SNRS

Securing Networks with Cisco Routers and Switches

Volume 1

Version 2.0

Student Guide

Editorial, Production, and Web Services: 02.06.07

Americas 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: 408 527-0883

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

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

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

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

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0701R)

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.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Students, this letter describes important course evaluation access information!

Welcome to Cisco Systems Learning. Through the Cisco Learning Partner Program, Cisco Systems is committed to bringing you the highest-quality training in the industry. Cisco learning products are designed to advance your professional goals and give you the expertise you need to build and maintain strategic networks.

Cisco relies on customer feedback to guide business decisions; therefore, your valuable input will help shape future Cisco course curricula, products, and training offerings. We would appreciate a few minutes of your time to complete a brief Cisco online course evaluation of your instructor and the course materials in this student kit. On the final day of class, your instructor will provide you with a URL directing you to a short post-course evaluation. If there is no Internet access in the classroom, please complete the evaluation within the next 48 hours or as soon as you can access the web.

On behalf of Cisco, thank you for choosing Cisco Learning Partners for your Internet technology training.

Sincerely,

Cisco Systems Learning

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Table of Contents

Volume 1

Course Introduction ***1***

Overview 1
Learner Skills and Knowledge 2
Course Goal and Objectives 3
Course Flow 4
Additional References 5
Cisco Glossary of Terms 5
Your Training Curriculum 6

Layer 2 Security ***1-1***

Overview 1-1
Module Objectives 1-1

Examining Company ABC Unsecured **1-3**

Overview 1-3
Objectives 1-3
Company ABC Unsecured 1-4
Attacks and Vulnerabilities 1-5
Vulnerabilities 1-5
Attacks 1-6
Attacks on Company ABC 1-7
Network Under Attack 1-7
Summary 1-8

Examining Layer 2 Attacks **1-9**

Overview 1-9
Objectives 1-9
Types of Layer 2 Attacks 1-10
CAM Table Overflow Attack 1-12
CAM Table Overflow 1-12
Mitigating CAM Table Overflow Attacks 1-13
Port Security 1-14
Secure MAC Addresses 1-15
Security Violations 1-16
Default Port Security Configuration 1-17
Port Security Configuration Guidelines 1-18
Compatibility with Other Features 1-19
Enabling and Configuring Port Security 1-20
Returning to Default Configuration 1-23
Enabling and Configuring Port Security Aging 1-24
Configuring Port Security Aging 1-24
Verifying Port Security 1-26
VLAN Hopping Attacks 1-29
Switch Spoofing 1-29
Mitigating VLAN Hopping Attacks 1-30
STP Vulnerabilities 1-31
Mitigating STP Vulnerabilities 1-32
MAC Spoofing: Man-in-the-Middle Attacks 1-33
PVLAN Vulnerabilities 1-34
Proxy Attack 1-34
Mitigating PVLAN Vulnerabilities 1-35
Summary 1-36
References 1-36

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Configuring DHCP Snooping 1-37
Overview 1-37
Objectives 1-37
DHCP Starvation and Spoofing Attacks 1-38
Understanding DHCP Snooping 1-39
Mitigating DHCP Attacks 1-41
DHCP Snooping Configuration Guidelines 1-42
Enabling and Configuring DHCP Snooping 1-43
Verifying DHCP Snooping 1-45
Examples 1-46
Summary 1-49
References 1-49
Module Summary 1-50
References 1-50

Trust and Identity 2-1

Overview 2-1
Module Objectives 2-1

Implementing Identity Management 2-3

Overview 2-3
Objectives 2-3
Cisco Secure ACS for Windows Overview 2-5
Additional Features in Cisco Secure ACS 4.0 for Windows 2-6
Authentication, Authorization, and Accounting 2-9
Common Cisco IOS AAA Configuration 2-10
Method Lists 2-10
Method Lists and Server Groups 2-11
Authentication 2-13
Common Cisco IOS AAA Authentication Configuration 2-14
AAA Authentication Login Methods 2-14
Authentication Example 2-15
Authorization 2-16
Common Cisco IOS AAA Authorization Configuration 2-17
Authorization Example 2-20
Accounting 2-21
Common Cisco IOS AAA Accounting Configuration 2-21
Accounting Example 2-25
TACACS+ 2-26
RADIUS 2-28
Client-Server Model 2-28
Network Security 2-28
Configuring AAA to Work with External AAA Servers 2-29
Cisco Secure ACS as a AAA Server 2-31
TACACS+ and RADIUS 2-32
Cisco Secure ACS for Microsoft Windows Architecture 2-34
CSAdmin 2-34
CSAuth 2-34
CSDBSync 2-34
CSLog 2-35
CSMonitor 2-35
CSTacacs and CSRadius 2-35
Administering Cisco Secure ACS 2-36
Layout 2-36
Installing Cisco Secure ACS 2-38
Preparation 2-38
Creating an Installation 2-42
Adding an Administrator 2-44
Access Policy 2-45
Session Policy 2-45

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Working in Cisco Secure ACS 2-46
User Setup 2-46
Group Setup 2-46
Shared Profile Components 2-46
Network Configuration 2-47
System Configuration 2-47
Interface Configuration 2-48
Administration Control 2-49
External User Databases 2-49
Posture Validation 2-50
Network Access Profiles 2-50
Reports and Activity 2-50
Network Access Profiles 2-52
Policies 2-53
Configuring Cisco Secure ACS NAPs 2-54
Common Configurations in Cisco Secure ACS 2-54
Creating a NAP 2-56
Templates 2-56
NAFs 2-57
Protocol Types 2-57
Advanced Filtering 2-57
Configuring Profile-Based Policies 2-59
Before You Begin 2-59
Authentication Rules 2-60
Posture Validation Rules 2-61
Authorization Rules 2-62
Troubleshooting Cisco Secure ACS 2-63
Reports 2-63
Cisco Secure ACS Command-Line Utility 2-63
Cisco IOS Debug 2-64
Summary 2-65
References 2-67

Implementing Cisco IBNS 2-69

Overview 2-69
Objectives 2-69
Cisco IBNS Overview 2-70
Features and Benefits 2-72
Port-Based Access Control 2-74
IEEE 802.1x 2-75
802.1x Components 2-76
How 802.1x Works 2-78
Authentication Initiation and Message Exchange 2-79
Ports in Authorized and Unauthorized States 2-80
IEEE 802.1x Host Mode 2-82
Selecting the Correct EAP 2-83
EAP Methods 2-85
EAP-MD5 2-86
EAP-TLS 2-87
PEAP with MS-CHAPv2 2-89
EAP-FAST 2-91
802.1x and Port Security 2-93
802.1x and VLAN Assignment 2-95
Configuring VLAN Assignment 2-96
802.1x and Guest VLANs 2-97
Configuring a Guest VLAN on a Port 2-98
802.1x and Restricted VLANs 2-99
Compatibility with Other Security Features 2-100
Configuring a Restricted VLAN 2-100
Configuring 802.1x 2-101
Guidelines 2-102

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

VLAN Assignment, Guest VLANs, and Restricted VLANs 2-102
Configuring AAA 2-103
RADIUS Communications 2-107
Enabling 802.1X 2-108
Periodic Reauthentication 2-110
Manually Reauthenticating a Client 2-111
Adjusting the Quiet Period 2-111
Adjusting the Switch-to-Client Retransmission Time 2-111
Adjusting Timers for DHCP 2-111
Summary 2-114
References 2-115
Module Summary 2-116
References 2-116

Cisco Network Foundation Protection 3-1

Overview 3-1
Module Objectives 3-1

Introducing Cisco NFP 3-3

Overview 3-3
Objectives 3-3
Cisco NFP Overview 3-4
Network Device Planes 3-4
Cisco IOS Tools for a Secure Infrastructure 3-5
Summary 3-6
References 3-6

Securing the Control Plane 3-7

Overview 3-7
Router Control Plane 3-8
Tools for Securing the Control Plane 3-9
Overview of CPPr 3-10
Prerequisites 3-10
Restrictions for CPPr 3-10
Benefits 3-11
CPPr Architecture 3-12
Aggregate Control Plane Services 3-12
Control Plane Interface and Subinterface 3-12
Port Filtering 3-13
Queue Thresholding 3-14
Configuring CPPr 3-15
Configuring CoPP 3-16
Defining Packet Classification Criteria for CoPP 3-16
Defining a CoPP Service Policy 3-18
Entering Aggregate Control Plane Configuration Mode 3-22
Applying a CoPP Service Policy to the Host Subinterface 3-23
Configuring a Port-Filter Policy 3-25
Defining Port-Filter Parameters 3-26
Defining Port-Filter Packet Classification Criteria 3-26
Defining a Port-Filter Service Policy 3-28
Applying a Port-Filter Service Policy to the Host Subinterface 3-28
Configuring a Queue-Threshold Policy 3-30
Restrictions 3-31
Configuring Queue Thresholding 3-31
Defining Queue-Threshold Packet Classification Criteria 3-32
Defining a Queue-Threshold Service Policy 3-33
Applying a Queue-Threshold Policy to the Host Subinterface 3-34
Verifying CPPr 3-35
Summary 3-39
References 3-39

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Securing the Management Plane 3-41
Overview 3-41
Objectives 3-41
The Management Plane 3-42
Tools for Securing the Management Plane 3-43
Cisco MPP Feature 3-44
Prerequisites 3-45
Restrictions 3-45
Securing the Management Plane 3-46
Configuring MPP 3-46
Verifying MPP 3-48
Summary 3-49
References 3-49
Securing the Data Plane 3-51
Overview 3-51
Objectives 3-51
Data Plane Attacks 3-52
Data Plane Protection 3-53
Flexible Packet Matching 3-54
Deployment 3-55
Protocol Header Definition File 3-55
Filter Description 3-57
Prerequisites and Restrictions 3-57
Configuring FPM 3-58
PHDFs and Traffic Classes 3-59
Loading a PDHF 3-59
Creating a Traffic Class 3-60
Creating a Traffic Policy 3-64
Applying the Service Policy 3-66
Verifying FPM 3-67
Troubleshooting FPM 3-71
Summary 3-72
References 3-72
Module Summary 3-73
References 3-73

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

SNRS

Course Introduction

Overview

Securing Networks with Cisco Routers and Switches (SNRS) v2.0 is a five-day, instructor-led, lab-intensive course that is delivered by Cisco Learning Partners. It is aimed at providing network specialists with the knowledge and skills needed to secure Cisco IOS router and switch networks. Successful graduates will be able to secure the network environment using existing Cisco IOS security features. This includes the ability to configure some of the primary components of the Cisco IOS Firewall feature set, which include the following:

- Cisco IOS classic firewall (formerly known as Context-Based Access Control [CBAC])
- Cisco IOS Intrusion Prevention System (IPS)
- Cisco IOS authentication proxy
- Cisco IOS zone-based policy firewall
- Application inspection and control

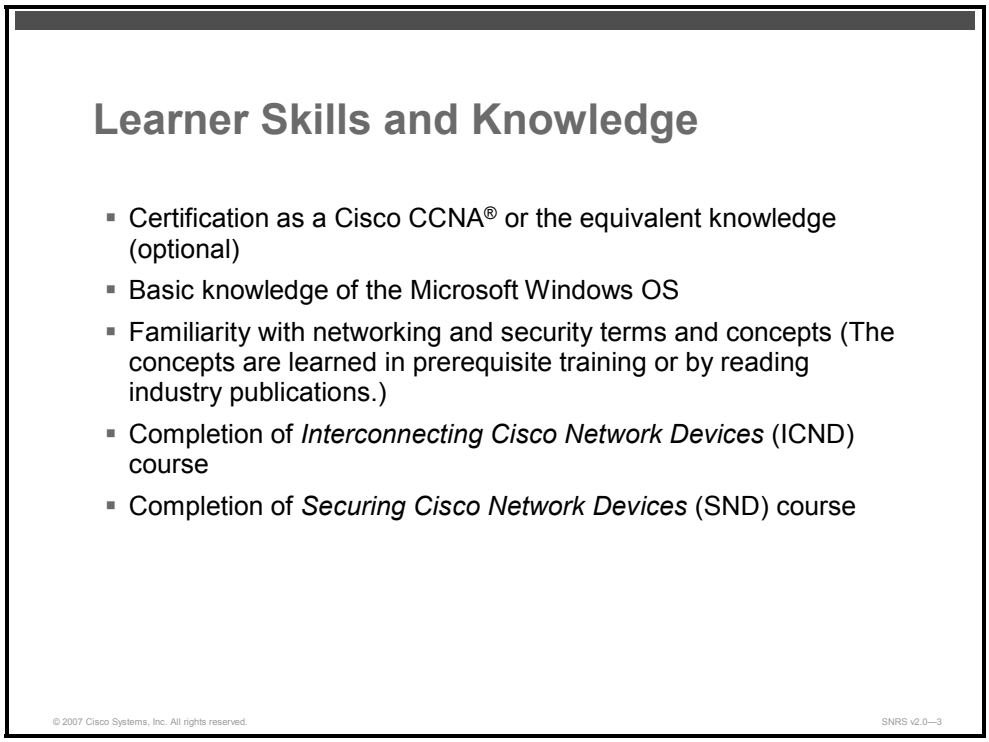
Learners will also have the ability to implement the following:

- Secure tunnels using generic routing encapsulation (GRE) and IP Security (IPsec) technology
- Basic Layer 2 switch security
- The Cisco Trust and Identity Management model to control network access
- Command-line Cisco Network Foundation Protection (NFP)

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Learner Skills and Knowledge

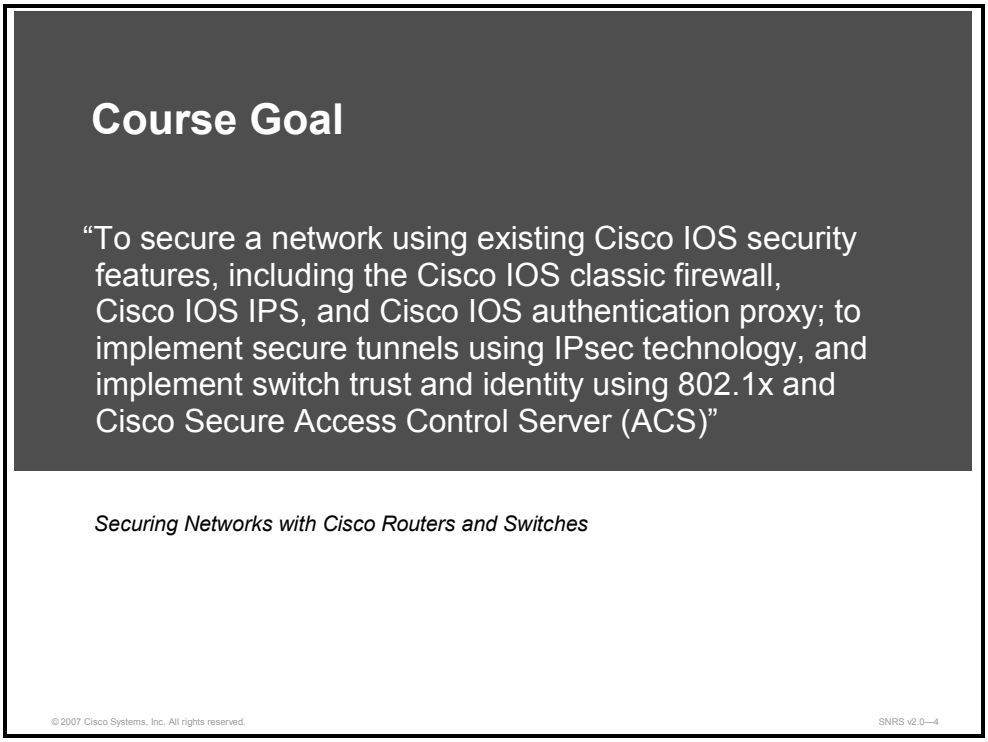
This subtopic lists the skills and knowledge that learners must possess to benefit fully from the course. The subtopic also includes recommended Cisco learning offerings that learners should first complete to benefit fully from this course.



The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Course Goal and Objectives

This topic describes the course goal and objectives.



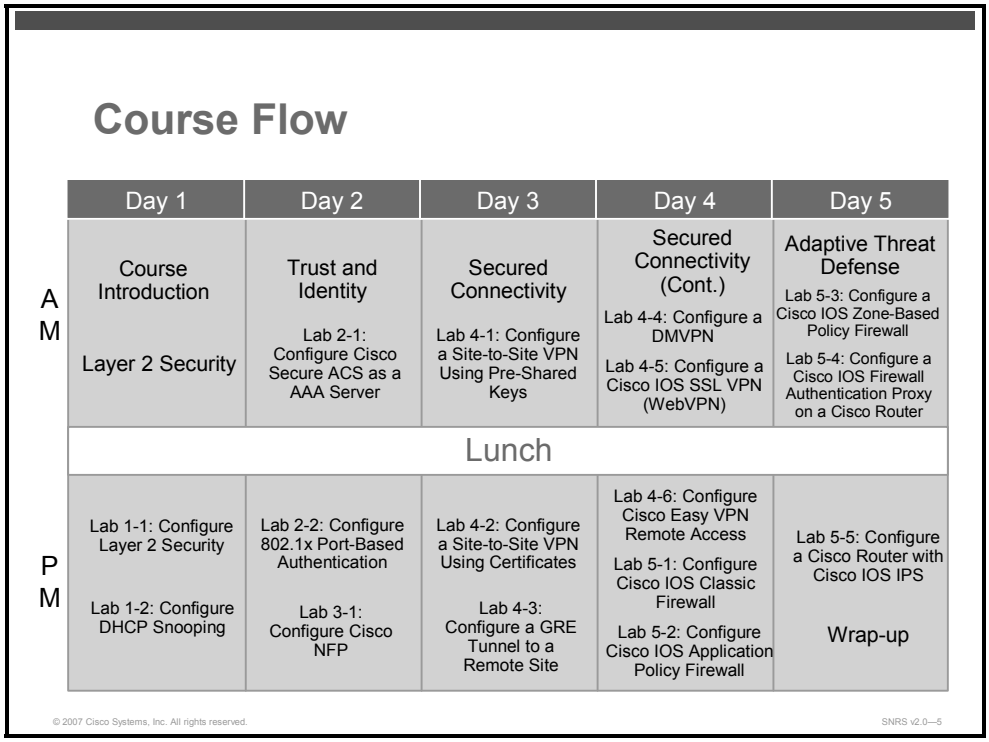
Upon completing this course, you will be able to meet these objectives:

- Implement Layer 2 security features
- Implement the Cisco Trust and Identity Management model to control network access
- Implement command-line Cisco NFP to protect infrastructure devices
- Implement secure IPsec VPNs and GRE tunnels using Cisco routers
- Install, configure, and troubleshoot Cisco IOS Firewall features, including IOS Classic Firewall, Cisco IOS Firewall authentication proxy, and Cisco IOS IPS on a Cisco router

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Course Flow

This topic presents the suggested flow of the course materials.

