udp any any eq 27444 log
R2(config)# access-list 190 deny udp any any eq 31335 log
R2(config)# interface e0/0
R2(config-if)# ip access-group 190 in
R2(config-if)# end
R2(config)# interface e0/1
R2(config-if)# ip access-group 190 in
R2(config-if)# end
```

Module Self-Check Answer Key

- Q1) proxy server
- Q2) packet filtering
- Q3) B, E
- Q4) TCP; specifically, the TCP handshake
- Q5) B, D, E, H
- Q6) A, C
- Q7) This standard ACL command statement allows traffic from all addresses in the range 192.168.3.0 to 192.168.3.255 into the network.
- Q8) This extended ACL command statement says that ACL 101 will permit traffic originating from any address on the 63.36.9.0/24 network to any destination host port 80 (HTTP).
- Q9) dynamic, time-based, reflexive, CBAC
- Q10) D
- Q11) A
- Q12) ACL 90 allows only hosts 16.2.1.3 and 16.2.1.2 to access router R2 using Telnet (port 23). The ACL denies Telnet access to router R2 by other hosts. This ACL also logs all successful and unsuccessful attempts to access router R2 using Telnet.
- Q13) ACL 105 permits only those packets that contain source addresses from the 16.2.1.0/24 network and denies all others. This ACL is applied inbound to the inside interface (e0/1) of router R2.
- Q14) The Trin00 attack sets up communications between clients, handlers, and agents using ports 1524 tcp, 27665 tcp, 27444 udp and 31335 udp. This ACL translates to “ACL number 190 will deny any TCP or UDP (as applicable) traffic going from any network to any network that has the port equivalent to TCP 1524 , TCP 27665 , 27444 udp, and 31335 udp, and this will be logged.”

Securing Networks with Cisco IOS IPS

Overview

In technology environments, Internet worms and viruses can spread across the world in a matter of minutes. Without the luxury of time to react, a network must possess the ability to instantaneously recognize and mitigate worm and virus threats. A networking architecture paradigm shift is required to defend against these fast moving attacks. It is no longer possible to contain the intrusions at a few points in the network. Intrusion prevention is required throughout the entire network to detect and stop an attack at every ingress and egress point in the network. The only scalable and cost-effective way to accomplish this is by integrating intrusion prevention systems (IPSs) into the access points of the network. This module provides the knowledge and skills to configure IPS on Cisco routers.

Module Objectives

Upon completing this module, you will be able to configure Cisco IOS IPS on Cisco network routers. This ability includes being able to meet these objectives:

- Describe the underlying IDS and IPS technology embedded in the Cisco HIPS and Cisco network IDS and IPS solutions
- Configure Cisco IOS IPS
- Describe the features and functions of the Cisco IPS product family

Lesson 1

Introducing IDS and IPS

Overview

Intrusion detection system (IDS) and intrusion prevention system (IPS) solutions form an integral part of a robust network defense solution. Maintaining secure network services is a key requirement of a profitable IP-based business. Using Cisco products and technologies as examples, this lesson defines IDS and IPS and how these systems work.

Objectives

Upon completing this lesson, you will be able to describe the underlying IDS and IPS technology embedded in the Cisco host IPS (HIPS) and Cisco network IDS and IPS solutions. This ability includes being able to meet these objectives:

- Describe the functions and operations of IDS and IPS
- Describe the types of IDS and IPS sensors
- Explain IPS technologies, attack responses, and monitoring options
- Describe HIPS and network IDS and IPS monitoring
- Explain how IDS and IPS signatures are used to detect malicious network traffic
- Explain how SDFs and signature micro-engines work together
- Describe the role of signature alarms in a Cisco IPS solution

Introducing IDS and IPS

This topic describes the functions and operations of IDS and IPS.

Defining IDS and IPS

- **Intrusion detection system**
 - An IDS analyzes copies of the traffic stream.
 - Network traffic is not slowed.
 - Some malicious traffic is allowed into the network.
- **Intrusion protection system**
 - Works in line in real time to monitor Layer 3 to Layer 7 traffic and content
 - Sensor needs to be able to handle network traffic
 - Prevents malicious traffic entering the network

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0-5-3

IDS and IPS work together to provide a network security solution. An IDS captures packets in real time, processes them, and can respond to threats but works on copies of data traffic to detect suspicious activity by using signatures. In the process of detecting malicious traffic, an IDS allows some malicious traffic to pass before the IDS can respond to protect the network. An IDS analyzes a copy of the monitored traffic rather than the actual forwarded packet. The advantage of operating on a copy of the traffic is that the IDS does not affect the packet flow of the forwarded traffic. The disadvantage of operating on a copy of the traffic is that the IDS cannot stop malicious traffic from single-packet attacks from reaching the target system before the IDS can apply a response to stop the attack. An IDS often requires assistance from other networking devices, such as routers and firewalls, to respond to an attack.

An IPS works in line in the data stream to provide protection from malicious attacks in real time. Unlike an IDS, an IPS does not allow packets to enter the trusted side of the network. An IPS monitors traffic on Layer 3 and Layer 4 and analyzes the contents and payload of the packets for more sophisticated embedded attacks that might include malicious data on Layer 3 to Layer 7. This deeper analysis lets the IPS identify, stop, and block attacks that would normally pass through a traditional firewall device. When a packet comes in through an interface on an IPS, that packet is not sent to the outbound or trusted interface until the packet has been determined to be clean. An IPS builds upon previous IDS technology; Cisco IPS platforms use a blend of detection technologies including profile-based intrusion detection, signature-based intrusion detection, and protocol analysis intrusion detection.

The key to differentiating an IDS from an IPS is that an IPS responds immediately and does not allow any malicious traffic to pass, whereas an IDS allows malicious traffic to pass before it can respond.

IDS and IPS Common Characteristics


- **IDS and IPS technology is deployed in a sensor. These are sensor options:**
 - A router configured with Cisco IOS IPS
 - An appliance specifically designed to provide dedicated IDS or IPS services
 - A network module installed in an adaptive security appliance, in a switch, or in a router
- **The network can be monitored by IDS and IPS technologies—Network IDS and Network IPS.**
- **Host computers can be monitored by HIPS.**
- **IDS and IPS technologies use a set of rules called a signature to detect typical intrusive activity.**
- **IDS and IPS technologies look for these patterns of misuse:**
 - An atomic pattern
 - A composite pattern

© 2006 Cisco Systems, Inc. All rights reserved.

SNID v2.0—5-4

IDS and IPS technologies share these characteristics:

- IDS and IPS technologies are deployed as sensors. An IDS or an IPS sensor can be any of these devices:
 - A router configured with Cisco IPS
 - An appliance specifically designed to provide dedicated IDS or IPS services
 - A network module installed in an adaptive security appliance, switch, or router
- IDS and IPS technologies typically monitor for malicious activities in these two spots:
 - Malicious activity is monitored at the network detecting attacks against a network, including attacks against hosts and devices, network IDS, and network IPS.
 - Malicious activity is monitored on a host detecting attacks launched from or on target machines; host-based attacks are detected by reading security event logs, checking for changes to critical system files, and checking system registries for malicious entries, and host IPS (HIPS).
- IDS and IPS technologies use signatures to detect patterns of misuse in network traffic. A signature is a set of rules that an IDS or IPS uses to detect typical intrusive activity. Signatures are usually chosen from a broad cross section of intrusion detection signatures, and can detect severe breaches of security, common network attacks, and information gathering.
- IDS and IPS technologies look for these patterns of misuse:
 - In an atomic pattern, an attempt is made to access a specific port on a specific host, and malicious content is contained in a single packet. An IDS is particularly vulnerable to an atomic attack because until it finds the attack, malicious single packets are being allowed into the network. An IPS prevents these packets from entering at all.
 - A composite pattern is a sequence of operations distributed across multiple hosts over an arbitrary period of time.



The figure shows a sensor deployed in IDS mode and a sensor deployed in IPS mode. In Step 1, an attack is launched on a network with a sensor deployed in IDS mode and the Cisco switch sends copies of all packets to the IDS sensor (configured in promiscuous mode) to analyze the packets. At the same time, the target machine experiences the malicious attack. In Step 2, the IDS sensor, using a signature, matches the malicious traffic to the signature and, in this example, sends the switch a command to deny access to the malicious traffic, and, in Step 3, sends an alarm to a management console for logging and other management purposes.

When an attack is launched on a network with a sensor deployed in IPS mode in Step 1, the IPS sensor, in Step 2 (configured in in-line mode), analyzes the packets as soon as they come into the IPS sensor interface. The IPS sensor, using signatures, matches the malicious traffic to the signature and the attack is stopped immediately. In Step 3, the IPS sensor can send an alarm to a management console for logging and other management purposes.

Comparing IDS and IPS Solutions

	Advantages	Disadvantages
IDS (Promiscuous mode)	• No impact on network (latency, jitter) • No impact on sensor failure • No network impact on sensor overload	• Response action cannot stop trigger packets • Correct tuning required for response actions • More vulnerable to network evasion techniques
IPS (In-line mode)	• Trigger packets stopped • Can use stream normalization techniques	• Sensor issues might affect network traffic • Sensor overloading impacts network • Some impact on network (latency, jitter)

© 2006 Cisco Systems, Inc. All rights reserved.

SNID v2.0—5-6

The table in the figure shows some of the advantages and disadvantages of an IDS in promiscuous mode and an IPS in in-line mode. This list here expands on the advantages and disadvantages of an IDS and an IPS solution.

The advantages of an IDS platform in promiscuous mode are as follows:

- Deploying the IDS sensor does not have any impact on the network (latency, jitter, and so on).
- The IDS sensor is not in line and, therefore, a sensor failure cannot impact network functionality.
- Overrunning the IDS sensor with data will not impact network traffic. It will affect the ability of the IDS to analyze the data.

The disadvantages of an IDS platform in promiscuous mode are as follows:

- IDS sensor response actions cannot stop the trigger packet and are not guaranteed to stop a connection. IDS response actions are typically better at stopping an attacker more than a specific attack itself.
- IDS sensor response actions are less helpful in stopping e-mail viruses and automated attackers such as worms.
- Users deploying IDS sensor response actions must have a well thought-out security policy combined with a good operational understanding of their IDS deployments. Users must spend time to correctly tune IDS sensors to achieve expected levels of intrusion detection.
- Being out of band (OOB), IDS sensors are more vulnerable to network evasion techniques and must expend significant resources attempting to penetrate resources, such as a flooding attack.


The advantages of an IPS platform in in-line mode are as follows:

- An IPS sensor can be configured to perform a packet drop that can stop the trigger packet, the packets in a connection, or packets from a source IP address.
- Being in line, an IPS sensor can use stream normalization techniques to reduce or eliminate many of the network evasion capabilities that exist.

The disadvantages of an IPS platform in in-line mode are as follows:

- An IPS sensor must be in line and, therefore, IPS sensor errors or failure can have a negative effect on network traffic.
- Overrunning IPS sensor capabilities with too much traffic will negatively affect the performance of the network.
- Users deploying IPS sensor response actions must have a well thought-out security policy combined with a good operational understanding of their IPS deployments.
- An IPS sensor will affect network timing because of latency, jitter, and so on. An IPS sensor must be appropriately sized and implemented so that time-sensitive applications such as VoIP are not negatively effected.

Placement of IDS and IPS Sensors



The figure shows the typical placement options for IDS and IPS sensors. The considerations listed here should be taken into account to determine the placement of IDS and IPS sensors.

■ Untrusted perimeter (outside firewall):

- In this placement, an IDS sensor will detect a large number of attacks and generate a lot of alarms. Place an IDS here so that you can analyze what kind of traffic is coming into the firewall and how attacks are executed.
- With an IDS sensor placed here, the alarms that are generated are not significant because the firewall will block most of the malicious traffic.

■ Network edge (inside firewall):

- By placing an IDS sensor, shown in the topology as an appliance-based sensor, the alarms that are generated detect firewall misconfigurations. The IDS sensor is detecting the malicious traffic that the firewall configuration has let in to the network.

Note Cisco platforms such as Cisco ASA 5500 Series Adaptive Security Appliances with the Cisco ASA Advanced Inspection and Prevention Security Services Module (AIP SSM) can act as an IDS and an IPS sensor.

- The appliance-based sensor placed inside the firewall is also configured as an IPS sensor. Placing an IPS sensor here allows for more focused application protection of the demilitarized zone (DMZ) applications and for additional attack protection. The IPS can immediately drop new attacks that the firewall was not configured to stop.

■ Host-specific (critical servers):

- The IDS implementation of a switch module-based IDS and IPS sensor in front of critical servers detects and analyzes attacks that are hitting the specific server.
- The IPS implementation of a switch module-based IDS and IPS sensor is configured here to block all traffic that should not hit the server.

Types of IDS and IPS Sensors

This topic describes the types of IDS and IPS sensors.

Types of IDS and IPS Sensors		
	Advantages	Disadvantages
Signature-Based	• Easy configuration • Fewer false positives • Good signature design	• No detection of unknown signatures • Initially a lot of false positives • Signatures must be created, updated, and tuned
Policy-Based	• Simple and reliable • Customized policies • Can detect unknown attacks	• Generic output • Policy must be created
Anomaly-Based	• Easy configuration • Can detect unknown attacks	• Difficult to profile typical activity in large networks • Traffic profile must be constant
Honey Pot-Based	• Window to view attacks • Distract and confuse attackers • Slow down and avert attacks • Collect information about attack	• Dedicated Honey pot server • Honey pot server must not be trusted

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0-5-8

The table in the figure summarizes the advantages and disadvantages of the various types of IDS and IPS sensors available. The list here describes these IDS and IPS sensors in more detail.

- **Signature-based:** A signature-based IDS or IPS sensor looks for specific, predefined patterns (signatures) in network traffic. It then compares the traffic to a database of known attacks and triggers an alarm or prevents communication if a match is found. The signature may be based on a single packet or a sequence of packets. New attacks that do not match a signature will not result in detection. For this reason, the signature database needs to be constantly updated.

Note Protocol analysis-based intrusion detection is similar to signature-based intrusion detection, but it performs a more in-depth analysis of the protocols specified in the packets.

Signature-based pattern matching is an approach that is rigid but simple to employ. In most cases, the pattern is matched against only if the suspect packet is associated with a particular service or, more precisely, destined to and from a particular port. This helps to lessen the amount of inspection done on every packet. However, it tends to make it more difficult for systems to deal with protocols that do not reside on well-defined ports and, in particular, Trojan horses and their associated traffic, which can usually be moved at will.

At the initial stage of incorporating signature-based IDS or IPS, before the signatures are tuned there can be a lot of false positives (traffic generating an alert which is no threat for the network). After the system is tuned and adjusted to the specific network parameters there will be fewer false positives than with the next approach, the policy-based approach.

- **Policy-based:** The IDS or IPS sensor is preconfigured based on the network security policy. You must create the policies used in a policy-based IDS or IPS. Any traffic detected outside the policy will generate an alarm or will be dropped. Creating a security policy requires detailed knowledge of the network traffic and is a time-consuming task. Policy-based signatures use an algorithm to determine if an alarm should be fired. Often policy-based signature algorithms are statistical evaluations of the traffic flow. For example, in a policy-based signature that is used to detect a port sweep, the algorithm issues an alarm when the threshold number of unique ports is scanned on a particular machine. Policy-based signature algorithms could be designed to only analyze a specific type of packets, for example, SYN packets. The policy itself may require tuning. For example, you might have to adjust the threshold level of certain types of traffic so that the policy conforms to the utilization patterns on the network that it is monitoring. Policies may be used to look for very complex relationships.
- **Anomaly-based:** Anomaly-based or profile-based signatures typically look for network traffic that deviates from what is seen “normally.” The biggest issue with this methodology is that you first need to define what “normal” is. Some systems have hard-coded definitions of normal traffic patterns and, in this case, they could be considered heuristic-based systems.

Other systems are built to learn normal traffic behavior; however, the challenge with these systems is in eliminating the possibility of improperly classifying abnormal behavior as normal. Also, if the traffic pattern being learned is assumed to be normal, the system must contend with how to differentiate between allowable deviations and those deviations not allowed or that represent attack-based traffic. Normal network traffic can be difficult to define.

- **Honey pot-based:** Honey pot systems use a dummy server to attract attacks. The purpose of the honey pot approach is to distract attacks away from real network devices. By staging different types of vulnerabilities in the honey pot server, you can analyze incoming types of attacks and malicious traffic patterns. You can use this analysis to tune your sensor signatures to detect new types of malicious network traffic.

Intrusion Prevention Technologies

This topic provides an explanation of IPS technologies, attack responses, and monitoring options.

Cisco IOS IPS Attack Responses

• Deny Attacker Inline	• Produce Alert
• Deny Connection Inline	• Produce Verbose Alert
• Deny Packet Inline	• Request Block Connection
• Log Attacker Packets	• Request Block Host
• Log Pair Packets	• Request SNMP Trap
• Log Victim Packets	• Reset TCP Connection

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—5-9

When an IPS sensor, configured with Cisco IOS IPS 5.0 or later, detects malicious activity, it can choose from any or all of these actions:

- **Deny Attacker Inline:** This action terminates the current packet and future packets from this attacker address for a specified period of time. The sensor maintains a list of the attackers currently being denied by the system. You can remove entries from the list or wait for the timer to expire. The timer is a sliding timer for each entry. Therefore, if attacker A is currently being denied, but issues another attack, the timer for attacker A is reset and attacker A remains on the denied attacker list until the timer expires. If the denied attacker list is at capacity and cannot add a new entry, the packet is still denied.
- **Deny Connection Inline:** This action terminates the current packet and future packets on this TCP flow.
- **Deny Packet Inline:** This action terminates the packet.
- **Log Attacker Packets:** This action starts IP logging on packets that contain the attacker address and sends an alert. This action causes an alert to be written to the Event Store, even if the Produce Alert action is not selected.
- **Log Pair Packets:** This action starts IP logging on packets that contain the attacker and victim address pair. This action causes an alert to be written to the Event Store, even if the Produce Alert action is not selected.
- **Log Victim Packets:** This action starts IP logging on packets that contain the victim address and sends an alert. This action causes an alert to be written to the Event Store, even if the Produce Alert action is not selected.
- **Produce Alert:** This action writes the event to the Event Store as an alert.

- **Produce Verbose Alert:** This action includes an encoded dump of the offending packet in the alert. This action causes an alert to be written to the Event Store, even if the Produce Alert action is not selected.
- **Request Block Connection:** This action sends a request to a blocking device to block this connection.
- **Request Block Host:** This action sends a request to a blocking device to block this attacker host.
- **Request SNMP Trap:** Sends a request to the notification application component of the sensor to perform Simple Network Management Protocol (SNMP) notification. This action causes an alert to be written to the Event Store, even if Produce Alert action is not selected.
- **Reset TCP Connection:** This action sends TCP resets to hijack and terminate the TCP flow.

Note	The Reset TCP Connection action can be used in conjunction with deny packet and deny flow actions. However, deny packet and deny flow actions do not automatically cause TCP reset actions to occur.

Event Monitoring and Management

- **There are two key functions of event monitoring and management:**
 - Real-time event monitoring and management
 - Analysis based on archived information (reporting)
- **Event monitoring and management hosted are on a single server or on separate servers for larger deployments.**
- **To decide how to implement your monitoring services consider this:**
 - It is recommended that a maximum of 25 well-tuned sensors can report to a single IDS management console.
- **Recommended approaches to implementing multiple IDS management consoles:**
 - Separate monitoring domain
 - Hierarchical monitoring structure

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-10

Event monitoring and management can be divided into the need for real-time event monitoring and management and the need to perform analysis based on archived information (reporting). These functions can be handled by a single server, or the functions can be placed on separate servers to scale deployment. The number of sensors that should be forwarding alarms to a single IDS management console is a function of the aggregate number of alarms per second generated by those sensors.

Experience with customer networks has shown that the number of sensors reporting to a single IDS management console should be limited to 25 or fewer. These customers use a mixture of default signature profiles and tuned signatures. The number of alarms generated by each sensor is determined by how sensitively the sensor is tuned; the more sensitive the tuning, the fewer the alarms generated and the larger the number of sensors that can report to a single IDS management console.

It is essential to tune out false positives to maximize the scalability of the network IDS deployment. Sensors that are expected to generate a large number of alarms, such as those sitting outside the corporate firewall, should log in to a separate IDS management console, because the number of false alarms raised increases the noise-to-signal ratio dramatically and makes it difficult to identify otherwise valid events.

When implementing multiple IDS management consoles, implement either separate monitoring domains or a hierarchical monitoring structure.


Separate Monitoring Domains

Separate security monitoring domains can be used when separate security operations groups are responsible for different geographic areas or business units. In this implementation, there is no ability to monitor real-time activity across the entire enterprise; sensors send alarms only to the IDS management console in their geographic area or business unit. To offset lost functionality, implement a separate, centralized system for trend analysis and reporting, and have all IDS management consoles forward events to this system. This architecture is similar to a Manager of Managers (MoM) architecture.

Hierarchical Monitoring Structure

An effective hierarchical monitoring structure requires an alarm or event policy to distinguish and identify those alarms requiring a local, regional, or corporate-wide response. Local alarms indicate a small-scale, localized attack against a branch network or remote office. Regional alarms indicate an attack against several branch networks, telecommuters, or remote office networks within a given geographic region. Regional incidents can be escalated to a more corporate-wide audience if it is determined that additional resources are necessary. Corporate-wide incidents represent a broad, enterprise-wide attack from one or more sources. In the latter situation, an enterprise security incident response team must coordinate the response. Local security personnel in regional networks may require direction to effectively coordinate resources to contain the various incidents and restore overall network integrity.

Two-Tier Hierarchical Cisco Security-MARS IPS Monitoring System



- The Cisco Security-MARS Global Controller monitors two or more local zones. Each zone consists of a cluster of monitored devices and is managed by a Cisco Security-MARS Local Controller.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-11

This figure shows a two-tiered hierarchical security monitoring system using the Cisco Security Monitoring, Analysis, and Response System (MARS) as an example. Cisco Security MARS can be deployed in a two-tiered architecture using a Cisco Security MARS Global Controller. The Cisco Security MARS Global Controller monitors two or more local zones. Each zone consists of a cluster of monitored devices (such as IPS sensors, firewalls, routers, and servers), and each zone is managed by a Cisco Security MARS Local Controller. The Cisco Security MARS Global Controller and Local Controller architecture has these advantages:

- The architecture allows for centralized, distributed management of network topology.
- A Cisco Security Global Controller manages multiple Cisco Security MARS Local Controllers (restricted by the license key), and each Cisco Security MARS Local Controller manages one zone.
- The architecture lets remote sites view their own data while keeping data private between Cisco Security MARS Global Controllers and Local Controllers.
- You can view the entire network from the Cisco Security MARS Global Controller.
- You can use multiple Cisco Security MARS Local Controllers to isolate departmental functions.

HIPS and Network IPS

This topic describes how HIPS and network IPS monitoring work.

HIPS Features


- **HIPS audits host log files, host file systems, and resources.**
- **Cisco Security Agent software is installed on each host.**
- **HIPS provides individual host detection and protection.**
- **HIPS does not require special hardware.**
- **Cisco Security Agent can stop attacks without any updates by identifying their behavior as malicious and responding to stop the attacks in real time.**

HIPS audits host log files, host file systems, and resources. A significant advantage of HIPS is that it can monitor operating system processes and protect critical system resources, including files that may exist only on that specific host. HIPS combines behavioral analysis and signature filters. HIPS can also combine the best features of antivirus, network firewalls, and application firewalls in one package.

A simple form of HIPS enables system logging and log analysis on the host. However, this approach can be extremely labor-intensive. When implementing HIPS, the Cisco Security Agent (CSA) software should be installed on each host to monitor all activity performed on and against the host. CSA performs the intrusion detection analysis and protects the host.

A Cisco HIPS deployment using CSA provides proactive security by controlling access to system resources. This approach avoids the race to update defenses to keep up with the latest exploit and protects hosts even on day zero of a new attack. For example, the Nimda and SQL Slammer worms did millions of dollars of damage to enterprises on the first day of their appearance, before updates were even available; however, a network protected with a CSA stopped these attacks without any updates by identifying their behavior as malicious.

HIPS Operation Details



1. An application calls for system resources.
2. HIPS checks the call against the policy.
3. Requests are allowed or denied.

- HIPS intercepts operation system and application calls.
- Rules control application and network stacks.
- Processor controls limit buffer overflow, registry updates, writes to the system directory, and the launching of installation programs.
- HIPS is behavior-based.

© 2006 Cisco Systems, Inc. All rights reserved.


SND v2.0—5-13

Recall that HIPS operates by detecting attacks occurring on a host on which it is installed. HIPS works by intercepting operating system and application calls, securing the operating system and application configurations, validating incoming service requests, and analyzing local log files for after-the-fact suspicious activity.

HIPS uses rules based on a combination of known attack signatures and a detailed knowledge of the operating system and specific applications running on the host. These rules enable HIPS to determine abnormal or out-of-bound activity and, therefore, prevent the host from executing commands that do not fit the correct behavior of the operating system or application.

HIPS improves the security of hosts and servers by using rules that control operating system and network stack behavior. Processor control limits activity such as buffer overflows, registry updates, writes to the system directory, and the launching of installation programs. Regulation of network traffic can help ensure that the host does not participate in accepting or initiating FTP sessions, can rate-limit when a denial of service (DoS) attack is detected, or can keep the network stack from participating in a DoS attack.

Cisco HIPS Deployment



The topology in the figure shows a typical Cisco HIPS deployment. CSA is installed on publicly accessible servers, corporate mail servers, application servers, and on user desktops. CSA reports events to a central console server located inside the corporate firewall. CSA is managed from a central management console.

The advantages and disadvantages of HIPS are as follows:

- **Advantages of HIPS:** The success or failure of an attack can be readily determined. A network IPS sends an alarm upon the presence of intrusive activity but cannot always ascertain the success or failure of such an attack. HIPS does not have to worry about fragmentation attacks or variable Time to Live (TTL) attacks because the host stack takes care of these issues. If the network traffic stream is encrypted, HIPS has access to the traffic in unencrypted form.
- **Disadvantages of HIPS:** There are two major drawbacks to HIPS. First, HIPS does not provide a complete network picture. Because HIPS examines information only at the local host level, HIPS has difficulty constructing an accurate network picture or coordinating the events happening across the entire network. Second, HIPS has a requirement to support multiple operating systems. HIPS needs to run on every system in the network. This requires verifying support for all different operating systems used.

NIPS Features

- **Sensors are connected to network segments. A single sensor can monitor many hosts.**
- **Sensors are network appliances tuned for intrusion detection analysis.**
 - The operating system is “hardened.”
 - The hardware is dedicated to intrusion detection analysis.
- **Growing networks are easily protected.**
 - New hosts and devices can be added without adding sensors.
 - New sensors can be easily added to new networks.



© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-15


Network IPS involves the deployment of monitoring devices, or sensors, throughout the network to capture and analyze the traffic. Sensors detect malicious and unauthorized activity in real time and can take action when required. Sensors are deployed at designated network points that enable security managers to monitor network activity while it is occurring, regardless of the location of the attack target.

Network IPS sensors are usually tuned for intrusion detection analysis. The underlying operating system of the platform on which the HIPS software is mounted is stripped of unnecessary network services and essential services are secured (that is, hardened). The hardware includes these components:

- **Network interface card (NIC):** Network IPS must be able to connect to any network (Ethernet, Fast Ethernet, Gigabit Ethernet).
- **Processor:** Intrusion detection requires CPU power to perform intrusion detection analysis and pattern matching.