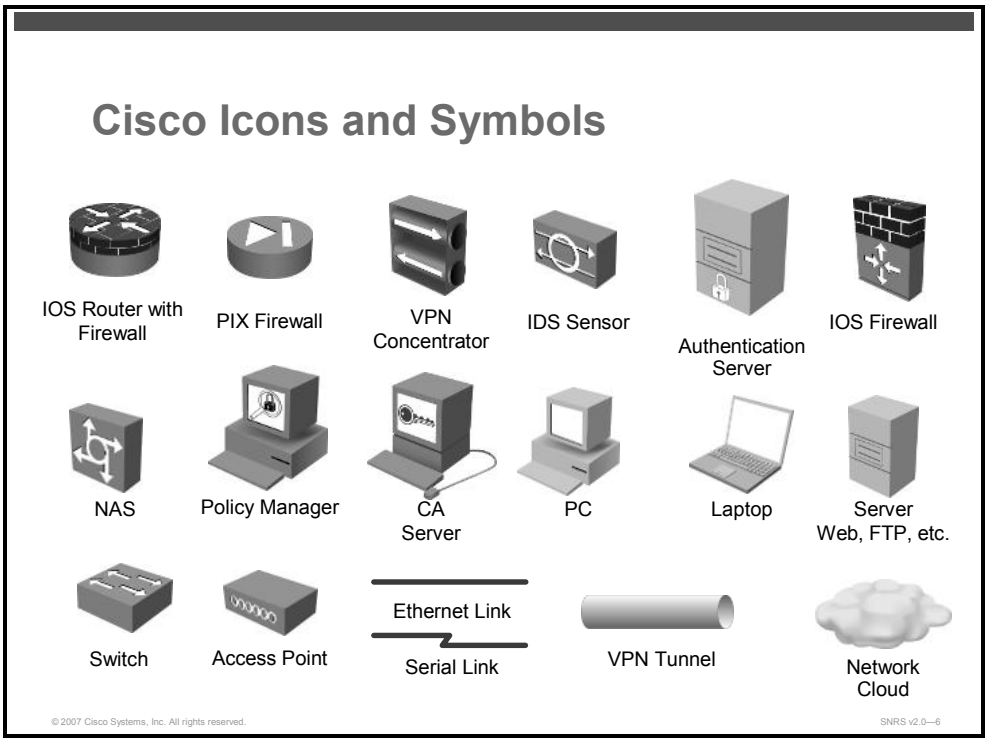


The schedule reflects the recommended structure for this course. This structure allows enough time for the instructor to present the course information and for you to work through the lab activities. The exact timing of the subject materials and labs depends on the pace of your specific class.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Additional References

This topic presents the Cisco icons and symbols that are used in this course, as well as information on where to find additional technical references.



Cisco Glossary of Terms

For additional information on Cisco terminology, refer to the *Cisco Internetworking Terms and Acronyms* glossary of terms at http://www.cisco.com/univercd/cc/td/doc/cisintwk/ita/index.htm.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

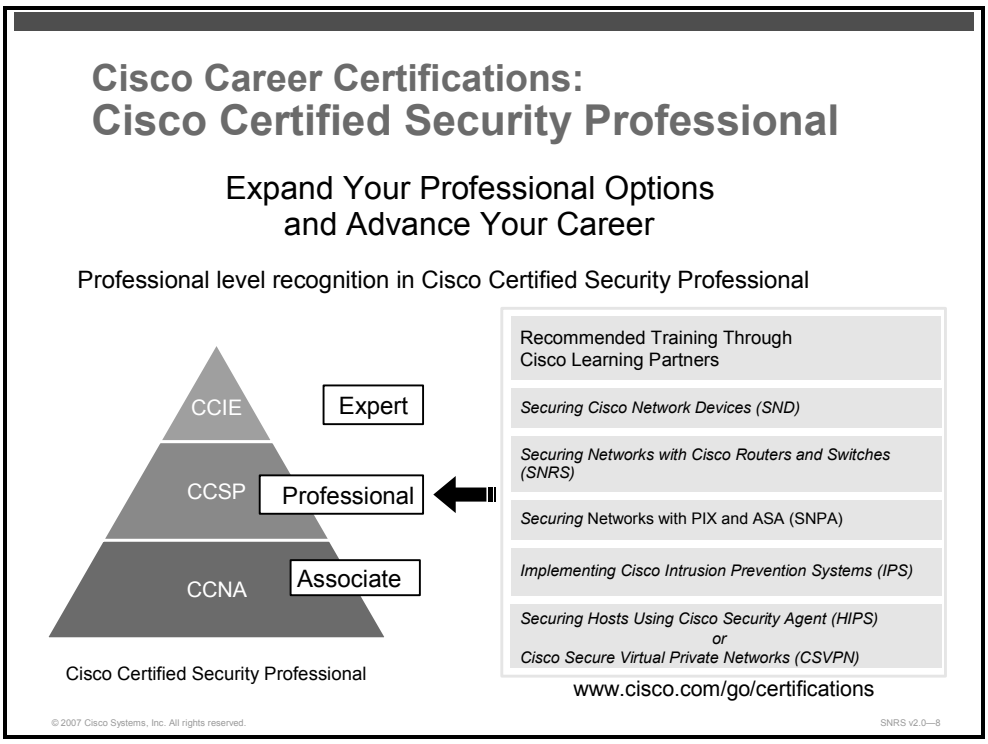
Your Training Curriculum

This topic presents the training curriculum for this course.



You are encouraged to join the Cisco Certification Community, a discussion forum open to anyone holding a valid Cisco Career Certification (such as Cisco CCIE®, CCNA®, CCDA®, CCNP®, CCDP®, CCIP®, CCVP™, or CCSP™). It provides a gathering place for Cisco certified professionals to share questions, suggestions, and information about Cisco Career Certification programs and other certification-related topics. For more information, visit www.cisco.com/go/certifications.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.



The figure above represents the Cisco Certified Security Professional (CCSP) certification path.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Module 1

Layer 2 Security

Overview

Layer 2 security is now available as a tool to use against the attacks faced by networks today. In this module, you will identify types of Layer 2 attacks and examine and configure Cisco Layer 2 security features. You will first be introduced to a hypothetical unsecured network belonging to Company ABC. You will then proceed to mitigate various attacks that are being experienced by that network.

Module Objectives

Upon completing this module, you will be able to implement Layer 2 security features using Cisco IOS commands. This ability includes being able to meet these objectives:

- Describe the network of Company ABC and examine the vulnerabilities and attacks that the company network experiences
- Describe the types of Layer 2 attacks and the strategies to mitigate them
- Implement port security on a Cisco Catalyst switch
- Implement DHCP snooping on a Cisco Catalyst switch

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Lesson 1

Examining Company ABC Unsecured

Overview

In this lesson, you will gain experience in examining an unsecured network for vulnerabilities and attacks that are being experienced.

Objectives

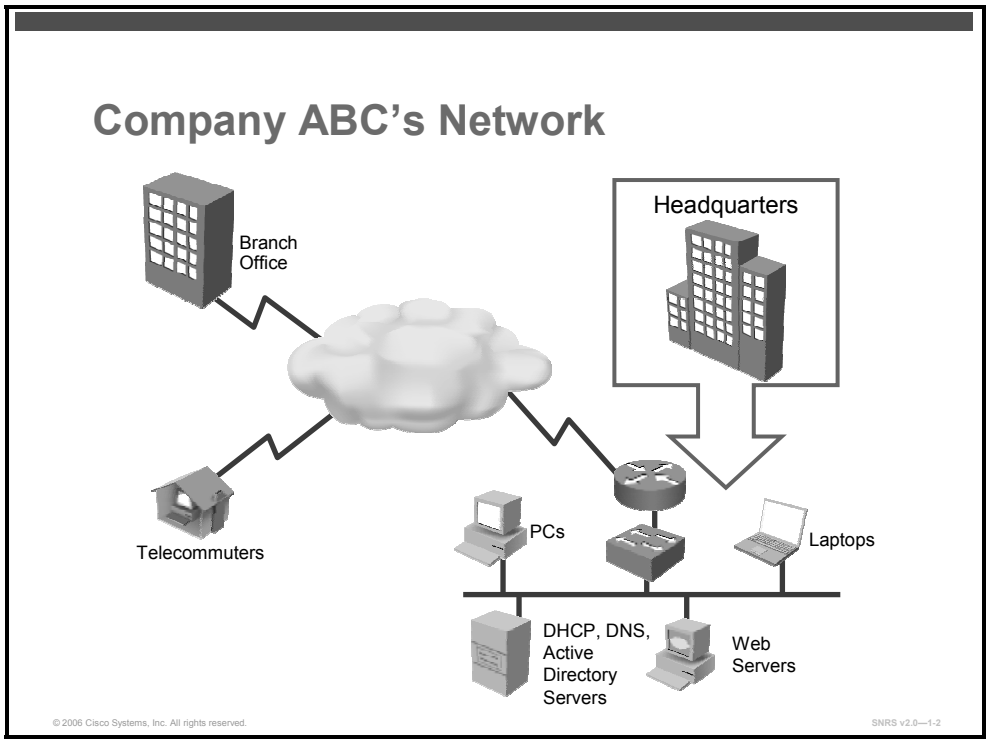
Upon completing this lesson, you will be able to describe the network of Company ABC and examine various vulnerabilities and attacks that the company network experiences. This ability includes being able to meet these objectives:

- Describe the network of Company ABC
- Describe some of the network attacks and the vulnerabilities that are being exploited
- Examine the attacks to which the network is exposed

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Company ABC Unsecured

This topic describes the network of Company ABC.



The figure represents a simplified view of the network architecture of Company ABC. The company headquarters contains the data center for the company. Company ABC also has a remote branch and some telecommuters in several cities.

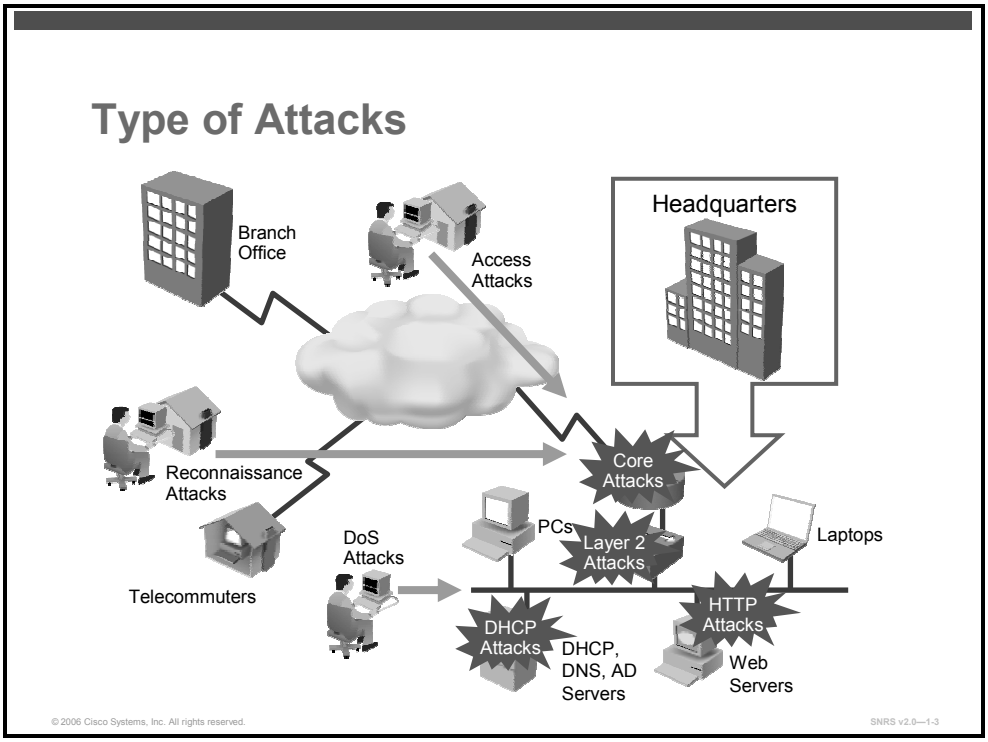
At this time, the network has no security policies and is experiencing some attacks that affect its performance. There has also been some data stolen and damage done because intruders have been able to gain access to the network data or the data flowing through the network.

Your job is to determine a proper security solution using the security features contained within the current Cisco IOS release for the Cisco router and switch used at the company.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Attacks and Vulnerabilities

This topic describes some of the network attacks and the vulnerabilities that are being exploited.



Because of either software- or network-related vulnerabilities, the network is exposed to several types of attacks.

Vulnerabilities

For an attack to take place, there must be some weakness to exploit. Here are some of these weaknesses:

- Missing network security policies
  - **No written policies:** Usually results in little to no security configurations
  - **No patch management:** Exposes network to OS vulnerabilities
- OS or application weaknesses
  - **65,535 ports that you can remotely connect to:** Easy access to known vulnerabilities
  - **Applications not written with security in mind:** Easy access to known vulnerabilities
- Protocol weaknesses
  - **Routing protocols**: Authentication features not commonly used
  - **TCP/IP:** Contains Internet Control Message Protocol (ICMP), which is commonly used in denial of service (DoS) and distributed DoS (DDoS) attacks

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

- **Microsoft networking:** Usually blocked at the perimeter to prevent remote access to known vulnerabilities

- Network configuration weaknesses
  - **Weak or unencrypted passwords:** Allows access to network devices
  - **Exposed services**: Allow for DoS attacks on known services
  - **Default configurations used:** Results in no security configurations on network devices
- Human weaknesses
  - Social engineering uses deception to gain the username and password or other credentials of a legitimate user.

Attacks

Attacks are generally broken into types:

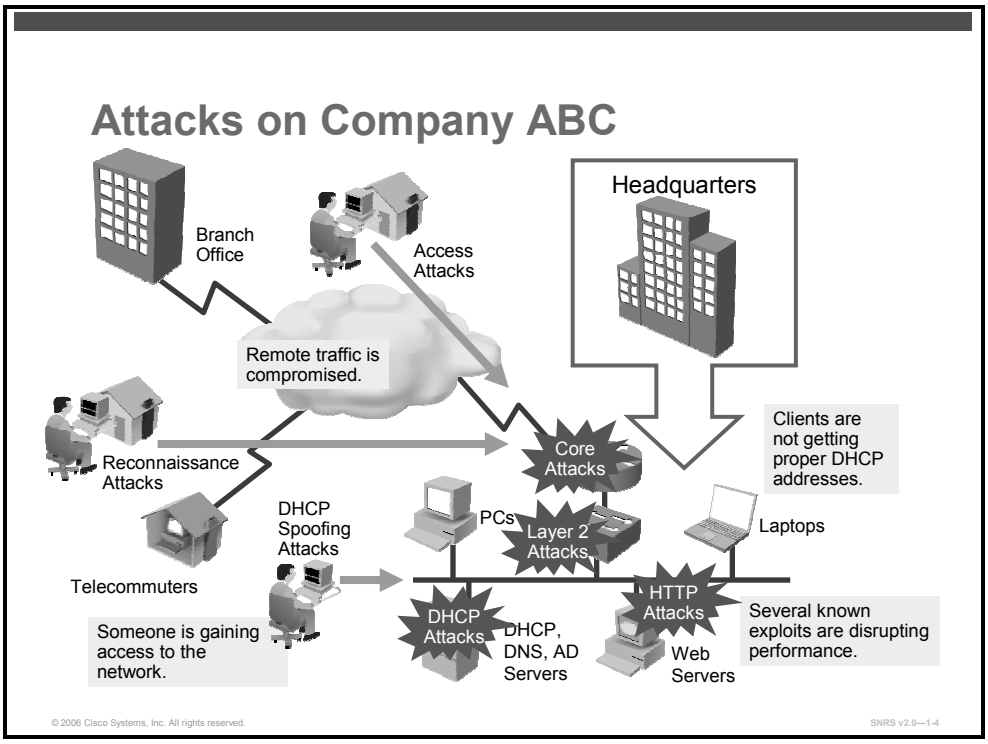
- **DoS attacks:** DoS attacks cause an interruption of access to a system or the network, usually by overloading network resources or infrastructure devices.
- **DDoS attacks:** DDoS attacks also cause an interruption of access to a system or the network by overloading network resources or infrastructure devices, but this time, the attack is executed from several different networks working in conjunction with each other.
- **Reconnaissance attacks:** These types of attacks are used to map a network and discover which hosts and services are exposed. Reconnaissance attacks do not cause immediate damage.
- **Access attacks:** Access attacks are used to illegally gain access to the network to install back doors, trojans, worms, and other exploits; destroy, download, or change data; change device configurations, and perform packet capture.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Attacks on Company ABC

This topic describes the attacks on the network.



Lack of a security policy has resulted in attacks on the network.

Network Under Attack

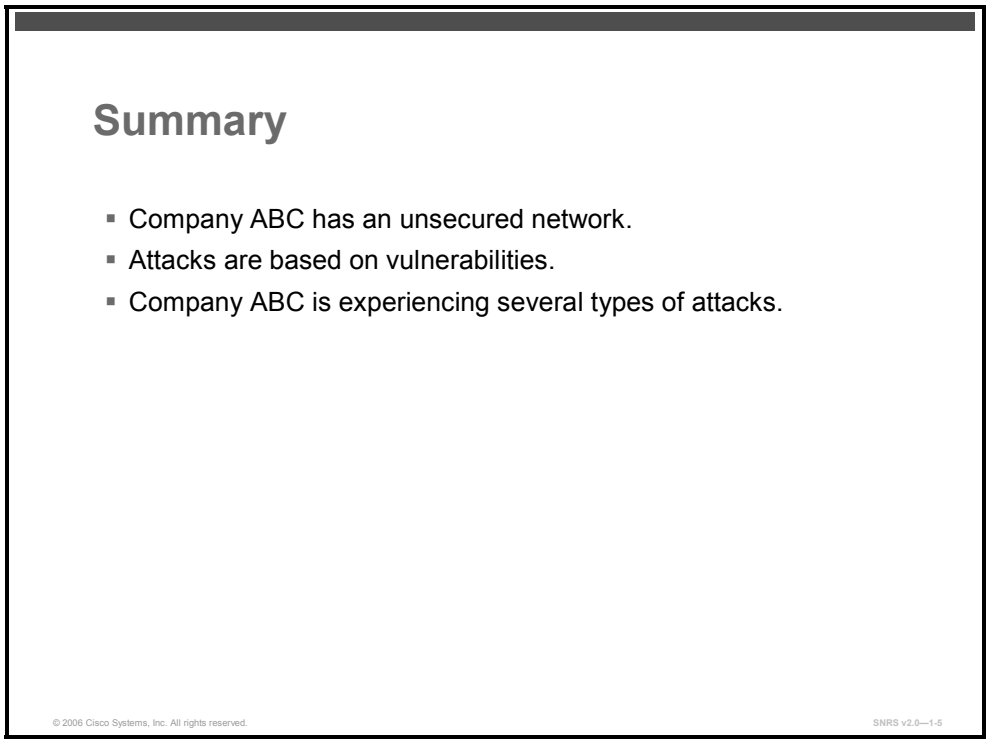
These are some of the attacks that the network is experiencing:

1. Several types of Layer 2 access attacks
   - — DHCP spoofing
   - — MAC spoofing
2. DoS attacks
   - — Known exploits
3. Core attacks
   - — Attacks on core devices
4. Man-in-the-middle attacks
   - — Interception of remote communications
   - — MAC spoofing

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Summary

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



The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Lesson 2

Examining Layer 2 Attacks

Overview

Access to switches is a convenient entry point for attackers who are intent on illegally gaining access to a corporate network. With access to a switch, an attacker can set up rogue DHCP servers, access points, and protocol analyzers, and launch all types of attacks from within the network. Attackers can even spoof the MAC and IP addresses of critical servers to do a great deal of damage. In this lesson, you will examine various Layer 2 attacks and strategies to mitigate them.

Objectives

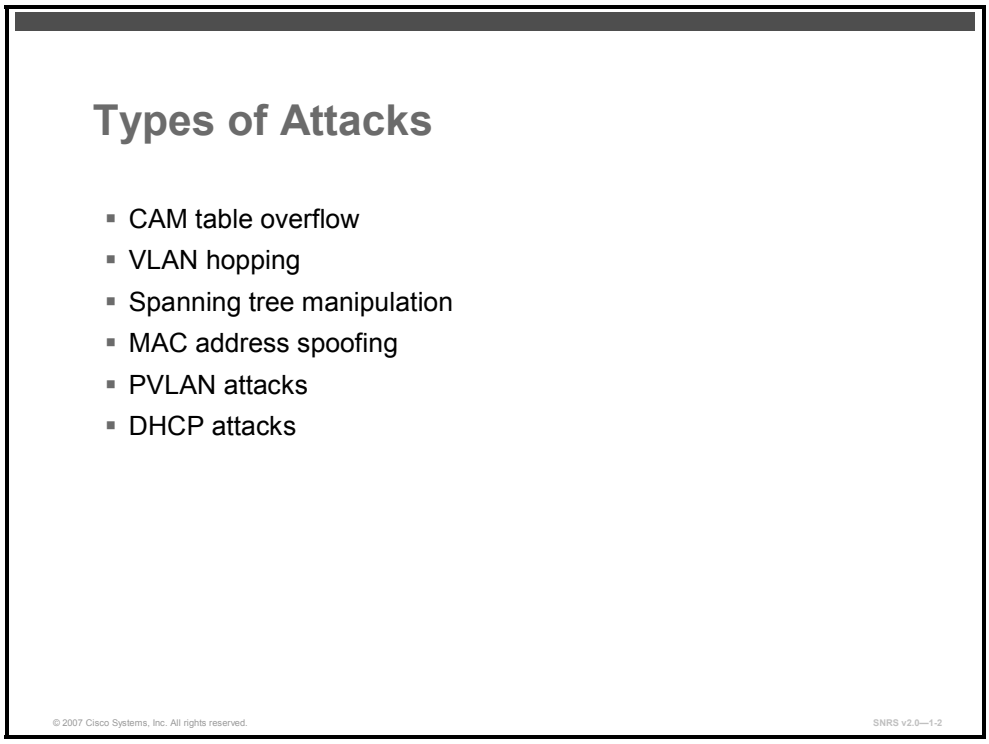
Upon completing this lesson, you will be able to describe the types of Layer 2 attacks and the strategies to mitigate them. This ability includes being able to meet these objectives:

- Describe the various types of Layer 2 network attacks
- Describe a CAM table overflow attack
- Describe the port security feature and how to configure it
- Describe some commands used to verify port security configuration and operations
- Describe VLAN hopping and a strategy to mitigate an attack
- Describe STP vulnerabilities and a strategy to mitigate an attack
- Describe MAC spoofing attacks and a strategy to mitigate an attack
- Describe PVLAN vulnerabilities and a strategy to mitigate an attack

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Types of Layer 2 Attacks

This topic describes various types of Layer 2 attacks.



Like routers, both Layer 2 and Layer 3 switches have their own sets of network security requirements. However, not as much public information is available about the network security risks in switches and what can be done to mitigate those risks. Switches are susceptible to many of the same Layer 3 attacks as routers. Switches, and Layer 2 of the Open Systems Interconnection (OSI) model in general, are subject to network attacks in unique ways, including the following:

- Content-addressable memory (CAM) table overflow

This attack involves an attacker who floods the switch with bogus MAC addresses in an effort to cause the switch to flood all packets out of all its ports.

- VLAN hopping

This attack involves an attacker who gains access to another VLAN other than the one they are assigned to.

- Spanning tree manipulation

This attack involves an attacker who wants to manipulate the Spanning Tree Protocol (STP)) in an attempt to change the root bridge of the network or subnet.

- MAC spoofing

This attack involves an attacker who falsifies their MAC address to execute a man-in-the-middle attack.

- Private VLAN (PVLAN) attacks

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

This attack involves an attacker who tries to gain access to data on a PVLAN. This attacks utilizes the IP address of the target but the MAC address of the router or Layer 3 device doing the routing for the target PVLAN subnet.

- DHCP attacks

These attacks can include the following.

- DHCP Starvation

This attack involves an attacker who floods the DHCP server with DHCP IP address requests in an attempt to use up all of the DHCP addresses and starve the rest of the clients of valid IP addresses.

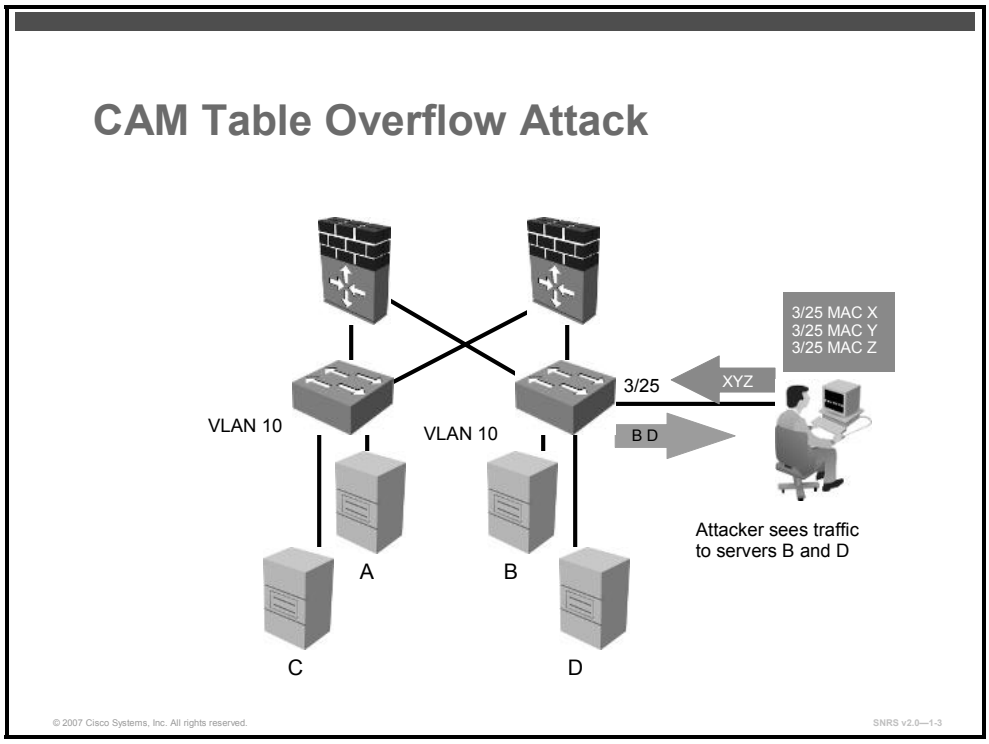
- Setting up a rogue DHCP server on the network

This attack involves an attacker who sets up a DHCP server on a network in order to hand out erroneous DHCP addresses.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

CAM Table Overflow Attack

This topic describes the CAM table overflow attack.



This diagram illustrates a CAM table overflow attack. In this figure, the attacker is sending out multiple packets with various source MAC addresses. Over a short period of time, the CAM table in the switch fills up until it cannot accept new entries. As long as the flood is left running, the CAM table on the switch will remain full. When this happens, the switch begins to broadcast all packets that it receives out of every port so that packets sent to and from server B and server D are also broadcast out of port 3/25 on the switch to which the attacker is attached. In the diagram, the machine of the attacker resides on VLAN 10. The attacker floods MAC addresses to port 3/25 on the switch. When the CAM table threshold is reached, the switch operates as a hub and simply floods traffic out all ports.

CAM Table Overflow

MAC flooding is the attempt to exploit the fixed hardware limitations of the switch CAM table. The Cisco Catalyst switch CAM table stores the source MAC address and the associated port of each device connected to the switch. The CAM table on the Cisco Catalyst 6000 Series Switch can contain 128,000 entries. These 128,000 entries are organized as 8 pages that can store approximately 16,000 entries.

CAM tables are limited in size. If enough entries are entered into the CAM table before other entries are expired, the CAM table fills up to the point that no new entries can be accepted. Typically, a network intruder will flood the switch with a large number of invalid source MAC addresses until the CAM table fills up. When that occurs, the switch will flood 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. If the intruder does not maintain the flood of invalid source MAC addresses, the switch will eventually time out older MAC address

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

entries from the CAM table and begin to act like a switch again. CAM table overflow only floods traffic within the local VLAN so the intruder will see only traffic within the local VLAN to which the intruder is connected.

The macof tool was released in May of 1999. Written in Perl code, the macof tool was later ported to C and incorporated into the dsniff package. This tool floods a switch with packets containing randomly generated source and destination MAC and IP addresses. When the CAM table of the switch fills up with these addresses, the switch begins to forward all frames that it receives to every port.

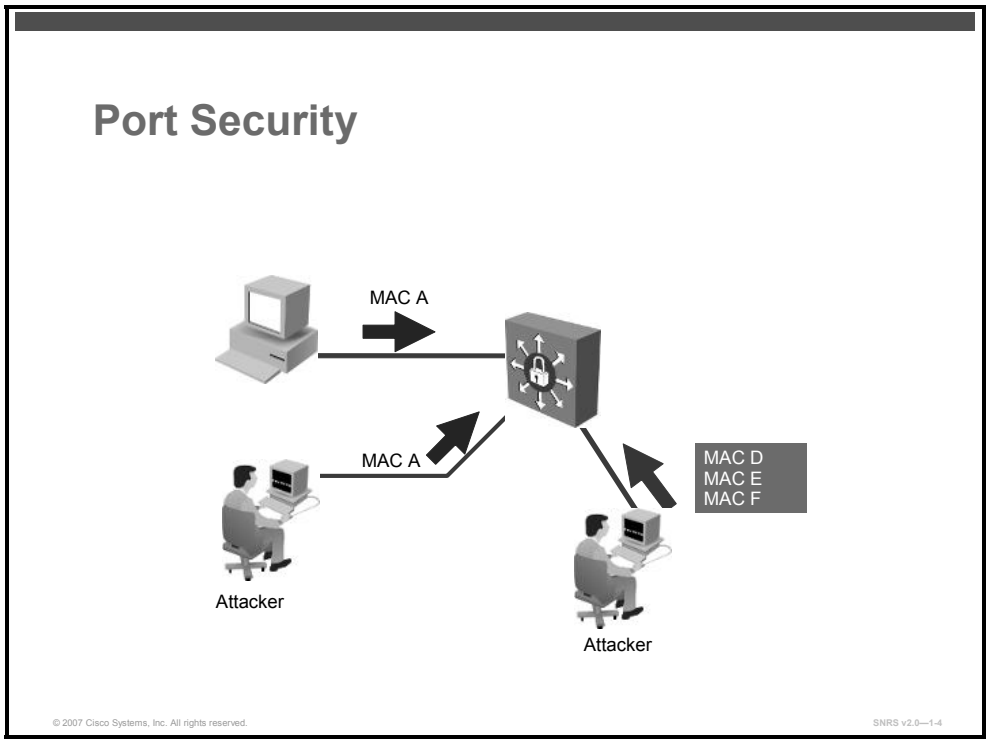
Mitigating CAM Table Overflow Attacks

You will use the **switchport port-security** command to mitigate a CAM table overflow attack by limiting the number of MAC addresses allowed on a switch port. The port security feature is described in the “Port Security” topic.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Port Security

This topic describes the port security feature and how to configure it.



You can use the port security feature to restrict input to an interface by limiting and identifying the MAC addresses of the stations 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.

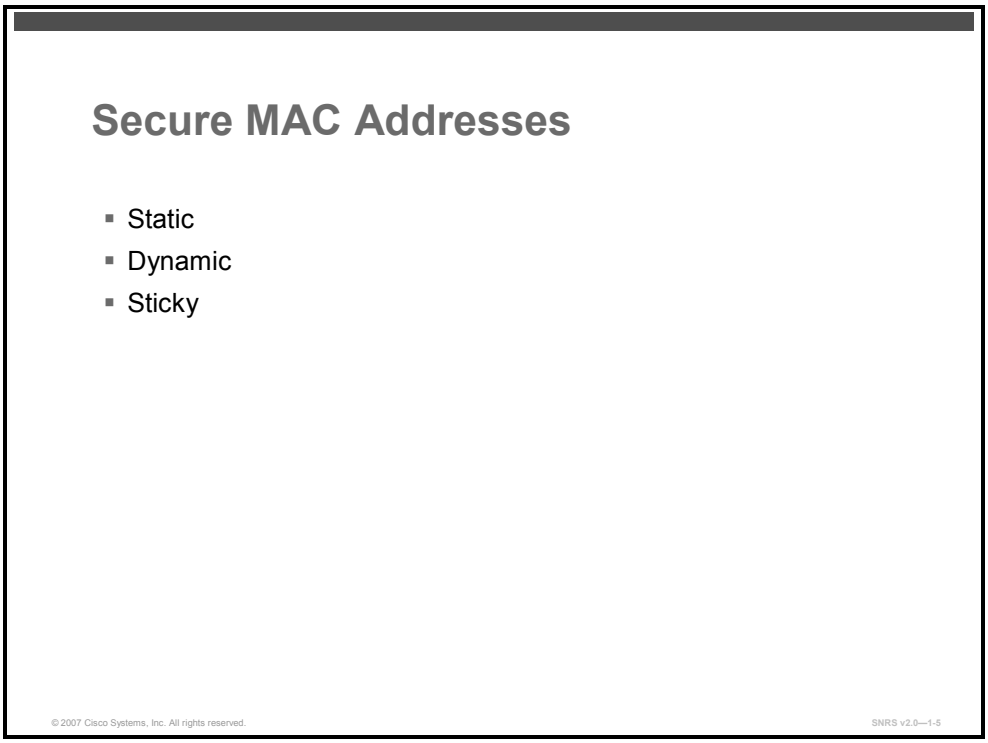
Port security allows you to specify MAC addresses for each port or to permit a limited number of MAC addresses. When a secure port receives a packet, the source MAC address of the packet is compared to the list of secure source addresses that were manually configured or autoconfigured (learned) on the port. If a MAC address of a device attached to the port differs from the list of secure addresses, the port either shuts down permanently (default mode) or drops incoming packets from the insecure host. The behavior of the port depends on how you configure it to respond to a security violator.

Cisco recommends that you configure the port security feature to issue a shutdown instead of dropping packets from insecure hosts through the restrict option. The restrict option may fail under the load of an attack, and the port will be disabled anyway.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Secure MAC Addresses

This section covers the types of secure MAC addresses and security violation mode actions.



A secure port can have 1 to 132 associated secure addresses. The total number of available secure addresses on the switch is 1024.

After you have set the maximum number of secure MAC addresses allowed on a port, you can add secure addresses to the address table by manually configuring them, by allowing the port to dynamically configure them, or by configuring some MAC addresses and allowing the rest to be dynamically configured.

Here are the three types of secure MAC addresses, which you can configure:

- **Static secure MAC addresses:** These MAC 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 MAC addresses are dynamically learned, stored only in the address table, and removed when the switch restarts. You can delete dynamic secure MAC addresses from the address table by entering the **clear port-security dynamic** privileged EXEC command.
- **Sticky secure MAC addresses:** These MAC addresses can be dynamically learned, stored in the address table, and added to the running configuration. If these addresses are saved in the configuration file, the interface does not need to dynamically relearn them when the switch restarts.

You can configure an interface to convert the dynamic MAC addresses to sticky secure MAC addresses and to add them to the running configuration by enabling sticky learning.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

You can enable sticky learning on an interface by using the **switchport port-security mac-address sticky** interface configuration command. When you enter this command, the interface converts all the dynamic secure MAC addresses, including those that were dynamically learned before sticky learning was enabled, to sticky secure MAC addresses.

The sticky secure MAC addresses do not automatically become part of the startup 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, those addresses are lost.

If you disable sticky learning, the sticky secure MAC addresses are converted to dynamic secure addresses and are removed from the running configuration.

You can delete a sticky secure MAC addresses from the address table by using the **clear port-security sticky** *mac-addr* privileged EXEC command. To delete all the sticky addresses on an interface, use the **clear port-security sticky** *interface-id* privileged EXEC command.

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 shut down** interface configuration commands.

Security Violations

A security violation occurs in these situations:

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

Here are the three violation modes:

- **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 increments.
- **Shutdown:** This is the default mode. In this mode, a port security violation causes the interface to immediately become error-disabled, and turns off the port LED. The switch also sends an SNMP trap, logs a syslog message, and increments 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.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Security Violation Mode Actions

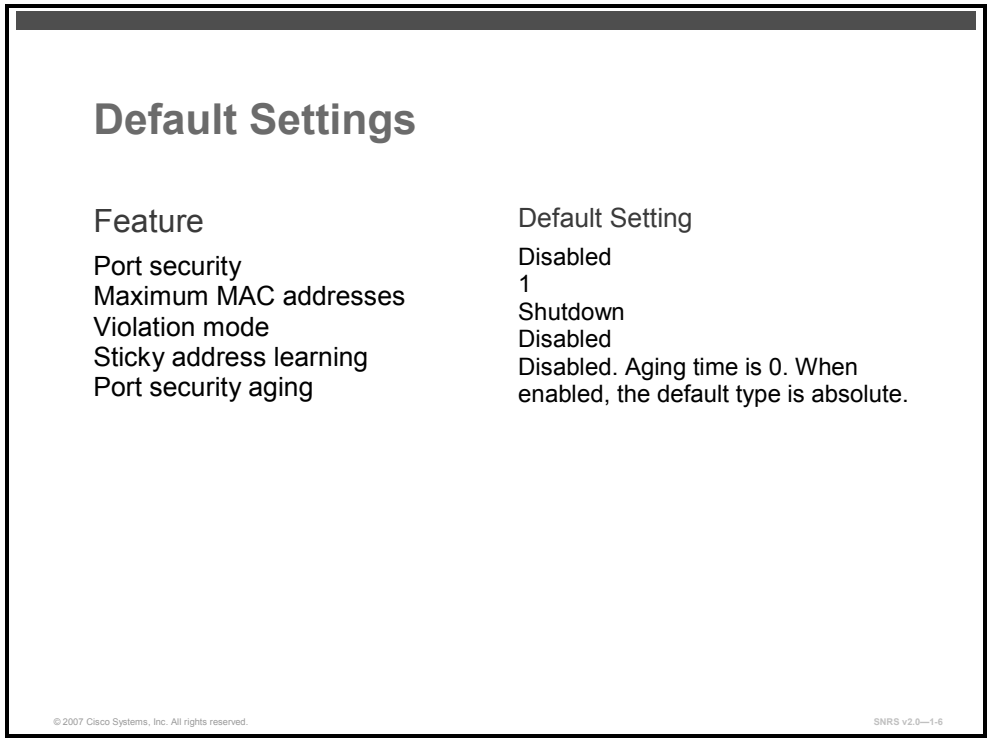
Violation Mode	Traffic Is Forwarded[1]	Sends SNMP Trap	Sends Syslog Message	Displays Error Message[2]	Violation Counter Increments	Shuts Down Port
Protect	No	No	No	No	No	No
Restrict	No	Yes	Yes	No	Yes	No
Shutdown	No	Yes	Yes	No	Yes	Yes

1. Packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses.

2. The switch will return an error message if you manually configure an address that would cause a security violation.

Default Port Security Configuration

This section describes the default port security interface configuration.



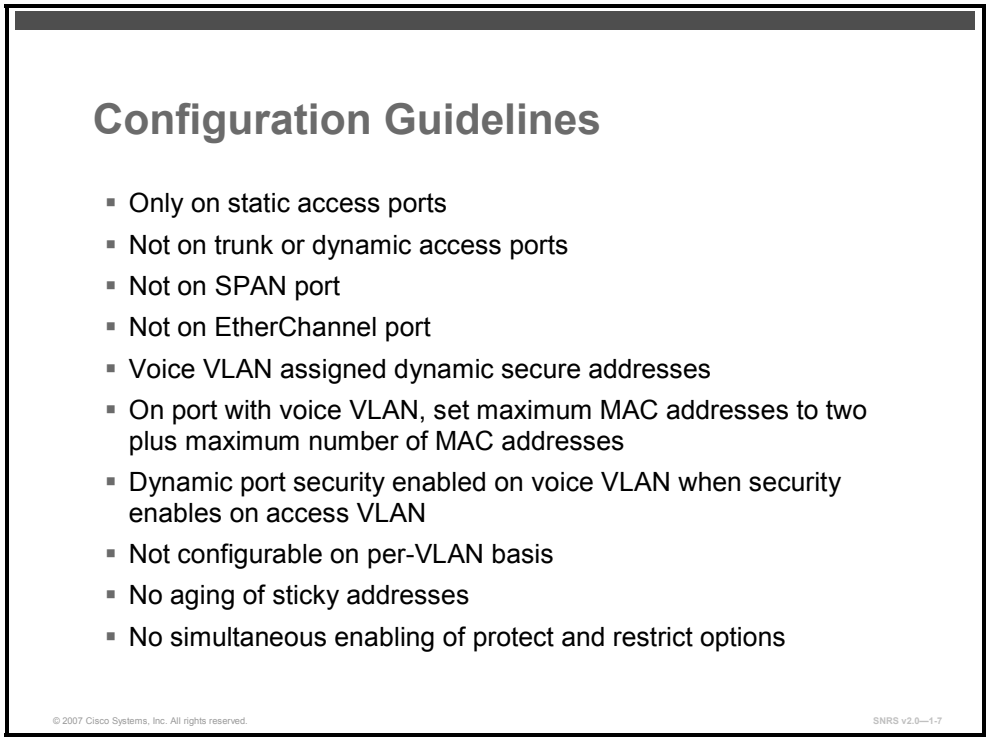
The default port security interface configuration settings are as follows:

- Ports security is **disabled**.
- Maximum MAC addresses setting is **1**.
- Violation mode is **shutdown**.
- Sticky address learning is **disabled**.
- Port security aging is **disabled**. Aging time is **0** and the default type is **absolute**.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Port Security Configuration Guidelines

This section describes some guidelines that you need to use when configuring port security on an interface.