- **Memory:** Intrusion detection analysis is memory-intensive. Memory directly impacts the ability of a network IPS to efficiently and accurately detect an attack.

Network IPS gives security managers real-time security insight into their networks regardless of network growth. Additional hosts can be added to protected networks without needing more sensors. When new networks are added, additional sensors are easy to deploy. Additional sensors are only required when their rated traffic capacity is exceeded, when their performance does not meet current needs, or when a revision in security policy or network design requires additional sensors to help enforce security boundaries.

Cisco NIPS Deployment



The figure shows a typical network IPS deployment. The key difference between this Network IPS deployment example and the previous HIPS deployment example is that there are no CSA agents on the various platforms. In this topology, the network IPS sensors are deployed at network entry points that protect critical network segments. The network segments have internal and external corporate resources. The sensors report to a central management and monitoring server located inside the corporate firewall.

The advantages and disadvantages of network IPS are as follows:

- **Advantages of network IPS:** A network-based monitoring system has the benefit of easily seeing attacks that are occurring across the entire network. Seeing the attacks against the entire network gives a clear indication of the extent to which the network is being attacked. Furthermore, because the monitoring system is only examining traffic from the network, it does not have to support every type of operating system that is used on the network.
- **Disadvantages of network IPS:** Encryption of the network traffic stream can essentially blind network IPS. Reconstructing fragmented traffic can also be a difficult problem to solve. Possibly the biggest drawback to network-based monitoring is that as networks become larger (with respect to bandwidth), it becomes more difficult to place network IPS at a single location in the network and successfully capture all of the traffic. Eliminating this problem requires the use of more sensors throughout the network. However, this solution increases costs.

Comparing HIPS and Network IPS


	Advantages	Disadvantages
HIPS	• Host-specific • Understands context of attack • Protects host after decryption • Application-level encryption protection	• Operating system dependent • Lower level network events not seen • Host is visible to attackers
Network IPS	• Cost-effective • Not visible on the network • Operating system independent • Lower level network events seen	• Cannot examine encrypted traffic • Does not understand context of an attack

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-17

The table compares HIPS and network IPS advantages and disadvantages.

HIPS and Network IPS Monitoring



The figure shows the range of features of a blended HIPS and network IPS implementation. HIPS and network IPS implementations complement one another. A host-based monitoring system examines information at the local host or operating system. Network-based monitoring systems examine packets that are traveling through the network for known signs of intrusive activity. As you move down the feature list toward network IPS, the features describe network-based monitoring features; application-level encryption protection is a HIPS feature, while DoS prevention is a network IPS feature.

Note Network-based monitoring systems do not assess the success or failure of the actual attacks. They only indicate the presence of intrusive activity.

Introducing Signatures

This topic explains how IDS and IPS signatures are used to detect malicious network traffic.

IPS Signature Operational Characteristics

- **A network IPS signature is a set of rules used to detect intrusive activity.**
- **Sensors scan network packets using existing signatures to detect known attacks and respond with predefined actions.**
- **You need to tune signatures to reduce false positives. Tune signatures by altering signature parameters.**
- **You cannot add or delete built-in signatures. Some built-in signatures can provide tuning information.**
- **Some signatures have subsignatures. Configuring a subsignature changes only that subsignature.**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-19

A signature is a set of rules that an IDS and an IPS use to detect typical intrusive activity, such as DoS attacks. Similar signatures are grouped together into SDFs, and a signature micro-engine is used to implement the signatures. SDFs and signature micro-engines are discussed in the “Examining SDFs and Signature Micro-Engines” topic. Signatures are easily installed using IDS and IPS management software such as the Cisco IDS Device Manager (IDM). Sensors allow you to modify existing signatures and define new ones.

As sensors scan network packets, they use signatures to detect known attacks and respond with predefined actions. A malicious packet flow has a specific type of activity and signature, and an IDS or IPS sensor will examine the data flow using many different signatures. When an IDS or IPS sensor matches a signature with a data flow, the sensor takes action, such as logging the event or sending an alarm to IDS or IPS management software such as the Cisco Router and Security Device Manager (SDM).

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 Internet Control Message Protocol (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 can tune built-in signatures (tuned signatures) by adjusting the many signature parameters. For example, signature 993 has these configuration parameters that can be tuned:

- **MpcTimeout in seconds $5 \leq \text{MpcTimeout} \leq 2500$ (default = 30)**
 - **MpcTimeout** is the interval between alarms.
- **MpcPercentThreshold in percent $0 \leq \text{MpcPercentThreshold} \leq 100$ (default = 0)**
 - **MpcPercentThreshold** is the percentage of missed packets that must be exceeded to trigger an alarm. A value of 100 percent disables this threshold.

Built-in signatures are included in the sensor software. You cannot add to or delete from the list of built-in signatures and you cannot rename them. Many built-in signatures are based on known attacks. Some built-signatures provide information about your sensor. For example, signature 993 (missed packet count) alerts you if the sensor is dropping packets. If the alarms show that there are no dropped packets or a very small percentage of dropped packets, this means that the sensor is still able to monitor the quantity of traffic being sent. If you see signature 993 alerts with a high percentage of dropped packets, your sensor is oversubscribed.

Some signatures have subsignatures; that is, the signature is divided into subcategories. When you configure a subsignature, changes made to the parameters of one subsignature apply only to that subsignature. For example, if you edit signature 3050 subsignature 1 and change the severity, the severity change applies to only subsignature 1, and not to signature 3050 subsignature 2.

IPS Signature Characteristics (Cont.)

- There are four types of signatures:
 - Exploit
 - Connection
 - String
 - DoS
 - State-tracking
- The type of signature used depends on these factors:
 - Network infrastructure
 - Protocols used
 - Operating systems
 - Services enabled
- The number of signatures available depends on the IPS sensor platform type.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-20

- Here are the four categories of signatures:
 - **Exploit:** Exploit-specific signatures seek to identify network activity or upper-layer protocol transactions that are unique to a specific exploit or attack tool. Consequently, each new exploit may require its own signature. Because a successful exploit can be created by slightly modifying the attack payload, updating exploit signatures is critical to a well-protected network. Fortunately, exploit signatures are often relatively easy to produce for simple protocols and attacks. The uniqueness of the exploit signature can be used by a security analyst to gain some insight into the methodology of an attacker, providing another way to identify and mitigate targeted vulnerabilities.
 - **Connection:** Connection signatures are attack signatures based on the TCP or UDP and port number of the packets being monitored. Connection signatures can be used for these activities:
 - Identifying misconfigured firewalls
(A misconfigured firewall is discovered if a signature matches a connection on the trusted side of the firewall when it had been assumed that the firewall has been configured to block this connection.)
 - Identifying services that are added to the network
 - Identifying back doors established on systems in the network
 - Verifying user access to specific ports
 - **String:** String pattern matching is based on looking for a fixed sequence of bytes in a single packet. Pattern matching is an approach that is not easy to program but is simple to employ. Usually the string pattern is matched against only if the suspect packet is associated with a particular service. This helps to lessen the amount of inspection done on every packet. However, it tends to make it more difficult for systems to deal with protocols that do not reside on well-defined ports and, in particular, Trojan horses and their associated traffic.

- **DoS:** DoS signatures indicate attempts by attack tools to consume bandwidth or computing resources to disrupt normal operations. Typical attack tools include Trin00, TFN, Stacheldracht, and TCP SYN floods.
 - **State-tracking:** A sensor that performs TCP state tracking, IP fragment reassembly, and detection of sweeps and floods must keep track of the state of the traffic on the network. Many of the signatures used in this type of sensor are based on a certain threshold of events that occur within a specified period of time. The only way to assess event thresholds is to keep a count of events until the time has expired or the threshold has been exceeded. A sensor must track information about both the source and destination of the packets, the state of TCP connections, and the type and number of packets going to or from the hosts. All this information must be stored somewhere within the sensor. The storage medium for this information is the state database.
- The type of signature used in an IDS and IPS environment depends on the network infrastructure, the protocols used, the operating systems found in the network, and the services enabled.
 - The number of signatures available depends on the IPS sensor platform type. For example, the Cisco IPS 4200 Series Sensors support more signatures than Cisco IOS IPS, which is further limited by the amount of RAM on the router.

Note	The current list of Cisco IOS IPS signatures can be found at http://www.cisco.com/en/US/partner/products/ps6634/products_white_paper0900aecd8039e2e4.shtml .

Attack Methods, IPS Signature Types, and Capabilities

Attack Method	Signature	
	Type	Capabilities
Attempt to connect from a reserved IP address	Connection	Sensor checks the source address field in an IP header.
Illegal TCP flag combination	Connection	Sensor compares the flags set in a TCP header against known good or bad flag combinations.
E-mail infected with a virus	Exploit	Sensor compares the subject of e-mail messages to the subject of known e-mail messages associated with the viruses, or it can look for a specific attachment.
DNS buffer overflow attempt contained in the payload of a query	String	The sensor can parse the DNS fields and check their length or look for exploit shell code sequences in the payload.
Denial of service attack on a server	DoS	The sensor signature keeps track of how many times the command is issued and sends an alert if that number exceeds the set threshold.
Unauthorized access to an FTP server	State-tracking	The sensor monitors FTP traffic for an authorized login. An alert would be sent if unauthorized commands were issued before the user had been properly authenticated.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-21

The table in the figure lists the typical attack methods with the corresponding signature type and signature intrusion prevention capabilities.

Examining SDFs and Signature Micro-Engines

This topic explains how SDFs and signature micro-engines work together.

Signature Definition Files

- An SDF contains all or a subset of the signatures supported by Cisco IPS.
- An IPS loads the signatures contained in the SDF and scans incoming traffic for matching signatures.
- The IPS enforces the policy defined in the signature action.
- Cisco IPS uses the SDF to populate internal tables with the information necessary to detect each signature.
- The SDF can be saved on the router flash memory.
- SDFs are downloaded automatically using Cisco services.
- Three prebuilt SDFs come with Cisco integrated services routers:
 - 256MB.sdf
 - 128MB.sdf,
 - attack-drop.sdf

An SDF contains all or a subset of the signatures supported by a sensor. The sensor reads the SDF, parses the file, and populates the internal tables of the sensor with the information necessary to detect each signature. The SDF can be saved on the router flash memory (recommended), or users can specify the location of the SDF on the router using a security management tool such as Cisco SDM.

The option of where to save the SDF gives customers the flexibility to choose from a broad set of signatures and engines that they want to load and activate on the router. Additionally, as new signatures are released, the SDF can be updated with more signatures and then saved into the router flash memory. This makes signature management in a sensor independent of the Cisco IOS software image version. The Cisco IOS software image does not need to be upgraded when more signatures are made available. Signature management is done using the Cisco SDM, the Cisco IDM, or CiscoWorks Management Center for IPS Sensors (CiscoWorks IPS MC).

Note Cisco IOS IPS has the ability to download IPS signatures without the need for a Cisco IOS software image update. Cisco IOS IPS currently supports more than 1500 signatures. Typically, new signatures are released every two weeks, with emergency signature updates posted as needed. The signatures are posted to Cisco.com.

SDM signature updates can be downloaded from <http://www.cisco.com/cgi-bin/tablebuild.pl/ios-sigup> (requires CCO login).

CiscoWorks IPS MC signature updates can be downloaded at <http://www.cisco.com/cgi-bin/tablebuild.pl/idsmc-ids4-sigup> (requires CCO login).

Three prebuilt SDFs—128MB.sdf, 256MB.sdf, and attack-drop.sdf—are available with the Cisco SDM package when the integrated services routers are shipped from the factory. These files contain high confidence-rated worm and attack signatures. Because they are rated as highly severe, signatures in attack-drop.sdf (only) are preconfigured to either send an alarm, drop the connection, or reset the connection—they represent activities that users would not want to see on their networks. Users can append signatures to the prebuilt SDFs according to their needs.

For example, the Nimda virus can be detected by loading and enabling these signatures from the attack-drop.sdf:

- Signature ID: 5081:0 WWW WinNT cmd.exe access
- Signature ID: 5114:2 WWW IIS Unicode attack
- Signature ID: 5326:0 root.exe access

The number of signatures that can be loaded on a router is a function of the amount of memory (DRAM) available.

Note	All of the signatures discussed here are set to alarm, drop, and reset the connection.


Memory Requirements of Pre-Built SDFs

Memory Available	Recommended SDF	Number of Signatures
256 MB or lower	256MB.sdf	~ 500
128 MB or lower	128MB.sdf	~ 300
64 MB or lower	attack-drop.sdf	~ 82

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--9-23

The number of signatures that can go on a router is completely dependent on memory. Cisco has developed some recommendations for choosing SDFs. The table in the figure shows the amount of memory required for a recommended SDF and the approximate number of signatures that can be supported by that amount of memory.



The Cisco Security MARS Distributed Threat Mitigation (DTM) with IPS is a collaborative solution that proactively identifies active network threats and distributes IPS signatures to mitigate them. Thus, it provides distributed and rapid threat mitigation using Cisco IOS IPS software and Cisco IPS 4200 Series Sensors. Here are the steps in the sequence and the flow of DTM:

- Step 1** The DTM feature of Cisco Security MARS works in conjunction with a Cisco IPS device to generate and publish current SDFs to the IPS function on Cisco IOS routers. The SDF is an Extensible Markup Language (XML) file that a Cisco IOS software IPS device reads and, based on the parsed data, it then populates its internal tables with the signatures against which it has to inspect each packet.
- Step 2** The Cisco IOS IPS software or IPS appliance detects one or more signatures and sends an alarm to Cisco Security MARS for event correlation and monitoring.
- Step 3** Following the intelligent analysis and correlation of signature detection events, and based on default and user-defined rules on Cisco Security MARS, signatures to be added to routers are identified, depending on when the last signature file update was sent.
- Step 4** Cisco Security MARS keeps track of the signature running on the routers and from there compiles the new .sdf file. To this file it adds the new signatures, depending on how the appliance has been configured and what it is firing on. Cisco Security MARS checks the amount of memory that each signature needs and the free memory on the router.
- Step 5** Cisco Security MARS sends the update to all or a group of branch routers (as configured). It is important to note that signatures can be added in “alarm only” mode at this stage.

- Step 6** Following the update, the branch office routers will also detect the same signature or signatures, start sending alarms to Cisco Security MARS, and verify the occurrence of active threats. Those new alarms, in turn, trigger a new rule in Cisco Security MARS to have those signatures drop the suspicious packets, thus stopping the attack at the branch gateway. A new .sdf file will be pushed to the routers with the new action for those rules following the schema.
- Step 7** You can define queries and reports around any of the DTM events, which are organized under the Miscellaneous and DTM event group. These events provide status around the DTM updates that are published to the target devices.

Benefits of DTM with Cisco IOS IPS Software

- Attempts to use the resources of the router for IPS only occur when needed.
- This solution provides (optionally) automated tuning of IPS signatures.
- Customers that turn on and use the IPS feature on their branch routers will not have to deal with too many (sometimes false) alarms.
- This solution increases the value of a company investment in network-based intrusion detection products.
- This solution helps the operator to quickly locate the source of the attacks.
- The Cisco IOS IPS system evaluates network device and security device events, correlating that data with anomalous traffic flow analysis to determine the fidelity of an IPS device signature alert before applying that signature.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-25

The primary benefits provided by DTM with the Cisco IOS software IPS solution are as follows:

- Attempts to use the resources of the router for intrusion detection and or protection only occur when (and as much as) needed.
- This solution provides (optionally) automated tuning of IPS signatures. Customers will not have to worry about which IPS signature set has to be loaded and active on their branch routers; additional signatures will be loaded and activated as attacks are detected at other parts of the network.
- Customers that turn on and use the IPS feature on their branch routers will not have to deal with too many (sometimes false) alarms, because a smaller set of signatures will be active, generating fewer alarms with a much lower ratio of false positives.
- This solution increases the value of a company investment in network-based intrusion detection products.
- This solution helps the user (operator) to narrow down or locate the source of the attack more quickly while dropping malicious traffic in line to protect other network segments. This solution also helps the user to quickly identify the branch or regional offices that are already affected by the attack and those that have not seen the attack yet.
- This combination of software creates the only in-line IPS system that evaluates network device and security device events and correlates that data with anomalous traffic flow analysis to determine the fidelity of an IPS device signature alert before applying that signature potentially enterprise-wide to quarantine an outbreak.

Signature Micro-Engines

- Cisco IPS relies on signature micro-engines to support IPS signatures.
 - All the signatures in a signature micro-engine are scanned in parallel.
- Each signature micro-engine does the following:
 - Categorizes a group of signatures (and each signature detects patterns of misuse in network traffic)
 - Is customized for the protocol and fields it is designed to inspect
 - Defines a set of legal parameters that have allowable ranges or sets of values
 - Uses router memory to compile, load, and merge signatures

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--9-26

A signature micro-engine is a component of an IDS and IPS sensor that supports a group of signatures in a common category (an SDF). Each engine is customized for the protocol and fields that it is designed to inspect and defines a set of legal parameters that have allowable ranges or sets of values. The signature micro-engines look for malicious activity in a specific protocol. Signatures can be defined for any of the supported signature micro-engines using the parameters offered by the supporting micro-engine. Packets are scanned by the micro-engines that understand the protocols contained in the packet.

Cisco signature micro-engines implement parallel scanning. All the signatures in a given signature micro-engine are scanned in parallel fashion, rather than serially. Each signature micro-engine extracts values from the packet and passes portions of the packet to the regular expression engine. The regular expression engine can search for multiple patterns at the same time (in parallel). To facilitate this parallel scanning, after loading a new SDF, the sensor builds the signature micro-engines by compiling all the signatures in each signature micro-engine for parallel scanning. Parallel scanning increases efficiency and results in higher throughput.

When IDS (promiscuous mode) or IPS (in-line mode) is enabled, a signature micro-engine is loaded (or built) on to the router. When a signature micro-engine is built, the router may need to compile the regular expression found in a signature. Compiling a regular expression requires more memory than the final storage of the regular expression. Be sure to determine the final memory requirements of the finished signature before loading and merging signatures.

Note	A regular expression is a systematic way to specify a search for a pattern in a series of bytes.

Note	For the list of currently supported signature micro-engines, refer to the "Lists of Supported Signature Engines" section in the <i>Cisco IOS Security Guide, Release 12.4</i> available at http://www.cisco.com/en/US/partner/products/ps6350/products_configuration_guide_chapter09186a00804453cf.html .

Supported Signature Micro-Engines

Signature Micro-Engines	Description
Atomic	Signatures that examine simple packets, such as ICMP and UDP
Service	Signatures that examine the many services that are attacked
String	Signatures that use regular expression-based patterns to detect intrusions
Multi-string	Supports flexible pattern matching and supports Trend Labs signatures
Other	Internal engine to handle miscellaneous signatures

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-27

The “Supported Signature Micro-Engines” table summarizes the types of signature micro-engines available in Cisco IOS Release 12.3(14)T.

Supported Signature Micro-Engines

Signature Micro-Engine	Description
ATOMIC.IP	Provides simple Layer 3 IP alarms
ATOMIC.ICMP	Provides simple ICMP alarms based on these parameters: type, code, sequence, and ID
ATOMIC.IPOPTIONS	Provides simple alarms based on the decoding of Layer 3 options
ATOMIC.UDP	Provides simple UDP packet alarms based on these parameters: port, direction, and data length
ATOMIC.TCP	Provides simple TCP packet alarms based on these parameters: port, destination, and flags
SERVICE.DNS	Analyzes the Domain Name System (DNS) service
SERVICE.RPC	Analyzes the remote procedure call (RPC) service
SERVICE.SMTP	Inspects Simple Mail Transfer Protocol (SMTP)
SERVICE.HTTP	Provides HTTP protocol decode-based string engine; includes antievasive URL deobfuscation
SERVICE.FTP	Provides FTP service special decode alarms
STRING.TCP	Offers TCP regular expression-based pattern inspection engine services
STRING.UDP	Offers UDP regular expression-based pattern inspection engine services
STRING.ICMP	Provides ICMP regular expression-based pattern inspection engine services
MULTI-STRING	Supports flexible pattern matching and supports Trend Labs signatures
Other	Provides internal engine to handle miscellaneous signatures

Examining Signature Micro-Engine and SDF Build Failures

Type of Failure	Default Response
Signature micro-engine build failure	Fail open
SDF load failure	Fail back to previously loaded SDF
SDF merge failure	Fail back to previously loaded SME
Unsupported signature or signature parameter	Print a syslog message

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-28

There are times when building a signature micro-engine it will fail. The signature micro-engine can fail for reasons such as attempting to load a corrupted SDF file or the signature micro-engine exceeding memory limitations of the router. The “Signature Micro-Engine Failure Types” table lists types of SDF and signature micro-engine failures, the default sensor responses, and a description of suggested responses and examples.

Signature Micro-Engine Failure Types

Type of Failure	Default Response	Description
Signature micro-engine build failure	Fail open	When a signature micro-engine build fails, you need to reconfigure the router to drop all packets destined for that signature micro-engine using the ip ips fail closed command at the router command-line interface (CLI). By default, when Cisco IOS IPS is building a signature micro-engine, it is designed to “fail open.” When Cisco IOS IPS fails to build a signature micro-engine, all packets that are destined for that particular signature micro-engine will pass traffic without scanning (that is, an “open” condition).
SDF load failure	Fail back to previously loaded SDF	If the Cisco IOS IPS is not able to load the attack-drop.sdf file onto a router, the router will revert to the previously loaded available signatures. In most cases, the previously loaded signatures are the Cisco IOS built-in signatures.
SDF merge failure	Fail back to previously loaded signature micro-engine	If an engine build fails when you are merging the attack-drop.sdf file with the built-in signatures, IPS will revert, by default, to the previously available engine (or engines).
Unsupported signature or signature parameter	Print a syslog message	If a signature or a signature parameter is not supported, Cisco IOS IPS will print a syslog message indicating that the signature or parameter is not supported.

Introducing Signature Alarms

This topic describes the role of signature alarms in a Cisco IPS solution.

Cisco Signature Alarm Types


- **False positive:**
 - Normal traffic or a benign action causes the signature to fire.
- **False negative:**
 - An actual attack is not detected.
- **True positive:**
 - An attack is detected as expected.
- **True negative:**
 - Normal traffic or a benign action does not cause an alarm.

© 2006 Cisco Systems, Inc. All rights reserved. SNO v2.0--5-29

The ability of IDS and IPS sensors to accurately detect an attack or a policy violation and generate an alarm is critical to the functionality of the sensors. Attacks can generate these types of alarms:

- **False positive:** A false positive is an alarm triggered by normal traffic or a benign action. Consider this scenario: A signature exists that generates alarms if the enable password of any network devices 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.
- **False negative:** A false negative occurs when a signature is not fired when offending traffic is detected. Offending traffic ranges from someone sending confidential documents outside of the corporate network to attacks against corporate web servers. False negatives are bugs in the IDS and IPS software and should be reported. A false negative should only be considered a software bug if the IDS and IPS have a signature that has been designed to detect the offending traffic.
- **True positive:** A true positive occurs when an IDS and IPS signature is correctly fired when offending traffic is detected and an alarm is generated. For example, consider a Unicode attack. Cisco IPS sensors have signatures that detect Unicode attacks against Microsoft Internet Information Services (IIS) web servers. If a Unicode attack is launched against Microsoft IIS web servers, the sensors detect the attack and generate an alarm.
- **True negative:** A true negative occurs when a signature is not fired when nonoffending traffic is captured and analyzed. In other words, the sensor does not fire an alarm when it captures and analyzes “normal” network traffic.

Support for SDEE and Syslog



© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-30

The figure shows how Cisco IPS alerts can be sent using the Security Device Event Exchange (SDEE) protocol and using a syslog-based approach. The sensor generates an alarm when an enabled signature is triggered. Alarms are stored on the sensor. A host can pull the alarms from the sensor. Pulling alarms from a sensor allows multiple hosts to subscribe to the event “feed” to allow a host or hosts to subscribe on an as-needed basis.

The support for SDEE and syslog in the Cisco IPS solution is as follows:

- Cisco IOS software supports the SDEE protocol.
- SDEE uses a pull mechanism. That is, requests come from the network management application, and the IDS and IPS router responds.
- SDEE becomes the standard format for all vendors to communicate events to a network management application.
- The Cisco IOS IPS router will still send IPS alerts via syslog.

Note The use of HTTP Secure (HTTPS) ensures that data is secured as it traverses the network.

Viewing SDEE Alarm Messages

The screenshot displays the Cisco SDM (Security Device Manager) interface for an Intrusion Prevention System (IPS). The left sidebar contains navigation options: Home, Configure, Monitor, Refresh, Save, Search, and Help. The main pane shows the 'Intrusion Prevention System (IPS)' configuration page, with 'SDEE Messages' selected under the 'Edit IPS' tab. A table of SDEE messages is displayed, showing a list of alarms with their timestamps, status, and descriptions. The messages include SDF_LOAD_SUCCESS, ENGINE_BUILDING, ENGINE_READY, ENGINE_BUILD_SKIPPED, and ENGINE_BUILDING_STRING_ICMP. The interface also shows a 'Cisco Systems' logo and a 'SND v2.0-5-31' version indicator.

Time	Type	Description
20:45:57 GMT+00:00 Tue Jan 17 2006	Status	SDF_LOAD_SUCCESS: SDF loaded successfully from flash:/12
20:45:57 GMT+00:00 Tue Jan 17 2006	Status	ENGINE_BUILDING: OTHER - 4 signatures - 1 of 15 engines
20:45:57 GMT+00:00 Tue Jan 17 2006	Status	ENGINE_READY: OTHER - 0 ms - packets for this engine will be
20:45:57 GMT+00:00 Tue Jan 17 2006	Status	ENGINE_BUILDING: MULTI-STRING - 0 signatures - 2 of 15 eng
20:45:57 GMT+00:00 Tue Jan 17 2006	Status	ENGINE_BUILD_SKIPPED: MULTI-STRING - there are no new s
20:45:57 GMT+00:00 Tue Jan 17 2006	Status	ENGINE_BUILDING_STRING_ICMP - 1 signatures - 3 of 15 engi
20:45:57 GMT+00:00 Tue Jan 17 2006	Status	ENGINE_READY: STRING_ICMP - 20 ms - packets for this engi
20:45:57 GMT+00:00 Tue Jan 17 2006	Status	ENGINE_BUILDING_STRING_UDP - 16 signatures - 4 of 15 engi

When you use Cisco SDM, you can keep track of alarms that are common in SDEE system messages including IPS signature alarms. Here is an example of an SDEE system alarm message:

```
%IPS-4-SIGNATURE:Sig:1107 Subsig:0 Sev:2 RFC1918 address  
[192.168.121.1:137 ->192.168.121.255:137]
```

Note

For a complete list of the Cisco IPS system messages, refer to the "Interpreting Cisco IPS System Messages" section in the *Cisco IOS Security Configuration Guide, Release 12.4* available at http://www.cisco.com/en/US/products/ps6350/products_configuration_guide_chapter09186a00804453cf.html.

Implementing Alarms in Signatures

- The level assigned to the signature determines the alarm severity level.
- The alarm severity levels are informational, low, medium, and high.
- Make the severity level of the signature the same as the severity level of the alarm.
- Tune your signatures to recognize intrusion patterns out of character with your network traffic patterns.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-32

Alarms fire when specific parameters are met. You must balance the number of incorrect alarms that you can tolerate with the ability of the signature to detect actual intrusions; too few alarms and you may be letting in more suspect packets but network traffic flows more quickly. If left with untuned signatures to use, IPS systems will produce many false positive alarms. Here are the factors that should be considered when implementing alarms used in a signature:

- The level assigned to the signature determines the alarm severity level.
- A Cisco IPS signature will be assigned one of these severity levels:
 - **Informational:** Activity that triggers the signature is not considered an immediate threat, but the information provided is useful information.
 - **Low:** Abnormal network activity is detected that could be perceived as malicious, but an immediate threat is not likely.
 - **Medium:** Abnormal network activity is detected that could be perceived as malicious, and an immediate threat is likely.
 - **High:** Attacks used to gain access or cause a DoS attack are detected, and an immediate threat is extremely likely.
- When turning a signature alarm, make the severity level of the signature the same as the severity level of the alarm.
- To minimize false positives, study your existing network traffic patterns and then tune your signatures to recognize intrusion patterns that are atypical (out of character) for your network traffic patterns. Do not base your signature tuning on traffic patterns that are based only on industry examples. Using an industry example as a starting point, determine what your own network traffic patterns are and use them in your signature alarm tuning efforts.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- **IDS technology is passive and monitors the network for suspicious activity and parsed system log files. IPS technology is reactive and is able to forward or drop packets based on what is detected. IDS and IPS can be implemented on the same sensor.**
- **There are four types of IDS and IPS sensors:**
 - Signature-based
 - Policy-based
 - Anomaly-based
 - Honey pot-based
- **When an IPS sensor, configured with Cisco IOS IPS software it can detect malicious activity and respond to protect a network in real-time. Using event monitoring and management tools malicious activity can be done in real-time or on archived information.**
- **HIPS and Network IPS implementations complement one another.**
 - HIPS examines local host or operating system information.
 - Network IPS examines network packets for instructive activity.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--9-33

Summary (Cont.)

- **Cisco IPS uses a signature to detect known intrusive activity and to respond with actions that you define.**
- **An SDF is a bundle of common signature files. Cisco IPS uses signature micro-engines to implement the signatures found in SDFs to detect malicious traffic.**
- **Configure alarms in a signature file by making the severity level of the alarm the same severity level as the signature. Minimize false positives and tune your signatures to recognize the intrusion patterns of your network.**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--9-34

Configuring Cisco IOS IPS

Overview

Configuring Cisco IOS Intrusion Prevention System (IPS) is a core competency for a network security administrator. This lesson describes how to configure Cisco IOS IPS on routers using the Cisco Router and Security Device Manager (SDM) IPS GUI. You will discover that the Cisco SDM GUI makes it easy to configure and manage Cisco IOS IPS on routers and security devices.

Objectives

Upon completing this lesson, you will be able to describe how to configure Cisco IOS IPS. This ability includes being able to meet these objectives:

- Describe the IPS features of Cisco IOS software
- Explain how to configure Cisco IOS IPS using the Cisco SDM GUI
- Describe how to navigate the Cisco SDM IPS home page
- Explain how to configure Cisco IPS rules using the Cisco SDM GUI
- Explain how to configure IPS signatures using the Cisco SDM GUI
- Explain how to configure IPS global settings using the Cisco SDM GUI
- Describe how to deliver a Cisco IPS configuration to a router using the Cisco SDM GUI

Cisco IOS IPS Features

This topic describes the IPS features of Cisco IOS software.

Cisco IOS IPS Intrusion Detection Technology

Cisco IOS IPS uses a blend of Cisco IDS and IPS products:

- Cisco IDS Series appliances
- Cisco Catalyst 6500 Intrusion Detection System Services Modules
- Cisco IDS Network Module

Cisco IOS IPS uses a blend of detection technologies:

- Profile-based
- Signature-based
- Protocol analysis-based

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0-5-3

Cisco Systems has implemented IPS functions into its internetwork operating system, Cisco IOS software. Cisco IOS IPS combines existing Cisco intrusion detection system (IDS) and IPS product features with three different intrusion detection techniques.

Cisco IOS IPS uses a blend of Cisco IDS and IPS products. Cisco IOS IPS uses technology from Cisco IDS and IPS sensor product lines, including Cisco IPS 4200 Series Sensors, Cisco Catalyst 6500 Intrusion Detection System Services Module, and Cisco IDS Access Router Network Module.

Cisco IOS IPS uses a blend of these detection techniques:

- **Profile-based intrusion detection:** Profile-based intrusion detection generates an alarm when activity on the network goes outside a defined profile. With anomaly detection, profiles are created for each user or user group on your system. These profiles are then used as a baseline to define normal user and network activity. A profile could be created to monitor web traffic.
- **Signature-based intrusion detection:** Signature-based intrusion detection is less prone to triggering a false alarm when detecting unauthorized activity. A signature is a set of rules pertaining to typical intrusion activity. Signature-based intrusion detection uses signatures based on values in IP, TCP, UDP, and Internet Control Message Protocol (ICMP) headers. Network engineers research known attacks and vulnerabilities and then develop signatures to detect these attacks and vulnerabilities on the network. These attack signatures encompass specific traffic or activity based on known intrusive activity.

A pattern matching approach searches for a fixed sequence of bytes in a single packet. Pattern matching is a rigid approach but is simple to employ. In most cases, the pattern is matched against a packet only if the suspect packet is associated with a particular service or, more precisely, destined to or from a particular port. For example, a signature might be based on a simple pattern matching approach such as the following:

If <the packet is IPv4 and TCP> and <the destination port is 2222> and <the payload contains the string “foo”> then <fire an alarm>.

Cisco IOS IPS implements signatures that can look at every packet going through the network and generate alarms when necessary. A Cisco IOS IPS generates alarms when a specific pattern of traffic is matched or a signature is triggered. You can configure a Cisco IOS IPS to exclude signatures and modify signature parameters to work optimally in your network environment.

- **Protocol analysis-based intrusion detection:** Protocol analysis-based intrusion detection is similar to signature-based intrusion detection, but it performs a more in-depth analysis of the protocols specified in the packets. A deeper analysis examines the payloads within TCP and UDP packets, which contain other protocols. For example, a protocol such as Domain Name System (DNS) is contained within TCP or UDP, which itself is contained within IP.

The first step of protocol analysis is to decode the packet IP header information and determine whether the payload contains TCP, UDP, or another protocol. For example, if the payload is TCP, some of the TCP header information within the IP payload is processed before the TCP payload is accessed (for example, DNS data). Similar actions are mapped for other protocols.

Protocol analysis requires that the IPS sensor knows how various protocols work so that it can more closely analyze the traffic of those protocols to look for suspicious or abnormal activity. For each protocol, the analysis is based not only on protocol standards, particularly the RFCs, but also on how things are implemented in the real world. Many implementations violate protocol standards. It is very important that signatures reflect common and accepted practice rather than the RFC-specified ideal; otherwise, false results can be reported.

Primary Benefits of the Cisco IOS IPS Solution

Cisco IOS IPS:

- **Uses the underlying routing infrastructure to provide an additional layer of security**
- **Denies malicious traffic from both the inside and outside network**
- **Works with Cisco IDS, Cisco IOS Firewall, VPN, and NAC solutions**
- **Is supported by Cisco SDM and CiscoWorks VMS**
- **Integrates smoothly into existing network infrastructure**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-4

These attributes describe the primary benefits of the Cisco IOS IPS solution:

- Cisco IOS IPS uses the underlying routing infrastructure to provide an additional layer of security with investment protection.
- Because Cisco IOS IPS is in line and supported on a broad range of routing platforms, attacks can be effectively mitigated to deny malicious traffic from both inside and outside the network.
- When used in combination with Cisco IDS, Cisco IOS Firewall, virtual private network (VPN), and Network Admission Control (NAC) solutions, Cisco IOS IPS provides superior threat protection at all entry points to the network.
- Cisco IOS IPS is supported by easy and effective management tools, such as Cisco SDM and CiscoWorks VPN/Security Management Solution (CiscoWorks VMS).
- Whether threats are targeted at endpoints, servers, or the network infrastructure, Cisco offers pervasive intrusion prevention solutions that are designed to integrate smoothly into the network infrastructure and to proactively protect vital resources.

Cisco IOS IPS Signature Features

Cisco IOS IPS Signature Feature	Description
Regular expression string pattern matching	Enables the creation of string patterns using regular expressions
Response actions	Enables the sensor to take an action when the signature is triggered
Alarm summarization	Enables the sensor to aggregate alarms; does this to limit the number of times an alarm is sent when the signature is triggered
Threshold configuration	Enables a signature to be tuned to perform optimally in a network
Antievasive techniques	Enables a signature to defeat evasive techniques used by an attacker

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-5

The table in the figure describes the features of Cisco IOS IPS-based signatures.

Configuring Cisco IOS IPS Using Cisco SDM

This topic explains how to configure Cisco IOS IPS using the Cisco SDM GUI.

Using Cisco SDM to Configure Cisco IOS IPS


1. Launch Cisco SDM.
2. Launch the IPS Rules Wizard.
3. Choose a router interface to apply the IPS rule.
4. Choose the traffic flow direction to be inspected by the IPS rules.
5. Specify where the router will find the SDFs.
6. Confirm status of interfaces and signature files.
7. Configure signature alarm severity, event actions, and parameters.
8. Save the Cisco IPS configuration to the router.

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0-5-6

The figure summarizes the steps used to configure Cisco IOS IPS using the Cisco SDM GUI.

Using the Cisco SDM GUI for IPS

This topic describes how to navigate to the Cisco SDM IPS home page.




The tasks associated with managing routers and security devices are displayed in a task pane on the left-hand side of the screen of the Cisco SDM home page. Choosing the Intrusion Prevention icon reveals the Cisco SDM IPS GUI. You will use the Cisco SDM IPS GUI to configure Cisco IOS IPS on routers and security devices.

The network topology in the figure shows a typical use scenario. To configure Cisco IOS IPS on the router or security device, click the **Launch IPS Rule Wizard** button. The wizard that is launched does more than just configure a rule; it performs all of the Cisco IOS IPS configuration steps.

Note In the Cisco SDM IPS GUI, when you see the words “the IPS rule configuration” substitute “the IPS signature configuration.”


Using Cisco SDM GUI to Edit Existing IPS Rules



The Edit IPS tab provides access to the main Cisco IOS IPS management and tuning functions buttons, including IPS Policies, Global Settings, SDEE Messages, and Signatures. All interfaces on the router or security device are prominently displayed in a list that includes an IPS status report.

Configuring IPS Rules

This topic explains how to configure Cisco IPS rules using the Cisco SDM GUI.



Using the Cisco SDM GUI to create a new rule on a Cisco router can be done manually through the Edit IPS tab, or automatically using the IPS rule wizard. The Cisco IOS IPS Deployment Guide recommends using the IPS rule wizard. The process is straightforward and includes these steps:

- Click the **Launch IPS Rule Wizard** button.
- Read about the need for the SDM to obtain Security Device Event Exchange (SDEE) messages and click **OK**.
- Choose a router interface to apply the IPS rule.
- Choose the traffic flow direction to be inspected by the IPS rule.
- Specify where the router will find the signature definition files (SDFs).

The wizard will present you with dialog boxes. One dialog box will show the progress of the configuration tasks, and another will displaying a signature micro-engine build report when the configuration is complete.

Confirming Cisco IOS IPS on Inbound and Outbound Interfaces


The screenshot shows the 'Edit IPS' tab in the Cisco IOS IPS configuration interface. The interface includes a sidebar with navigation options: 'Create IPS', 'Edit IPS', 'IPS Policies', 'Global Settings', 'SDEE Messages', and 'Signatures'. The main area displays a table of interfaces with their status for inbound and outbound traffic.

Interface Name	IP	Inbound	Outbound	VFR status	Description
FastEthernet0/0	128.107.245.35	Disabled	Disabled	on	
FastEthernet0/1	172.16.1.1	Enabled	Enabled	on	

The screen capture in the figure shows the Edit IPS tab form. The list includes the IPS status of the router showing that the inbound and outbound ports on the interface named “FastEthernet0/1” are now enabled with Cisco IOS IPS.

Configuring IPS Signatures

This topic explains how to configure IPS signatures using the Cisco SDM GUI.



IPS signatures are loaded as a part of the procedure used to create a Cisco IPS rule using the IPS rule wizard. The figure shows a screen capture of all currently loaded signatures. Because signatures are what optimize your configuration, confirm that all of the correct signatures are loaded on the router or security device.

Configuring Signatures Using Cisco SDM (Cont.)

Import Select by: All Signatures Criteria: --N/A--

Select All Add Edit Delete Enable Disable

Enabled	I	Sig ID	SubSig ID	Name	
☒		3157	0	FTP PASV Port Spoof	ala
☒		11002	0	Gnutella Server Reply	ala
☒		5088	0	vWWW Akopia MiniVend access	ala
☒		3153		Address	ala
☒		3129			ala
☒		5084			ala
☒		5084		NSDB Help	ala
☒		11212	0	Yahoo Messenger Through HT	ala
☒		5080	0	WWW IRM WebSphere access	ala

Actions...

- Set Severity To
 - high
 - informational
 - low
 - medium
- Restore Defaults
- NSDB Help

Signature Alarm Severity

Assign Actions

You can specify actions the sensor should perform when it detects the selected signature(s). To assign an action, click the checkbox next to the action. A checkmark indicates the action will be performed. No checkmark indicates the action will not be performed. A gray checkmark indicates the action is assigned to some, but not all of the signatures you selected.

☒ alarm
☐ denyAttackerInline
☐ denyFlowInline
☐ drop
☐ reset

All None

OK Cancel Help

Signature Event Actions

A signature configuration can be tuned using the Cisco SDM GUI. The screen capture in the figure shows the popup menu used to change the alarm severity and event actions of a signature.

Configuring Signatures Using Cisco SDM (Cont.)

Signature Parameters


Name	Value
SIOID:	3153
SigName:	FTP Improper Address
SubSig:	0
AlarmInterval:	
AlarmSeverity:	medium
AlarmThrottle:	FpsAB
AlarmTraits:	
BadPortCmdAddress:	True
BadPortCmdPort:	Undefined
BadPortCmdShort:	Undefined
ChokeThreshold:	100
Direction:	ToService
Enabled:	True
EventAction:	IDSmon

■ Parameter uses the Default Value. Click the icon to edit the value.
◆ Parameter uses a User-Defined Value. Click the icon to restore the default value.

OK Cancel Help

The figure shows a screen capture of the dialog box that you will use to configure signature parameters. To access and configure signature parameters, choose the signature and then click the **Edit** button in the Cisco SDM Configure Signatures window.

Importing Signature Definition Files



This series of screen captures shows how to update the IPS signatures with the latest SDF. To update an SDF from a PC, follow these steps:


- Step 1** Navigate to the **Edit IPS > Signatures** form and click the **Import** menu button and choose the **From PC** menu item.
- Step 2** The Import dialog box opens. Go to the directory and choose the SDF, then click the **Open** button.
- Step 3** When the SDF file has been imported, the IPS Import window appears.

The signature list in the IPS Import window displays the signatures available in the SDF. Review the signatures and choose the ones that you want to import. If you want to import all of the signatures, click the **Select All** button. In the signature list area, the Name column lists the name of the signature; for example, Ping of Death. The Deployed column displays either Yes or No to signal if the signature is already loaded on the router. The Import column displays a check box for each signature so that you can decide which group of signatures to import.

The other option offered in the IPS Import window is the option to merge signature files. You can import only one category at a time; if you switch to another category, the current import selection is lost.

Configuring Global Settings


This topic explains how to configure IPS global settings using the Cisco SDM GUI.



The screen capture shows the global features that you can configure using the Cisco SDM GUI. To access and configure a particular global feature, choose the item name and click the **Edit** button.

Delivering the Configuration to the Router

The topic describes how to deliver a Cisco IPS configuration to a router using the Cisco SDM GUI.



The figure shows two menu selections found in the File menu option. Once you have configured Cisco IPS on a router or security device, you can either write the new IPS configuration to the starting configuration of the router or save the running configuration to the PC to use later.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Cisco IOS IPS uses the underlying routing infrastructure to provide an additional layer of security to deny malicious traffic from inside and outside the network. Cisco IOS IPS works with Cisco IDS, Cisco IOS Firewall, VPN, and NAC solutions, and it is supported by Cisco SDM and CiscoWorks device management software.
- Using the Cisco SDM software there are basically eight steps to configure Cisco IOS IPS on a router.
- The Cisco SDM GUI is browser-based and easy to use.
- Use the Cisco SDM GUI wizard to configure IPS rules.
- Use the Cisco SDM GUI to configure and tune IPS signatures.
- Use the Cisco SDM GUI to configure global settings.
- Using the Cisco SDM GUI, you can save your configuration to the router or to a file on your PC.

© 2006 Cisco Systems, Inc. All rights reserved.

SDM v2.0-5.17

Defending Your Network with the Cisco IPS Product Family

Overview

When you need to defend your network, the Cisco Intrusion Prevention System (IPS) product family provides a comprehensive suite of routers, switches, and appliance and network modules from which to choose. In this lesson, you will learn about the relative positioning of Cisco IDS and IPS sensor platforms and modules. The Cisco host IPS (HIPS) solution feature of Cisco Security Agent (CSA) will be examined, as will how the complementary Cisco Guard Distributed Denial of Service (DDoS) Mitigation Appliances (Cisco Guard) and Cisco Traffic Anomaly Detectors defend against DDoS attacks. You will learn to choose the correct Cisco IPS solution for your application and learn policies and best practices to help you defend your network using the Cisco IPS product family.

Objectives

Upon completing this lesson, you will be able to describe the features and functions of the Cisco IPS product family. This ability includes being able to meet these objectives:

- Describe the relative positioning of Cisco IDS and IPS sensor platforms and modules
- Explain the Cisco HIPS solutions
- Select appropriate Cisco IPS solutions
- Describe IPS best practices

Network IPS Solutions

This topic describes the relative positioning of Cisco IDS and IPS sensor platforms and modules.



Cisco IPS solutions run on a variety of platforms. Here is a brief description of the available platforms:

- **Cisco ASA 5500 Series Adaptive Security Appliances:** The Cisco ASA 500 Series Adaptive Security Appliances offer a purpose-built, high-performance security solution. These appliances integrate the technologies from Cisco PIX 500 Series Security Appliances, Cisco IPS 4200 Series Sensors, and Cisco VPN 3000 Series Concentrators. The Cisco ASA 5500 Series Adaptive Security Appliances are a key component of the Cisco Adaptive Threat Defense (ATD) strategy. The Cisco ASA 5500 Series Adaptive Security Appliances combines a wide range of security and virtual private network (VPN) technologies to provide rich application security, anti-X defenses, network containment and control, and secure connectivity.
- **Cisco ASA 5500 SSM:** The Cisco ASA Advanced Inspection and Prevention Security Services Module (AIP SSM) uses advanced inspection and prevention technology to provide high-performance security services, such as intrusion prevention services and advanced anti-X services. The Cisco AIP SSM products include an AIP-SSM-10 module with a 1-GB memory and an SSM-AIP-20 module with a 2-GB memory.
- **Cisco IPS 4200 Series Sensors:** Cisco IPS 4200 Series Sensors offer significant protection to your network by helping to detect, classify, and stop threats, including worms, spyware and adware, network viruses, and application abuse. Using Cisco IPS Sensor Software Version 5.1, the Cisco IPS solution combines in-line intrusion prevention services with innovative technologies that improve accuracy. As a result, more threats can be stopped without the risk of dropping legitimate network traffic. Cisco IPS Sensor Software Version 5.1 includes enhanced detection capabilities and improved scalability, resiliency, and performance features.

- **Cisco Catalyst 6500 Intrusion Detection System Services Module (IDSM-2):** The Cisco IDSM-2 is part of the Cisco IPS solution. It works in concert with the other components to efficiently protect your data infrastructure. With the increased complexity of security threats, achieving efficient network intrusion security solutions is critical to maintaining a high level of protection. Vigilant protection ensures business continuity and minimizes the effect of costly intrusions.
- **Cisco Intrusion Detection System (IDS) Access Router Network Module (Cisco IDS Network Module):** The Cisco IDS Network Module—for the Cisco 2600 Series Multiservice Access Routers, Cisco 2800 Series Integrated Services Routers, Cisco 3600 Series Multiservice Platforms, Cisco 3700 Series Multiservice Access Routers, and Cisco 3800 Series Integrated Services Routers—is part of the Cisco IDS and IPS sensor family portfolio and Cisco IPS. The Cisco IDS Network Module is designed to be integrated with branch office routing to reduce the complexity of securing WAN links and to reduce operational costs.

Throughput on Cisco Routers That Support Cisco IOS IPS

Cisco Platform Tested	Maximum Throughput
Cisco 1841 Integrated Services Router	60 Mbps
Cisco 2801 Integrated Services Router	65 Mbps
Cisco 2811 Integrated Services Router	70 Mbps
Cisco 2821 Integrated Services Router	200 Mbps
Cisco 2851 Integrated Services Router	250 Mbps
Cisco 3825 Integrated Services Router	325 Mbps
Cisco 3845 Integrated Services Router	425 Mbps

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-4

The table in the figure lists the maximum throughput obtained for various router platforms with Cisco IOS IPS enabled. Maximum throughput numbers change often. The numbers presented in the table provide a good comparison of the relative performances of the various Cisco routers. Check Cisco.com for the latest maximum throughput numbers. Use these numbers to help you determine what the performance impact might be when implementing Cisco IOS IPS on your network. These numbers are derived from tests performed in isolated or ideal conditions. Your results may vary depending on your network architecture, traffic patterns, and the services running on the router.

Note To determine what the expected performance of Cisco IOS IPS could be Cisco used the Spirent Communications Avalanche (Avalanche) and Spirent Communications Reflector (Reflector) test system. Avalanche was programmed to rapidly establish multiple HTTP sessions to Reflector, with the router running Cisco IOS IPS between them.

Avalanche opens the TCP session to Reflector and requests a single 64-KB file from Reflector. Reflector sends the file, and then Avalanche receives the file and closes the TCP and HTTP session. Avalanche is programmed to increase the number of such transfers, increasing the total throughput load offered to the router. The load is increased until sessions start failing because of the performance limitations of the router.

Performance and Limitations of Platforms

Cisco IDS or IPS	Cisco IDS 4215 Sensor	Cisco IDS 4250 XL Sensor	Cisco IPS 4240 Sensor	Cisco IPS 4255 Sensor
Inline (IPS) Ready	Yes	Yes	Yes	Yes
Performance (Mbps)	65	800	250	500
Standard Monitoring Interface	10/100 BASE-TX	10/100/1000 Dual BASE-SX	Four 10/100/1000 BASE-TX	Four 10/100/1000 BASE-TX
Standard Command and Control Interface	10/100 BASE-TX	10/100/1000 BASE-TX	10/100 BASE-TX	10/100 BASE-TX
Optional Interface	Four 10/100 BASE-TX (4-FE)	None	Four 10/100/1000 BASE-TX (4-FE) Four 10/100/1000 BASE-SX (future)	Four 10/100/1000 BASE-TX (4-FE) Four 10/100/1000 BASE-SX (future)

FE = Fast Ethernet

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.6—5-5

The table in the figure shows the performance and interface limitations of the Cisco 4200 Series IDS and IPS platforms running as an in-line IPS sensor.

Note The performance numbers for the Cisco platforms listed in the table are slightly higher if these sensors are running in IDS promiscuous mode.

Performance and Limitations of Cisco ASA 5500 Series Platforms

ASA Performance with the Security Service Module

Cisco ASA 5500 Series Adaptive Security Appliance	Cisco ASA 5510 AIP SSM-10	Cisco ASA 5520 AIP SSM-20	Cisco ASA 5540 AIP SSM-20
Firewall + anti-X (Mbps)	150	375	450
Maximum VLANs	0 (10 sec+)	25	100
Interfaces (10/100)	3+Out-of-Band	1	1
Interfaces (10/100/1000)	—	4	4


© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-6

The table in the figure shows the performance and interface limitations of the Cisco ASA 5500 Series Adaptive Security Appliance platform. Refer to the Cisco ASA 5500 Series Adaptive Security Appliance Platform and Module Datasheet for an in-depth discussion on the performance metrics. This datasheet can be found at http://www.cisco.com/en/US/products/ps6120/products_data_sheet0900aecd802930c5.html.

Note The Cisco AIP SSM-10 can also run on the Cisco ASA 5520 Adaptive Security Appliance platform.

Relative Positioning of Cisco IPS Sensors



The diagram shows the relative positioning of some of the Cisco IDS and IPS sensors. Use this chart as a guide to select the Cisco IDS and IPS sensor platform with the correct performance and media support for your application.

Note For the complete line of Cisco IPS 4200 Series Sensors refer to <http://www.cisco.com/en/US/products/hw/vpndevc/ps4077/index.html>.

Cisco IPS Management Software

Available to help with IPS solutions:

- Cisco Security MARS
- Cisco ICS
- Cisco SDM
- Cisco IDM
- CiscoWorks SIMS
- Cisco Security Manager
- CiscoWorks VMS
- CiscoWorks IPS MC

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-8

An IPS solution can be configured using the CLI, but configuration is simpler with a GUI-based device manager. The following describes the Cisco device management software available to help you manage an IPS solution:

- **Cisco Security Monitoring, Analysis, and Response System (MARS):** Cisco Security MARS is an appliance-based, all-inclusive solution that allows network and security administrators to monitor, identify, isolate, and counter security threats. This family of high-performance appliances enables organizations to make more effective use of their network and security resources.

Cisco Security MARS can monitor security events and information from a wide variety of sources, including third-party devices and hosts. With its correlation engine, vector analysis, and hotspot identification, Cisco Security MARS can not only identify anomalous behavior and security threats, but can also recommend precision removal of those elements, leading to rapid threat mitigation. In addition, Cisco Security MARS incorporates a comprehensive reporting engine that provides easy access to information for compliance reporting.


- **Cisco Incident Control System (ICS):** The Cisco ICS prevents new worm and virus outbreaks from affecting businesses by enabling the network to rapidly adapt and provide a distributed response.

Because the time that it takes a worm or virus outbreak to spread around the world has decreased from days to minutes, a proactive response within minutes after an outbreak, regardless of its location, is necessary to help ensure the safety of business networks. The Cisco ICS solution meets that need by delivering a network-wide defense within minutes of an outbreak anywhere on the globe. Using the global monitoring capability of Trend Labs, Cisco ICS collaborates with existing Cisco Systems network and security devices to rapidly distribute worm and virus immunization capabilities throughout the network. This fast, proactive approach prevents worms and viruses from becoming entrenched, thus helping ensure network availability and decreasing the costs associated with damage cleanup.

- **Cisco SDM:** Cisco SDM is a web-based device management tool for Cisco routers that can improve the productivity of network managers, simplify router deployments, and help troubleshoot complex network and VPN connectivity issues. SDM supports a wide range of Cisco IOS software releases and is available free on Cisco router models including the Cisco 830 Series Secure Broadband Routers and the Cisco 7301 Router.
- **Cisco IDS Device Manager (IDM):** Cisco IDM is a web-based configuration tool for network IDS appliances. It is shipped at no additional cost with the IDS sensor code. Cisco IDM implements a web-based GUI that is similar to the CiscoWorks Management Center for IPS Sensors software. For additional scalability, customers may consider the CiscoWorks Management Center for IPS Sensors (CiscoWorks IPS MC) software.
- **CiscoWorks Security Information Management Solution (CiscoWorks SIMS):** An important element of the Cisco Self-Defending Network, CiscoWorks SIMS integrates, correlates, and analyzes security event data from the enterprise network to improve visibility and provide actionable intelligence for strengthening the security of an organization.
- **Cisco Security Manager:** Cisco Security Manager is a powerful but very easy-to-use solution to centrally provision all aspects of device configurations and security policies for Cisco firewalls, VPNs, and IPS. The solution is effective for managing even small networks consisting of fewer than 10 devices, but also scales to efficiently manage large-scale networks composed of thousands of devices. Scalability is achieved through intelligent policy-based management techniques that can simplify administration.