Here are some guidelines to use when configuring port security:

- Port security can only be configured on static access ports.
- A secure port cannot be a dynamic access port or a trunk port.
- A secure port cannot be a destination port for Switched Port Analyzer (SPAN).
- A secure port cannot belong to a Fast EtherChannel or Gigabit EtherChannel port group.
- You cannot configure static secure or sticky secure MAC addresses on a voice VLAN.
- When you enable port security on an interface that is also configured with a voice VLAN, set the maximum allowed secure addresses on the port to two. When the port is connected to a Cisco IP phone, the IP phone requires one MAC address. The Cisco IP phone address is learned on the voice VLAN, but is not learned on the access VLAN. If you connect a single PC to the Cisco IP phone, no additional MAC addresses are required. If you connect more than one PC to the Cisco IP phone, you must configure enough secure addresses to allow one for each PC and one for the phone.
- If any type of port security is enabled on the access VLAN, dynamic port security is automatically enabled on the voice VLAN.
- When a voice VLAN is configured on a secure port that is also configured as a sticky secure port, all addresses seen on the voice VLAN are learned as dynamic secure addresses, and all addresses seen on the access VLAN (to which the port belongs) are learned as sticky secure addresses.
- You cannot configure port security on a per-VLAN basis.
- The switch does not support port security aging of sticky secure MAC addresses.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Compatibility with Other Features

The table includes other switch features that are compatible with port security configured on a port.

Type of Port	Compatible with Port Security
Dynamic Trunking Protocol (DTP) port[1]	No
Trunk port	No
Dynamic access port[2]	No
SPAN source port	Yes
SPAN destination port	No
EtherChannel	No
Protected port	Yes
IEEE 802.1x port	Yes
Voice VLAN port[3]	Yes

1. A port configured with the **switchport mode dynamic** interface configuration command.

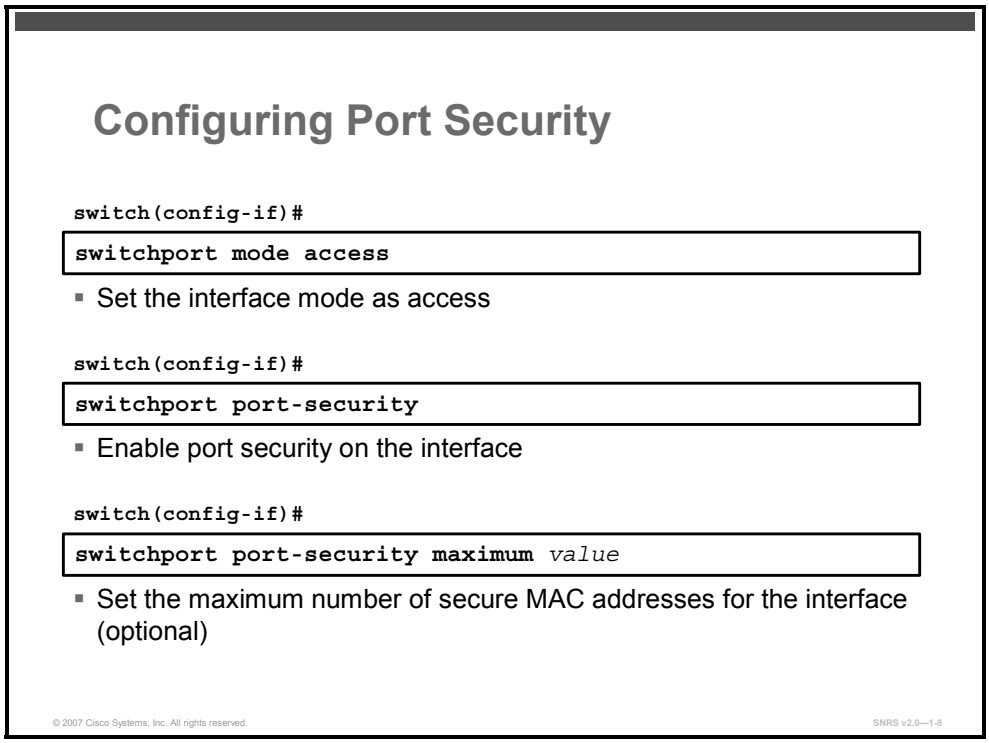
2. A VLAN Query Protocol (VQP) port configured with the **switchport access vlan dynamic** interface configuration command.

3. You must set the maximum allowed secure addresses on the port to two plus the maximum number of secure addresses allowed on the access VLAN.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Enabling and Configuring Port Security

This section describes how to enable and configure Port Security.



Complete these steps to configure port security on an interface.

Step 1 Enter interface configuration mode.

```
switch(config)# interface FastEthernet 0/8
```

Step 2 Configure the interface as an access interface.

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

Note	With an interface in the default mode (dynamic desirable) it cannot be configured as a secure port.

Step 3 Enable port security on the interface.

```
switch(config-if)# switchport port-security [mac-address mac-address] | [mac-address sticky [mac-address]] | [maximum value] | [violation {protect | restrict | shutdown}]
```

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Syntax Description

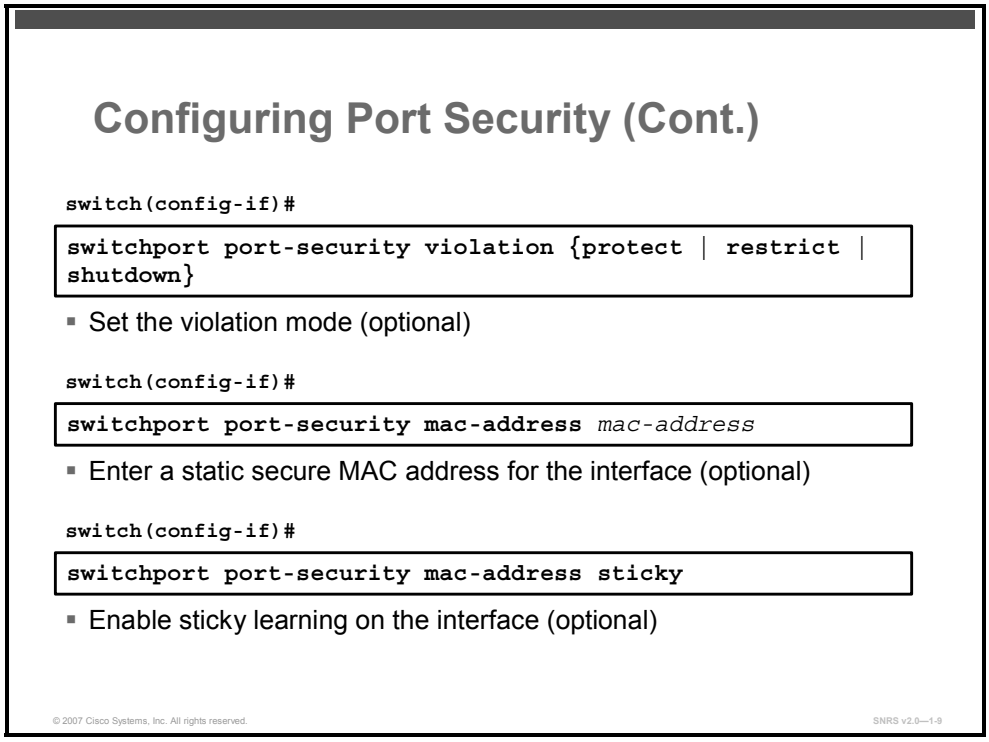
mac-address *mac-address*	(Optional) Specify a secure MAC address for the port by entering a 48-bit MAC address. You can add additional secure MAC addresses up to the maximum value configured.
mac-address sticky [*mac-address*]	(Optional) Enable the interface for *sticky learning* by entering only the **mac-address sticky** keywords. When sticky learning is enabled, the interface adds all secure MAC addresses that are dynamically learned to the running configuration and converts these addresses to sticky secure MAC addresses. Specify a sticky secure MAC address by entering the **mac-address sticky** *mac-address* keywords. **Note** Although you can specify a sticky secure MAC address by entering the **mac-address sticky** *mac-address* keywords, we recommend using the **mac-address** *mac-address* interface configuration command to enter static secure MAC addresses.
maximum *value*	(Optional) Set the maximum number of secure MAC addresses for the interface. The range is 1 to 132. The default is 1.
violation	(Optional) Set the security violation mode or the action to be taken if port security is violated. The default is **shutdown**.
protect	(Optional) Set the security violation protect mode. 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	(Optional) Set the security violation restrict mode. 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, an SNMP trap is sent, a syslog message is logged, and the violation counter increments.
shutdown	(Optional) Set the security violation shutdown mode. In this mode, a port security violation causes the interface to immediately become error-disabled and turns off the port LED. It also sends an SNMP trap, logs a syslog message, and increments 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 shut down** interface configuration commands.

Step 4 (Optional) Set the maximum number of secure MAC addresses for the interface.

```
switch(config-if)# switchport port-security maximum
value
```

Note The range is 1 to 132; the default is 1.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.



Step 5 (Optional) Set the violation mode. This is the action to be taken when a security violation is detected:

```
switch(config-if)# switchport port-security violation
{protect | restrict | shutdown}
```

Note 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.

Step 6 (Optional) Enter a static secure MAC address for the interface with this command:

```
switch(config-if)# switchport port-security mac-address
mac-address
```

Note Repeat this command as many times as necessary for each secure MAC address.

Step 7 (Optional) Enable sticky learning on the interface with this command:

```
switch(config-if)# switchport port-security mac-address sticky
```

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Returning to Default Configuration

Use the **no switchport port-security** interface configuration command to return the interface to the default condition as being not a secure port. The sticky secure addresses remain part of the running configuration.

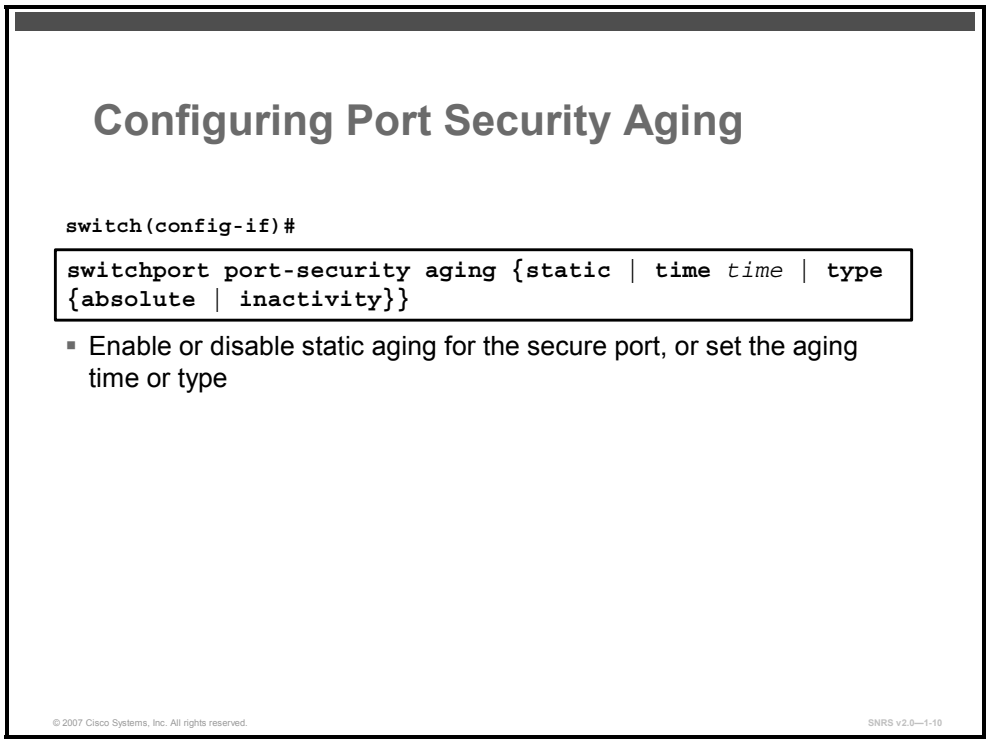
Use the **no switchport port-security maximum** *value* interface configuration command to return the interface to the default number of secure MAC addresses.

Use the **no switchport port-security violation** {**protect** | **restrict**} interface configuration command to return the violation mode to the default condition (shutdown mode).

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Enabling and Configuring Port Security Aging

This section describes how to enable and configure port security aging.



You can use port security aging to set the aging time for static and dynamic secure addresses on a port.

Here are the two types of aging supported per port:

- **Absolute:** The secure addresses on the port are deleted after the specified aging time.
- **Inactivity:** The secure addresses on the port are deleted only if the secure addresses are inactive for the specified aging time.

You can use this feature to remove and add secure MAC addresses on a secure port without manually deleting the existing secure MAC addresses and to still limit the number of secure addresses on a port. Also, you can enable or disable the aging of statically configured secure addresses on a per-port basis.

Configuring Port Security Aging

Use the **switchport port-security aging** command to enable or disable static aging for the secure port, or set the aging time or type.

```
switch(config-if)# switchport port-security aging {static |
time time | type {absolute | inactivity}}
```

Enter **static** to enable aging for the statically configured secure addresses on this port.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

For time, specify the aging time for this port. The valid range is from 0 to 1440 minutes. If the time is equal to 0, aging is disabled for this port.

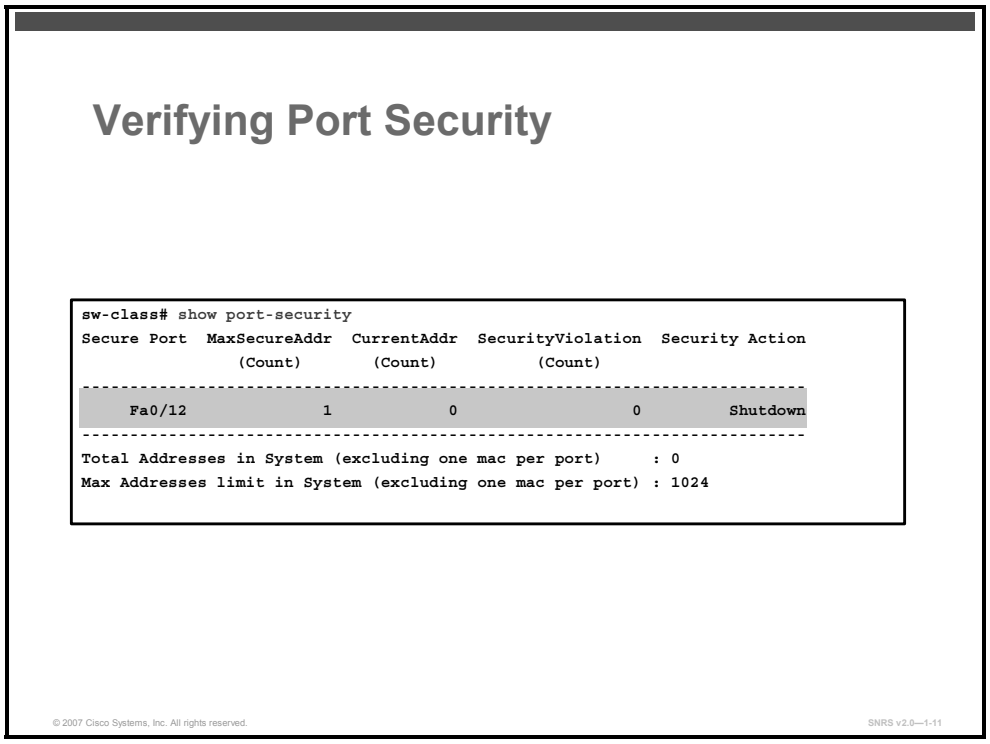
For type, choose one of these keywords:

- **absolute:** All the secure addresses on this port age out after the specified time (minutes) lapses and are removed from the secure address list.
- **inactivity:** The secure addresses on this port age out only if there is no data traffic from the secure source addresses for the specified time period.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Verifying Port Security

This topic describes some commands used to verify port security configuration and operations.



Use the **show port-security** command to view port security settings for the switch including violation count, configured interfaces, and security violation actions.

```
          sw-class# show port-security
Secure Port  MaxSecureAddr  CurrentAddr  SecurityViolation  Security Action
               (Count)       (Count)          (Count)
---------------------------------------------------------------------------
     Fa0/12              1            0                  0         Shutdown
---------------------------------------------------------------------------
Total Addresses in System (excluding one mac per port)     : 0
Max Addresses limit in System (excluding one mac per port) : 1024
```

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Verifying Port Security (Cont.)

```
sw-class# show port-security interface fa0/12
Port Security              : Enabled
Port Status                : Secure-down
Violation Mode             : Shutdown
Aging Time                 : 0 mins
Aging Type                 : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 1
Total MAC Addresses        : 1
Configured MAC Addresses   : 1
Sticky MAC Addresses       : 0
Last Source Address        : 0000.0000.0000
Security Violation Count   : 0
```

© 2007 Cisco Systems, Inc. All rights reserved.

Use the **show port-security** [**interface** *interface-id*] command to view port security settings for the specified interface, including the maximum allowed number of secure MAC addresses for each interface, the number of secure MAC addresses on the interface, the number of security violations that have occurred, and the violation mode.

```
sw-class# show port-security interface fa0/12
Port Security              : Enabled
Port Status                : Secure-down
Violation Mode             : Shutdown
Aging Time                 : 0 mins
Aging Type                 : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 1
Total MAC Addresses        : 1
Configured MAC Addresses   : 1
Sticky MAC Addresses       : 0
Last Source Address        : 0000.0000.0000
Security Violation Count   : 0
```


The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Verifying Port Security (Cont.)

```
sw-class# show port-security address
          Secure Mac Address Table
-----------------------------------------------------------------
Vlan    Mac Address       Type                Ports   Remaining Age
                                                        (mins)
----    -----------       ----                -----   -------------
   1    0000.ffff.aaaa    SecureConfigured    Fa0/12       -
-----------------------------------------------------------------
Total Addresses in System (excluding one mac per port)     : 0
Max Addresses limit in System (excluding one mac per port) : 1024
```

© 2007 Cisco Systems, Inc. All rights reserved.

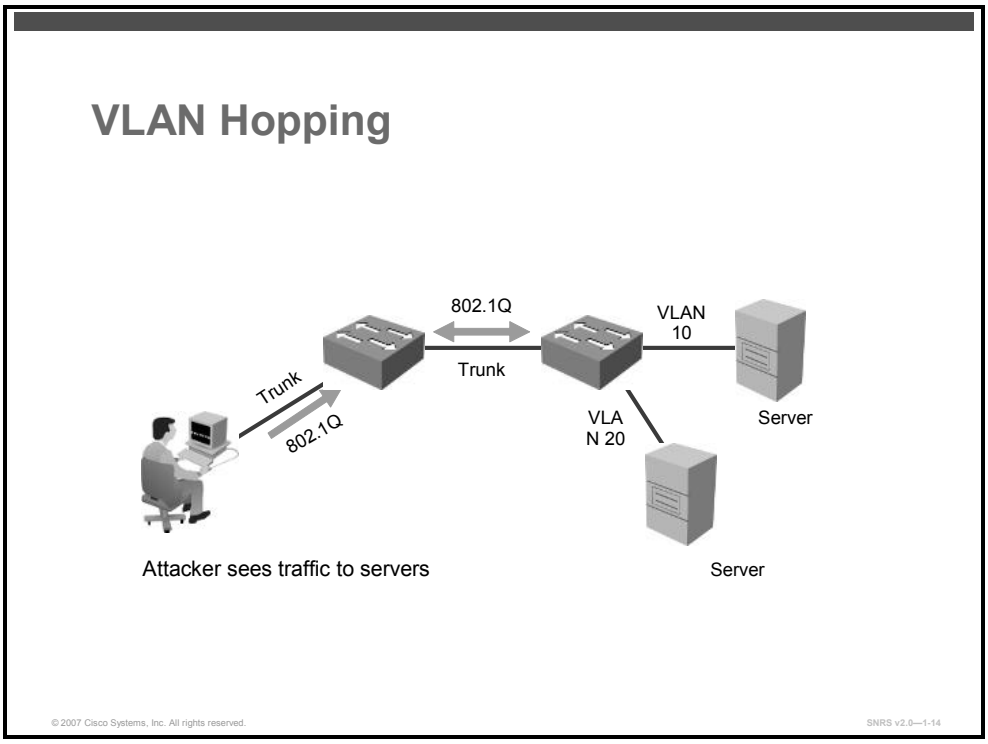
Use the **show port-security** [**interface** *interface-id*] **address** command to view all secure MAC addresses configured on all switch interfaces or on a specified interface with aging information for each address.

```
sw-class# show port-security address
          Secure Mac Address Table
-----------------------------------------------------------------
Vlan    Mac Address       Type                Ports   Remaining Age
                                                        (mins)
----    -----------       ----                -----   -------------
   1    0000.ffff.aaaa    SecureConfigured    Fa0/12       -
-----------------------------------------------------------------
Total Addresses in System (excluding one mac per port)     : 0
Max Addresses limit in System (excluding one mac per port) : 1024
```


The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

VLAN Hopping Attacks

This topic describes VLAN hopping attacks.