- **CiscoWorks VMS:** CiscoWorks VMS is an integral element of the Cisco SDM strategy and contributes to organizational productivity by combining web-based tools for configuring, monitoring, and troubleshooting VPNs, firewalls, network IDS, and HIPS. CiscoWorks VMS also includes network device inventory, change audit, and software distribution features.
- **CiscoWorks IPS MC:** CiscoWorks IPS MC is a tool for configuring Cisco network sensors, switch IPS sensors, IPS network modules for routers, and in-line intrusion prevention software in routers. The tool allows administrators to save time by using group profiles to configure multiple sensors concurrently. It also provides a powerful signature management feature that increases the accuracy and specificity of detecting possible network intrusions.

HIPS Solutions


This topic explains the Cisco HIPS solutions.



The CSA architecture model consists of these elements:

- **CSA Management Center (CSA MC):** The administrator divides network hosts into groups by function and security requirements and then configures security policies for those groups. The CSA MC maintains a log of security violations and sends alerts through e-mail or a pager.
- **CSA:** Software installed in the host systems 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.
- **Administration workstation:** Any workstation can be connected securely to the CSA MC using a web interface with enabled Secure Sockets Layer (SSL). This workstation then applies the behavioral requirements to internal system resource calls.

Application, Kernel, and Interceptors



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, CSA takes these actions:

- The request is blocked; it is 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 other calls 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 by the file system interceptor and allowed or denied based on the security policy.
- **Network interceptor:** The network interceptor changes and controls network driver interface specifications and clears network connections through the security policy by port and IP address pairs. The number of network connections allowed within a specified time can also be limited to prevent denial of service (DoS) attacks.
- **Configuration interceptor:** Read or write requests to the registry on Microsoft Windows or to rc files on UNIX are intercepted by the configuration interceptor. 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 the dynamic run-time environment of each application. It detects and blocks requests to write to memory that are not owned by the requesting application. Attempts by one application to inject code, such as a shared library or dynamic link library (DLL), into another application are also detected and blocked by this interceptor. Buffer overflow attacks are detected by this interceptor. 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

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	?

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-11

When intercepting communications between applications and the underlying operating system, CSA combines the functionality of these traditional security approaches:

- **Distributed firewall:** The network interceptor performs the duties of a host firewall.
- **HIPS:** The network interceptor teams with the execution space interceptor to provide the alerting capability of HIPS with the proactive enforcement of a 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. A combination of the file system, configuration, and the execution space interceptors provides this security service.
- **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.

CSA Features

CSA features:

- Supports real-time enterprise-class protection decisions
- Provides defense in-depth approach
- Deploys and manages easily
- Supports many platforms and operating systems
- Provides enforce rule and detect rule organization
- Supports internationalization and localization for Microsoft Windows agents
- Integrates with the Cisco Trust Agent

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--9-12

Here are the key features of the CSA:

- CSA provides protection in real time before attacks have a chance to enter the network. Real-time correlation at the CSA level and at the enterprise level reduces false positives and allows the CSA to adapt to new threats across an enterprise. Here is a description of the CSA enterprise-level features:
 - A CSA can scan a network over multiple systems within a configured time period creating logs during significant events.
 - When a CSA discovers a worm event on multiple systems, the CSA will instruct all computer systems to quarantine the contaminated files.
 - The CSA can create network termination event logs and virus scanner logs that can be correlated across an entire enterprise.
- Here is how the CSA controls access to host system resources to provide a defense-in-depth solution:
 - The CSA intercepts communication between applications and the kernel.
 - The CSA protects a host system at these locations:
 - Network
 - File system
 - Configuration
 - Execution space

- The CSA can be installed, configured, and be on line with default policies quickly and allows easy configuration of custom policies. The CSA eases administration because there is no need for a constant review of logs; the CSA proactive defense approach minimizes the need for constant administrator involvement. There are no updates, and the CSA is always analyzing and interpreting traffic flows for malicious activity. When there is a need to administer the CSA, you can use a web browser to securely manage your CSA solution. The CSA provides centralized event management using e-mail, pager, Simple Network Management Protocol (SNMP) alerts controlled at the CSA MC, and logging and report-generating capabilities.
- The CAS supports Microsoft Windows or UNIX (Solaris and Linux)-based servers and desktops.
- With the CSA, combined rule lists are organized as a combination of enforce and detect rules. Enforce rules are primarily access control rules that allow, deny, or terminate an action. Detect rules are monitoring, logging, and tagging rules. In rule display lists, enforce rules are shown at the top of the list and detect rules are shown at the bottom. These rule types work together to monitor actions, build application classes, and protect systems.
- The CSA now accepts and displays query text characters appropriately for the selected language type. It also displays events in non-ASCII characters so that events generated by systems in other parts of the world can be incorporated.
- The CSA is a supported configuration for the Cisco Trust Agent feature. The Cisco Trust Agent is a core component of the Network Admission Control (NAC) solution. NAC is a Cisco led, multipartner program designed to limit damage caused by viruses and worms. This Cisco Trust Agent client software must be installed on hosts whose host policy state is to be validated before permitting network access. Cisco Trust Agent allows NAC to determine if the CSA or antivirus software is installed and current, and can determine current operating system and patch levels.

Positioning IPS Solutions

This topic explains how to select appropriate Cisco IPS solutions.

Cisco IPS Selection Considerations

- **Network media**
- **Intrusion detection analysis performance**
- **Network environment**
- **Number of sensors**
- **Sensor placement**
- **Management and monitoring options**
- **External sensor communication**

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0--5-13

Several factors affect the decisions that you make when selecting sensors for a Cisco IPS solution. Here are the technical factors to consider:

- **Network media:** Sensor selection is affected by the network media and environment. Cisco IPS sensor network interface cards (NICs) range from Ethernet to Gigabit Ethernet.
- **Intrusion detection analysis performance:** The performance of sensors is rated by the number of bits per second (bps) that can be captured and accurately analyzed. Cisco IPS sensor performance ranges from 45 Mbps to 1000 Mbps.
- **Network environment:** Cisco IPS sensors are suited for networks that have network speeds ranging from 10/100BASE-T Ethernet to Gigabit Ethernet.
- **Number of sensors:** Knowledge of your network topology helps you determine how many IDS and IPS appliances are required, the hardware configuration for each IDS and IPS appliance (for example, the size and type of network interface cards), and how many IDS and IPS management workstations are needed. The IDS and IPS appliance monitors all traffic across a given network segment. Given these facts, you should consider all of the connections to the network that you want to protect. Before you deploy and configure your IDS and IPS appliances, you should understand these factors of your network:
 - The size and complexity of your network
 - Connections between your network and other networks, including the Internet
 - The amount and type of network traffic on your network

- **Sensor placement:** It is recommended that sensors be placed at network entry and exit points to provide sufficient intrusion detection coverage. Determine network entry and exit points in your network to determine which segments of the network you want to monitor. Keep in mind that each IDS and IPS appliance maintains a security policy configured for the segment that it is monitoring. The security policies can be standard across the organization or unique for each IDS and IPS appliance. You may consider changing your network topology to force traffic across a given monitored network segment. When you are finished, you should have an approximate number of IDS and IPS appliances required to protect the desired network. You can place an IDS and IPS appliance in front of or behind a firewall. Each position has its benefits and drawbacks, which are summarized in the “IPS Best Practices” topic.
- **Management and monitoring options:** Review the management and monitoring options to select those that are most appropriate for your network. Keep in mind that the number of sensors that you deploy correlates directly to the type of management console that you select.
- **External sensor communication:** Traffic on the communication port between sensors and external systems must be allowed through firewalls to ensure functionality.

IPS Best Practices

This topic describes IPS best practices.

IPS Configuration Best Practices

- **When setting up a large deployment of sensors, automatically update signature packs rather than manually upgrading every sensor.**
- **Place the signature packs on a dedicated FTP server within the management network.**
- **Stagger the time of day when the sensors check the FTP server for new signature packs.**
- **Group IPS sensors together under a few larger profiles.**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-14

This discussion covers some configuration best practices that will improve IPS efficiency.

When setting up a large deployment of sensors, automatically update signature packs rather than manually upgrading every sensor. Security operations personnel will then have more time to analyze events. When new signature packs are available, download the new signature packs to a secure server within the management network.

Place the signature packs on a dedicated FTP server within the management network. If a signature update is not available, a custom signature may be created to detect and mitigate a specific attack. The FTP server should be configured to allow read-only access to the files within the directory on which the signature packs are placed, and then only from the account that the sensors will use. The sensors can then be configured to automatically check the FTP server periodically, such as once a week on a certain day, to look for the new signature packs and to update the sensors. A IPS can be used to protect this server from attack by an outside party.

Stagger the time of day when the sensors check the FTP server for new signature packs, perhaps through a predetermined change window. This will prevent multiple sensors from overwhelming the FTP server by asking for the same file at the same time. The need to upgrade sensors with the latest signature packs must be balanced against the momentary downtime—and, therefore, the vulnerability to attack—incurred while upgrading them. Finally, the signature levels supported on the management console must remain synchronized with the signature packs on the sensors themselves.

Group IPS sensors together under a few larger profiles. Every signature upgrade requires that all new signatures be appropriately tuned on every sensor. Tuning signatures for groups of sensors rather than for each sensor on the network significantly reduces configuration time. This administrative advantage must be balanced against the ability to finely tune sensor configuration by establishing a separate profile for each sensor.

Accommodating Network Growth

- Network growth can occur by adding additional hosts or new networks.
 - Additional hosts added to protected networks are covered without adding new sensors.
 - Additional sensors can easily be deployed to protect the new networks.
- Some of the factors that influence the addition of sensors are as follows:
 - Exceeded traffic capacity
 - Performance capabilities of the sensor
 - Network implementation

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-15

Network IPS gives security managers real-time security insight into their networks regardless of network growth. Network growth can occur by adding either additional hosts or new networks. Additional hosts added to protected networks are covered without adding any new sensors. Additional sensors can easily be deployed to protect the new networks. Some of the factors that influence the addition of sensors are as follows:

- **Exceeded traffic capacity:** For example, the addition of a new gigabit network segment requires a high-capacity sensor.
- **Performance capabilities of the sensor:** The current sensor may not be able to perform with the new traffic capacity.
- **Network implementation:** The security policy or network design may require additional sensors to help enforce security boundaries.

Scaling HIPS Systems

- **Deploy a central management console to maintain a database of policies and system nodes.**
- **HIPS agents installed on similar systems should be grouped together.**
- **Ensure that you place common HIPS hosts into groups based on your security plan.**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--9-16

HIPS implementations scale in a similar way as network IDS and IPS implementations. Here are some best practices when scaling a HIPS system:

- Deploy a central management console that is used to maintain a database of policies and system nodes. Each system node will have a HIPS agent installed.
- To streamline the process of assigning policies on many HIPS systems, the HIPS agents installed on similar systems should be grouped together. Servers that perform mission-critical roles benefit from being grouped together even if they perform different functions; their value to the enterprise is great, and they may, therefore, warrant special policies that are not applicable on other systems
- Ensure that you have a security plan in place before placing common HIPS hosts into groups. Grouping hosts together has the added benefit of applying a consistent set of policies across multiple host systems.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- The Cisco IPS solution runs on network modules, purpose-built appliances, and routers, and it is implemented in software.
- The CSA solution consists of the CSA MC, the CSA software, and an administration workstation. The CSA intercepts operating system calls. It then determines if the call should be passed to the kernel for execution or if the suspicious nature of the call warrants an action.
- Use these factors to select the best Cisco IPS solution for your needs:
 - Network media
 - Intrusion detection analysis performance
 - Network environment
 - Number of sensors
 - Sensor placement
 - Management and monitoring options
 - External sensor communication
- IPS best practices support IPS policies. The key is to reduce the effort required to manage your sensors while maximizing their ability to defend your network.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—5-17

Module Summary

This topic summarizes the key points that were discussed in this module.

Module Summary

- **IDS technology is passive; it monitors the network for suspicious activity and parsed system log files. IPS technology is reactive; it is able to forward or drop packets based on what is detected.**
- **Use the Cisco SDM to configure Cisco IOS IPS on Cisco network routers.**
- **Cisco IPS sensor platforms are used together to provide intrusion protection. Cisco IOS IPS is used in combination with the Cisco IDS, Cisco IOS Firewall, VPN, and NAC Cisco product. It provides superior threat protection at all entry points into the network.**

© 2006 Cisco Systems, Inc. All rights reserved.

SDM v2.0—5.1

This module described how intrusion detection system (IDS) and intrusion prevention system (IPS) technology embedded in Cisco host- and network-based IDS and IPS solutions fight Internet worms and viruses in real time. This module presented how you can easily configure a Cisco IOS IPS solution using the Cisco SDM GUI-based software. You have learned that the features and functions of the Cisco IPS product family provide a scalable and cost-effective way to integrate IPS into the access points of your network.

References

For additional information, refer to these resources:

- Cisco Systems, Inc. Cisco ASA 5500 Series Adaptive Security Appliance Platform and Module Datasheet.
http://www.cisco.com/en/US/products/ps6120/products_data_sheet0900aecd802930c5.html.
- Cisco Systems, Inc. Cisco Incident Control System.
http://www.cisco.com/en/US/products/ps6542/products_data_sheet0900aecd8033185b.html.
- Cisco Systems, Inc. Cisco IOS Intrusion Prevention System (IPS): Cisco IOS IPS Supported Signature List.
http://www.cisco.com/en/US/products/ps6634/products_white_paper0900aecd8039e2e4.shtml.
- Cisco Systems, Inc. Cisco IOS Intrusion Prevention System (IPS): Cisco IOS IPS Deployment Guide.
http://www.cisco.com/en/US/products/ps6634/products_white_paper0900aecd80327257.shtml.
- Cisco Systems, Inc. Cisco IOS IPS (Q&A).
http://www.cisco.com/en/US/products/ps6634/products_qanda_item0900aecd803137cc.shtml.
- Cisco Systems, Inc. Cisco IPS 4200 Series Sensors: Cisco IOS IPS.
http://www.cisco.com/en/US/products/hw/vpndevc/ps4077/products_data_sheet0900aecd80200749.html.
- Cisco Systems, Inc. Cisco Security Management Suite.
http://www.cisco.com/en/US/netsol/ns647/networking_solutions_package.html.
- Cisco Systems, Inc. Cisco Trust Agent.
<http://www.cisco.com/en/US/products/ps5923/index.html>.
- Cisco Systems, Inc. CiscoWorks Management Center for IPS Sensors 2.2.
<http://www.cisco.com/en/US/products/ps6680/index.html>.
- Cisco Systems, Inc. Configuring Cisco IOS Firewall Intrusion Detection System.
http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800ca7c6.html.
- Cisco Systems, Inc. Configuring Cisco IOS Intrusion Prevention System (IPS).
http://www.cisco.com/en/US/products/ps6350/products_configuration_guide_chapter09186a00804453cf.html.
- Cisco Systems, Inc. IDS Device Manager Configuration Tasks.
http://www.cisco.com/en/US/customer/products/sw/secursw/ps2113/products_installation_and_configuration_guide_chapter09186a008014a218.html.
- Cisco Systems, Inc. *Creating Custom Signature*.
http://www.cisco.com/en/US/products/hw/vpndevc/ps4077/products_configuration_guide_chapter09186a00803eb043.html.
- Cisco Systems, Inc. Cisco Router and Security Device Manager End-User Guides.
http://www.cisco.com/en/US/products/sw/secursw/ps5318/products_user_guide_list.html.

- Cisco Systems, Inc. Working with Signature Engines.
http://www.cisco.com/en/US/products/hw/vpndevc/ps4077/module_installation_and_configuration_guides_chapter09186a00801a0c28.html#wp787211.
- National Security Agency and Central Security Service. Security Configuration Guides.
<http://www.nsa.gov/snac/index.cfm?MenuID=scg10.3.1>.

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) What is the key difference between an IDS and an IPS? (Source: Introducing IDS and IPS)
- A) An IPS works in line in the data stream and in real time, while an IDS monitors copies of traffic off line.
 - B) An IPS allows some malicious traffic to enter the network, while an IDS stops malicious traffic before it can enter the network.
 - C) An IPS is passive monitoring technology, while an IDS is reactive monitoring technology.
- Q2) For each description of an IDS and IPS type listed in the table in the figure, identify the correct IDS and IPS sensor type. (Source: Introducing IDS and IPS) The instructor will use an animated slide to support this practice item.

Introducing IDS and IPS	
Description	Type
- Simple and reliable - Customized policies - Can detect unknown attacks	Policy-based
- Fewer false positives - Needs fine tuning	Signature-based
- Window to view attacks - Distract and confuse attackers - Slow down and avert attacks	Honey Pot-based
- Traffic profile must be constant - Can detect unknown attacks	Anomaly-based

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—5-2

- Q3) Which responses are the correct Cisco IOS IPS software-based IPS sensor attack responses? (Source: Introducing IDS and IPS)
- A) Attacker Alarm, Reset, Deny
 - B) Deny Inline, Deny TCP Promiscuous, Deny UDP Inline
 - C) Drop, Deny Attacker Inline, Deny Flow Inline

- Q4) Which attack response creates an ACL that denies all traffic from the IP address source of the attack? (Source: Introducing IDS and IPS)
- A) Attacker Alarm
 - B) Deny TCP Promiscuous
 - C) Deny Flow Inline
 - D) Deny Attacker Inline
- Q5) Describe the differences between HIPS and network IPS monitoring. (Source: Introducing IDS and IPS)

- Q6) What is a signature micro-engine used for? (Source: Introducing IDS and IPS)
- A) A signature micro-engine is used to implement the Cisco IOS IPS signatures found in SDFs.
 - B) A signature micro-engine is used by SDFs to apply Cisco IOS IPS signatures.
 - C) A signature micro-engine is an IDS-specific tool used to support policy-based signature monitoring.
- Q7) Describe the signature alarm types in the table. (Source: Introducing IDS and IPS) The instructor will use an animated slide to support this practice item.

Signature Alarm Type	Description
False positive	
False negative	
True positive	
True negative	

- Q8) Identify the enhanced Cisco IOS security services and describe the key feature for each. (Source: Configuring Cisco IOS IPS) The instructor will use an animated slide to support this practice item.

Cisco IOS IPS Signature Feature	Description
Regular expression string pattern matching	
Response actions	
Alarm summarization	
Threshold configuration	
Antievasive techniques	

- Q9) Which HIPS interceptor intercepts write requests to the Microsoft Windows registry? (Source: Defending Your Network with the Cisco IPS Product Family)

- A) file system interceptor
- B) configuration interceptor
- C) registry interceptor
- D) Microsoft Windows registry interceptor

- Q10) Describe some of the key features of the CSA. (Source: Defending Your Network with the Cisco IPS Product Family)

Module Self-Check Answer Key

Q1) A

Q2) The figure describes the sensor location for each type of network protection listed in the table.

Introducing IDS and IPS	
Description	Type
• Simple and reliable • Customized policies • Can detect unknown attacks	Policy-based
• Fewer false positives • Needs fine tuning	Signature-based
• Window to view attacks • Distract and confuse attackers • Slow down and avert attacks	Honey Pot-based
• Traffic profile must be constant • Can detect unknown attacks	Anomaly-based

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—6.2

Q3) C

Q4) D

Q5) The summary should touch on these points:

- HIPS operates by detecting attacks occurring on a host that it is installed on. HIPS works by intercepting operating system and application calls, securing the operating system and application configurations, validating incoming service requests, and analyzing local log files for after-the-fact suspicious activity.
- Network IPS involves the deployment of monitoring devices, or sensors, throughout the network to capture and analyze the traffic. Sensors detect malicious and unauthorized activity in real time and can take action when required. Sensors are deployed at designated network points that enable security managers to monitor network activity while it is occurring, regardless of the location of the attack target.

Q6) A

Q7) The figure describes the signature alarm types.

Introducing IDS and IPS (Cont.)	
Signature Alarm Type	Description
False positive	An alarm is triggered by normal traffic or a benign action.
False negative	A signature is not fired when offending traffic is detected.
True positive	A signature is correctly fired when offending traffic is detected and an alarm is generated.
True negative	A signature is not fired when nonoffending traffic is captured and analyzed.

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—5-3

Q8) The figure identifies and describes the enhanced Cisco IOS IPS signature features.

Defending Your Network with Cisco IOS IPS	
Cisco IOS IPS Signature Feature	Description
Regular expression string pattern matching	Enables the creation of string patterns using regular expressions
Response actions	Enables the sensor to take an action when the signature is triggered
Alarm summarization	Enables the sensor to aggregate alarms, to limit the number of times an alarm is sent when the signature is triggered
Threshold configuration	Enables a signature to be tuned to perform optimally in a network
Antievasive techniques	Enables a signature to defeat evasive techniques used by an attacker

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—5-4

Q9) B

Q10) The description should touch on these points:

- The CSA provides protection in real time before attacks have a chance to enter the network.
- The CSA controls access to host system resources to provide a defense-in-depth solution by intercepting communication between applications and the kernel.
- The CSA protects a host system at these locations:
 - Network
 - File system
 - Configuration
 - Execution space
- The CSA can be installed, configured, and be on line with default policies in 30 minutes and allows easy configuration of custom policies.
- The CAS supports Microsoft Windows or UNIX (Solaris and Linux)-based servers and desktops.
- With the CSA, combined rule lists are organized as a combination of enforce and detect rules.
- The CSA can incorporate events generated by systems in other parts of the world.
- The CSA is a supported configuration for the Cisco Trust Agent feature.

Building IPsec VPNs

Overview

An IPsec virtual private network (VPN) uses the Internet to connect branch offices, remote employees, and business partners to the resources of your company. It is a reliable way to maintain your company privacy while streamlining operations, reducing costs, and allowing for flexible network administration. This module describes the fundamental concepts, technologies, and terms used with VPNs. The module describes how to configure a site-to-site IPsec VPN and how to configure a site-to-site IPsec VPN with preshared key authentication using the Cisco Router and Security Device Manager (SDM). In this module, you will learn how to configure Cisco Easy VPN Server and Cisco Easy VPN Remote using the Cisco SDM. You will also learn about how to select the appropriate mix of Cisco security products to support your IPsec VPN solution. This module explains the best practices for deploying the hardware and software components of the Cisco VPN product family.

Module Objectives

Upon completing this module, you will be able to configure point-to-point and remote-access VPNs using Cisco IOS features. This ability includes being able to meet these objectives:

- Describe the fundamental concepts, technologies, and terms used with VPNs
- Describe how to configure a site-to-site IPsec VPN
- Explain how to configure a site-to-site IPsec VPN with preshared key authentication using Cisco SDM
- Describe how to configure Cisco Easy VPN Server and Cisco Easy VPN Remote
- Explain the best practices for deploying the hardware and software components of the Cisco VPN product family

Introducing IPsec VPNs

Overview

The IPsec virtual private network (VPN) is an essential tool for providing a secure network for business communication. This lesson explains how IPsec VPNs work by exploring the underlying technologies that support the tools themselves. IPsec protocol and Internet Key Exchange (IKE) protocol concepts and functions are described. This lesson also describes how IPsec provides message authentication and integrity checks, symmetrical and asymmetrical encryption methods, and public key infrastructure (PKI).

Objectives

Upon completing this lesson, you will be able to describe the fundamental concepts, technologies, and terms used with IPsec VPNs. This ability includes being able to meet these objectives:

- Describe the IPsec protocol, its basic functions, and the advantages of IPsec VPNs versus other types of VPNs
- Explain the IKE protocols
- Describe the additional functionality available within IKE
- Describe the ESP and AH protocols and the transport and tunnel modes used for IPsec
- Describe message authentication and integrity check
- Explain the difference between and the functionality of symmetric and asymmetric encryption algorithms
- Describe the PKI environment

IPsec Overview

This topic describes the IPsec protocol, its basic functions, and the advantages of IPsec VPNs versus other types of VPNs.

Introducing IPsec

IPsec has these features:

- It is an IETF standard (RFC 2401-2412).
- It defines how a VPN can be set up using the IP addressing protocol.
- It determines how the interface appears to the encryption protocol, not which type of encryption is used.
- It provides these essential functions:
 - Confidentiality
 - Integrity
 - Authentication

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0-4-3

IPsec is an Internet Engineering Task Force (IETF) standard (RFC 2401-2412) that defines how a VPN can be set up using the IP addressing protocol. The IPsec protocol determines how the interface on a router appears to the encryption protocol, not which type of encryption is used. IPsec provides these essential security functions:

- **Confidentiality:** IPsec ensures confidentiality by using encryption. Data encryption prevents third parties from reading the data, especially data that is transmitted over public networks or over wireless networks.
- **Integrity:** IPsec ensures that data arrives unchanged at the destination; that is, that the data is not manipulated at any point along the communication path. IPsec ensures data integrity by using checksums. A checksum is a simple redundancy check. The IPsec protocol adds up the basic components of a message (typically the number of bytes) and stores the total value. IPsec performs a checksum operation on received data and compares the result to the authentic checksum. If the sums match, the data is considered unmanipulated.
- **Authentication:** Authentication ensures that the connection is actually made with the desired communication partner. IPsec authenticates users (people) and devices that can carry out communication independently. These types of authentication are used by IPsec:
 - Username and password
 - One-time password
 - Biometric
 - Preshared keys
 - Digital certificates

Internet Key Exchange

This topic explains the IKE protocol.

Internet Key Exchange

- IPsec uses the IKE protocol to authenticate a peer computer and to generate encryption keys.
- The IKE protocol automates the key exchange process by:
 - Negotiating SA characteristics
 - Automatically generating keys
 - Automatically refreshing keys
 - Allowing manual configuration
- The IKE protocol uses these modes to secure communications:
 - Main mode
 - Aggressive mode
 - Quick mode

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—6.4

IPsec implements a VPN solution using an encryption process that involves the periodic changing of encryption keys. IPsec uses the IKE protocol to authenticate a peer computer and to generate encryption keys. IKE negotiates a Security Association (SA), which is an agreement between two peers engaging in an IPsec exchange, and consists of all required parameters necessary to establish successful communication.

IPsec uses the IKE protocol to provide these functions:

- Negotiation of SA characteristics
- Automatic key generation
- Automatic key refresh
- Manageable manual configuration

To establish a secure communication channel between two peers, the IKE protocol uses these three modes of operation:

- **Main mode:** In main mode, an IKE session begins with one computer (the initiator) sending a proposal or proposals to another computer (the responder). The proposal sent by the initiator defines what encryption and authentication protocols are acceptable, how long keys should remain active, and whether perfect forward secrecy should be enforced. Here is a description of the three exchanges typical of the main mode:
 - The first exchange between the initiator and the responder establishes the basic security policy. The responder chooses a proposal that is best suited to the security situation and then sends that proposal to the initiator.

- The next exchange passes Diffie-Hellman (DH) public keys between the two users. DH key exchange is a cryptographic protocol that allows two parties that have no prior knowledge of each other to establish a shared secret key over an insecure communications channel. All further negotiation is encrypted within the IKE SA.
- The third exchange authenticates an Internet Security Association and Key Management Protocol (ISAKMP) session. Once the IKE SA is established, IPsec negotiation (quick mode) begins.
- **Aggressive mode:** Aggressive mode compresses the IKE SA negotiation phases described thus far into three packets. In aggressive mode, the initiator passes all data required for the SA. The responder sends the proposal, key material, and ID and authenticates the session in the next packet. The initiator replies by authenticating the session. Negotiation is quicker, and the initiator and responder IDs pass in clear text.