A VLAN hopping attack occurs when an attacker sends out packets destined for a system on a different VLAN that cannot normally be reached by the attacker. This traffic is tagged with a different VLAN ID (VID) to which the attacker belongs. Or, the attacking system may be trying to behave like a switch and negotiate trunking so that the attacker can send and receive traffic between other VLANs.

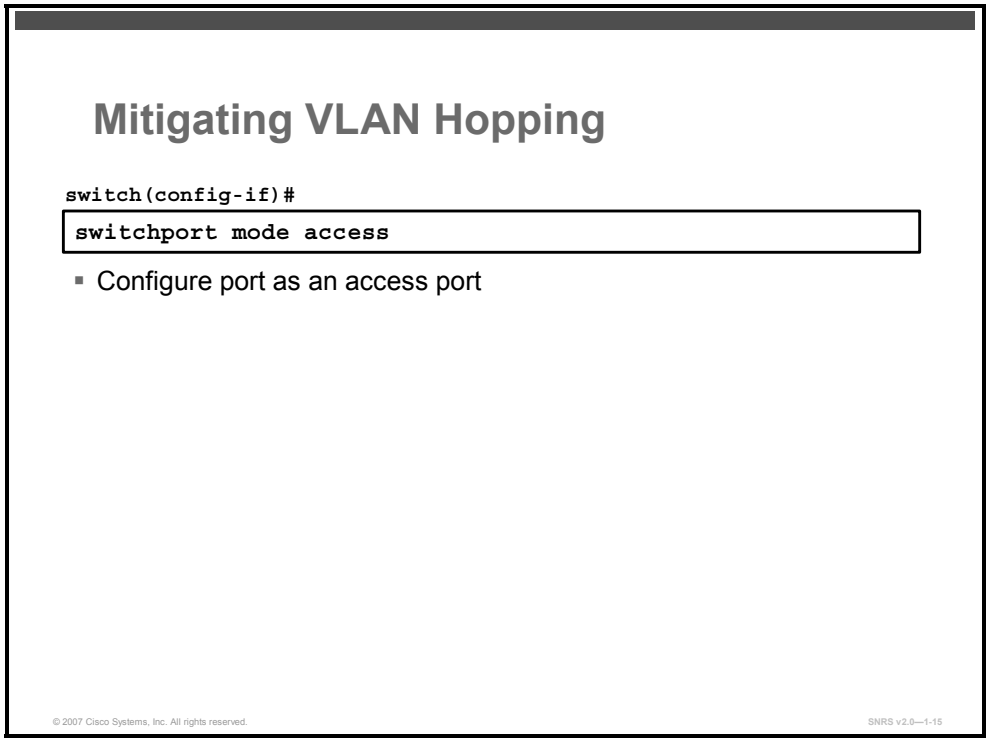
Switch Spoofing

In a VLAN hopping attack, the network attacker configures a system to spoof itself as a switch. This requires that the network attacker be capable of emulating either Inter-Switch Link (ISL) or 802.1Q signaling along with DTP signaling. Using this method, a network attacker can make a system appear to be a switch with a trunk port. If successful, the attacking system then becomes a member of all VLANs.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Mitigating VLAN Hopping Attacks

This subtopic describes a strategy to mitigate a VLAN hopping attack.



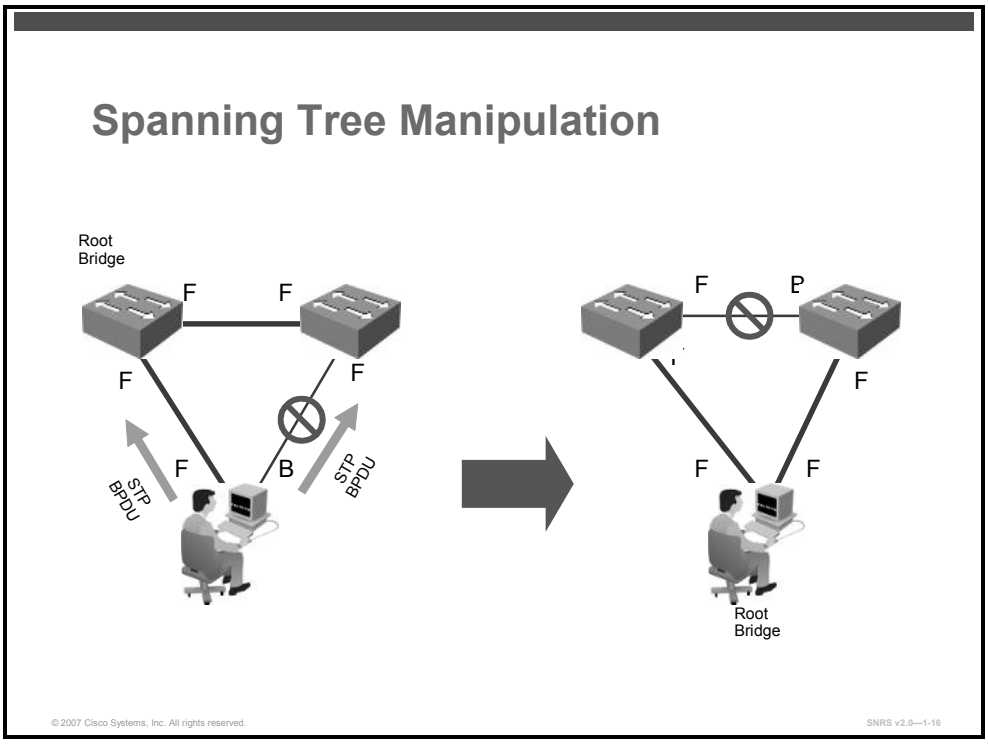
You can mitigate VLAN hopping attacks by putting all user ports into access mode using the **switchport mode access** command. Several other modifications to the VLAN configuration are also recommended. One of the more important elements is to use dedicated VLAN IDs as the active (allowed) VLANs for all trunk ports. Also, disable all unused switch ports and place them in an unused VLAN.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

STP Vulnerabilities

This topic describes spanning tree manipulation.



The diagram illustrates how a network attacker can use STP to change the topology of a network so that it appears that the network attacker host is a root bridge with a higher priority.

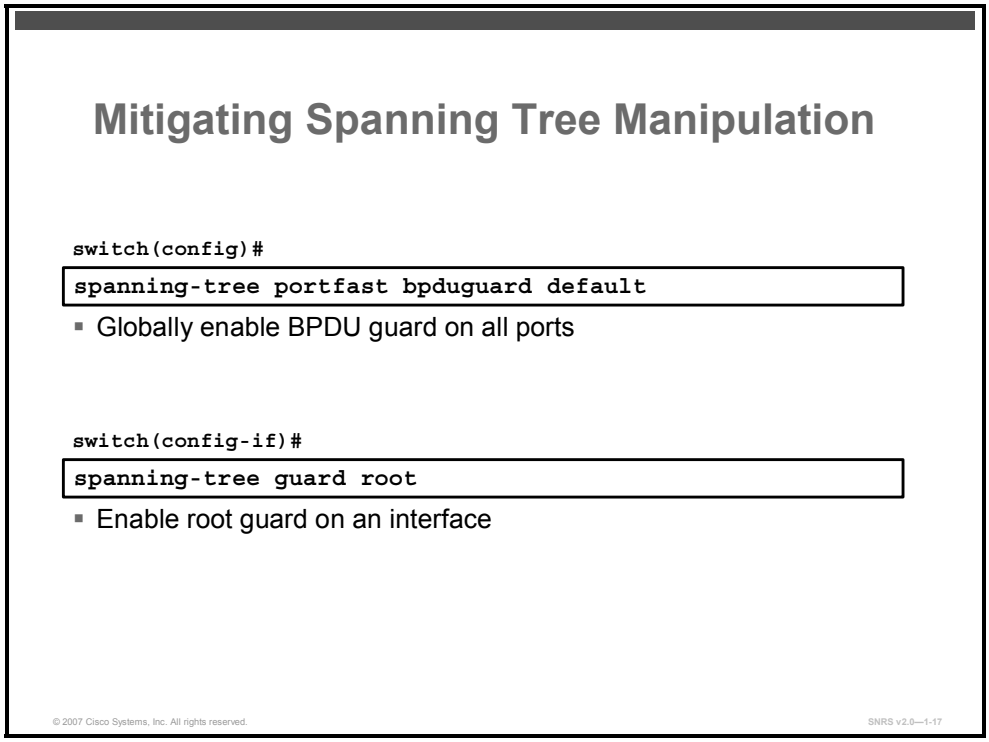
One attack against switches involves intercepting traffic by attacking the STP. This protocol is used in switched networks to prevent the creation of bridging loops in an Ethernet network topology. Upon bootup, the switches begin a process of determining a loop-free topology. The switches identify one switch as a root bridge and block all other redundant data paths.

By manipulating the STP root bridge determination calculations, network attackers hope to spoof their system as the root bridge in the topology. To do this, the network attacker broadcasts out STP configuration and topology change bridge protocol data units (BPDUs) in an attempt to force spanning-tree recalculations. The BPDUs sent out by the system of the network attacker announce that the attacking system has a lower bridge priority. If successful, the network attacker becomes the root bridge and can see a variety of frames. By transmitting spoofed STP frames, the network attacker causes the switches to initiate spanning-tree recalculations that then result in having all of the interfaces in the system of the network attacker to be in the forwarding mode.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Mitigating STP Vulnerabilities

This subtopic describes a strategy to mitigate an STP attack.

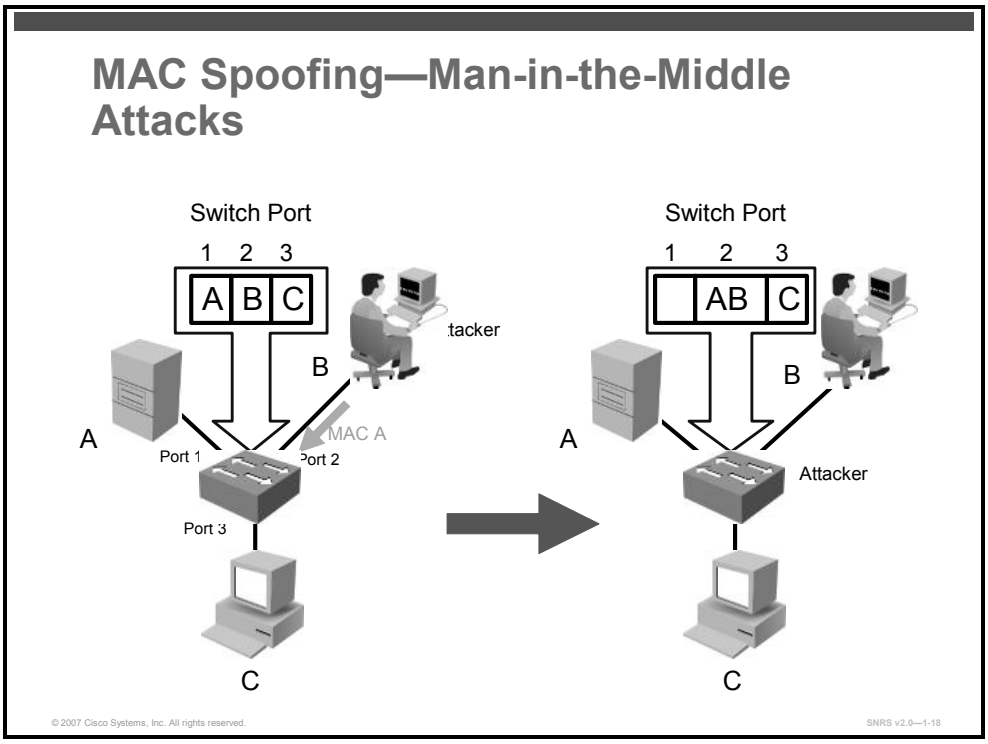


To mitigate STP manipulation, use the **root guard** and the **bpdu-guard** enhancement commands to enforce the placement of the root bridge in the network and enforce the STP domain borders. The root guard feature is designed to provide a way to enforce the root-bridge placement in the network. The STP BPDU guard is designed to allow network designers to keep the active network topology predictable. While BPDU guard may seem unnecessary given that the administrator can set the bridge priority to zero, there is still no guarantee that it will be elected as the root bridge. This is because there might be a bridge with priority zero and a lower bridge ID (BID). BPDU guard is best deployed toward user-facing ports to prevent rogue switch network extensions by an attacker.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

MAC Spoofing: Man-in-the-Middle Attacks

This topic describes MAC spoofing—one type of man-in-the-middle attacks.



MAC spoofing attacks involve the use of a known MAC address of another host to attempt to make the target switch forward frames destined for the remote host to the network attacker. By sending a single frame with the source Ethernet address of the other host, the network attacker overwrites the CAM table entry so that the switch forwards packets destined for the host to the network attacker. Until the host sends traffic, it will not receive any traffic. When the host sends out traffic, the CAM table entry is rewritten once more so that it moves back to the original port.

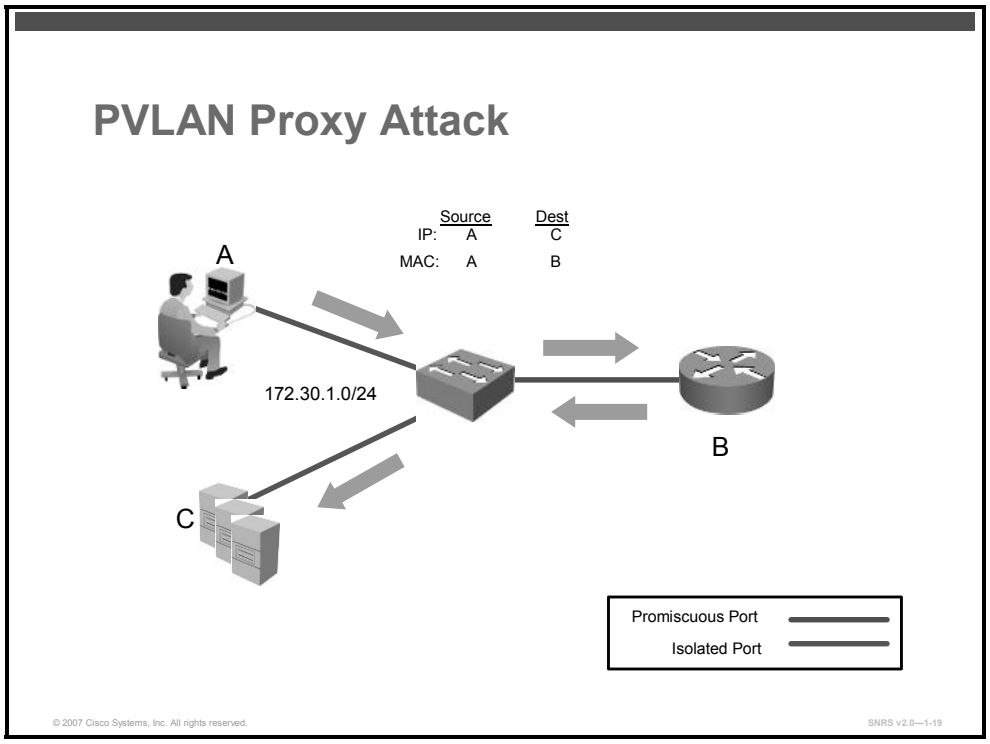
The diagram shows how MAC spoofing works. In the beginning, the switch has learned that host A is on port 1, host B is on port 2, and host C is on port 3. Host B sends out a packet identifying itself as the address of host B but with the MAC address of host A or another packet with the same IP address and MAC address combination. This traffic causes the switch to move 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.

This attack can also be mitigated using port security.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

PVLAN Vulnerabilities

This topic describes PVLAN attacks.



Even though PVLANs are a common mechanism to restrict communications between systems on the same logical IP subnet, they are not always 100 percent secure. PVLANs work by limiting the ports within a VLAN that can communicate with other ports in the same VLAN. Isolated ports within a VLAN can communicate only with promiscuous ports. Community ports can communicate only with other members of the same community and promiscuous ports. Promiscuous ports can communicate with any port. One network attack capable of bypassing the network security of PVLANs involves the use of a proxy to bypass access restrictions to a PVLAN.

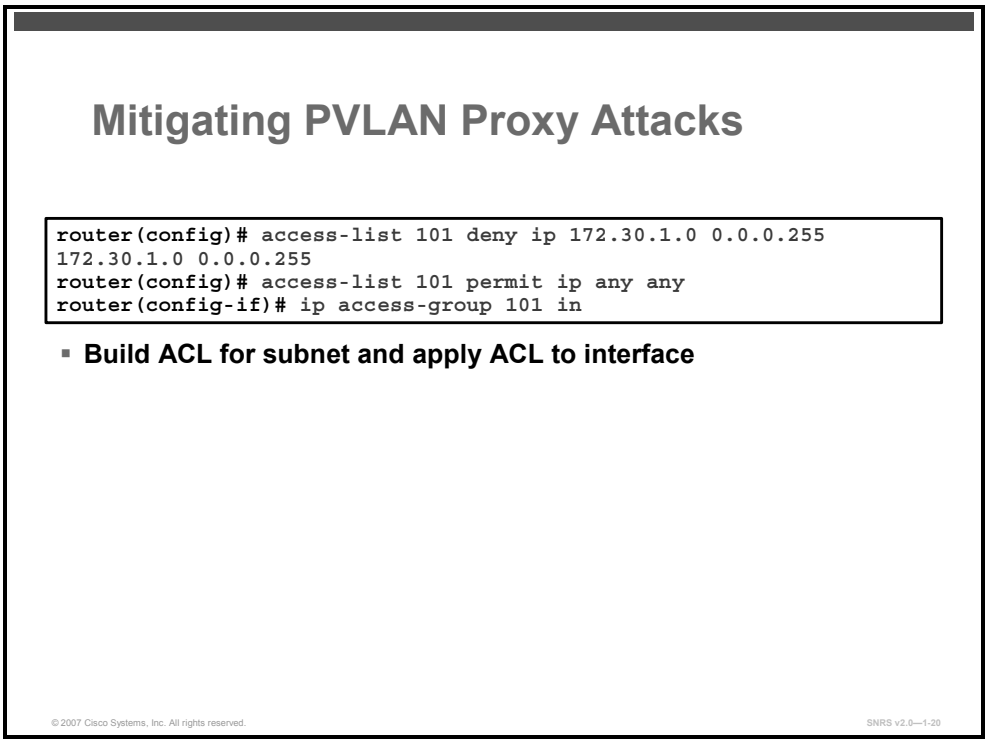
Proxy Attack

In this network attack against private VLANs, frames are forwarded to a host on the network connected to a promiscuous port such as a router. In the diagram, the network attacker sends a packet with the source IP and MAC address of attacker device, a destination IP address of the target system, but a destination MAC address of the router. The switch forwards the frame to the switch port of the router. The router routes the traffic, rewrites the destination MAC address as that of the target, and sends the packet back out. Now the packet has the proper format and is forwarded to the target system. This network attack allows only for unidirectional traffic because any attempt by the target to send traffic back will be blocked by the PVLAN configuration. If both hosts are compromised, static Address Resolution Protocol (ARP) entries could be used to allow bidirectional traffic. This scenario is not a PVLAN vulnerability because all of the rules of PVLANs were enforced; however, the network security was bypassed.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Mitigating PVLAN Vulnerabilities

This subtopic describes a strategy to mitigate PVLAN attacks.