- **Quick mode:** Quick mode IPsec negotiation is similar to aggressive mode IKE negotiation, except that negotiation is protected within an IKE SA. Quick mode negotiates the SA for the data encryption and manages the key exchange for that IPsec SA.

IKE Communication Negotiation Phases

IKE uses these phases to secure a communication channel between two peers:

- **IKE Phase 1:** Transform sets, hash methods, and other parameters are determined.
- **IKE Phase 1.5 (optional):** XAUTH protocol can be used to provide user authentication of IPsec tunnels within the IKE protocol to provide additional authentication of the VPN clients.
- **IKE Phase 2:** SAs are negotiated by ISAKMP, where quick mode is used. In this phase, the IPsec SAs are unidirectional.

© 2006 Cisco Systems, Inc. All rights reserved.

SNID v2.0—6-5

To establish a secure communication channel between two peers, the IKE protocol executes these phases:

- **IKE Phase 1:** In IKE Phase 1, two IPsec peers perform the initial negotiation of SAs. In this phase, the SA negotiations are bidirectional; data may be sent and received using the same encryption key. In IKE Phase 1, the transform sets, hash methods, and other parameters are determined. Optionally, IKE Phase 1 can include authentication in which each peer in the SA negotiation is able to verify the identity of the other. Even if the SA negotiation data stream between the two IPsec peers is compromised, there is little chance that the encryption keys could be decrypted.
- **IKE Phase 1.5 (optional):** Using the Extended Authentication (XAUTH) protocol, user authentication of IPsec tunnels within the IKE protocol can be used to provide additional authentication of the VPN clients. It is also possible to exchange other parameters between the two peers in the IPsec negotiation. For example, mode configuration can be used to provide parameters such as the IP address, the NetBIOS Name Service (NBNS), and the Domain Name System (DNS) address to one of the clients.
- **IKE Phase 2:** In IKE Phase 2, SAs are negotiated by the IKE process ISAKMP on behalf of other services, such as IPsec, that need encryption key material for operation. Quick mode is used for IKE Phase 2 SA negotiations. In this phase, the SAs used by IPsec are unidirectional; therefore, a separate key exchange is required for each data flow.

IKE: Other Functions

This topic describes additional functionality available within IKE.

IKE: Other Functions

These IKE functions are also available:


- NAT traversal
- NAT detection
- NAT traversal decision
- UDP encapsulation of IPsec packets
- UDP encapsulated process for software engines: Transport mode and tunnel mode ESP encapsulation
- Mode configuration option
- Extended Authentication

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0-6-6

The IKE protocol provides a number of functions that are used to verify if the peer device is still active, to pass IPsec through Network Address Translation (NAT) devices, or to exchange additional configuration parameters.

IKE: Other Functions (Cont.)

IPsec and NAT: The Problem




Without the IKE NAT traversal feature, a standard IPsec VPN tunnel with one or more NAT or Port Address Translation (PAT) points in the delivery path of an IPsec packet will not work. The reason is that there are no port numbers in the IPsec headers that can be used to create and maintain translation tables; IPsec encrypts the Layer 4 port information.

IKE: Other Functions (Cont.)

Need NAT traversal with IPsec over TCP and UDP

- NAT traversal detection
- NAT traversal decision
- UDP encapsulation of IPsec packets
- UDP encapsulated process for software engines



© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0-6-8

The solution is to use the IKE protocol NAT traversal feature with IPsec. The diagram shows how, by encapsulating IPsec packets in a UDP wrapper, the NAT traversal function enables IPsec traffic to travel through NAT or PAT devices in the network. NAT traversal works using these features:

- **NAT detection:** During IKE Phase 1 negotiation, two types of NAT detections occur before IKE quick mode begins: NAT support and NAT existence along the network path. To detect NAT support, the vendor ID string is exchanged with the remote peer. The remote peer sends a vendor ID string payload to its peer to indicate that this version supports NAT traversal. Thereafter, NAT existence along the network path can be determined.

NAT traversal enables an IPsec device to find any NAT device between two IPsec peers. To detect whether a NAT device exists along the network path, the peers send a payload with hashes of the IP address and port of both the source and destination address from each end of the communication path. The hashes are sent as a series of NAT discovery payloads. If, upon receipt, both ends recalculate the hashes and the hashes match the payload hash, each peer knows that no NAT device exists on the network path between them. If the payload hash and recalculated hashes do not match (that is, someone translated the address or port), each peer needs to perform NAT traversal to enable the IPsec packet to go through the network.

- **NAT traversal decision:** While IKE Phase 1 detects NAT support and NAT existence along the network path, IKE Phase 2 decides whether the peers at both ends will use NAT traversal. Quick mode SA payload is used for NAT traversal negotiation.


- **UDP encapsulation of IPsec packets:** In addition to allowing IPsec packets to traverse across NAT devices, UDP encapsulation addresses these incompatibility issues between IPsec and NAT and PAT:
 - **Incompatibility between IPsec Encapsulating Security Payload (ESP) and PAT:** If PAT finds a legislative IP address and port, it drops the ESP packet. To prevent this action from happening, UDP encapsulation is used to hide the ESP packet behind the UDP header. In this way, PAT treats the ESP packet as a UDP packet and processes the ESP packet as a normal UDP packet.
 - **Incompatibility between checksums and NAT:** In the new UDP header, the checksum value is always assigned as zero. This value prevents an intermediate device from validating the checksum against the packet checksum and resolves the checksum issue because NAT changes the IP source and destination addresses.
 - **Incompatibility between fixed IKE destination ports and PAT:** PAT changes the port address in the new UDP header for translation and leaves the original payload unchanged.
- **UDP encapsulated process for software engines—Transport mode and tunnel mode ESP encapsulation:** After the IPsec packet is encrypted by a hardware accelerator or a software cryptographic engine, a UDP header and a non-IKE marker (which is 8 bytes in length) are inserted between the original IP header and ESP header. The total length, protocol, and checksum fields are changed to match this modification.

Note	NAT keepalives can be used to keep the dynamic NAT mapping alive during a connection between two peers. NAT keepalives are UDP packets with an unencrypted payload of 1 byte. By default, there are no NAT keepalives sent.

- **IKE mode configuration option:** Using the IKE mode configuration option, it is possible to send system parameters such as an IP address, the DNS server name, the NBNS, and split tunnel attributes to one of the peers in an IPsec VPN tunnel. Typically, this peer is the client using a remote access VPN.
- **XAUTH:** XAUTH is based on the IKE protocol. XAUTH allows authentication, authorization, and accounting (AAA) methods to perform user authentication in an optional phase (IKE Phase 1.5). IKE Phase 1.5 occurs after the IKE authentication Phase 1 exchange. XAUTH does not replace the IKE protocol. The IKE protocol provides device authentication, and XAUTH provides user authentication. User authentication occurs after IKE protocol-based device authentication.

ESP and AH Protocols, Transport and Tunnel Modes

This topic describes the ESP and Authentication Header (AH) protocols and the transport and tunnel modes used for IPsec.




Here is a description of the two IP protocols used in the IPsec standard, the ESP and AH protocols:


- The ESP header (IP protocol 50) forms the core of the IPsec protocol. The ESP protocol used in conjunction with an encryption method or transform set makes the data flow difficult to decrypt. The diagram shows how the ESP protocol protects the data portion of the packet and, optionally, provides authentication of the protected data.
- The AH protocol (IP protocol 51) provides connectionless integrity and data origin authentication to IP packets but does not encrypt the IP data. The diagram shows the AH protocol protecting the entire data packet, including the address fields of the IP header. Used alone, the AH protocol provides weak protection to the IP header. This is because as the packet travels through the network, the AH protocol cannot stop the values in the IP header field from changing; the sender cannot tell which fields may have been changed. Consequently, the AH protocol is used with the ESP protocol to provide data encryption and tamper-aware security features.

Transport and Tunnel Mode

Transport Mode



Tunnel Mode



The IPsec protocol can be configured to use the transport mode or the tunnel mode to forward data across a network. These descriptions explain the transport and tunnel modes:

- **Transport mode:** IPsec transport mode is used when packet expansion during the forwarding of small packets is a concern. The diagram shows how the transport mode inserts the ESP header between the IP header and the next protocol or the transport layer of the packet. Using the transport mode here requires no additional IP header, so there is minimal packet expansion.


In transport mode, the IP addresses of the two network nodes whose traffic is being protected by IPsec are visible. The transport mode can be susceptible to traffic analysis. Transport mode can be deployed with either ESP or AH or with both protocols.

Note Transport mode works well with generic routing encapsulation because generic routing encapsulation already hides the addresses of the end stations by adding its own IP header.

- **Tunnel mode:** IPsec tunnel mode is used when packet expansion is not a large concern. Tunnel mode works by encapsulating and protecting an entire IP packet. The diagram shows a new IP header being added to the packet so that the packet can be successfully forwarded through the network. Tunnel mode encapsulates or hides the IP header of the original packet. Tunnel mode adds approximately an extra 20 bytes to the packet by adding a new IP header. The encrypting devices themselves own the IP addresses used in this new header. These IP addresses can be specified in the configuration of Cisco IOS routers. Tunnel mode may be employed with either ESP or AH or with both protocols.

Message Authentication and Integrity Check

This topic describes message authentication and integrity check.



For message authentication and integrity checking, the IPsec protocol uses a Hashed Message Authentication Code (HMAC). HMAC can be used with any iterative cryptographic hash function in combination with a secret shared key. Examples of an iterative cryptographic hash function include Message Digest 5 (MD5) and Secure Hash Algorithm 1 (SHA-1). The cryptographic strength of HMAC depends on the properties of the underlying hash function. However, HMAC also uses a secret key for calculation and verification of the message authentication values.

The diagram shows a sender using a private key and a message with a hash to create a HMAC. The message and the HMAC are sent over an insecure channel. The receiver uses the message, the private key, and a hash function to produce a hash output and tries to match the hash output with the HMAC received with the message.

MD5 and SHA-1

- MD5 produces a 128-bit message digest.
- SHA-1 produces a 160-bit message digest.
- IPsec protocol uses only the first 96 bits of the SHA-1 message digest.
- SHA-1 is computationally slower than MD5, but more secure.

© 2006 Cisco Systems, Inc. All rights reserved.


SND v2.0--6-12

The IPsec protocol uses the MD5 and SHA-1 algorithms to produce hashes. These hash functions take a variable-length input message and create a fixed-length message digest. The MD5 algorithm takes a variable-length message and produces a 128-bit message digest or hash. The MD5 algorithm is intended for digital signature application, where a large file must be compressed in a secure manner before being encrypted with a private key using a public key cryptosystem such as Rivest, Shamir, and Adleman (RSA). The SHA-1 algorithm can take a message less than 2^{64} bits as input, and produce a 160-bit output message digest. The SHA-1 algorithm is based on principles similar to those used in Message Digest 4 (MD4) and though computationally slower than the improved MD5 algorithm, SHA-1 is considered more secure than the MD5 algorithm.

Note	Only the first 96 bits of the message digest are used by the IPsec protocol.

Symmetric vs. Asymmetric Encryption Algorithms

This topic explains the difference between and the functionality of symmetric and asymmetric encryption algorithms.



The figure shows the differences between symmetric and asymmetric encryption. In symmetric encryption, the sender and the receiver are using the same secret key to encrypt and decrypt the message. The weakness in symmetric encryption is the secret key. Any user can obtain the secret key to crack the code. In asymmetric encryption, one key is used for encryption and another key is used for decryption. In this case, you must obtain both keys to crack the code.

Symmetric vs. Asymmetric Encryption Algorithms (Cont.)

Comparing key lengths required for asymmetric keys and symmetric keys

Symmetric Key Length	Asymmetric Key Length
80	1024
112	2048
128	3072
192	7680
256	15360

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--6-14

One way to compare the strength of two encryption methods is to examine the length of the keys that each method uses. The information in the table in the figure compares symmetric key length with the equivalent asymmetric key length. A symmetric algorithm using a 256-bit key is comparable to an asymmetric algorithm using a 15,360-bit key.

Symmetric vs. Asymmetric Encryption Algorithms (Cont.)

Comparing security levels of cryptographic algorithms

Security Level	Work Factor	Algorithms
Weak	$O(2^{40})$	DES, MD5
Legacy	$O(2^{64})$	RC4, SHA-1
Baseline	$O(2^{80})$	3DES
Standard	$O(2^{128})$	AES-128, SHA-256
High	$O(2^{192})$	AES-192, SHA-384
Ultra	$O(2^{256})$	AES-256, SHA-512

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6-15

The table in the figure shows the security level and the amount of computational work that it takes to crack a variety of symmetric encryption algorithms. The Work Factor column lists estimates of the number of hash computations that a computer would need to make to decipher the symmetric key. For example, a work factor of 2^{69} hash computations takes a computer with 10,000 custom application-specific integration circuits, each performing two billion hash operations per second, about one year to crack an RC4 or SHA-1 algorithm. It is an interesting comment on the resources of malicious attackers that a work factor of 2^{80} (Triple-Data Encryption Standard [3DES] cryptosystem, described in the next discussion) is considered a baseline level of encryption security.

Note The "O" in the Work Factor column is referred to as "big O" notation.

Symmetrical Key Encryption Algorithms

- **DES**
 - Uses a 56-bit key
 - Is considered outmoded and insecure
- **Triple-DES**
 - Uses a 168-bit key
 - Only provides baseline encryption protection
- **AES**
 - The 126-bit key version is deemed acceptable by the NSA for U.S. government nonclassified data.

The three important symmetric key encryption algorithms used by the IPsec protocol are as follows:

- Data Encryption Standard (DES) is a cipher selected in 1976 as an official Federal Information Processing Standard (FIPS) for the United States. DES use is widespread. The DES algorithm is a symmetric key encryption algorithm. When the DES algorithm was introduced, it had classified design elements and a relatively short key length. The secrecy surrounding DES created intense academic scrutiny and lead to a new understanding of block ciphers and block cipher cryptanalysis. Because of the 56-bit key size, DES is now considered to be insecure for many applications. DES keys have been broken in less than 24 hours. The DES cipher has been superseded by the Advanced Encryption Standard (AES). In some documentation, DES is referred to as the data encryption algorithm.
- 3DES (also DESede) is another symmetric key encryption algorithm based on the DES algorithm. 3DES uses the DES algorithm three times by performing a DES encryption, then a DES decryption, and then a DES encryption again. 3DES has a key length of 168 bits (formed from three 56-bit DES keys). Recalling the table comparing various symmetric key encryption algorithms, 3DES is considered to provide baseline encryption security.
- The AES symmetric key encryption algorithm is a block cipher adopted as an encryption standard by the U.S. government. AES has a fixed block size of 128 bits and a key size of 128, 192, or 256 bits. At the time of publication of this student guide, AES remains a secure encryption. The National Security Agency (NSA) reviewed commercial implementations of AES and has declared it secure enough for U.S. government nonclassified data. AES is also known as the Rijndael algorithm.

DH and RSA Asymmetric Encryption Algorithms

Diffie-Hellman key agreement protocol:

- The first practical method for establishing a shared secret over an unprotected communications channel
- Vulnerable to a man-in-the-middle attack because there is no requirement to authenticate the sender and receiver

RSA cryptosystem:

- Most popular asymmetric encryption system available
- Provides encryption and digital signatures for authentication
- RSA keys are typically 1024–2048 bits long

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6-17

Of the asymmetric algorithms available for use with the IPsec protocol, the DH key agreement and the RSA algorithm highlight how asymmetric algorithms work.


The DH key agreement protocol was the first practical method for establishing a shared secret over an unprotected communications channel. This protocol allows two users to exchange a secret key over an insecure medium without any prior secrets. The key pairs of the parties may be regenerated at each run of the protocol. The public keys may be certified so that the parties can be authenticated and so that you can specify a combination of available attributes.

The original DH key agreement protocol is vulnerable to a man-in-the-middle attack because there is no requirement to authenticate the sender and receiver. In a man-in-the-middle attack, an attacker intercepts the public value of the sender and sends its own public value to the recipient. The attacker intercepting the public key value of the sender must now remain “in the middle,” so the attacker can substitute its own public key when the recipient sends its key back to the sender. Once the attacker has the keys of both the sender and the recipient, the attacker can then decrypt messages. Newer versions of the DH key agreement protocol include the use of digital signatures and other protocol variants.

Ronald Rivest, Adi Shamir, and Leonard Adleman developed the RSA system in 1977. The RSA cryptosystem is the most popular public key cryptosystem available that offers both encryption and digital signatures (authentication). The security of the RSA system is based on the reasonable assumption that discovering the private key by factoring is too difficult for most attackers. Because RSA keys are typically 1024 to 2048 bits long, RSA remains a secure encryption protocol.

PKI Environment

This topic describes the PKI environment.



PKI provides a scalable, secure mechanism for distributing, managing, and revoking encryption and identity information in a secured data network. A PKI is composed of these entities:

- Computers (peers) communicating on a secure network
- At least one certificate authority (CA) (A CA grants and maintains certificates.)
- Digital certificates (A digital certificate can contain the certificate validity period, peer identity information, encryption keys used for secure communications, and the signature of the issuing CA.)
- An optional registration authority (RA) to offload the CA by processing enrollment requests
- A distribution mechanism using, for example, the Lightweight Directory Access Protocol (LDAP) or HTTP for certificate revocation lists

Every entity (either a person or a device) participating in a secured communication event is enrolled in the PKI process. In the PKI process, an entity generates an RSA private key and an RSA public key (a key pair) and has its identity validated by a trusted entity (also known as a CA or trust point). After each entity enrolls in a PKI, every peer in a PKI is granted a digital certificate that has been issued by a CA. When peers want to communicate securely, they negotiate a secured communication session and exchange digital certificates. Using the information in the digital certificate, each peer validates the identity of the other peer. At this point, the peers can establish an encrypted session using the RSA public keys contained in the digital certificate to decrypt their messages.

PKI Certificates



- **A PKI uses a CA to:**
 - **Manage certificate requests and issue certificates**
 - **Provide a centralized trusted source for key management**
 - **Provide a trusted source to validate identities and to create digital certificates**
- **The CA starts by generating its own public key pair and creates a self-signed CA certificate. Then the CA can sign certificate requests and begin peer enrollment for the PKI.**
- **Use a third-party CA vendor, or use the Cisco IOS certificate server for your own CA-signed certificates.**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6.19

A CA, or trust point, manages certificate requests and issues certificates to network devices participating in a PKI environment. A CA provides a centralized trusted source for key management for devices participating in a PKI. A CA becomes an explicitly trusted source for the receiver of communications from unknown sources to validate identities and to create digital certificates. Before any PKI operations can begin, the CA generates its own public key pair and creates a self-signed CA certificate. Thereafter, the CA can sign certificate requests and begin peer enrollment for the PKI. You can use a CA provided by a third-party CA vendor, or you can use an “internal” CA, such as the Cisco IOS certificate server.

Hierarchical CA Frameworks

A PKI allows a hierarchical CA framework supporting multiple CAs with these features:

- The root CA holds a self-signed certificate and an RSA key pair.
- Subordinate CAs enroll with either the root CA or with another subordinate CA.
- Each enrolled peer can validate the certificate of another enrolled peer.
- Multiple CAs provide users with added flexibility and reliability.
- A subordinate CA can be placed in a branch office, and the root CA can be placed at office headquarters.
- One CA can automatically grant certificate requests, while another CA can require only manually granted certificate requests.

© 2006 Cisco Systems, Inc. All rights reserved.


SND v2.0--6-20

A PKI can be set up in a hierarchical framework to support multiple CAs. At the top of the hierarchy is a root CA, which holds a self-signed certificate. The trust within the entire hierarchy is derived from the RSA key pair of the root CA. The subordinate CAs within the hierarchy can be enrolled with either the root CA or with another subordinate CA. Using these enrollment options, multiple tiers of CAs can be configured. If the peers within a hierarchical PKI share a trusted root CA certificate or a common subordinate CA, each enrolled peer can validate the certificate of another enrolled peer.

Multiple CAs provide users with added flexibility and reliability. For example, subordinate CAs can be placed in branch offices while the root CA is at office headquarters. Also, different certificate granting policies can be implemented. For example, one CA can automatically grant certificate requests, while another CA within the hierarchy requires that each certificate request be manually granted. A minimum of a two-tier CA hierarchy is recommended in these situations:

- A two-tier CA hierarchy is recommended for large and very active networks in which a large number of certificates are revoked and reissued. A multiple-tier CA helps to control the size of the certificate revocation list (CRL), a list of revoked or cancelled certificates.
- When online enrollment protocols are used, the root CA can be kept off line with the exception of issuing subordinate CA certificates. This scenario provides added security for the root CA.

PKI Certificates



© 2006 Cisco Systems, Inc. All rights reserved.


SND v2.0—6-21

Securely exchanging secret keys among users is only practical for small networks. Certificates are used to implement public key cryptography in most applications. A certificate is an electronic document that, typically, binds a user identity with a public key. A certificate is issued by a CA. The diagram shows a schematic of a X.509 v3 certificate. X.509 certificates are defined by the ITU (formerly known as the International Telegraph and Telephone Consultative Committee) as a standard certificate structure. The structure of an X.509 v3 digital certificate is shown in the diagram.

A certificate may be revoked if it is discovered that its related private key has been compromised, or if the relationship between an entity and a public key embedded in the certificate is discovered to be incorrect or has changed. To check the validity of a certificate, the certificate is compared to the CRL. Ensuring that the CRL is up to date and accurate is an essential function in a centralized PKI. Because checking the validity of a certificate is the responsibility of the user, this is a key point of failure in the PKI system. To mitigate this weakness, people who rely on digital certificates must know that they have ready access to an up-to-date CRL.

Another way to check certificate validity is to use the Online Certificate Status Protocol (OCSP). If the application supports OCSP, OCSP allows a client to request information about the validity of certificates from a CA. The information about the validity of a certificate itself will be digitally signed by the CA. The CA, using OCSP, provides real-time status information to the user with only the information about the validity of the certificate that the user requested. HTTP is the protocol used to transport the request and the response between a client and the CA.

PKI Message Exchange



In the sample PKI message exchange in the figure, when two people (Alice and Bob) want to engage in secure communications using PKI, they must obtain a CA-signed certificate. This is how the enrollment occurs between Alice, the CA, and Bob:

- Step 1** Alice generates an RSA key pair and requests the CA public key.
- Step 2** The CA sends its public key to Alice.
- Step 3** Alice generates a certificate request and forwards it to the CA (or the RA, if applicable). The CA receives the certificate enrollment request, and, depending on the network configuration, one of these options occurs:
 - Manual intervention is required to approve the request.
 - Alice is configured to automatically request a certificate from the CA, and human intervention is not required after the enrollment request is sent to the CA server.
- Step 4** After the request is approved, the CA uses the private key belonging to Alice and a hash algorithm to create a message digest. The message digest is signed using the CA private key.
- Step 5** The CA returns the completed CA-signed certificate to Alice. Alice stores the CA-signed certificate on a computer for safe keeping and on her router for efficient distribution of CA-signed certificates.
- Step 6** Now, when Alice sends a message to Bob, Bob can query the CA using the CA-signed certificate belonging to Alice and can verify the identity of Alice.

PKI Credentials

Storing PKI credentials:

- RSA keys and certificates
- NVRAM or eToken storage

eToken prerequisites:

- Cisco 871 Integrated Service Router; Cisco 1800, 2800, or 3800 Series Integrated Service Routers
- Cisco IOS Release 12.3(14)T image
- USB eToken supported by Cisco
- A Cisco K9 image

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6-23

PKI credentials, such as RSA keys and CA-signed certificates, are typically stored on a router (in NVRAM). There are other places PKI credentials can be stored. Selected Cisco platforms support smartcard technology implemented in hardware as a Universal Serial Bus (USB) key (eToken key). The eToken can securely store any type of file within its 32 KB of storage space. Configuration files that are stored on the eToken can be encrypted and accessed only via a user PIN. The router will not load the configuration file unless the proper PIN has been configured for secure deployment of router configuration files.

To use an eToken key, you plug it into the USB port on the router and log in to the eToken. After you have successfully logged into the eToken, you can copy files from the router on to the eToken via the **copy** command. By default, after the eToken is removed from the router, all associated RSA keys are removed and IPsec tunnels are not torn down until the next IKE negotiation period. These Cisco platform requirements are needed to use an eToken key:

- A Cisco 871 Integrated Services Router; Cisco 1800, Cisco 2800, or Cisco 3800 Series Integrated Services Routers
- At least a Cisco IOS Release 12.3(14)T image running on any of the supported platforms
- A USB eToken supported by Cisco
- A Cisco K9 image, which is a Cisco IOS security feature set

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- IPsec is an IETF standard that defines how a VPN can be set up using the IP addressing protocol. IPsec provides confidentiality, integrity, and authentication security functions.
- IPsec relies on the IKE protocol to provide the negotiation of SA characteristics, automatic key generation, the automatic refreshing of keys, and a way to manage the manual configuration of keys.
- The IKE protocol supports the verification of peer device activity, the passing of IPsec packets through NAT devices, and the exchange of additional configuration parameters between peer devices.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0-6-24

Summary (Cont.)

- Together the ESP and AH protocols provide an undecipherable data flow and a tamper-evident seal. The ESP and AH protocols can use the IPsec transport mode when packet size is a concern or the IPsec tunnel mode when packet expansion is not a concern.
- The IPsec protocol uses HMAC to provide an iterative cryptographic hash function. The strength of HMAC depends on the properties of the underlying hash function.
- IPsec uses symmetric and asymmetric encryption. In symmetric encryption the sender and the receiver use the same secret key; in asymmetric encryption, one key is used for encryption and another key is used for decryption.
- PKI provides a scalable, secure mechanism for distributing, managing, and revoking encryption and identity information in a secured data network.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0-6-25

Building a Site-to-Site IPsec VPN Operation

Overview

Building a site-to-site IPsec virtual private network (VPN) operation is an essential part of a plan to meet the security requirements of customers. In this lesson, you will learn how to configure a site-to-site IPsec VPN to securely connect two or more subnets over the Internet or an intranet.


Objectives

Upon completing this lesson, you will be able to describe how to configure a site-to-site IPsec VPN. This ability includes being able to meet these objectives:

- Describe the five steps of IPsec operation
- Explain the procedure to configure IPsec
- Describe the configuration of the ISAKMP parameters
- Describe the configuration to define the cryptographic access list
- Describe the configuration to apply the cryptographic map to the interface
- Describe the configuration to apply to the interface access list

Site-to-Site IPsec VPN Operations

This topic describes the five steps of IPsec operation.



IPsec VPN negotiation can be broken down into five steps, including Phase 1 and Phase 2 of Internet Key Exchange (IKE):

- Step 1** An IPsec tunnel is initiated by interesting traffic when host A sends “interesting” traffic to host B. Traffic is considered interesting when it travels between the IPsec peers.
- Step 2** In IKE Phase 1, the IPsec peers (routers A and B) negotiate the established IKE Security Association (SA) policy. Once the peers are authenticated, a secure tunnel is created using Internet Security Association and Key Management Protocol (ISAKMP).
- Step 3** In IKE Phase 2, the IPsec peers use the authenticated and secure tunnel to negotiate IPsec SA transforms. The negotiation of the shared policy determines how the IPsec tunnel is established.
- Step 4** The IPsec tunnel is created and data is transferred between the IPsec peers based on the IPsec parameters configured in the IPsec transform sets.
- Step 5** The IPsec tunnel terminates when the IPsec SAs are deleted or when their lifetime expires.

Configuring IPsec

This topic explains the procedure to configure IPsec.

Site-to-Site IPsec Configuration

- **Step 1: ISAKMP policy**
- **Step 2: IPsec transform set**
- **Step 3: Cryptographic access list**
- **Step 4: Create and apply the cryptographic map**
- **Step 5: Interface access list**

© 2006 Cisco Systems, Inc. All rights reserved.


SND v2.0—6-4

Here are the five steps used to configure a site-to-site IPsec VPN:

- Step 1** Establish the ISAKMP parameters that will define how the tunnel will be established by configuring the ISAKMP policy.
- Step 2** Define the IPsec transform set. The definition of the transform set defines the parameters used for the IPsec tunnel, and can include the encryption and integrity algorithms used.
- Step 3** Create a cryptographic access list. The cryptographic access list defines which traffic should be sent through the IPsec tunnel.
- Step 4** Create and apply a cryptographic map. The cryptographic map maps the previously configured parameters together and defines the IPsec peer devices. The cryptographic map is applied to the outgoing interface of the VPN device.
- Step 5** Configure the interface access list. Usually, there are some restrictions on the interface that is used for VPN traffic. For example, block all traffic that is not IPsec or IKE.

Site-to-Site IPsec Configuration—Phase 1


This topic describes the configuration of the ISAKMP parameters.



The first part of site-to-site IPsec configuration is to configure the ISAKMP parameters. In the topology, router 1 and router 2 are configured with a policy, policy 1, which employs preshared authentication using the Secure Hash Algorithm (SHA) hash function and Advanced Encryption Standard (AES) 128-bit encryption. In the example, you see that preshared authentication is used with the secret “SeCrEt” to the IPsec peer.

Site-to-Site IPsec Configuration—Phase 2

This topic describes the configuration to define the cryptographic access list.



The next part of the configuration defines the cryptographic access list. This access list contains a “permit” entry for the traffic that should be sent into the IPsec tunnel. If packets do not match the access list, they are not encrypted but they are not dropped. All IP traffic passing through the interface where the cryptographic map is applied is evaluated against the applied cryptographic map set. If a cryptographic map entry sees outbound IP traffic that should be protected and the cryptographic map specifies the use of IKE, an SA is negotiated with the remote peer according to the parameters included in the cryptographic map entry.

After the parameters are defined, they are mapped with the cryptographic map configuration. The cryptographic map (for example, VPN_To_R2) maps the configured access list with the transform set (IPsec parameters). Additionally, the cryptographic map defines the IP address of the IPsec peer.


Cryptographic map entries created for IPsec combine the needed configuration parameters of IPsec security associations, including these parameters:

- What traffic should be protected by IPsec (per a cryptographic access list)
- 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)
- (Optional) 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)

Cryptographic map entries with the same cryptographic map name (but different map sequence numbers) are grouped into a cryptographic map set.

Site-to-Site IPsec Configuration—Apply VPN Configuration

This topic describes the configuration to apply the cryptographic map to the interface.



In this step in configuring the site-to-site IPsec VPN, the cryptographic map is applied to the outgoing interface of the VPN tunnel. Configure the routing information needed to send packets into the tunnel.

All IP traffic passing through the interface where the cryptographic map is applied is evaluated against the applied cryptographic map set. If a cryptographic map entry sees outbound IP traffic that should be protected and the cryptographic map specifies the use of IKE, a SA is negotiated with the remote peer according to the parameters included in the cryptographic map entry.

Site-to-Site IPsec Configuration—Interface Access List

This topic describes the configuration to apply to the interface access list.

Site-to-Site IPsec Configuration—Interface Access List

- **When using only IPsec VPN on a router interface, block all traffic except the traffic that you want.**
- **Block unwanted traffic by enabling the IPsec protocol 50 for ESP, or protocol 51 for AH, and enable IKE to configure UDP on port 500.**
- **To pass IPsec traffic through a NAT or PAT device or both, be sure to permit UDP port 4500 or the correct TCP port.**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6-8

Finally, if you are using only IPsec VPN on a router interface, you need to block all unwanted traffic and allow the traffic that you want. To block unwanted traffic, define an access list and apply it to all incoming packets on your IPsec interface. To do this, enable the IPsec protocol (protocol 50 for Encapsulating Security Payload [ESP] or 51 for Authentication Header [AH]) and IKE (UDP port 500). If you need to pass IPsec traffic through a Network Address Translation (NAT) device or a Port Address Translation (PAT) device (or both), be sure to permit UDP port 4500 or the correct TCP port. IPsec NAT transversal is accomplished by encapsulating IPsec packets with a UDP header or TCP header.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- **IPsec VPN negotiation can be broken down into five steps:**
 - **Step 1: Interesting traffic**
 - **Step 2: IKE Phase 1**
 - **Step 3: IKE Phase 2**
 - **Step 4: IPsec tunnel created**
 - **Step 5: IPsec tunnel terminates**
- **The five steps to configure a site-to-site IPsec VPN are as follows:**
 - **Step 1: Configure the ISAKMP policy**
 - **Step 2: Define the IPsec transform set**
 - **Step 3: Create a crypto access list**
 - **Step 4: Create and apply a cryptographic map**
 - **Step 5: Configure the interface access list**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6-9

Summary (Cont.)

- **In the site-to-site IPsec configuration IKE Phase 1, ISAKMP parameters are configured.**
- **In the site-to-site IPsec configuration IKE Phase 2, the cryptographic access list is configured to permit entry of the traffic that you want. Packets that do not match the access list pass through the router unencrypted.**
- **When you apply the VPN configuration in the site-to-site IPsec configuration, the crypto map is applied to the outgoing interface of the VPN tunnel.**
- **Be sure to define an access list and apply it to all incoming packets on your IPsec interface. To pass IPsec traffic through a NAT or PAT device, be sure to permit UDP port 4500 or the correct TCP port.**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6-10

Configuring IPsec Site-to-Site VPNs Using Cisco SDM

Overview

This lesson describes the configuration steps that you use to implement an IPsec site-to-site virtual private network (VPN) using the Cisco Router and Security Device Manager (SDM).


Objectives

Upon completing this lesson, you will be able to describe the procedure to configure a site-to-site IPsec VPN with preshared key authentication using Cisco SDM. This ability includes being able to meet these objectives:

- Describe how to navigate the Cisco SDM site-to-site VPN wizard interface
- Describe the components that you configure using the Cisco SDM site-to-site VPN wizard
- Explain how to launch the Cisco SDM site-to-site VPN wizard
- Explain how to configure the site-to-site VPN tunnel connection settings using Cisco SDM
- Explain how Cisco SDM sets IKE policies
- Explain how to select a transform set and associate additional transform sets as required using Cisco SDM
- Explain how to define the traffic that the VPN protects using Cisco SDM
- Explain how to complete the site-to-site VPN configuration using Cisco SDM

Introducing the Cisco SDM VPN Wizard Interface

This topic describes how to navigate the site-to-site VPN wizard interface.



To select and start a VPN wizard, follow these steps:

- Step 1** Click the **Configure** icon in the top horizontal navigation bar to enter the configuration page.
- Step 2** Click the **VPN** icon in the left vertical navigation bar to open the VPN page.
- Step 3** Choose the **Site to Site VPN** wizard in the list in the middle section of the window.
- Step 4** Click the **Create Site to Site VPN** radio button in the section on the right side of the window and choose the VPN implementation subtype.
- Step 5** Click the **Launch the Selected Task** button to start the wizard.

Site-to-Site VPN Components

This topic describes the components that will be configured by the Cisco SDM site-to-site VPN wizards.

Site-to-Site VPN Components

- **VPN wizards use two sources to create a VPN connection:**
 - User input during step-by-step wizard process
 - Preconfigured VPN components
- **Cisco SDM provides some default VPN components:**
 - Two IKE policies
 - IPsec transform set for quick setup wizard
- **Other components are created by the VPN wizards.**
- **Some components (such as PKI) must be configured before the wizards can be used.**

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—6-4

The VPN wizards of the Cisco SDM use these two sources to create a VPN connection:

- User input during step-by-step wizard process
- Preconfigured VPN components


Here are the default VPN components that Cisco SDM provides:

- Two Internet Key Exchange (IKE) policies
- IPsec transform set for the quick setup wizard

Other components are created by the VPN wizards during the step-by-step configuration process. Some components must be configured before the wizards can be used (for example, public key infrastructure [PKI]).

Site-to-Site VPN Components (Cont.)

- **Two main components:**
 - IPsec
 - IKE
- **Two optional components:**
 - **Group Policies for Cisco Easy VPN Server functionality**
 - **Public Key Infrastructure for IKE authentication using digital certificates**




The figure illustrates the VPN navigation bar, which contains two major sections.

- These VPN wizards at the top:
 - Site to Site VPN
 - Easy VPN Remote
 - Easy VPN Server
 - Dynamic Multipoint VPN
- These individual IPsec components below:
 - Main components:
 - IPsec
 - IKE
 - Optional components:
 - Group Policies (for easy VPN server functionality)
 - Public Key Infrastructure (for IKE authentication using digital certificates)
 - The VPN Key Encryption settings window appears if the Cisco IOS image on your router supports type 6 encryption, also referred to as VPN key encryption. You can use this window to specify a master key to use when encrypting VPN keys, such as preshared keys, Cisco Easy VPN keys, and Extended Authentication (XAUTH) keys. When encrypted, these keys will not be readable by someone viewing the router configuration file.

The VPN wizards are used to simplify the configuration of individual VPN components. On the other hand, the individual IPsec components section can be used later to modify some parameters that may have been misconfigured during the VPN wizard step-by-step configuration.

Launching the Site-to-Site VPN Wizard

This topic explains how to launch the Cisco SDM site-to-site VPN wizard.




1. Click the **Launch the Selected Task** button after selecting the VPN type.

Use a web browser to connect to an HTTP server of a router. Select the VPN wizard by choosing **Configure > VPN > Site to Site VPN**. Follow these steps to create and configure a classic site-to-site VPN:

- Step 1** Click the **Create a Site to Site VPN** radio button and click the **Launch the Selected Task** button.

Launching the Site-to-Site VPN Wizard (Cont.)




2. Choose the wizard mode:
 - a. Quick setup uses predefined IKE and IPsec policies.
 - b. Step-by-step setup includes IKE and IPsec policy configuration steps.
3. Proceed to the configuration of parameters.

- Step 2** A window will pop up asking you which wizard mode to use:
- The “Quick setup” option uses Cisco SDM-default IKE policies and IPsec transform sets.
 - The “Step by step wizard” option allows the administrator to specify all of the details.
- Step 3** Click the **Next** button to configure the parameters of the VPN connection.

Quick Setup

The quick setup requires a single window to complete the configuration of the VPN.



- Configure all parameters on one page.

The quick setup includes these parameters:

- Outside interface
- IP address of the peer
- Authentication method:
 - Preshared keys (specify the secret)
 - Digital certificates (choose a certificate that should have been created beforehand)
- Traffic to protect:
 - Coming from IP subnet configured on the selected source interface
 - Going to defined remote IP subnet

Step-by-Step Setup

The step-by-step wizard requires multiple steps to configure the VPN connection.

Step-by-Step Setup

Multiple steps used to configure the VPN connection:

- **Defining connection settings:** Outside interface, peer address, authentication credentials
- **Defining IKE proposals:** Priority, encryption algorithm, HMAC, authentication type, Diffie-Hellman group, lifetime
- **Defining IPsec transform sets:** Encryption algorithm, HMAC, mode of operation, compression
- **Defining traffic to protect:** Single source and destination subnets, ACL
- **Reviewing and completing the configuration**

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0-6.9

The step-by-step wizard includes these parameters:

- **Connection settings:** Outside interface, peer address, and authentication credentials
- **IKE proposals:** IKE proposal priority, encryption algorithm (Data Encryption Standard [DES], Triple-DES[3DES], Advanced Encryption Standard [AES], or Software Encryption Algorithm [SEAL]), Hashed Message Authentication Code (HMAC), Secure Hash Algorithm 1 (SHA-1) or Message Digest 5 (MD5), IKE authentication method (preshared secrets or digital certificates), Diffie-Hellman (DH) group (1, 2, or 5), and IKE lifetime
- **IPsec transform sets:** Encryption algorithm (DES, 3DES, AES, or SEAL), HMAC (SHA-1 or MD5), mode of operation (tunnel or transport), and compression
- **Traffic to protect:** Defining single source and destination subnets or an access control list (ACL) for more complex VPNs

The last task of the step-by-step wizard is reviewing and completing the configuration.

Connection Settings

This topic explains how to configure the site-to-site VPN tunnel connection settings using Cisco SDM.

Connection Settings

1. Choose the outside interface toward the IPsec peer over the untrusted network.
2. Specify the IP address of the peer.
3. Choose the authentication method and specify credentials.
4. Proceed to the next task.

© 2006 Cisco Systems, Inc. All rights reserved. SDM v2.0—6-10

The first task in the step-by-step wizard is setting the connection settings. Follow these steps:

- Step 1** Choose the outside interface toward the IPsec peer over the untrusted network.
- Step 2** Specify the IP address of the peer.
- Step 3** Choose the authentication method and specify credentials. Use long and random preshared keys to prevent brute-force and dictionary attacks against IKE.
- Step 4** Click the **Next** button to proceed to the next task.

IKE Proposals

This topic explains how Cisco SDM sets IKE policies.

IKE Proposals

1. Use the IKE proposal predefined by Cisco SDM.

2. Add a custom IKE proposal.

3. Proceed to the next task.

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0-6.11

The second task in the step-by-step wizard is configuring IKE proposals. Follow these steps:

- Step 1** You can use the IKE proposal predefined by Cisco SDM.
- Step 2** If you want to use a custom IKE proposal, define it by clicking the Add button and specifying these required parameters:
- IKE proposal priority
 - Encryption algorithm
 - HMAC
 - IKE authentication method
 - DH group
 - IKE lifetime
- Step 3** When you are finished with adding IKE policies, click the **Next** button to proceed to the next task.

Transform Set

This topic explains how to select a transform set and associate additional transform sets to the VPN connection using Cisco SDM.

Transform Set

1. Use the IPsec transform set predefined by Cisco SDM.

2. Add a custom transform set.

3. Proceed to the next task.

© 2006 Cisco Systems, Inc. All rights reserved. SDM v2.0—6-12

The third task in the step-by-step wizard is configuring a transform set. Follow these steps:

- Step 1** You can use the IPsec transform set predefined by Cisco SDM.
- Step 2** If you want to use a custom IPsec transform set, define it by clicking the Add button and specifying these parameters:
- Transform set name
 - Encryption algorithm
 - HMAC
 - Mode of operation
 - Optional compression
- Step 3** When finished, click the **Next** button to proceed to the next task.

Defining What Traffic to Protect

This topic explains the two options you have to define the traffic that the VPN protects using Cisco SDM.

Option 1: Single Source and Destination Subnet

To define the traffic that needs protection, you can use the simple mode, allowing protection of traffic between one pair of IP subnets.

Option 1: Single Source and Destination Subnet

• Use the top option if protecting traffic between a single pair of subnets.

To protect traffic between a particular pair of IP subnets, follow these steps:

- Step 1** Choose the **Protect All Traffic Between the Following Subnets** option.
- Step 2** Define the IP address and subnet mask of the local network where IPsec traffic originates.
- Step 3** Define the IP address and subnet mask of the remote network where IPsec traffic is sent.

Option 2: Using an ACL

Alternatively, you can define a more complex ACL to identify which traffic to protect.

Option 2: Using an ACL

- Use the bottom option if protecting traffic between multiple subnets.


© 2006 Cisco Systems, Inc. All rights reserved. SMD v2.0—6.14

To specify an IPsec rule that defines the traffic types to be protected, follow these steps:

- Step 1** Choose the **Create/Select an Access-List for IPsec Traffic** option.
- Step 2** Click the ... button on the bottom-right side of the window to choose an existing ACL or to create a new one.
- Step 3** If an ACL you would like to use already exists, choose the **Select an Existing Rule (ACL)** option. If you would like to create a new ACL, choose the **Create a New Rule (ACL) and Select** option.

Option 2: Using an ACL (Cont.)

1. Create an IPsec rule.
2. Add entries to the rule.



When creating a new ACL to define traffic that needs protection, you will be presented with a window listing the created access rule entries. To create a new rule, follow these steps:

- Step 1** Give the access rule a name and description.
- Step 2** Click the **Add** button to start adding rule entries.

Option 2: Using an ACL (Cont.)

1. Choose an action and add a description.
2. Define source and destination networks or addresses.
3. (Optional) Define the protocol and port numbers.

© 2006 Cisco Systems, Inc. All rights reserved. SMD v2.0—6-16

Follow these steps to configure a new rule entry:

- Step 1** Choose an action and write a description of the rule entry.
- Step 2** Each rule entry defines one pair of source and destination addresses or networks.

Note You must use wildcard bits instead of subnet masks.

- Step 3** Optionally, you can provide protection for individual Open Systems Interconnection (OSI) Layer 4 protocols by choosing the required protocol radio box (TCP or UDP) and the required port numbers. If the rule applies to all IP traffic, leave the default radio box setting (IP).

Completing the Configuration

This topic explains how to complete the site-to-site VPN configuration using Cisco SDM.

Review the Generated Configuration

Hash	DH Group	Authentication	Encryption
SHA_1	group5	PRI_SHARE	AES_256
SHA_1	group2	PRI_SHARE	3DES

- Click Back to modify the configuration.
- Click Finish to complete the configuration.

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—6-17

At the end of the configuration, the wizard will present a summary of all the configured parameters. You can go back to modify the configuration in case you have made a mistake. Click the **Finish** button to complete the configuration.

Testing the Tunnel Configuration and Operation

This subtopic explains how to see the status of the site-to-site tunnel that you created.

Test Tunnel Configuration and Operation

• Check the VPN configuration and status.

You can click the Test Tunnel button to run the test to determine the configuration of the tunnel. You can also click the Generate Mirror button to generate a mirroring configuration that is required on the other end of the tunnel. This is typically useful if the other router does not have Cisco SDM and if you have to use the command-line interface (CLI) to configure the tunnel.

Monitoring Tunnel Operation

This subtopic describes the monitoring page that can be used to display the status of the tunnel.

Monitor Tunnel Operation

1

2

3

VPN Status

IPsec Tunnels DMVPN Tunnels Easy VPN Server IKE SAs

Each row represents one IPsec Tunnel

Local IP	Remote IP	Peer	Tunnel Status	Encapsulation f	Decapsulation f	Send Error Pack	Received Error
172.31.235.21	10.2.1.3	10.2.1.3:500	Up	4	105	0	0

Test Tunnel Update

Monitor the router 23:28:51 UTC Fri May 01 2002

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0—6-19

- List all IPsec tunnels, their parameters, and status.

To see all IPsec tunnels, their parameters, and status, follow these steps:

- Step 1** Click the **Monitor** icon in the top horizontal navigation bar.
- Step 2** Click the **VPN Status** icon in the left vertical navigation bar.
- Step 3** Click the **IPsec Tunnels** tab.

Advanced Monitoring

The basic Cisco IOS web interface also allows administrators to use the web interface to enter Cisco IOS CLI commands to monitor and troubleshoot the router. This subtopic explains two useful show commands.

router#

show crypto isakmp sa

- Lists active IKE sessions

router#

show crypto ipsec sa

- Lists active IPsec security associations

Central_site /level/15/exec/- - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Address http://10.1.1.1/level/15/exec/-

Central_site

Home Exec Configure

Command show crypto isakmp sa

Output

Command base-URL was: /level/15/exec/-

Complete URL was: /level/15/exec/-

Exec commands:

access-enable

Create a temporary Access-List entry

access-profile

Apply user-profile to interface

access-template

Create a temporary Access-List entry

archive

manage archive files

auto

Exec level Automation

bfe

For manual emergency modes setting

Done

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6-20

The “Show Commands” table lists two of the most useful **show** commands to determine the status of IPsec VPN connections.

Show Commands

Command	Description
show crypto isakmp sa	To display all current IKE Security Associations (SAs), use the show crypto isakmp sa command in EXEC mode. QM_IDLE status indicates an active IKE SA.
show crypto ipsec sa	To display the settings used by the current SAs, use the show crypto ipsec sa command in EXEC mode. Nonzero encryption and decryption statistics can indicate a working set of IPsec SAs.

Troubleshooting

This subtopic explains troubleshooting and debugging practices.

Troubleshooting

```
router#  
debug crypto isakmp
```

- Debugs IKE communication
- Advanced troubleshooting performed using the Cisco IOS CLI
- Requires knowledge of Cisco IOS CLI commands

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0-6-21

You should use a terminal to connect to the Cisco IOS router if you want to use debugging commands to troubleshoot VPN connectivity.

The **debug crypto isakmp** EXEC command displays detailed information about the IKE Phase 1 and IKE Phase 2 negotiation processes.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Cisco SDM is a GUI for simplified management of security mechanisms on Cisco IOS routers.
- By default the Cisco SDM provides, two IKE policies and an IPsec transform set for the quick setup wizard.
- To launch the Cisco SDM site-to-site VPN wizard use a web browser to connect to an HTTP server of the appropriate router. Then using the Cisco SDM GUI, select the VPN wizard by choosing **Configure > VPN > Site to Site VPN**.
- Other components for a VPN connection are created by the VPN wizards during the step-by-step configuration process. Cisco SDM can be used to configure a simple site-to-site VPN in three ways:
 - Using the quick setup wizard
 - Using the step-by-step wizard
 - Configuring individual VPN components

© 2006 Cisco Systems, Inc. All rights reserved.

SDM v2.0—6-22

Summary (Cont.)

- The Cisco SDM step-by-step wizard helps you to configure the outside interface toward the IPsec peer over the untrusted network, the IP address of the peer and to configure the router to use long and random preshared keys; all necessary settings for a VPN tunnel connection.
- The Cisco SDM configures, or sets, IKE proposals using the step-by-step wizard. The Cisco SDM GUI allows you to select a predefined IKE proposal or to create a custom IKE proposal by specifying the individual IKE parameters.
- Define the traffic to protect using the simple mode and allow protection of traffic between one pair of IP subnets, or use an ACL to define a more complex set of proxy identities.
- Upon completing the configuration, the Cisco SDM converts the configuration into the Cisco IOS CLI format.

© 2006 Cisco Systems, Inc. All rights reserved.

SDM v2.0—6-23

Building Remote-Access VPNs

Overview

The Cisco Router and Security Device Manager (SDM) virtual private network (VPN) GUI interface makes the task of building and managing VPN servers and VPN remote clients straightforward. This lesson describes the Cisco Easy VPN Server and Cisco Easy VPN Remote client and explains how to configure and manage them.

Objectives

Upon completing this lesson, you will be able to describe how to configure Cisco Easy VPN Server and Cisco Easy VPN Remote solutions. This ability includes being able to meet these objectives:

- Describe the functions of the two components of Cisco Easy VPN
- Describe how to configure Cisco Easy VPN Server using the Cisco SDM wizard
- Describe how to manage Cisco Easy VPN Server connections using Cisco SDM
- Describe how to configure a router as a Cisco Easy VPN Remote client using Cisco SDM

Cisco Easy VPN

This topic describes the functions of the two components of Cisco Easy VPN.

Cisco Easy VPN Components

- **Cisco Easy VPN Server** enables Cisco IOS routers, Cisco PIX Security Appliances, and Cisco VPN 3000 Series Concentrators to act as VPN headend devices in site-to-site or remote-access VPNs. In this configuration, remote office devices use the Cisco Easy VPN Remote feature.
- **Cisco Easy VPN Remote** enables Cisco IOS routers, Cisco PIX Security Appliances, and Cisco VPN 3002 Hardware Clients or software clients to act as remote VPN clients.

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0-6.3


Cisco Easy VPN consists of these two components:

- **Cisco Easy VPN Server:** Cisco Easy VPN Server enables Cisco IOS routers, Cisco PIX Security Appliances, and Cisco VPN 3000 Series Concentrators to act as VPN headend devices in site-to-site or remote-access VPNs, where the remote office devices are using the Cisco Easy VPN Remote feature. Using this feature, security policies defined at the headend are sent to the remote VPN device, ensuring that connections have up-to-date policies in place before the connection is established.

A Cisco Easy VPN Server-enabled device can terminate IPsec tunnels initiated by mobile remote workers running VPN client software on PCs. This flexibility makes it possible for mobile and remote workers to access their headquarter intranet.

- **Cisco Easy VPN Remote:** Cisco Easy VPN Remote enables Cisco IOS routers, Cisco PIX Security Appliances, and Cisco VPN 3002 Hardware Clients to act as remote VPN clients. These devices can receive security policies from a Cisco Easy VPN Server, minimizing VPN configuration requirements at the remote location. This is a cost-effective solution that is ideal for remote offices with little information technology support or for large sites where it is impractical to individually configure multiple remote devices. This feature makes VPN configuration as easy as entering a password. This ease increases productivity and lowers costs because the need for local information technology support is minimized.

Remote Access Using Cisco Easy VPN



The figure shows a Cisco IOS router VPN gateway running the Cisco Easy VPN Server feature. A variety of remote Cisco IOS routers, Cisco VPN Software Clients can connect to the Cisco IOS router Easy VPN Server for access to the corporate intranet.

Here are the restrictions that apply when using Cisco Easy VPN Remote:

- **Required Cisco Easy VPN Server:** The Cisco Easy VPN Remote feature requires that the destination peer be a Cisco IOS Easy VPN Server or VPN concentrator that supports the Cisco Easy VPN Server feature. For a current list of the Cisco products supporting the Cisco Easy VPN Server feature, refer to <http://www.cisco.com/en/US/partner/products/sw/secursw/ps5299/index.html>.
- **Only Internet Security Association and Key Management Protocol (ISAKMP) policy group 2 supported on Cisco Easy VPN Server:** The Cisco VPN Client client/server protocol supports only ISAKMP policies that use Diffie-Hellman (DH) group 2 (1024-bit DH) Internet Key Exchange (IKE) negotiation; therefore, the Cisco Easy VPN Server being used with the Cisco Easy VPN Remote feature must be configured for a group 2 ISAKMP policy. The Cisco Easy VPN Server cannot be configured for ISAKMP group 1 or group 5 when being used with a Cisco Easy VPN client.
- **Transform sets supported:** To ensure a secure tunnel connection, the Cisco Easy VPN Remote feature does not support transform sets that provide encryption without authentication (**esp-des** and **esp-3des** sets) or transform sets that provide authentication without encryption (**esp-null esp-sha-hmac** and **esp-null esp-md5-hmac** sets).

Note The Cisco VPN Client client and server protocol does not support Authentication Header (AH) authentication, but it does support Encapsulating Security Payload (ESP).

- **Dial backup for Cisco Easy VPN Remote:** Line status-based backup is not supported in this feature.
- **Network Address Translation (NAT) interoperability support:** NAT interoperability is not supported in client mode with split tunneling.

Cisco Easy VPN Remote Connection Process

1. **Device authentication via ISAKMP**
2. **User authentication using IKE XAUTH**
3. **VPN policy push (using mode configuration)**
4. **IPsec Security Association creation**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--6-5

The figure lists the key steps used to initiate a connection between a Cisco Easy VPN Remote device and a gateway configured with Cisco Easy VPN Server.

Configuring Cisco Easy VPN Server

This topic describes how to configure Cisco Easy VPN Server using the Cisco SDM wizard.