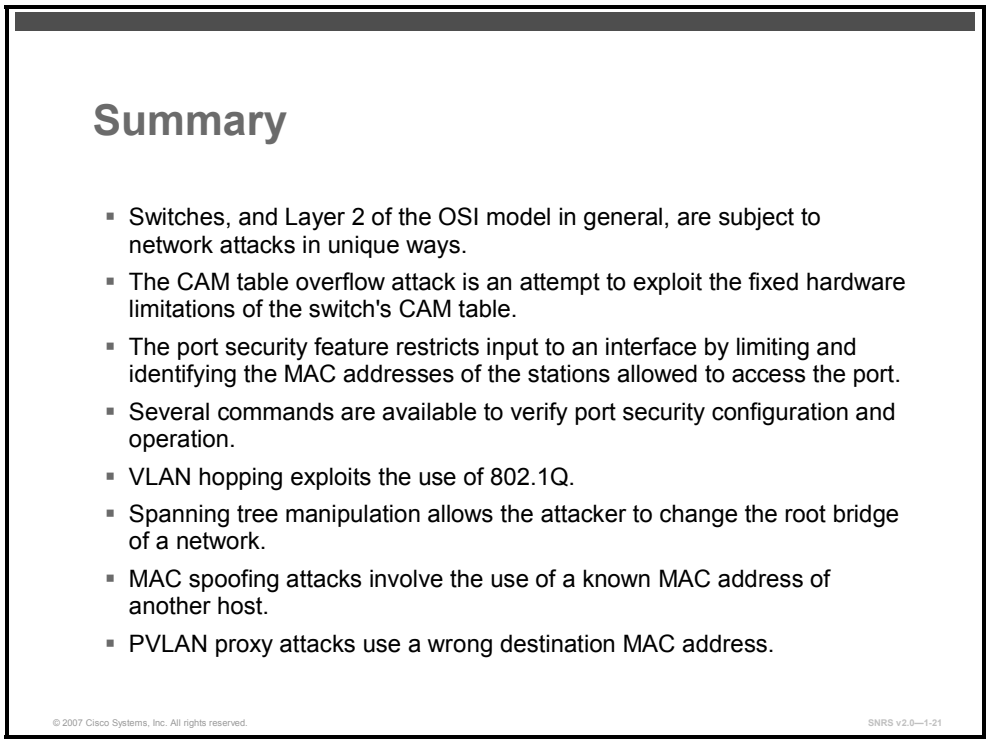


Configure access control lists (ACLs) on the router port to mitigate PVLAN attacks. An example of using ACLs on the router port is if a server farm segment existed on subnet 172.30.1.0/24 and target C was in the server farm, then configuring the ACL (in the previous diagram) on the default gateway would mitigate the PVLAN proxy attack.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Summary

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



References

For additional information, refer to these resources:

- *SAFE Layer 2 Security In-depth Version 2:* http://www.cisco.com/en/US/partner/netsol/ns340/ns394/ns171/ns128/networking_solutions_white_paper09186a008014870f.shtml#wp1002210.
- *Catalyst 2950 and Catalyst 2955 Switch Command Reference, 12.1(22)EA7:* http://www.cisco.com/en/US/products/hw/switches/ps628/products_command_reference_chapter09186a0080647ba5.html#wp1862599.
- *Catalyst 2950 and Catalyst 2955 Switch Software Configuration Guide, 12.1(22)EA7:* http://www.cisco.com/univercd/cc/td/doc/product/lan/cat2950/12122ea7/scg/index.htm.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Lesson 3

Configuring DHCP Snooping

Overview

The DHCP spoofing and DHCP starvation attacks are also types of Layer 2 attacks. In this lesson, you will be introduced to the Cisco IOS DHCP snooping feature. You will get an overview of a DHCP starvation attack and then examine DHCP snooping and how to configure it to mitigate this type of attack.

Objectives

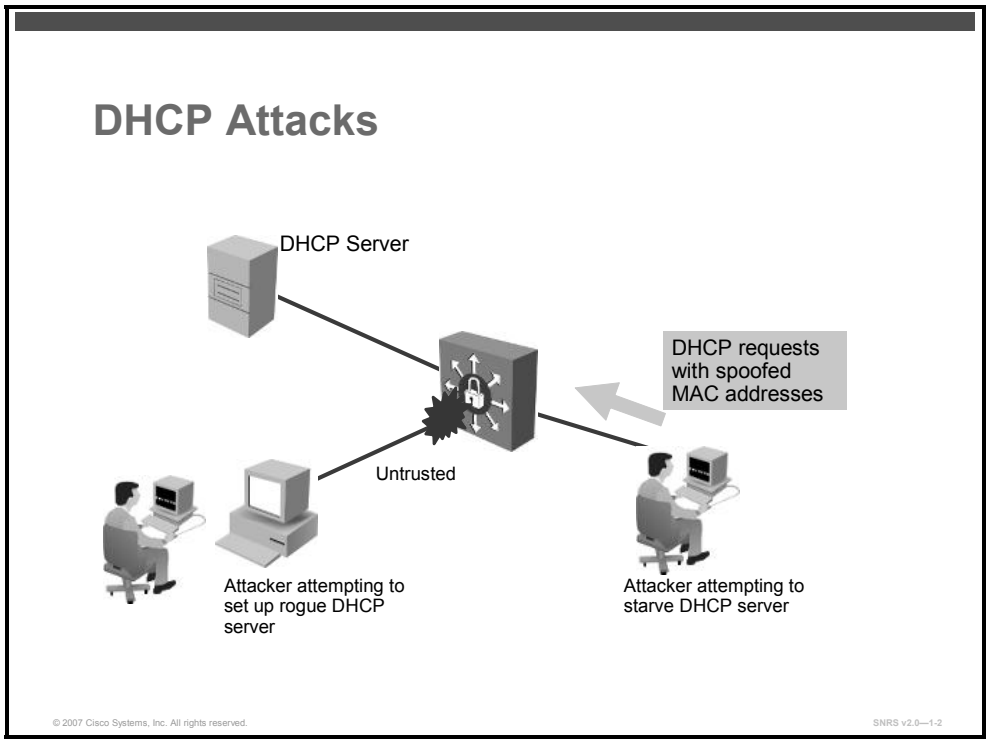
Upon completing this lesson, you will be able to implement DHCP snooping on a Cisco Catalyst switch. This ability includes being able to meet these objectives:

- Describe DHCP spoofing and starvation attacks
- Describe the Cisco DHCP snooping feature
- Describe a method used to mitigate DHCP spoofing and starvation attacks
- Describe the guidelines for DHCP snooping configuration
- Enable and configure DHCP snooping
- Verify DHCP snooping operation

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

DHCP Starvation and Spoofing Attacks

This topic describes the DHCP spoofing and starvation attacks.



A DHCP starvation attack works by broadcasting DHCP requests with spoofed MAC addresses. This is easily achieved with attack tools such as "the gobbler." If enough requests are sent, the network attacker can exhaust the address space available to the DHCP servers for a period of time. This is a simple resource starvation attack just like a synchronization (SYN) flood is a starvation attack. Network attackers can then set up a rogue DHCP server on their system and respond to new DHCP requests from clients on the network. Exhausting all of the DHCP addresses is not *required* to introduce a rogue DHCP server, though, as stated in RFC 2131:

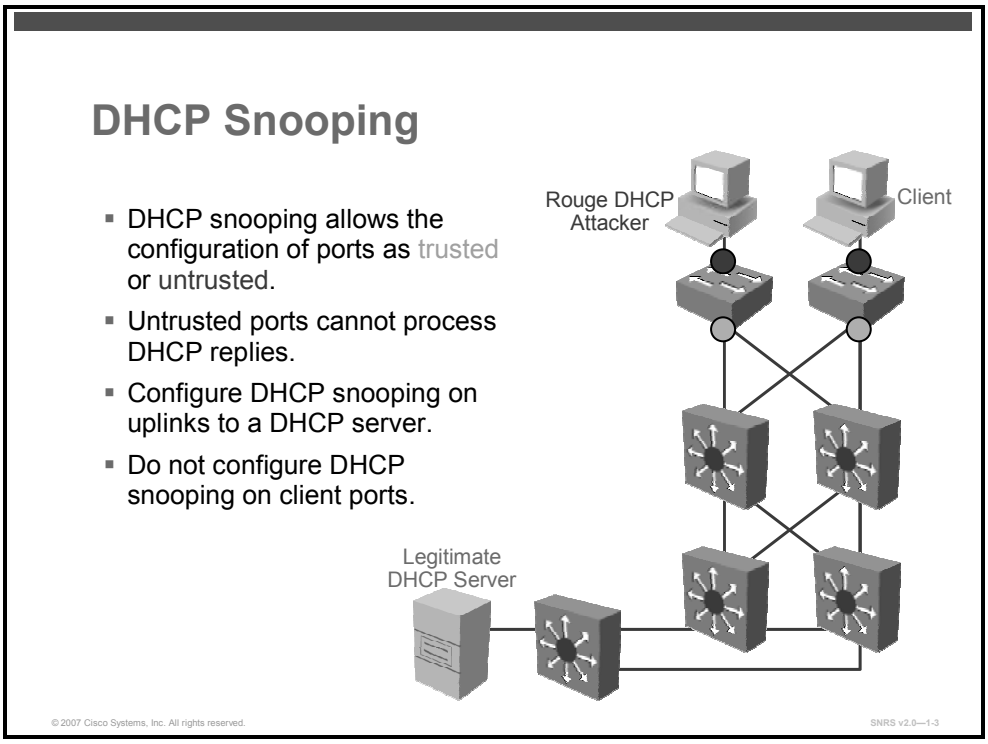
> "The client collects DHCPOFFER messages over a period of time, selects one DHCPOFFER message from the (possibly many) incoming DHCPOFFER messages (for example, the first DHCPOFFER message or the DHCPOFFER message from the previously used server) and extracts the server address from the `server identifier' option in the DHCPOFFER message. The time over which the client collects messages and the mechanism used to select one DHCPOFFER are implementation dependent."

By placing a rogue DHCP server on the network, a network attacker can provide clients with addresses and other network information. Because DHCP responses typically include default gateway and Domain Name System (DNS) server information, network attackers can supply their own system as the default gateway and DNS server resulting in a man-in-the-middle attack.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Understanding DHCP Snooping

This topic describes the Cisco IOS DHCP snooping feature.



DHCP snooping is a DHCP security feature that provides network security by filtering untrusted DHCP messages and by building and maintaining a DHCP snooping binding database, which is also referred to as a DHCP snooping binding table.

DHCP snooping is a Cisco Catalyst switch 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, while 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 the 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 DHCPOFFER, DHCPACK, or DHCPNAK.

Note	For DHCP snooping to function properly, all DHCP servers must be connected to the switch through trusted interfaces.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

When a switch receives a packet on an untrusted interface and the interface belongs to a VLAN in which DHCP snooping is enabled, the switch compares the source MAC address and the DHCP client hardware address. If addresses match (the default), the switch forwards the packet. If the addresses do not match, the switch drops the packet.

The switch drops a DHCP packet when one of these situations occurs:

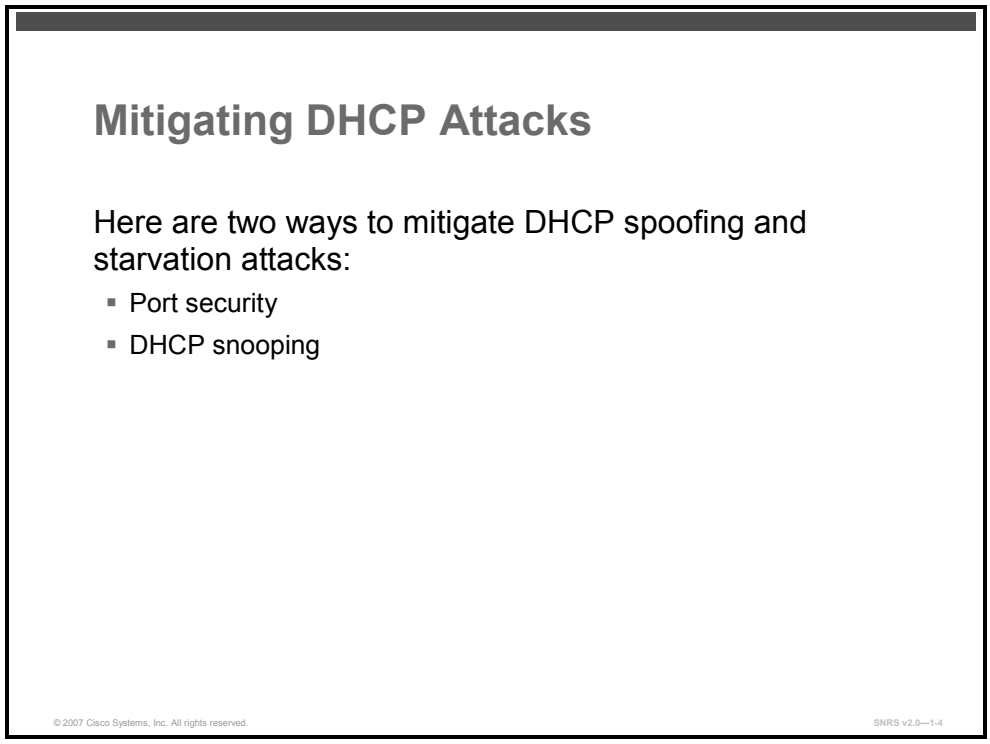
- A packet from a DHCP server, such as a DHCPOFFER, DHCPACK, DHCPNAK, or DHCPLEASEQUERY packet, is received from an untrusted interface.
- A packet is received on an untrusted interface, and the source MAC address and the DHCP client hardware address do not match.
- The switch receives a DHCPRELEASE or DHCPDECLINE broadcast message that contains a MAC address in the DHCP snooping binding table, but the interface information in the binding table does not match the interface on which the message was received.
- A DHCP relay agent forwards a DHCP packet that includes a relay agent IP address that is not 0.0.0.0, or the relay agent forwards a packet that includes DHCP option 82 information to an untrusted port.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Mitigating DHCP Attacks

This topic describes strategies to mitigate DHCP spoofing and starvation attacks.



In this lesson, the following two ways to mitigate DHCP spoofing and starvation attacks are discussed.

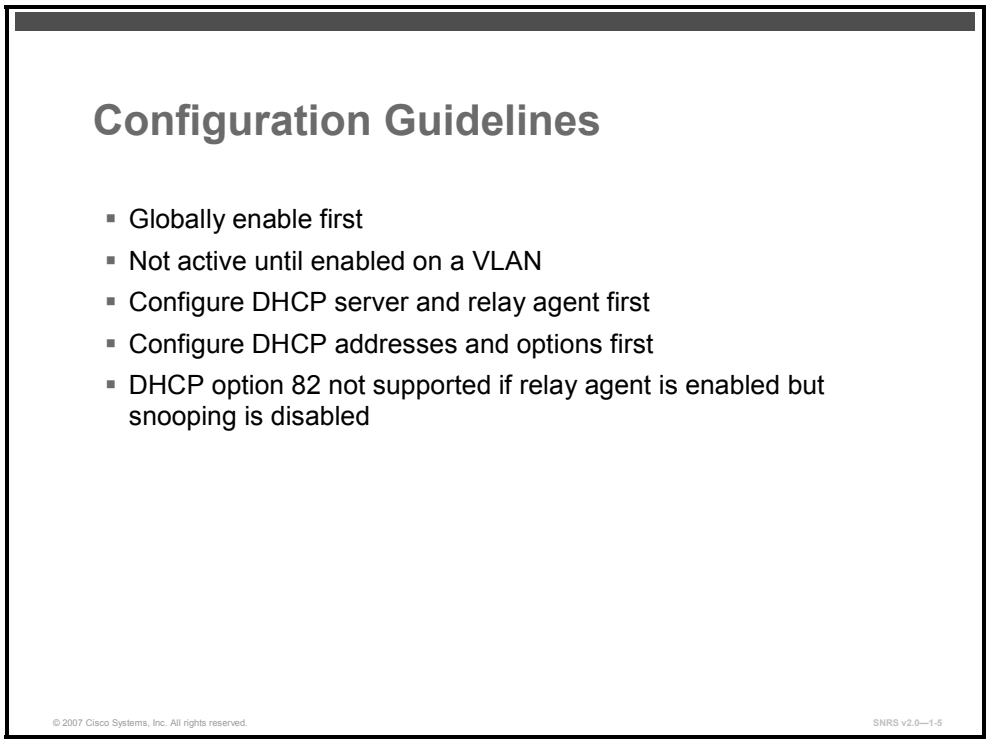
- **Port security:** The techniques that mitigate CAM table flooding also mitigate DHCP starvation by limiting the number of MAC addresses on a switch port. You would use the **port-security** command to set the MAC address of a valid DHCP server on a switch port to prevent any other device from connecting to that trusted port.
- **DHCP snooping:** DHCP snooping is a security feature that filters untrusted DHCP messages and builds and maintains a DHCP snooping binding table.

Because covered port security has already been covered, the next sections will cover the configuration of DHCP snooping.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

DHCP Snooping Configuration Guidelines

This topic describes some DHCP snooping configuration guidelines.



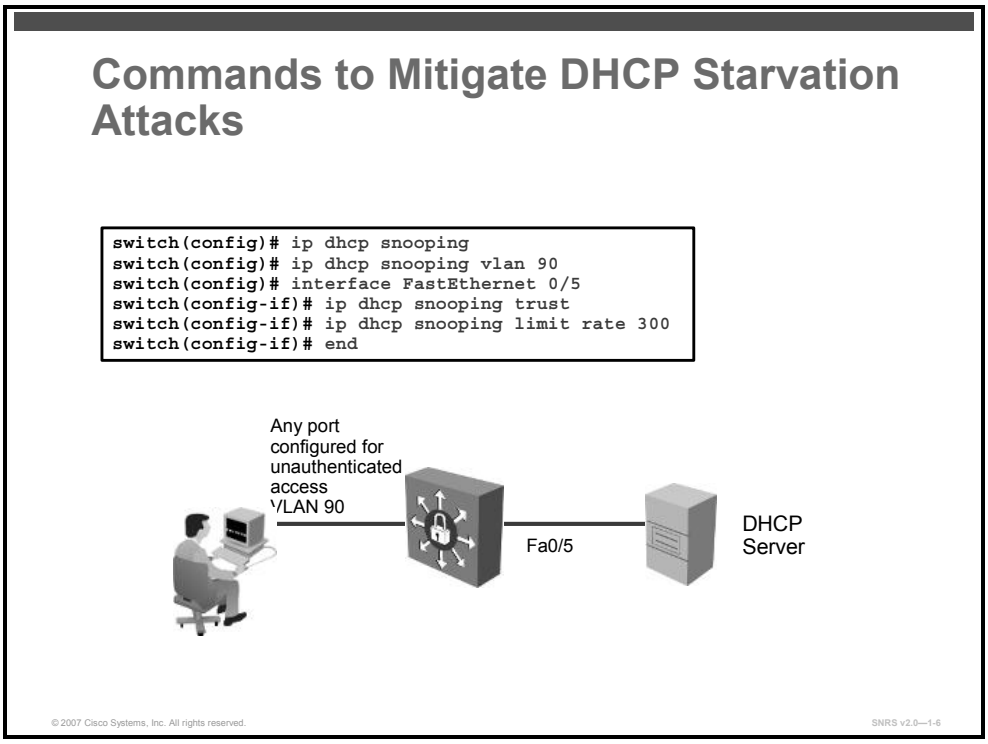
Here are some guidelines to use when configuring DHCP snooping:

- You must globally enable DHCP snooping on the switch.
- DHCP snooping is not active until DHCP snooping is enabled on a VLAN.
- Before globally enabling DHCP snooping on the switch, make sure that the devices acting as the DHCP server and the DHCP relay agent are configured and enabled.
- Before configuring the DHCP snooping information option on your switch, be sure to configure the device that is acting as the DHCP server. For example, you must specify the IP addresses that the DHCP server can assign or exclude, or you must configure DHCP options for these devices.
- If the DHCP relay agent is enabled but DHCP snooping is disabled, the DHCP option 82 data insertion feature is not supported.
- If a switch port is connected to a DHCP server, configure a port as trusted by entering the **ip dhcp snooping trust** interface configuration command.
- If a switch port is connected to a DHCP client, configure a port as untrusted by entering the **no ip dhcp snooping trust** interface configuration command.
- Do not enter the **ip dhcp snooping information option allow-untrusted** command on an aggregation switch to which an untrusted device is connected. If you enter this command, an untrusted device might spoof the DHCP option 82 information.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Enabling and Configuring DHCP Snooping

This topic describes how to enable and configure DHCP snooping.



To enable and configure DHCP snooping, follow these steps:

Step 1 Globally enable DHCP snooping.

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

Step 2 Enable DHCP snooping on a VLAN or range of VLANs.

```
switch(config)# ip dhcp snooping vlan vlan-range
```

Syntax Description

vlan *number*	(Optional) Single VLAN number or a range of VLANs; valid values are from 1 to 4094

DHCP snooping is enabled on a VLAN only if both the global snooping and the VLAN snooping are enabled.

Step 3 Change to interface configuration mode.

```
switch(config)# interface interface-id
```

Step 4 Configure the interface as trusted where a DHCP server is connected to the switch.

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

Step 5 (Optional; *see Note*) Configure the number of DHCP packets per second that an interface can receive.

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


The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Note	Typically, the rate limit applies to the untrusted interfaces. If you want to set up rate limiting for the trusted interfaces, note that the trusted interfaces aggregate all DHCP traffic in the switch, and you will need to adjust the rate limit of the interfaces to a higher value.

Syntax Description

rate	Number of DHCP messages that a switch can receive per second

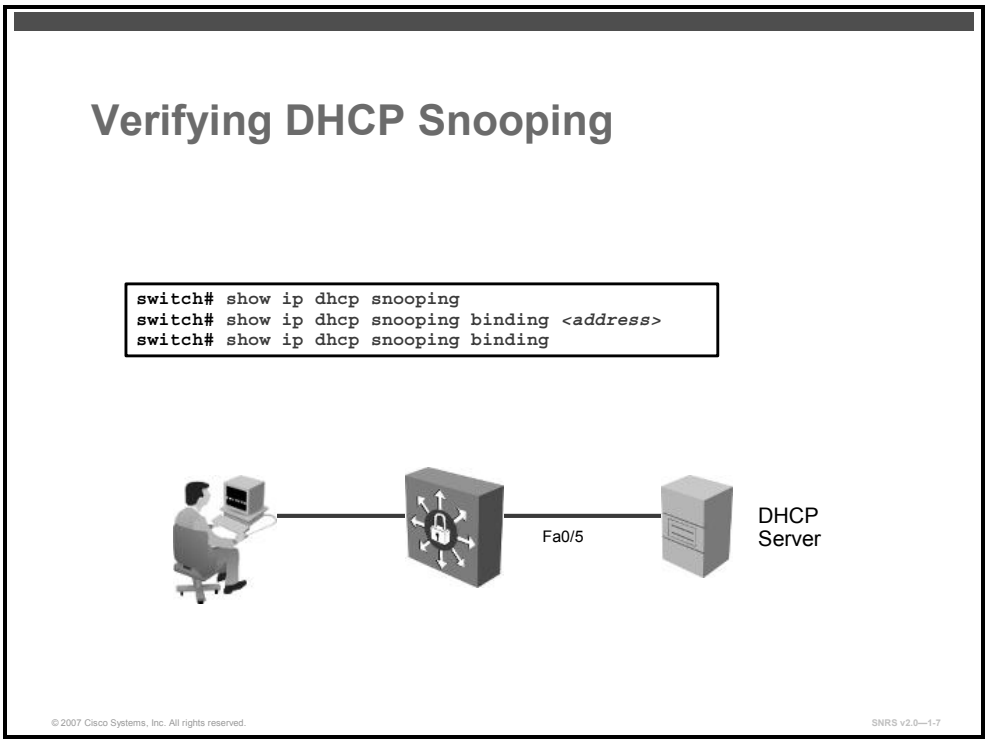
To disable DHCP snooping, use the **no ip dhcp snooping** global configuration command. To disable DHCP snooping on a VLAN or range of VLANs, use the **no ip dhcp snooping vlan** *vlan-id* global configuration command.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Verifying DHCP Snooping

This topic describes how to verify DHCP snooping configuration and operations.



The following commands are available to verify DHCP snooping configuration and operation.

- To display the DHCP snooping configuration for a switch, use the **show ip dhcp snooping** command.

```
switch# show ip dhcp snooping
```

- To display only the dynamically configured bindings in the DHCP snooping binding database, use the **show ip dhcp snooping binding** command.

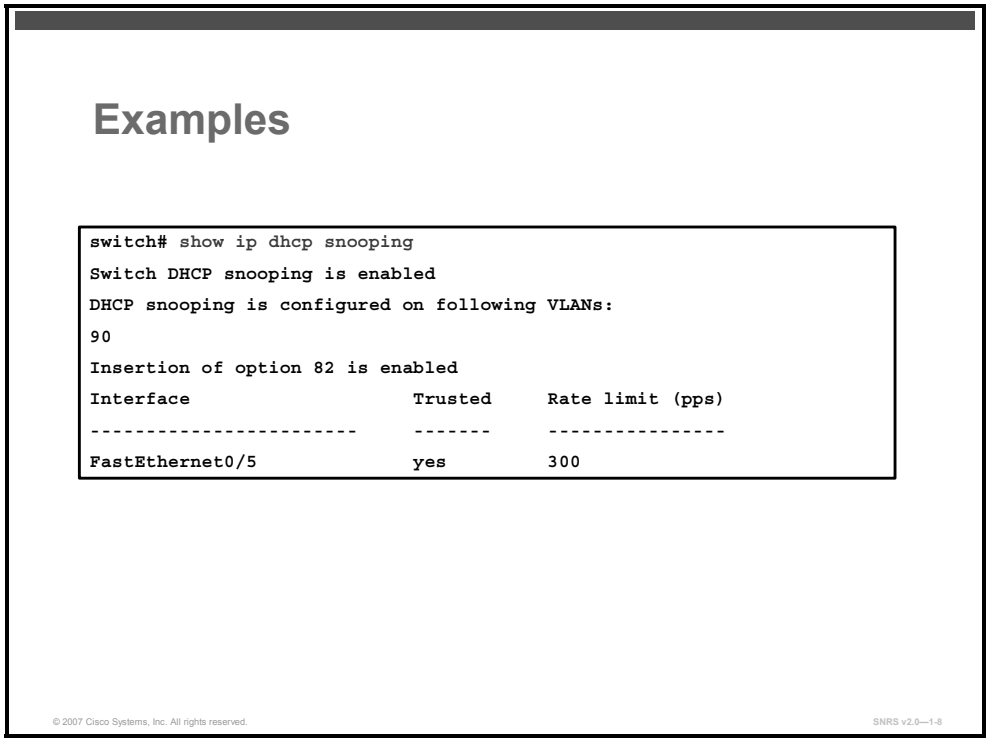
```
switch# show ip dhcp snooping binding
```

This command is used to display DHCP binding information for IP address assignment and subnet allocation. If the address is not specified, all address bindings are shown. Otherwise, only the binding for the specified client is displayed. The output from this command displays binding information for individual IP address assignment and allocated subnets. The output that is generated for DHCP IP address assignment and subnet allocation is almost identical, except that subnet leases display an IP address followed by the subnet mask (which shows the size of the allocated subnet). Bindings for an individual IP address only display an IP address and are not followed by a subnet mask.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Examples

Here are some examples of DHCP snooping verification.

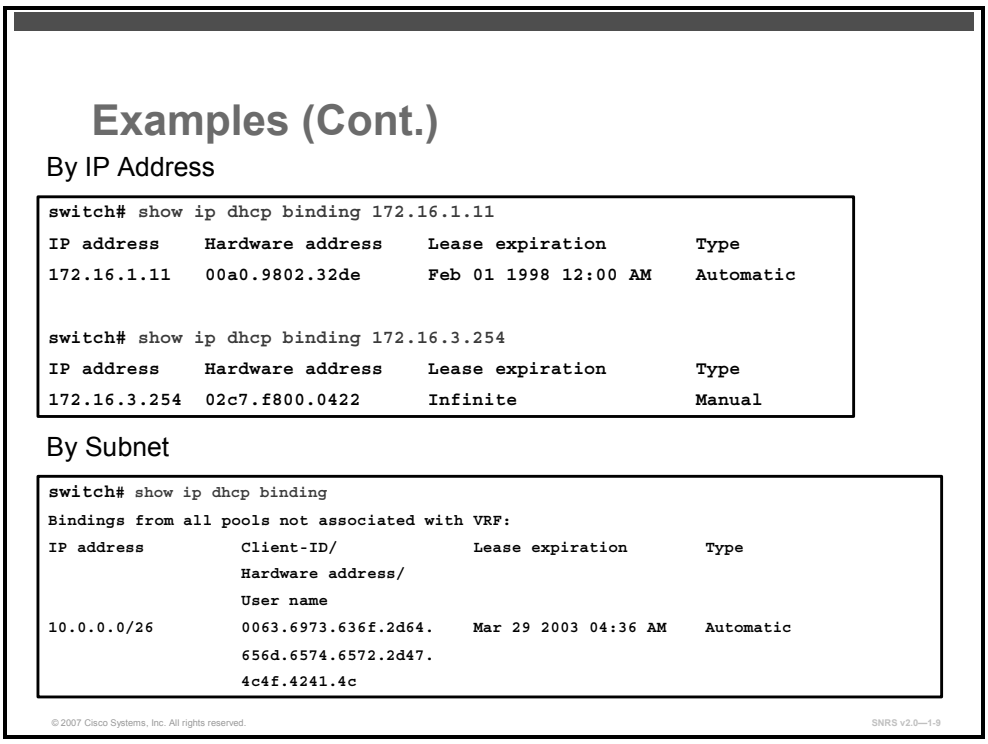


This example displays DHCP snooping configuration information on the switch.

The next example shows the output of the **show ip dhcp binding** command.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.



This example displays the DHCP bindings by IP address and subnet.

IP Address Assignment Example

The example shows the DHCP binding address parameters, including an IP address, an associated MAC address, a lease expiration date, and the type of address assignment that have occurred. The table describes the significant fields shown in the display.

show ip dhcp binding <*address*> Field Descriptions

Field	Description
IP address	The IP address of the host as recorded on the DHCP server
Hardware address	The MAC address or client identifier of the host as recorded on the DHCP server
Lease expiration	The lease expiration date and time of the IP address of the host
Type	The manner in which the IP address was assigned to the host

Subnet Allocation Example

The example shows the subnet lease to MAC address mapping, the lease expiration, and the lease type (subnet lease bindings are configured to be automatically created and released by default). The output that is generated for DHCP IP address assignment and subnet allocation is almost identical, except that subnet leases display an IP address followed by the subnet mask (which shows the size of the allocated subnet) in classless interdomain routing (CIDR) bit-count notation. Bindings for an individual IP address only display an IP address and are not followed by a subnet mask. The next table describes the significant fields shown in the display.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

show ip dhcp binding Field Descriptions

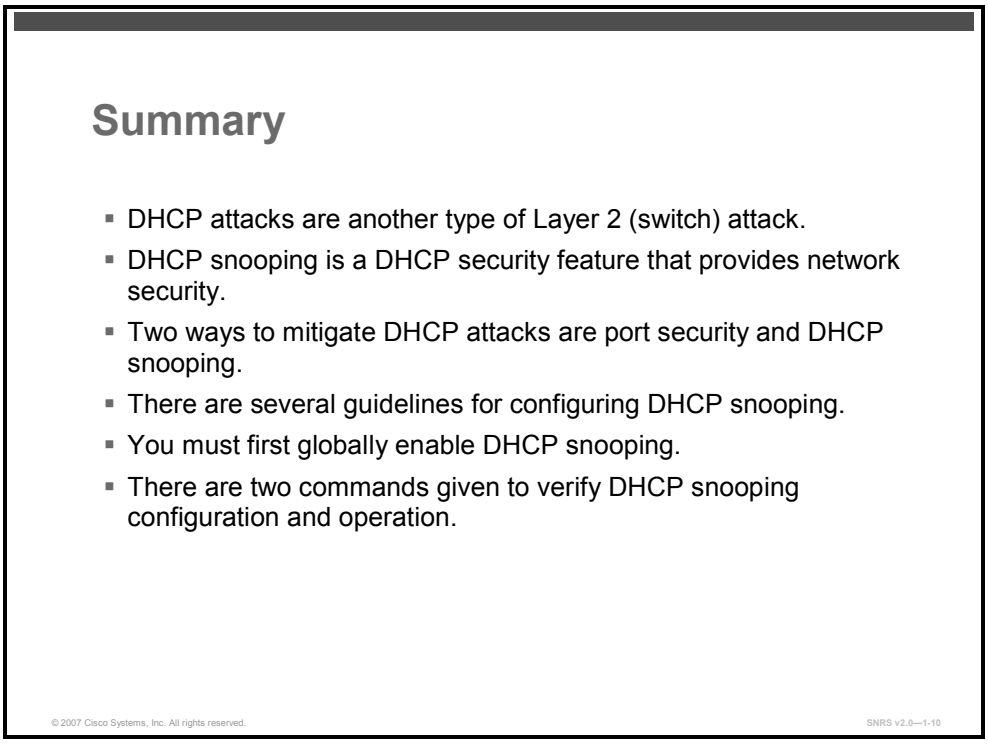
Field	Description
IP address	The IP address of the host as recorded on the DHCP server The subnet that follows the IP address (/26) in the example defines this binding as a subnet allocation binding.
Hardware address	The MAC address or client identifier of the host as recorded on the DHCP server
Lease expiration	The lease expiration date and time of the IP address of the host
Type	The manner in which the IP address was assigned to the host

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Summary

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



References

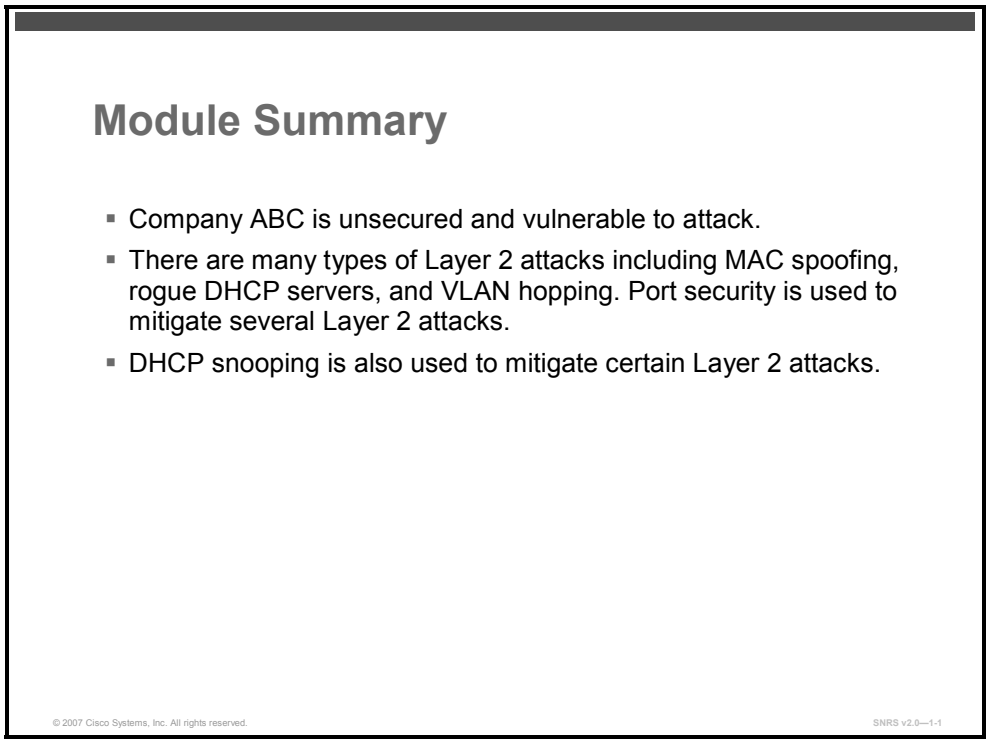
For additional information, refer to these resources:

- *SAFE Layer 2 Security In-depth Version 2:* http://www.cisco.com/en/US/partner/netsol/ns340/ns394/ns171/ns128/networking_solutions_white_paper09186a008014870f.shtml#wp1002210.
- *Catalyst 2950 and Catalyst 2955 Switch Software Configuration Guide, 12.1(22)EA7: http*://www.cisco.com/univercd/cc/td/doc/product/lan/cat2950/12122ea7/scg/index.htm.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Module Summary

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



In this module, you were introduced to an unsecured network belonging to Company ABC. Layer 2 attacks were examined and various types of mitigation strategies were discussed. Port security and DHCP snooping were among the strategies used to mitigate these Layer 2 attacks on the network. You were given an example of a common network topology in the business environment of today, and you will use this course to build a secure network using Cisco IOS security features.

References

For additional information, refer to these resources:

- *SAFE Layer 2 Security In-depth Version 2:* http://www.cisco.com/en/US/partner/netsol/ns340/ns394/ns171/ns128/networking_solutions_white_paper09186a008014870f.shtml#wp1002210
- *Catalyst 2950 and Catalyst 2955 Switch Command Reference, 12.1(22)EA7:* http://www.cisco.com/en/US/products/hw/switches/ps628/products_command_reference_chapter09186a0080647ba5.html#wp1862599
- *Catalyst 2950 and Catalyst 2955 Switch Software Configuration Guide, 12.1(22)EA7:* http://www.cisco.com/univercd/cc/td/doc/product/lan/cat2950/12122ea7/scg/index.htm

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Module 2

Trust and Identity

Overview

Businesses need to effectively and securely manage *who* and *what* can access the network, as well as *when*, *where*, and *how* that access can occur. Cisco Trust and Identity Management comprises Cisco identity management, Cisco Identity-Based Networking Services (IBNS), and Cisco Network Admission Control (NAC). In this module, you will be introduced to two of the three technologies. NAC is beyond the scope of this course and is covered in another course (*Implementing Cisco Network Admission Control* [NAC]).

Module Objectives

Upon completing this module, you will be able to implement the Cisco Trust and Identity Management model to control network access. This ability includes being able to meet these objectives:

- Describe and configure the Cisco Secure ACS as a AAA server
- Describe the IBNS model and how to configure 802.1x authentication on a Cisco Catalyst switch

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Lesson 1

Implementing Identity Management

Overview

This lesson will allow you to examine the Cisco Secure Access Control Server (ACS) and configure the Cisco Secure ACS to implement identity management services within a network.