Cisco Easy VPN Server Configuration Tasks for the Cisco Easy VPN Server Wizard

The Cisco Easy VPN Server Wizard includes these tasks:

- **Choosing the interface on which to terminate IPsec**
- **IKE policies**
- **Group policy lookup method**
- **User authentication**
- **Local group policies**
- **IPsec transform set**


© 2006 Cisco Systems, Inc. All rights reserved.

SDM v2.0—6-6

The Cisco SDM Easy VPN Server Wizard guides the administrator through a set of steps that include the configuration of these parameters:

- Choosing the interface on which to terminate IPsec tunnels
- Setting these IKE policies:
 - Encryption algorithm
 - Hashed Message Authentication Code (HMAC)
 - Priority
 - Lifetime
 - DH group
- Choosing the IPsec transform set (that is, encryption algorithm, HMAC, and operation mode)
- Choosing local, RADIUS, or TACACS+ for the group policy lookup method
- Choosing the local or RADIUS user authentication method
- Setting the local group policies (for example, name, preshared secret, Domain Name System [DNS] servers, Microsoft Windows Internet Name Service [WINS] servers, split tunneling, and the IP address pool used to allocate an internal IP address to the client)


Starting the Cisco Easy VPN Server Wizard



The screen capture shows the Cisco SDM VPN configuration GUI. To launch the Cisco SDM Easy VPN Server Wizard, first choose the Easy VPN Server VPN type and then click the **Launch Easy VPN Server Wizard** button.


Note Authentication, authorization, and accounting (AAA) must be enabled to launch the Easy VPN Server Wizard.

Choosing an Interface for Terminating IPsec



In step 3, choose the outside interface which faces toward the IPsec peer over the untrusted network.


Configuring IKE Policies



To configure IKE policies, use the IKE proposal predefined by Cisco SDM or add a custom IKE proposal specifying these required parameters:

- IKE proposal priority
- DH group (1, 2, or 5)
- Encryption algorithm (Data Encryption Standard [DES], Triple-Data Encryption Standard [3DES], Advanced Encryption Standard [AES], or Software Encryption Algorithm [SEAL])
- HMAC (Secure Hash Algorithm 1 [SHA-1] or Message Digest 5 [MD5])
- IKE lifetime

Configuring IPsec Transform Sets



To configure the IPsec transform set, use the IPsec transform set that is predefined by Cisco SDM or add a custom IPsec transform set specifying these parameters:

- Transform set name
- Encryption algorithm (DES, 3DES, AES, or null)
- HMAC (SHA-1 or MD5)
- Optional compression
- Operation mode (tunnel or transport)


Configuring a Group Policy Configuration Location: Local Router Configuration



The Cisco Easy VPN Server allows you to group the remote users and set policies for all of the users in the same group. To configure the group policies, you have these three options that each configure where Cisco Easy VPN group policies will be stored:

- **Local:** This option means that all of the group configurations will be in the router configuration in NVRAM.
- **RADIUS:** This option means that the router will access a RADIUS server to look up group policies.
- **RADIUS and Local:** This option means that the router will also be able to look up group policies stored in a AAA server database reachable via RADIUS.
- **Select an existing AAA method list:** This option means that the router will be able to use a predefined AAA method list as a source for your group policies.

Configuring a Group Policy Configuration Location: External Location via RADIUS




The screen captures show the configuration steps needed to define a RADIUS server. In Step iii you configure these parameters:

- Specify the server IP address.
- Define RADIUS authentication and authorization ports (1645 and 1646 for Cisco Secure Access Control Server [ACS]; use 1812 and 1813 for other RADIUS servers).
- Cisco recommends that you use a key to authenticate individual RADIUS messages.

GUIs used to define a TACACS+ server are identical to those used to configure a RADIUS server. The parameters you configure in Step iii differ. These TACACS+ parameters are configured:

- Specify the server IP address.
- Cisco recommends that you enable the single-connection option if you are using Cisco Secure ACS to improve TACACS+ performance.
- Cisco recommends that you use a key to authenticate and encrypt the TACACS+ sessions.


Configuring a Local User Database: User Authentication



There are three options available to configure the location where the user records for Extended Authentication (XAUTH) (user authentication) will be stored:

- **Local Only:** This option means that all usernames and passwords will be in the router configuration in NVRAM. In Step 7a in the screen capture, if a local user database is used, you can immediately start adding users (Step i). In Step ii, use the default privilege level 1 for VPN users.
- **RADIUS and Local Only:** This option means that the router will also be able to authenticate users via RADIUS.
- **Select an Existing AAA Method List:** Alternatively, a previously configured AAA template can be selected for user authentication.

Configuring Local Group Policies



The screen capture shows where to configure the group policies. From this page, you can add a new group, edit an existing group, copy (clone) a group, or delete an existing group.

To edit a group policy, choose the desired group policy then click the **Edit** button.

Configuring Local Group Policy Parameters

Add Group Policy

General | DNS/WINS | Split Tunneling | Client Settings | XAuth Options

Group Information

Please enter a name of this group:

☒ **Pre-Shared Key**

Specify the key that will be used to authenticate the clients associated to this group.

Current Key: <None>

Enter new pre-shared key:

Re-enter new pre-shared key:

☒ **Pool Information**

Configure a local pool containing a range of addresses that will be used to allocate an internal IP address to a client.

☒ **Create a new pool**

IP Address Range: --


☐ **Select from an existing pool**

This figure shows the Add Group Policy form (similar to the Edit Group Policy form). These local group policy parameters can be configured:

- **General:** These parameters are the minimum required parameters for a functional group policy:
 - Group name
 - Group preshared secret
 - IP address pool to take addresses from to assign to clients
- **DNS/WINS:** Cisco recommends that you specify any internal DNS servers that may be required by clients to resolve host names that are only reachable inside the VPN. The same applies to WINS servers.
- **Split Tunneling:** Cisco recommends that you keep split tunneling disabled (the default setting) to prevent any compromised client PC from becoming a proxy between the Internet and the VPN. If split tunneling is required, be sure to complete one of these configuration options:
 - Define protected networks so that all other destinations will be reachable by bypassing the tunnel.
 - Use an existing access control list (ACL) or create a new ACL to configure split tunneling.
- **Client Settings:** Use this tab to define backup servers that will be available to clients. Using the options in this tab, you can enable polling for firewall status and allow local LAN access even if split tunneling is disabled.
- **XAUTH Options:** Use this tab to configure these user authentication parameters:
 - **Group lock:** Use this option to statically tie a user to a VPN group where users will have to use a group name as part of the XAUTH username.
 - **Enable save password:** This option means that the user saves a password in the VPN client.


- **Maximum logins:** This field defines the maximum number of concurrent logins. Use a maximum login number to prevent multiple users from sharing the same account at the same time.

Confirming Configuration Settings



At the end of configuration, the wizard will present a summary of all the configured parameters. You can go back to correct the configuration if you have made a mistake.


Testing the Cisco Easy VPN Server Configuration



Once the Cisco Easy VPN Server configuration is created, choose the VPN server name and click the **Test VPN Server** button. In the VPN Troubleshooting dialog box, click the **Start** button and observe the test results.

Managing Cisco Easy VPN Server Connections

This topic describes how to manage Cisco Easy VPN Server connections using Cisco SDM.



The screen captures show the Cisco SDM GUI used to view, add, edit, or delete Cisco Easy VPN Server connections.

Managing Cisco Easy VPN Server Connections (Cont.)

The image shows two side-by-side screenshots of Cisco SDM windows. The left window is titled 'Add Easy VPN Server Connection' and the right window is titled 'Edit Easy VPN Server Connection'. Both windows have a similar layout with the following fields:

- Select an Interface:** A dropdown menu. In the 'Add' window, it shows '-Select an entry'. In the 'Edit' window, it shows 'FastEthernet0/0'.
- Select an IPsec Policy:** A dropdown menu with an 'Add...' button. In the 'Add' window, it shows '-Select an entry'. In the 'Edit' window, it shows 'SDM_CMAP_1'.
- Select a Method List for Group Policy Lookup:** A dropdown menu. In the 'Add' window, it shows '-Select an entry'. In the 'Edit' window, it shows 'sdm_vpn_group_ml_1'.
- Enable User Authentication:** A checkbox that is checked in both windows.
- Select a Method List for User Authentication:** A dropdown menu. In the 'Add' window, it shows '-Select an entry'. In the 'Edit' window, it shows 'sdm_vpn_xauth_ml_1'.
- Mode Configuration:** Two radio buttons: 'Initiate' and 'Respond'. 'Respond' is selected in both windows.


At the bottom of each window are three buttons: 'OK', 'Cancel', and 'Help'.

Add Easy VPN Server

Edit Easy VPN Server


When you click the Add or Edit buttons, the options you can change are the same options that you used to create a Cisco Easy VPN Server connection using the Cisco SDM wizard.

Editing, Cloning, or Deleting Group Policies




The screen captures show the Cisco SDM GUI used to view, add, edit, clone, or delete Cisco Easy VPN Server group policies.

Creating or Editing a Local Pool for IP Addresses



The screen capture shows the dialog box used to edit the IP address pool information within the Cisco Easy VPN Server group policy.

Cisco VPN Client Software




Simple to deploy and operate, the Cisco VPN Client allows organizations to establish end-to-end, encrypted VPN tunnels for secure connectivity for mobile employees or teleworkers. This thin design, IPsec implementation is compatible with all Cisco VPN products. The Cisco VPN Client supports Microsoft Windows 98, Me, NT 4.0, 2000, and XP; Linux (Intel); Solaris (UltraSPARC 32- and 64-bit); and Mac OS X 10.2, 10.3, and 10.4. The Cisco VPN Client is compatible with these Cisco products:

- Cisco VPN 3000 Series Concentrator Software Version 3.0 and later
- Cisco IOS Release 12.2(8)T and later
- Cisco PIX Security Appliance Software Version 7.1 and later
- Cisco ASA 5500 Series Software Version 7.0 and higher

Note Customers with Cisco SMARTnet support contracts and encryption entitlement may download the Cisco VPN Client from the Cisco Software Center at no additional cost from <http://www.cisco.com/kobayashi/sw-center/vpn/client/>.

Configuring Cisco Easy VPN Remote


This topic describes how to configure a router as a Cisco Easy VPN Remote client using Cisco SDM.



The Cisco SDM GUI provides another wizard to configure a router as a Cisco Easy VPN Remote client. To start the wizard, click the **Launch Easy VPN Remote Wizard** button. The wizard steps you through these Cisco Easy VPN Remote client configuration tasks:

- **Configure the server information:** Provide the connection name and IP addresses of up to two Cisco Easy VPN Servers and choose a client or network extension operation mode.
- **Configure the device and user authentication:** For device authentication, provide the authentication method and enter a user group name and a secret key. User authentication configuration requires that you choose how you want to supply the XAUTH credentials each time that a user connects to the VPN server.
- **Specify interface and connection settings:** In this configuration step, you will choose the local interface that will connect to the networks behind a Cisco Easy VPN Server tunnel. For the connection settings, you will choose how you want the remote client to establish a VPN connection with the server. There are three options: automatically, manually, or when there is traffic from local networks (called interesting traffic).
- **Confirmation:** The wizard presents you with a dialog box to confirm your configuration.
- **Testing:** To ensure that you have correctly configured a Cisco Easy VPN Remote client, you should test your connection. This is done using the same procedure used to test a Cisco Easy VPN Server. Navigate to the **Edit Easy VPN Remote** tab, choose the new connection, click the **Test Tunnel** button, and click the **Start** button in the VPN troubleshooting dialog box.

Managing Cisco Easy VPN Remote Connections



The screen capture shows the Cisco SDM GUI used to view Cisco Easy VPN Remote connections and add, edit, or delete Cisco Easy VPN Server group policies. When you click the Add or Edit buttons, the options that you can change are the same options that you used to create a Cisco Easy VPN Remote connection using the Cisco SDM wizard. To reset a Cisco Easy VPN Remote connection, choose the connection and click the **Reset Connection** button.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- **Cisco Easy VPN consists of two components: Cisco Easy VPN Server and Cisco Easy VPN Remote.**
- **The Cisco Easy VPN Server Wizard easily configures the Cisco SDM Easy VPN Server.**
- **The Cisco SDM GUI manages Cisco Easy VPN Server connections.**
- **The Cisco Easy VPN Remote Wizard easily configures Cisco SDM Easy VPN Remote.**

© 2006 Cisco Systems, Inc. All rights reserved.

SDM v2.0—6-25

Defending Your Network with the Cisco VPN Product Family

Overview

A well-designed Cisco virtual private network (VPN) solution needs to provide private, ubiquitous communications to the locations and users that require it. This lesson describes the Cisco products that are available to support IPsec VPN solutions. The lesson covers the features of the product portfolio supporting Cisco VPN and how to select the appropriate products for your requirements. This lesson presents the VPN security features of the Cisco ASA 5500 Series Adaptive Security Appliances. In this lesson, the IPsec VPN features of the Cisco Router and Security Device Manager (SDM) are described. The lesson concludes with a number of IPsec VPN best practices.


Objectives

Upon completing this lesson, you will be able to explain the best practices for deploying the hardware and software components of the Cisco VPN product family. This ability includes being able to meet these objectives:

- Describe how secure connectivity is provided by VPNs
- Describe the security features of Cisco VPN products
- Describe optimum product positioning for a range of VPN requirements
- Describe IPsec VPN configuration best practices

Secure Connectivity—VPN Solutions


This topic describes how secure connectivity is provided by VPNs.




The type of VPN required is an important factor when deciding what kind of device best fits the needs of the VPN deployment. Here are the two applications for Cisco VPN-enabled devices:

- **Site-to-site VPN:** Site-to-site VPNs allow businesses to extend their network resources to branch offices, home offices, and business partner sites. All traffic sent between the sites is encrypted using IPsec, which provides network layer encryption for sensitive data passing across the VPN tunnel. This scenario includes the use of firewall-based VPN solutions using a router with a firewall and a VPN and also using a Cisco firewall with a security appliance.
- **Remote-access IPsec VPN:** IPsec VPN provides remote users with a robust remote-access environment by extending almost any data, voice, or video application available in the office to remote working locations, helping to create a user experience that emulates working in the main office location. Cisco IOS WebVPN provides Secure Sockets Layer (SSL) VPN-based remote-access connectivity from almost any Internet-enabled location using only a web browser and its native SSL encryption. Remote-access IPsec VPN enables companies to securely extend their enterprise networks to any authorized user by providing remote-access connectivity to corporate resources from any Internet-enabled location.

Site-to-Site VPNs



Site-to-site VPNs can be used to connect corporate sites. With Internet access, leased lines and Frame Relay lines can be replaced with site-to-site VPNs for network connection. VPN can support company intranets and business partner extranets. A site-to-site VPN is an extension of the classic WAN.




Remote-access VPNs are targeted to mobile users and home telecommuters. In the past, corporations supported remote users via dial-in networks, and gaining access to the corporate network often necessitated a toll or toll-free call. With the advent of VPNs, mobile users can use a dial-up or broadband connection to their Internet service provider (ISP) and then use IPsec to access the corporation via the Internet. Remote-access VPNs support the needs of telecommuters, mobile users, extranet consumer-to-business, and so on. The ubiquity of the Internet, combined with VPN technologies, allows organizations to cost-effectively and securely extend the reach of their networks to anyone, in any place, at any time.

VPNs have become the logical solution for remote-access connectivity because they provide these benefits:

- Secure communications with access rights tailored to individual users, including employees, contractors, and partners
- Enhanced productivity by extending corporate networks and applications
- Reduced communications costs and increased flexibility

Security Appliance-Based VPN Solutions



A firewall-based VPN solution is based on the capabilities of existing firewalls that can support both remote-access and site-to-site VPN requirements. Firewall-based VPN solutions are based more on management issues than on technical issues. The difference in the solution is in who manages the VPN network (either the owner or the ISP. If corporate security manages the VPN network, a firewall-based VPN may be the VPN solution of choice. Corporations can enhance their existing firewall systems to support VPN services.

Building Cisco IPsec VPNs

Product Choice	Remote-Access VPN	Site-to-Site VPN
Cisco VPN 3000 Series Concentrators	Primary role	Secondary role
Cisco VPN-Enabled Router	Secondary role	Primary role
Cisco PIX 500 Series security appliances	Enhance your existing Cisco PIX Security Appliance with the VPN remote-access solution	Security organization owns the VPN solution
Cisco ASA 500 Series Adaptive security appliances	Supports Cisco VPN 3000 features	Security organization owns the VPN solution

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0-6-7


Three product groups support VPN technology and are shown in the left column of the table in the figure. The top row of the matrix shows the two VPN applications. You can select the most appropriate Cisco product for your application using this matrix. For example, if your primary requirement is for a site-to-site VPN that allows for some remote access, a Cisco VPN-enabled router is the appropriate product choice. Similarly, if your primary need is to provide remote-access VPN with some site-to-site connectivity, a Cisco VPN 3000 Series Concentrator is the product of choice. The “VPN Products” table provides details about the available product choices.

VPN Products

VPN Application	Appropriate Cisco Product Choice
Dedicated VPN	Cisco VPN 3000 Series Concentrators for remote access Cisco 7200 Series Routers
VPN-enabled routers series	Cisco Small Office/Home Office (SOHO) 70 Series Router and Cisco 800 Series Routers Cisco 1700 Series Integrated Services Routers and Cisco 2600 Series Multiservice Access Routers Cisco 3700 Series Multiservice Access Routers and Cisco 3600 Series Routers Cisco 1800 Series, Cisco 2800 Series, and Cisco 3800 Series Integrated Services Routers Cisco 7200 Series Routers and Cisco 7400 Series Routers Cisco Catalyst 6500 Series Switches or Cisco 7600 Series Routers
Firewall VPN	Cisco PIX 500 Series Security Appliances

Secure Connectivity—Cisco VPN Product Family

This topic describes the security features of Cisco VPN products.



The portfolio of the Cisco VPN product family includes remote and site-to-site Cisco IOS VPN and firewall security routers, Cisco Catalyst 6500 Series Switches with VPN service modules (not shown), Cisco PIX security appliances, and Cisco ASA 5500 Series Adaptive Security Appliances.

Cisco VPN Product Family (Cont.)

VPN products:

- Cisco VPN-Enabled routers and switches
- Cisco VPN 3000 Series Concentrators
- Cisco ASA 5500 Series Adaptive Security Appliances
- Cisco PIX 500 Series Security Appliances

Hardware acceleration:

- AIM
- Cisco IPsec VPN SPA
- SEP
- VAC

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0-6-9

Here are more details on the Cisco VPN product family:

- **Cisco VPN-enabled routers and switches:** Cisco VPN security routers and switches represent the best options for customers of all sizes looking to take advantage of their existing network infrastructures to deploy VPNs and security while integrating all services in a single device with the widest selection of WAN and LAN interfaces.
- **Cisco VPN 3000 Series Concentrators:** Cisco VPN 3000 Series Concentrators are the most feature-rich remote-access VPN platform from Cisco, offering solutions for the most diverse remote-access deployment scenarios. Cisco VPN 3000 Series Concentrators offers both IPsec and SSL VPN connectivity on a single platform without the expense of individual feature licensing. Customers can achieve significant cost savings while experiencing the advanced features required by contemporary remote-access VPN deployments.
- **Cisco ASA 5500 Series Adaptive Security Appliances:** The Cisco ASA 5500 Series Adaptive Security Appliances are all-in-one security appliances that deliver enterprise-class security and IPsec VPN to small and medium-sized businesses and large enterprise networks in a modular, purpose-built appliance. Cisco ASA 5500 Series Adaptive Security Appliances incorporate a wide range of integrated security services, including firewall, intrusion prevention system (IPS), and VPN in an easy-to-deploy, high-performance solution. By integrating VPN and security services, the Cisco ASA 5500 Series Adaptive Security Appliances provide secure VPN connectivity and communications.

- **Cisco PIX 500 Series Security Appliances:** Cisco PIX 500 Series Security Appliances provide robust, enterprise-class, integrated network security services, including stateful inspection firewall, deep protocol and application inspection, IPsec VPN, multivector attack protection, and rich multimedia and voice security. Cisco PIX 500 Series Security Appliances are ideal for clients looking for the best-of-breed firewall combined with comprehensive VPN support. Cisco PIX 500 Series Security Appliances are also an excellent option for organizations whose security policies recommend separate management of the security infrastructure, setting a clear demarcation between security and network operation.

To enhance the performance and offload the encryption task to specialized hardware, the Cisco VPN family of devices offers hardware acceleration modules.

- **Advanced integration module (AIM):** A broad range of Cisco routers can be equipped with AIM. The AIM modules are installed inside the router chassis and offload encryption tasks from the router CPU.
- **Cisco IPsec VPN Shared Port Adapter (Cisco IPsec VPN SPA):** The Cisco IPsec VPN SPA delivers scalable and cost-effective VPN performance for Cisco Catalyst 6500 Series switches and Cisco 7600 Series Routers. Using the Cisco 7600 Series/Catalyst 6500 Series Services SPA Carrier-400 (Cisco Services SPA Carrier-400), each slot of the Cisco Catalyst 6500 switch or Cisco 7600 Series Router can support up to two Cisco IPsec VPN SPAs.
- **Scalable Encryption Processing (SEP):** Cisco VPN 3000 Series Concentrators can be upgraded with Enhanced Scalable Encryption Processing (SEP-E) modules. The modules perform hardware encryption of Data Encryption Standard (DES) 0, Triple-Data Encryption Standard (3DES), and Advanced Encryption Standard (AES) traffic.
- **Cisco PIX Security appliance VPN Accelerator Card+ (VAC+):** The VAC+ delivers hardware acceleration up to 425 Mbps of DES, 3DES, or AES IPsec encryption throughput.

Cisco IOS VPN Enabled Routers

- **V3PN**
 - Quality of service
 - IP telephony and video
- **IPsec**
 - IPsec stateful failover
- **DMVPN**
- **IPsec and MPLS integration**
- **Cisco Easy VPN**



© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6-10

With Cisco routers running Cisco IOS software, organizations can easily deploy and scale site-to-site VPNs of any topology—from hub-and-spoke VPNs to the more complex, fully meshed VPNs. In addition, the Cisco IOS security features combine the VPN feature set with firewall, intrusion prevention, and extensive Cisco IOS capabilities, including quality of service (QoS), multiprotocol, multicast, and advanced routing support.


The Cisco IOS feature sets incorporate these VPN features:

- Voice and Video Enabled VPN (V3PN) integrates IP telephony, QoS, and IPsec, providing an end-to-end VPN service that helps ensure the timely delivery of latency-sensitive applications such as voice and video.
- IPsec stateful failover provides fast and scalable network resiliency for VPN sessions between remote and central sites. With both stateless and stateful failover solutions available, options such as dead peer detection (DPD), Hot Standby Router Protocol (HSRP), Reverse Route Injection (RRI), and Stateful Switchover (SSO) help ensure maximum uptime of mission-critical applications.
- Dynamic Multipoint VPN (DMVPN) enables auto-provisioning of site-to-site IPsec VPNs, combining three Cisco IOS software features: Next Hop Resolution Protocol (NHRP), multipoint generic routing encapsulation, and IPsec VPN. This combination eases the provisioning challenges for customers and provides secure connectivity between all locations.
- IPsec and Multiprotocol Label Switching (MPLS) integration enables ISPs to map IPsec sessions directly into an MPLS VPN. This solution can be deployed on collocated edge routers that are connected to a Cisco IOS software MPLS provider edge network. This approach enables the ISP to securely extend its VPN service beyond the boundaries of the MPLS network by using the public IP infrastructure that securely connects enterprise customer remote offices, telecommuters, and mobile users from anywhere to the corporate network.

- Cisco Easy VPN simplifies VPN deployment for remote offices and teleworkers. The Cisco Easy VPN solution centralizes VPN management across all Cisco VPN devices, thus reducing the management complexity of VPN deployments.

Cisco VPN 3000 Series Concentrators

- Customized application access
- Cisco Secure Desktop
- Fully clientless Citrix support
- Integrated web-based management
- Clustering and load-balancing capabilities
- Broad user authentication support



© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6-11

Cisco VPN 3000 Series Concentrators are ideal for organizations that require advanced and flexible remote-access VPN technology and that prefer the operational simplicity and management segregation of a focused-function VPN device.

Here are some of the features of the Cisco VPN 3000 Series Concentrator platform:

- Customized application access with the Cisco WebVPN feature in Cisco VPN 3000 Series Concentrator Software v4.7 delivering clientless, thin client, and SSL tunneling client access methods.
 - Clientless access with Cisco WebVPN allows users to connect to a corporate network with few requirements beyond a basic web browser.
 - Thin client access with the Cisco WebVPN feature in Cisco VPN 3000 Series Concentrator Software v4.7 is achieved through a port forwarding mechanism enabled by a small Java applet download. Port forwarding relays data requested by the port on the local machine to the corresponding application port on the network side. This access method grants the user access to more applications and network resources than a web browser offers.
 - The Cisco SSL VPN Client for the Cisco WebVPN feature is a lightweight, centrally configured, and easy-to-support SSL VPN software client that allows access to virtually any application.
- The Cisco Secure Desktop is an endpoint security solution offering advanced endpoint security and data theft prevention.
- The Cisco VPN 3000 Series Concentrator platform gives fully clientless Citrix Systems support for terminal service environments without the need for any Cisco SSL VPN Client software. This support increases application performance and reduces endpoint software compatibility issues, providing users with rapid and highly stable system access, regardless of browser or security settings.
- The integrated web-based management system for the platform enables corporations to easily install, configure, and monitor their remote-access VPNs.

- The integrated clustering and load-balancing capabilities of the platform enable customers to scale their Cisco VPN 3000 Series Concentrator deployments to tens of thousands of users with low operational expense.
- The platform is capable of broad user authentication support, including single-use passwords, RADIUS, Active Directory, Security Dynamics International (SDI) Secure ID, digital certificates, and many others.

Cisco ASA 5500 Series Adaptive Security Appliances

Features of the Cisco PIX 500 Series Security Appliance plus advanced VPN features include:



- **Resilient clustering**
- **Cisco Easy VPN**
- **Cisco VPN Client updates**
- **Cisco IOS WebVPN**
- **VPN infrastructure for converged networks**
- **Integrated web-based management**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6-12

For VPN services, the Cisco ASA 5500 Series Adaptive Security Appliances offer flexible technologies that deliver tailored solutions to suit remote-access and site-to-site connectivity requirements. The Cisco ASA 5500 Series Adaptive Security Appliances provide easy-to-manage IPsec and SSL VPN-based remote-access and network-aware site-to-site VPN connectivity, enabling businesses to create secure connections across public networks to mobile users, remote sites, and business partners.

Cisco ASA Features 5500 Series Adaptive Security Appliances

The Cisco ASA 5500 Series Adaptive Security Appliances are multifunction security appliances with:

- **Adaptive identification and mitigation architecture**
- **Adaptive Threat Defense capabilities**
 - Application security
 - Anti-X defenses
 - Containment and control
- **Secure connectivity capabilities**
- **Remote access and Site-to-Site connectivity**
- **Converged security and VPN management**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--6-13


The Cisco ASA 5500 Series Adaptive Security Appliances form a high-performance, multifunction security appliance family delivering converged firewall, IPS, network antivirus, and VPN services. As a key component of the Cisco Self-Defending Network, Cisco ASA 5500 Series Adaptive Security Appliances provide proactive threat mitigation that stops attacks before they spread through the network, control network activity and application traffic, and deliver flexible VPN connectivity while remaining cost-effective and easy to manage.

Compared to Cisco PIX 500 Series Security Appliances, Cisco ASA 5500 Series Adaptive Security Appliances offer additional services, such as intrusion prevention, Cisco IOS WebVPN and AIM to enhance the processing capabilities of the appliances. Here are some of the features of Cisco ASA 5500 Series Adaptive Security Appliances:

- **Flexible platform:** Cisco ASA 5500 Series Adaptive Security Appliances offer both IPsec and SSL VPN on a single platform, eliminating the need to provide parallel solutions. In addition to VPN services, Cisco ASA 5500 Series Adaptive Security Appliances offer application inspection firewall and intrusion prevention services.
- **Resilient clustering:** Cisco ASA 5500 Series Adaptive Security Appliances allow remote-access deployments to scale cost-effectively by evenly distributing VPN sessions across all Cisco ASA 5500 Series Adaptive Security Appliances and Cisco VPN 3000 Series Concentrators without requiring any user intervention.
- **Cisco Easy VPN:** Cisco ASA 5500 Series Adaptive Security Appliances deliver uniquely scalable, cost-effective, and easy-to-manage remote-access VPN architecture. Cisco ASA 5500 Series Adaptive Security Appliances dynamically push the latest VPN security policies to remote VPN devices and clients, making sure that those endpoint policies are up to date before a connection is established.
- **Automatic Cisco VPN Client updates:** Cisco ASA 5500 Series Adaptive Security Appliances provide VPN client software auto update capabilities that enable automated version upgrades for Cisco VPN Client software operating on remote desktops.
- **Cisco IOS WebVPN** Cisco ASA 5500 Series Adaptive Security Appliances offer Cisco IOS WebVPN with clientless and thin client Cisco IOS WebVPN capabilities.

- **VPN infrastructure for contemporary applications:** Cisco ASA 5500 Series Adaptive Security Appliances provide a VPN infrastructure capable of converged voice, video, and data across a secure IPsec network by combining robust site-to-site VPN support with rich inspection capabilities, QoS, routing, and stateful failover features, allowing businesses to take advantages of the many benefits that converged networks deliver.
- **Integrated web-based management:** Cisco ASA 5500 Series Adaptive Security Appliances are managed via the integrated web-based Cisco Adaptive Security Device Manager (ASDM). Cisco ASDM manages all security and VPN functions of the appliances.

Positioning the Cisco ASA 5500 Series Adaptive Security Appliance Platforms



The graph shows the positioning of the Cisco ASA 5500 Series Adaptive Security Appliance platforms. This list shows which Cisco ASA 5500 Series Adaptive Security Appliance models are available and some of the features included:

- Cisco ASA 5510 Adaptive Security Appliance (3 Fast Ethernet interfaces, 50 VPN peers, 3DES, and AES)
- Cisco ASA 5510 Adaptive Security Appliance with Cisco ASA Advanced Inspection and Prevention Security Services Module (AIP SSM-10) (3 Fast Ethernet interfaces, 50 VPN peers, AIP-SSM-10, 3DES, and AES)
- Cisco ASA 5510 Security Plus Appliance (5 Fast Ethernet interfaces, 150 VPN peers, 3DES, and AES)
- Cisco ASA 5520 Adaptive Security Appliance (4 Gigabit Ethernet interfaces plus 1 Fast Ethernet interface, 300 VPN peers, 3DES, and AES)
- Cisco ASA 5520 Adaptive Security Appliance with AIP SSM-10 (4 Gigabit Ethernet interfaces plus 1 Fast Ethernet interface, 300 VPN peers, AIP-SSM-10, 3DES, and AES) interfaces
- Cisco ASA 5520 Adaptive Security Appliance with AIP SSM-20 (4 Gigabit Ethernet interfaces plus 1 Fast Ethernet interface, 300 VPN peers, AIP SSM-20, 3DES, and AES)
- Cisco ASA 5540 Adaptive Security Appliance (4 Gigabit Ethernet interfaces plus 1 Fast Ethernet interface, 500 VPN peers, 3DES, and AES)
- Cisco ASA 5540 Adaptive Security Appliance with AIP SSM-20 (4 Gigabit Ethernet interfaces plus 1 Fast Ethernet interface, 500 VPN peers, AIP SSM-20, 3DES, and AES)

Cisco ASA 5500 Series Adaptive Security Appliance Platforms

	5510 Security Plus	5520 VPN Plus	5540 VPN Plus	5540 VPN Premium
Simultaneous Web VPN (clientless) users	150	750	1250	2500
Site-to-site tunnels and remote access server (RAS) VPN peers	150	750	2000	5000
Encrypted throughput (Mbps)	170	225	325	325
Firewall throughput	300	450	650	650
Hardware encryption	Yes	Yes	Yes	Yes

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6-15

The table shows how the performance of Cisco ASA 5500 Series Adaptive Security Appliances depends on the platform feature license used. Here are the available licenses:

- Cisco ASA 5510 Adaptive Security Appliance: Base license and Security Plus license
- Cisco ASA 5520 Adaptive Security Appliance: Base license with VPN Plus add-on license
- Cisco ASA 5540 Adaptive Security Appliance: Base license with VPN Plus or VPN Premium add-on license

Cisco ASA 5500 Series Adaptive Security Appliances provide these feature, encryption, and platform licensing options:

- **Feature licenses:** These licenses are used to enable additional features such as security contexts and General Packet Radio Service (GPRS) tunneling protocol (GTP) inspection. The available feature licenses are as follows:
 - Security context licenses
 - GTP inspection license
- **Encryption licenses:** These licenses are used to extend the encryption capabilities to 3DES and AES. By default, there are 56-bit DES, 56-bit RC4, 512-bit Rivest, Shamir, and Adleman (RSA), and 512-bit Directory System Agent (DSA) encryptions available.
- **Platform licenses:** Various platform licenses are available based on the Cisco security appliance used.
 - **Cisco ASA 5510 Security Plus license:** This license increases port density on the platform by enabling the fourth Fast Ethernet port and removing restriction on the out-of-band (OOB) management port so that the port can be repurposed to a general traffic port if desired. Integration into switched network environments is simplified with this license, because support for up to 10 VLANs is enabled. Furthermore, this upgrade license enables active and standby high availability services and triples VPN capacity by supporting up to 150 concurrent VPN connections.

- **Cisco ASA 5520 VPN Plus license:** This license more than doubles the platform VPN capacity to support up to 750 concurrent VPN connections from mobile users, remote sites, and business partners.
- **Cisco ASA 5540 VPN Plus and VPN Premium licenses:** With a VPN Plus license, businesses quadruple the platform base VPN capacity to support up to 2000 concurrent IPsec VPN and 1250 Cisco IOS WebVPN connections from mobile users, remote sites, and business partners. The Cisco VPN Premium license maximizes the platform VPN capacity and offers 10 times the capacity of the base platform, supporting up to 5000 concurrent IPsec VPN and 2500 Cisco IOS WebVPN connections.

Cisco PIX 500 Series Security Appliances

- Spoke-to-spoke VPN support
- VPN NAT transparency
- Cisco VPN Client security posture enforcement
- Cisco VPN Client blocking by operating system and type
- OSPF dynamic routing Over VPN
- VPN hardware acceleration



© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6-16

The Cisco PIX 500 Series Security Appliances offer extensive features to deploy VPN service.


The features of Cisco PIX 500 Series Security Appliances with Cisco IOS Release 7.0 and above include these capabilities:

- **Enhanced spoke-to-spoke VPN support:** Support for spoke-to-spoke (and client-to-client) VPN communications allows encrypted traffic to enter and leave the same interface. Furthermore, split tunnel remote-access connections can now be terminated on the outside interface for the firewall allowing Internet-destined traffic from remote-access user VPN tunnels to leave on the same interface that it arrived on.
- **VPN Network Address Translation (NAT) transparency:** Cisco PIX 500 Series Security Appliances support Cisco TCP and UDP NAT traversal methods as methods complementary to existing support for the Internet Engineering Task Force (IETF) UDP wrapper mechanism for safe traversal through NAT and Port Address Translation (PAT) boundaries.
- **Cisco VPN Client security posture enforcement:** Cisco PIX 500 Series Security Appliances have the ability to perform Cisco VPN Client security posture checks when a VPN connection is initiated. Capabilities include enforcing use of authorized host-based security products (such as the Cisco Security Agent [CSA]) and verifying its version number, policies, and status (enabled or disabled).
- **Cisco VPN Client blocking by operating system and type:** Cisco PIX 500 Series Security Appliances can restrict different types of Cisco VPN Clients (Cisco VPN Software Client, , Cisco VPN 3002 Concentrator, and Cisco PIX security appliances) that are allowed to connect based on the type of client, operating system version installed, and Cisco VPN Client software version used. When noncompliant users attempt to connect, they can be directed to a group that specifically allows connections from noncompliant users.

- **Open Shortest Path First (OSPF) dynamic routing over VPN:** OSPF neighbors are supported across an IPsec VPN tunnel. This allows the CSA to support dynamic routing updates across a VPN tunnel to other OSPF peers. OSPF "hello packets" are unicast and encrypted for transport down the tunnel to an identified neighbor in an RFC-compliant manner.
- **VPN hardware acceleration:** Certain Cisco PIX 500 Series Security Appliance models have integrated hardware VPN acceleration capabilities. The Cisco VAC+ delivers up to 495 Mbps of DES, 3DES, or AES IPsec encryption throughput.

Secure Connectivity—VPN Product Positioning

This topic describes the optimum product positioning for a range of VPN requirements.



This figure shows the product placement considerations for Cisco VPN devices. In most networks, you will find some devices already in place. In this case, it is important to verify if interoperability between the different devices is possible. In a customer network, there may be a Cisco PIX 500 Series Security Appliance at one site and a Cisco router at another. A VPN tunnel can be established between the Cisco PIX 500 Series Security Appliance and the router as long as the software is at a minimum revision. The site-to-site VPN interoperability is given by choosing the following software releases (or later) for the example mentioned: Cisco IOS Release 12.2(8)T, Cisco PIX Security Appliance Version 7.2 and Cisco VPN 3000 Series Concentrator Software Version 3.x.

Cisco VPN Product Positioning

Cisco VPN Product Matrix

Site-to-Site VPN	IPsec Remote-Access VPN	SSL Remote-Access VPN
Cisco VPN-enabled router	Cisco ASA 5500 Series Adaptive Security Appliances	Cisco ASA 5500 Series Adaptive Security Appliances
Cisco ASA 5500 Series Adaptive Security Appliances	Cisco VPN 3000 Series Concentrators	Cisco VPN 3000 Series Concentrators
Cisco PIX 500 Series Security Appliances	Cisco VPN-enabled router	Cisco VPN-enabled router
Cisco VPN 3000 Series Concentrators	Cisco PIX 500 Series Security Appliances	

Products are ranked top to bottom

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--6-18

The table shows the products ranked from top to bottom based on performance, with the best product for the application at the top of each column.

Cisco VPN Product Positioning (Cont.)

	Remote Access	Site-to-Site	Cisco PIX 500 Series Security Appliance based	Cisco ASA 5500 Series Adaptive Security Appliance based
Large enterprise	Cisco VPN 3060 and 3080 Concentrators	Cisco Catalyst 6500, 7600 Series Switches Series Routers	PIX 535 Security Appliance	
Medium enterprise	Cisco VPN 3030 Concentrator	3700 Multiservice Access Routers, 3800 Series Integrated Service Routers, 7000 Series Routers	PIX 515E, 525 Security Appliances	ASA 5540, ASA 5520
Small business or remote office with branch office	Cisco VPN 3005 and 3015 Concentrators	1700, 1800; 2600 Series Multiservice Access Routers, 2800 Integrated Service Routers	PIX 506 Firewall, 515E Security Appliance	ASA 5510
SOHO market	Cisco VPN software and hardware Client	800 Series Routers, 1700 Series Integrated Services	PIX 501 Security Appliance, 506 Firewall	ASA 5510

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6-19

This table shows which Cisco device is best suited to the size of the network and the type of VPN application used.

Cisco VPN Best Practices

This topic describes Cisco VPN configuration best practices.

Cisco VPN Design Objectives

A Cisco IPsec VPN should emulate the functional requirements of your network. These design objectives should guide your decision making:

- **Secure connectivity**
- **Reliability, performance, and scalability**
- **Options for high availability**
- **Authentication of users and devices in VPN secure management**
- **Security and attack mitigation before and after IPsec**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--6-20

First and foremost, a well-designed Cisco VPN solution needs to provide private, ubiquitous communications to the locations and users that require it. It must do this in a secure manner while maintaining as many of the characteristics of traditional private WAN connections as possible. It must integrate with existing network designs based on Cisco Self-Defending Network security architecture. Although VPN design differs greatly with the size of enterprises, the underlying best practices remain virtually the same.

A Cisco IPsec VPN should emulate, as closely as possible, the functional requirements of your network. Implementation decisions vary, depending on the network functionality required. These design objectives, listed in order of priority, should guide your decision-making process:

- Secure connectivity
- Reliability, performance, and scalability
- Options for high availability
- Authentication of users and devices in the VPN secure management
- Security and attack mitigation before and after IPsec

Identity and IPsec Access Control Best Practices

- **Preshared keys**
 - **Group preshared keys are applicable only to remote access.**
 - **Do not use wildcard preshared keys for site-to-site device authentication.**
- **Digital certificates**
 - **Scale better than unique preshared keys**
 - **Use if the network of the VPN grows beyond 20 devices**
 - **Ensure that devices have the correct time of day**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6-21

In site-to-site and remote-access VPNs today, it is important that devices are identified in a secure and manageable way. In remote-access VPNs, user authentication and device authentication occur. When the remote device is authenticated, some level of access control needs to be in place to permit only permitted traffic over the tunnel. Device authentication uses either a preshared key or digital certificate to provide the identity of a device. The list that follows provides some best practices for managing identity and IPsec access control:

- **Preshared keys**
 - Group preshared keys are tied to a group name identity; currently, these are applicable only to remote access.
 - Wildcard preshared keys should not be used for site-to-site device authentication. When using wildcard preshared keys, every device in the network uses the same key. Therefore, if a single device in your network is compromised and the wildcard preshared key has been determined, a hacker can establish a tunnel with any device in the network.
- **Digital certificates**
 - Digital certificates scale better than unique preshared keys because they allow any device to authenticate to any other device, but digital certificates do not have the security properties of wildcard keys. Digital certificates are not tied to IP addresses; instead, they are tied to unique, signed information on the device that is validated by the certificate authority (CA) of the enterprise.
 - Consider using digital certificates if the size of the VPN grows beyond 20 devices, or sooner if there are requirements for strong device authentication.
 - Ensure that devices generating digital certificates or validating received certificates during tunnel authentication and establishment have the correct time of day configured, preferably Coordinated Universal Time (UTC).

Identity and IPsec Access Control Best Practices (Cont.)

- **Certificate revocation lists**
 - **Enable checking CRLs on remote and headend devices when digital certificates are deployed.**
 - **Consider a third-party managed CA when deploying an extranet VPN.**
- **Consider using a hardware-based solution to protect digital certificates and preshared key material.**
- **Use inbound ACLs on the VPN devices for site-to-site traffic.**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--6-22

- **Certificate revocation lists (CRLs)**
 - Checking CRLs should always be enabled on remote and headend devices when digital certificates are deployed.
 - Using a third-party-managed CA versus an enterprise-managed CA may help to facilitate deploying an extranet VPN.
- **Consider using a hardware-based solution for protecting the digital certificates and preshared key material because they are generally considered more secure than software implementations**
- **Use inbound access control lists (ACLs) on the VPN devices for site-to-site traffic. For remote-access traffic filtering, access control occurs dynamically by loading the per-user granular authorization information when the user successfully authenticates via Extended Authentication (XAUTH).**

IPsec Best Practices

- Use both encryption and integrity.
- Do not use single DES for data encryption.
- Use 3DES or AES for data encryption.
- Use SHA.
- Strong encryption algorithms cannot be exported to some countries or some customers.
- Do not change the SA lifetimes or to enable PFS unless the sensitivity of the data mandates it.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6-23

IPsec provides numerous security features. Here are IPsec best practices:

- Cisco highly recommends using both encryption and integrity.
- Cisco recommends that you do not use DES for data encryption.
- Cisco recommends the use of 3DES.
- Cisco recommends the use of Secure Hash Algorithm (SHA) because the increased security outweighs the slight cost of increased processor use. SHA is sometimes faster than Message Digest 5 (MD5) in certain hardware implementations.

Note The use of strong encryption algorithms outside the United States is sometimes regulated by local import and use laws. These strong encryption algorithms cannot be exported to some countries or to some customers. For more information, go to <http://www.cisco.com/www/export/crypto>.

- Regarding Security Association (SA) key aging, Cisco does not recommend changing SA lifetimes or enabling perfect forward secrecy (PFS) unless the sensitivity of the data mandates it. Changing the SA lifetimes or enabling PFS increases the level of security but at the same time increases processor overhead.

NAT Best Practices

- **Avoid the application of NAT to VPN traffic**
- **Use address ranges for your sites that do not overlap with other devices that you will connect via IPsec.**
- **When address translation occurs, make sure that a protocol-aware device carries out the address translation.**
- **Do not hide the public peer addresses of the VPN devices.**
- **When a remote-access client is not able to connect because of NAT-related issue, consider enabling NAT traversal mode.**
- **Use ESP tunnel mode and avoid NAT whenever possible.**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--6-24

These are some best practices for NAT implementation with an IPsec VPN:

- Use address ranges for your sites and remote-access VPN client virtual address pools that do not overlap with the addresses of other devices that you will connect to via IPsec. If this is not possible, use NAT only in this scenario to allow for connectivity.
- When address translation occurs, make sure that a protocol-aware device carries out the address translation, not only in the IP header but also in the data segment of the packet.
- Do not “address hide” the public peer addresses of the VPN devices, because this provides no real added security value and may cause connectivity problems.
- If you believe that NAT is involved when a remote-access client is not able to successfully establish a tunnel or send packets over an established tunnel, consider enabling NAT transparency mode.
- NAT transparency mode will not resolve the connection problems associated with client applications that are not NAT friendly.

This list explains the attributes of using NAT after IPsec:

- Applying NAT after IPsec encryption for address hiding provides no benefit because the actual IP addresses of the devices using the tunnel for transport are hidden via the encapsulation.
- Encapsulating Security Payload (ESP) and Authentication Header (AH) are higher-layer protocols on top of the IP that do not use ports.
- Use ESP tunnel mode and avoid NAT whenever possible.
- For remote access specifically, use NAT transparency mode when PAT is occurring.

The only scenario in which the application of NAT before IPsec is recommended is when two networks are connected via IPsec and the address ranges overlap. Using NAT before IPsec overcomes this restriction by translating one set of the overlapping networks into a unique network address range that will not interfere with the IPsec tunnel establishment.

Single-Purpose vs. Multipurpose Devices— Selection Considerations

- **When deciding which option to select, weigh your decision based on the capacity and functionality available on the appliance versus the functionality advantage of the integrated device.**
- **IPsec is a demanding function. As the size of the network increases, so does the likelihood that a VPN appliance needs to be selected over an integrated router or firewall.**

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6-25

At many points in the network design process you need to choose between using integrated functionality in a networking or security device versus using the specialized functions of a VPN appliance. The list that follows presents some factors to take into consideration when trying to select a single-purpose versus multipurpose device for your IPsec VPN solution:

- When deciding which option to select, weigh your decision based on the capacity and functionality available on the appliance versus the functionality advantage of the integrated device. For example, sometimes you can choose an integrated higher-capacity Cisco IOS router with IPsec encryption software as opposed to a smaller Cisco IOS router and an associated VPN device.
- IPsec is a demanding function. As the size of the network increases, so does the likelihood that a VPN appliance needs to be selected over an integrated router or firewall.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Cisco has a range of products to support site-to-site VPN, remote-access VPN, and remote-access web-based VPN solutions.
- The product portfolio supporting VPN consists of Cisco VPN-enabled routers, Cisco VPN 3000 Series Concentrators, Cisco PIX 500 Series Security Appliances, Cisco ASA 5500 Series Security Appliances, and Cisco Catalyst 6500 Series Switches.
- Placement of a VPN device depends on the functionality, the intended use, the supported features, and the required performance.
- A well-designed Cisco VPN solution needs to provide private, ubiquitous communications to the locations and users that require it.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0--6-26

Module Summary

This topic summarizes the key points that were discussed in this module.

Module Summary

- An IPsec VPN is a collection of protocols that help you to hook up your company private network to the public network. Ensure that you have proper security measures in place so that you select the best mix of encryption and authentication to support your IPsec VPN security policies.
- These five steps used to configure a site-to-site VPN:
 - Step 1 Establish the ISAKMP parameters
 - Step 2 Define the IPsec transform set.
 - Step 3 Create a cryptographic access list.
 - Step 4 Create and apply a cryptographic map.
 - Step 5 Configure the interface access list.
- Use the Cisco SDM wizard to help you to configure a site-to-site IPsec VPN with preshared key authentication using Cisco SDM.
- Use the Cisco SDM wizard to configure Cisco Easy VPN_servers and clients.
- Cisco has a range of products to support site-to-site VPN, remote-access VPN, and remote-access web-based VPN solutions. Placement of a VPN device depends on the functionality, the intended use, the supported features, and the required performance.

© 2006 Cisco Systems, Inc. All rights reserved. SND v2.0-4-1

Most businesses want remote locations to have full access to the company network for seamless communication and increased productivity. Virtual private networks (VPNs) offer a great opportunity for small and medium-sized businesses to improve operations and save money, but like all technology, VPNs need to be properly implemented. In this module, you learned about the fundamental concepts, technologies, and terms used with VPNs. You learned how to configure a site-to-site IPsec VPN using the Cisco Router and Security Device Manager (SDM). Cisco has a number of security products that are designed to support an IPsec VPN, and, in this module, you learned how to select the appropriate mix of Cisco products to support your requirements. This module provided you with a number of IPsec VPN best practices to ensure that you can implement a Cisco VPN solution to provide private, ubiquitous communications to the locations and users that require it.

References

For additional information, refer to these resources:

- Cisco Systems, Inc. Easy VPN.
http://www.cisco.com/en/US/partner/products/ps6659/products_ios_protocol_option_home.html.
- Mason, A. *Cisco Secure Virtual Private Networks*. San Jose, California: Cisco Press; 2002.
- Cisco Systems, Inc. Site to Site VPN Solution: Deploying IPsec Virtual Private Networks.
http://www.cisco.com/en/US/partner/netsol/ns340/ns394/ns171/ns142/networking_solutions_white_paper09186a0080117919.shtml.
- RSA Laboratories Inc. RSA Laboratories' Frequently Asked Questions About Today's Cryptography, Version 4.1. <http://www.rsasecurity.com/rsalabs/node.asp?id=2152>.
- Cisco Systems, Inc. *Voice and Video Enabled IPsec VPN (V3PN) Solution Reference Network Design*.
http://www.cisco.com/application/pdf/en/us/guest/netsol/ns241/c649/ccmigration_09186a00801ea79c.pdf.

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) What does the IPsec protocol determine? (Source: Introducing IPsec VPNs)
- A) how the interface on the router appears to the encryption protocol
 - B) which type of encryption is used (DES, 3DES, AES)
 - C) which type of message digest is used (MD1, MD5, SHA-1)
- Q2) Which mode of operation does the IKE protocol use when the responder in the process to establish secure communications sends the proposal, key material and ID, and authenticates the session in one packet? (Source: Introducing IPsec VPNs)
- A) main mode
 - B) aggressive mode
 - C) quick mode
- Q3) In which phase in the process to establish a secure communication using the IKE protocols is the SA negotiation unidirectional? (Source: Introducing IPsec VPNs)
- A) IKE Phase 1
 - B) IKE Phase 1.5
 - C) IKE Phase 2
- Q4) Explain why the IKE protocol NAT traversal feature is used with IPsec.(Source: Introducing IPsec VPNs)

- Q5) If packet expansion during the forwarding of small packets is a concern, which IPsec protocol packet forwarding mode can be used?(Source: Introducing IPsec VPNs)
- A) tunnel mode
 - B) channel mode
 - C) quick mode
 - D) transport mode

Q6) Which two hash functions does the IPsec protocol use? (Source: Introducing IPsec VPNs)

- A) MD5 and SHA-1
- B) MD4 and SHA
- C) RSA and MD5
- D) 3DES-CBC and AES

Q7) Describe the components of the PKI environment. (Source: Introducing IPsec VPNs)

Q8) The table in the figure shows the unordered steps used to configure a site-to-site IPsec VPN. Place the steps in the correct order used to configure a site-to-site IPsec VPN. (Source: Building a Site-to-Site IPsec VPN Operation). The instructor will use an animated slide to support this practice item.

Building a Site-to-Site IPsec VPN Operation

Step Number	Site-to-Site Configuration Step
1	Configure the ISAKMP policy required to establish an IKE tunnel.
2	Define the IPsec transform set.
3	Create a cryptographic access-list.
4	Create and apply a cryptographic map.
5	Configure the interface access list.

© 2006 Cisco Systems, Inc. All rights reserved.

SND v2.0—6.2

- Q9) Which type of VPN networks is the best choice when corporate security manages the VPN network? (Source: Defending Your Network with the Cisco VPN Product Family)
- A) remote-access VPN network
 - B) site-to-site VPN network
 - C) firewall-based VPN network
 - D) IPsec VPN
- Q10) If the primary role is to perform as a remote-access VPN with a few site-to-site connections, which product is the best choice? (Source: Defending Your Network with the Cisco VPN Product Family)
- A) VPN-enabled router
 - B) Cisco PIX 500 Series Security Appliance
 - C) Cisco VPN 3000 Series Concentrator
 - D) Cisco VPN 3002 Hardware Client
- Q11) Describe the primary roles of Cisco VPN concentrators and VPN-enabled routers. (Source: Defending Your Network with the Cisco VPN Product Family)

- Q12) Describe the design objectives that should guide your decision making when you are designing a Cisco IPsec VPN. (Source: Defending Your Network with the Cisco VPN Product Family)

Module Self-Check Answer Key

- Q1) A
- Q2) B
- Q3) C
- Q4) The explanation should cover these points: Without the IKE NAT traversal feature, a standard IPsec VPN tunnel with one or more NAT or PAT points in the delivery path of an IPsec packet will not work. The reason is that there are no port numbers in the IPsec headers that can be used to create and maintain translation tables; IPsec encrypts the Layer 4 port information. By encapsulating IPsec packets in a UDP wrapper, the NAT traversal function enables IPsec traffic to travel through NAT or PAT devices in the network.
- Q5) D
- Q6) A
- Q7) A PKI is composed of these components:
- Computers (peers) communicating on a secure network
 - At least one CA (A CA grants and maintains certificates.)
 - Digital certificates (A digital certificate can contain the certificate validity period, peer identity information, encryption keys used for secure communications, and the signature of the issuing CA.)
 - An optional RA to offload the CA by processing enrollment requests
 - A distribution mechanism using, for example, LDAP or HTTP for CRLs
- Q8) The steps in the table should be listed in this order: Step 1, Step 3, Step 2, Step 5, Step 4
- Q9) C
- Q10) C
- Q11) The description should cover these points: VPN concentrators can be configured to provide site-to-site VPNs; they are best suited to support remote-access VPNs. Site-to-site VPN requirements are best met using VPN-enabled routers.
- Q12) A Cisco IPsec VPN should emulate the functional requirements of your network. These design objectives should guide your decision making:
- Secure connectivity
 - Reliability, performance, and scalability
 - Options for high availability
 - Authentication of users and devices in VPN secure management
 - Security and attack mitigation before and after IPsec