Objectives

Upon completing this lesson, you will be able to configure a Cisco Secure ACS on a Microsoft Windows server. This ability includes being able to meet these objectives:

- Describe the function and features of Cisco Secure ACS for Windows
- Define AAA and describe the commands used to implement it on a NAD
- Describe AAA authentication
- Describe AAA authorization
- Describe AAA accounting
- Describe TACACS+
- Describe RADIUS
- Describe how to configure AAA services to work with external TACACS and RADIUS servers
- Describe how Cisco Secure ACS may function as a AAA server for NADs
- Describe the Cisco Secure ACS architectural components
- Describe the Cisco Secure ACS administrative web interface
- Describe how to install Cisco Secure ACS on a Microsoft Windows server
- Describe how to add an administrator to Cisco Secure ACS
- Give a brief overview of each section of the HTML interface
- Describe the use of NAPs

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

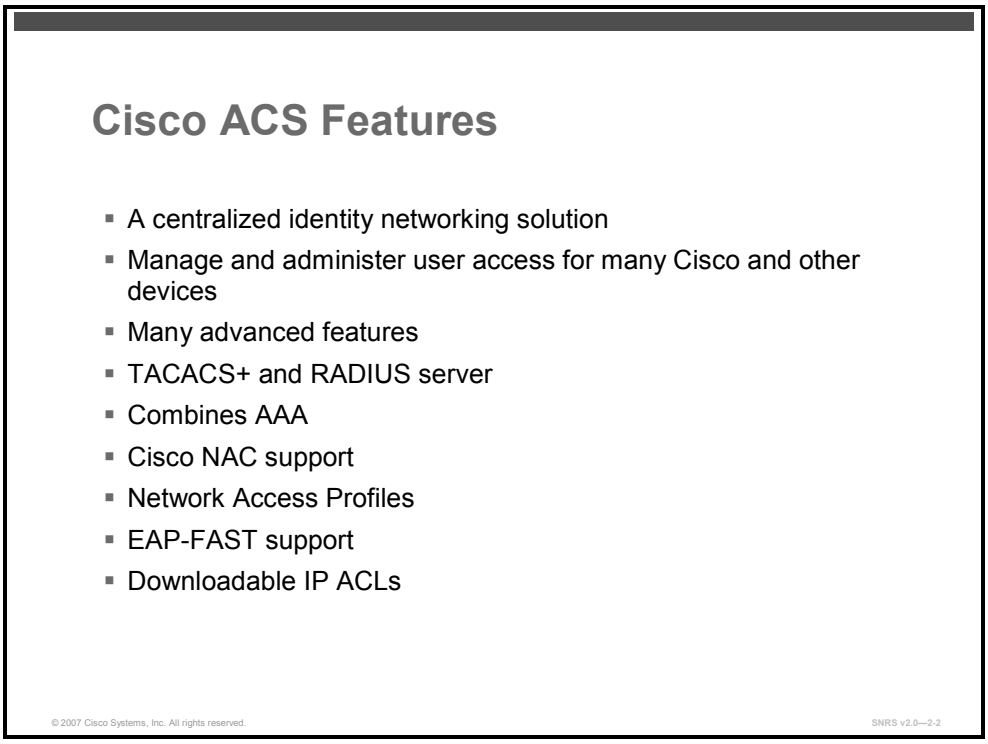
- Describe how to configure NAPs
- Describe how to create a NAP
- Describe how to configure profile-based policies
- Describe three techniques for troubleshooting network and device access in a Cisco Secure ACS environment

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Cisco Secure ACS for Windows Overview

This topic describes the function and features of Cisco Secure ACS for Windows.



Cisco Secure ACS for Windows provides a centralized identity networking solution and simplified user management experience across all Cisco devices and security management applications. Cisco Secure ACS helps to ensure enforcement of assigned policies by allowing network administrators to control who can log into the network, the privileges a user has on the network, and access to the administrative web interface for each configuration administrator. Cisco Secure ACS also helps to document security audits or account billing information.

With Cisco Secure ACS, you can manage and administer user access for Cisco IOS routers, virtual private networks (VPNs), firewalls, dialup and DSL connections, cable access solutions, storage, content, VoIP, Cisco wireless solutions, and Cisco Catalyst switches using IEEE 802.1x access control. Cisco Secure ACS is also an important component of the Cisco NAC, an industry initiative sponsored by Cisco Systems that uses the network infrastructure to enforce security policy compliance on all devices seeking to access network computing resources. Cisco Secure ACS 4.0 for Windows acts as a policy decision point in NAC deployments, evaluating credentials, determining the state of the host, and sending out per-user authorization to the network access devices.

Here are some of the advanced features that are included in Cisco Secure ACS 4.0 for Windows:

- Automatic service monitoring, database synchronization, and importing of tools for large-scale deployments
- Lightweight Directory Access Protocol (LDAP) and Open Database Connectivity (ODBC) user authentication support

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

- Flexible 802.1x authentication type support, including Extensible Authentication Protocol Transport Layer Security (EAP-TLS), Protected EAP (PEAP), Cisco Lightweight EAP (LEAP), EAP-Flexible Authentication via Secure Tunneling (EAP-FAST), and EAP-Message Digest 5 (EAP-MD5)
- Downloadable access control lists (ACLs) for any Layer 3 device, including Cisco routers, Cisco PIX Firewalls, and Cisco VPNs
- Device command set authorization
- Network access restrictions
- User and administrative access reporting
- Dynamic quota generation
- Restrictions, such as time of day and day of week
- User and device group profiles

Cisco Secure ACS is a scalable, high-performance RADIUS and TACACS+ security server. As the centralized control point for managing enterprise network users, network administrators, and network infrastructure resources, Cisco Secure ACS provides a comprehensive identity-based NAC solution for Cisco intelligent information networks.

Cisco Secure ACS extends network access security by combining traditional authentication, authorization, and accounting (AAA; pronounced “triple A”) with policy control. Cisco Secure ACS enforces a uniform network access security policy for network administrators and other network users.

Cisco ACS supports a broad variety of Cisco and other network access devices (NADs), also known as AAA clients, including the following:

- Wired and wireless LAN switches and access points
- Edge and core routers
- Dialup and broadband terminators
- Content and storage devices
- VoIP
- Firewalls
- VPNs

Additional Features in Cisco Secure ACS 4.0 for Windows

Cisco Secure ACS 4.0 for Windows provides the following additional features:

- **Cisco NAC support:** Cisco Secure ACS 4.0 for Windows acts as a policy decision point in NAC deployments. Using configurable policies, it evaluates and validates the credentials received from the Cisco Trust Agent (posture), determines the state of the host, and sends a per-user authorization to the NAD: ACLs, a policy-based ACL, or a private VLAN assignment. Evaluation of the host credentials can enforce many specific policies, such as OS patch level and antivirus **d**igital **a**udio **t**ape (DAT) file version. Cisco Secure ACS records the policy evaluation result for use with monitoring systems. Cisco Secure ACS 4.0 for Windows also allows hosts without the appropriate agent technology to be audited by third-party audit vendors before granting network access. Cisco Secure ACS policies can be extended with external policy servers to which Cisco Secure ACS forwards posture

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

credentials. For example, credentials specific to an antivirus vendor can be forwarded to the antivirus policy server of the vendor, and audit policy requests can be forwarded to third-party audit products.

- **Scalability improvements:** Cisco Secure ACS 4.0 for Windows has been upgraded to use an industry-standard relational database management system (RDBMS), improving the number of devices (AAA clients) by tenfold and the number of users by threefold. There have also been significant improvements in performance (transactions per second) across the protocol portfolio that Cisco Secure ACS supports.
- **Network Access Profiles (NAPs):** Cisco Secure ACS 4.0 for Windows provides a new feature, Network Access Profiles, which allow administrators to classify access requests according to network location, membership in a network device group (NDG), protocol type, or other specific RADIUS attribute values sent by the NAD through which the user connects. You can map AAA policies to specific profiles. For example, you can apply a different access policy for wireless access and remote (VPN) access.
- **Extended replication components:** Cisco Secure ACS 4.0 for Windows provides improved and enhanced replication. Administrators now can replicate NAPs, and all related configurations, including the following:
  - Posture validation settings
  - AAA clients and hosts
  - External database configuration
  - Global authentication configuration
  - NDGs
  - Dictionaries
  - Shared-profile components
  - Additional logging attributes
- **EAP-FAST enhanced support:** EAP-FAST is a new, publicly accessible IEEE 802.1x EAP type that Cisco developed to support customers who cannot enforce a strong password policy, or, who want to deploy an 802.1x EAP type that has these characteristics:
  - Does not require digital certificates
  - Supports a variety of user and password database types
  - Supports password expiration and change
  - Is flexible, and is easy to deploy and manage
- **Downloadable IP ACLs:** Downloadable IP ACLs extend per-user ACL support to any Layer 3 network device that supports this feature, such as Cisco PIX Firewalls, Cisco VPN solutions, and Cisco IOS routers. You can define sets of ACLs that can be applied per user or per group. This feature complements NAC support by enforcing the correct ACL policy. When used in conjunction with network access filters (NAFs), you can apply downloadable ACLs differently per device. You can, therefore, tailor ACLs uniquely per user, per access device.
- **Certificate revocation list (CRL) comparison:** This feature supports certificate revocation by using the X.509 CRL profile. A CRL is a time-stamped list identifying revoked certificates; the list is signed by a certificate authority (CA) or CRL issuer, and made freely available in a public repository.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

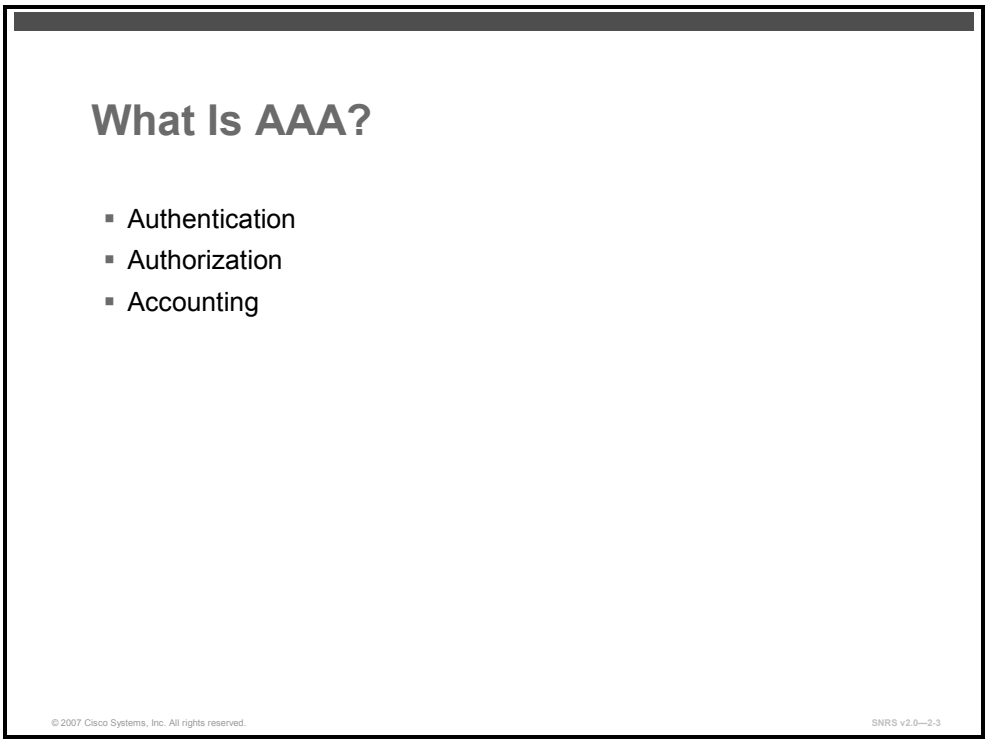
- **Machine access restrictions (MARs):** This feature includes MARs as an enhancement of Microsoft Windows machine authentication. When Microsoft Windows machine authentication is enabled, you can use MARs to control authorization of EAP-TLS, EAP-FASTv1a, and Microsoft PEAP users who authenticate with a Microsoft Windows external user database. Users who access the network with a computer that has not passed machine authentication within a configurable length of time are given the authorizations of a user group that you specify and that you can configure to limit authorization as needed. Alternatively, you can deny network access altogether.
- **NAFs:** NAFs are a new type of shared profile component. NAFs provide a flexible way to apply network access restrictions and downloadable ACLs on network device names, NDGs, or their IP address. NAFs applied by IP addresses can use IP address ranges and wildcards. This feature introduces granular application of network access restrictions (NARs) and downloadable ACLs, which previously supported only the use of the same access restrictions or ACLs to all devices. You can use NAFs to define flexible network device restriction policies to be defined, a requirement that is common in large environments.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Authentication, Authorization, and Accounting

This topic describes AAA.



AAA is a compilation of network security services that provides the framework through which you set up NAC. AAA provides a modular way of performing authentication, authorization, and accounting services. These services will be discussed further.

AAA uses protocols such as RADIUS, TACACS+, or Kerberos for its security functions. From the view of a network access server (NAS), AAA is the means through which the NAS establishes communication between itself and the RADIUS, TACACS+, or Kerberos security server.

AAA is designed to enable you to dynamically configure the type of authentication and authorization you want on a per-line (per-user) or per-service (for example, IP, Internetwork Packet Exchange [IPX], or virtual private dialup network [VPDN]) basis. You define the type of authentication and authorization you want by creating method lists, then applying those method lists to specific services or interfaces.

Cisco IOS Software also provides additional features for simple access control, such as local username authentication, line password authentication, and enable password authentication. However, these features do not provide the same degree of access control that is possible by using AAA.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Common Cisco IOS AAA Configuration

Use the **aaa new-model** command to enable AAA.

```
router(config)# aaa new-model
```

To disable AAA, use this command:

```
router(config)# no aaa new-model
```

To configure security on a Cisco router or access server using AAA, follow these steps:

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

Step 2 If you decide to use a separate security server, configure security protocol parameters, such as RADIUS, TACACS+, or Kerberos.

Step 3 Define the method lists for authentication by using an AAA authentication command.

Step 4 Apply the method lists to a particular interface or line, if required.

Step 5 (Optional) Configure authorization using the **aaa authorization** command.

Step 6 (Optional) Configure accounting using the **aaa accounting** command.

Method Lists

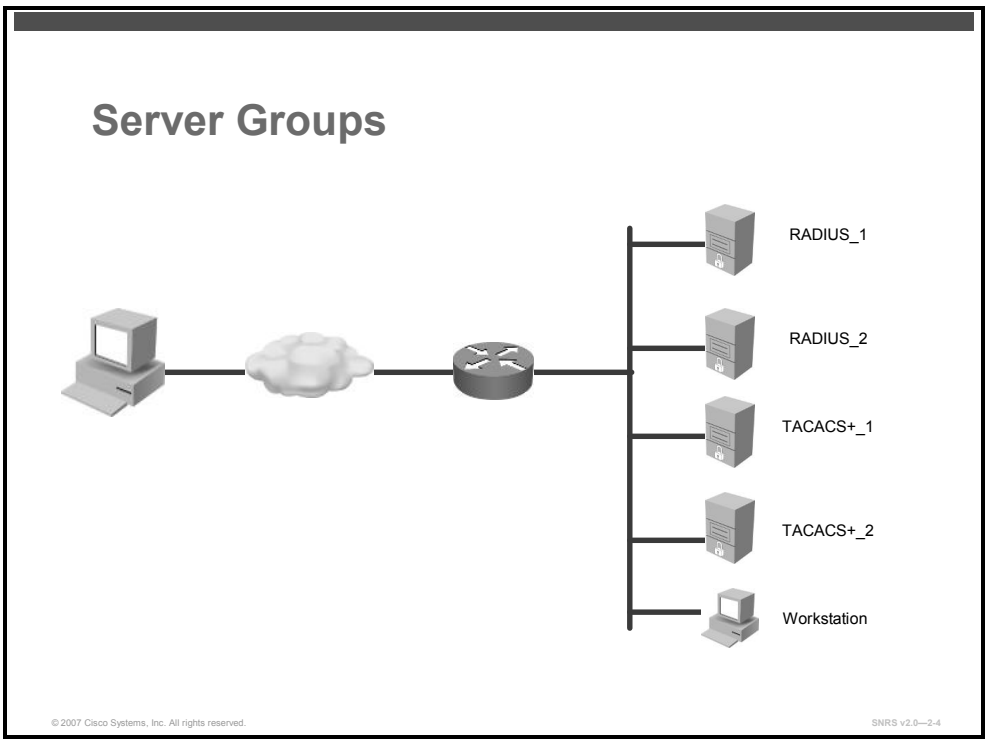
A method list is a sequential list that defines the authentication methods used to authenticate a user. Method lists enable you to designate one or more security protocols to be used for authentication, thus ensuring a backup system for authentication in case the initial method fails. Cisco IOS Software uses the first method listed to authenticate users; if that method does not respond, Cisco IOS Software selects the next authentication method in the method list. This process continues until there is successful communication with a listed authentication method or the authentication method list is exhausted, in which case authentication fails.

Note Cisco IOS Software attempts authentication with the next listed authentication method only when there is no response from the previous method. If authentication fails at any point in this cycle (for example, denied access), the authentication process stops and no other authentication methods are attempted. If there is an error reaching the security server, the process goes on to the next server if one is available.

A method list must be applied to a specific interface before any of the defined authentication methods will be performed. The only exception is the default method list (which is named "default"). The default method list is automatically applied to all interfaces except those that have a named method list explicitly defined. A defined method list overrides the default method list.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.



Method Lists and Server Groups

A server group is a way to group existing RADIUS or TACACS+ server hosts for use in method lists.

Using server groups, you can specify a subset of the configured server hosts and use them for a particular service. For example, server groups allow you to define Radius_1 and Radius_2 as a server group, and define TACACS+_1 and TACACS+_2 as a separate server group. For example, you can specify Radius_1 and TACACS+_1 in the method list for authentication login, while specifying Radius_2 and TACACS+_2 in the method list for PPP authentication.

Server groups can include multiple host entries for the same server, as long as each entry has a unique identifier. The combination of an IP address and a UDP port number creates a unique identifier, allowing different ports to be individually defined as RADIUS hosts providing a specific AAA service. In other words, this unique identifier enables RADIUS requests to be sent to different User Datagram Protocol (UDP) ports on a server at the same IP address. If two different host entries on the same RADIUS server are configured for the same service—for example, authentication—the second host entry configured acts as failover backup to the first one. Using this example, if the first host entry fails to provide accounting services, the NAS will try the second host entry configured on the same device for accounting services. The RADIUS host entries will be tried in the order in which they are configured.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Use this command to create server groups:

```
router(config)# aaa group server {radius | tacacs+} group-name
```

- The **aaa group server** command defines the AAA server group with a group name. All members of a group must be the same type; that is, RADIUS or TACACS+. This command puts the router in server group subconfiguration mode.

```
router(config-sg)# server ip-address [auth-port port-number]
[acct-port port-number]
```

- The **server** *ip-address* command associates a particular RADIUS server with the defined server group. Each security server is identified by its IP address and UDP port number.

Repeat this step for each RADIUS server in the AAA server group.

Note Each server in the group must be defined previously using the **radius-server host** or **tacacs-server host** command.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Authentication

This topic describes authentication.



Authentication is used to identify users before they gain access to the network and network services. This can include a login and password dialog, challenge and response, messaging support, and, depending on the security protocol that you select, encryption.

Authentication is configured by defining a named list of authentication methods, and then applying that list to an interface. The method list defines the types of authentication to be performed and the sequence in which they will be performed; it must be applied to a specific interface (for example, a vty or console) before any of the defined authentication methods will be performed. The only exception is the default method list (which is named “default”). The default method list is automatically applied to all interfaces if no other method list is defined. A defined method list overrides the default method list.

All authentication methods, except for local, line password, and enable authentication, must be defined through AAA.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Common Cisco IOS AAA Authentication Configuration

To enable AAA authentication and create a local authentication list, use the **aaa authentication login** command.

```
dev(config)# aaa authentication login {default | list-name}
password-expiry method1 [method2...]
```

Syntax Description

default	Uses the listed authentication methods that follow this keyword as the default list of methods when a user logs in.
list-name	Character string used to name the list of authentication methods activated when a user logs in.
password-expiry	Enables password aging on a local authentication list.
method1 **[***method2...***]**	Identifies the list of methods that the authentication algorithm tries in the given sequence. You must enter at least one method; you may enter up to four methods. Method keywords described in the table "AAA Authentication Login Methods".

The *list-name* is a character string used to name the list you are creating. The *method* argument refers to the actual method that the authentication algorithm tries. The additional methods of authentication are used only if the previous method returns an error, not if it fails. To specify that the authentication should succeed even if all methods return an error, specify **none** as the final method in the command line.

AAA Authentication Login Methods

The table lists AAA authentication login methods.

AAA Authentication Login Methods

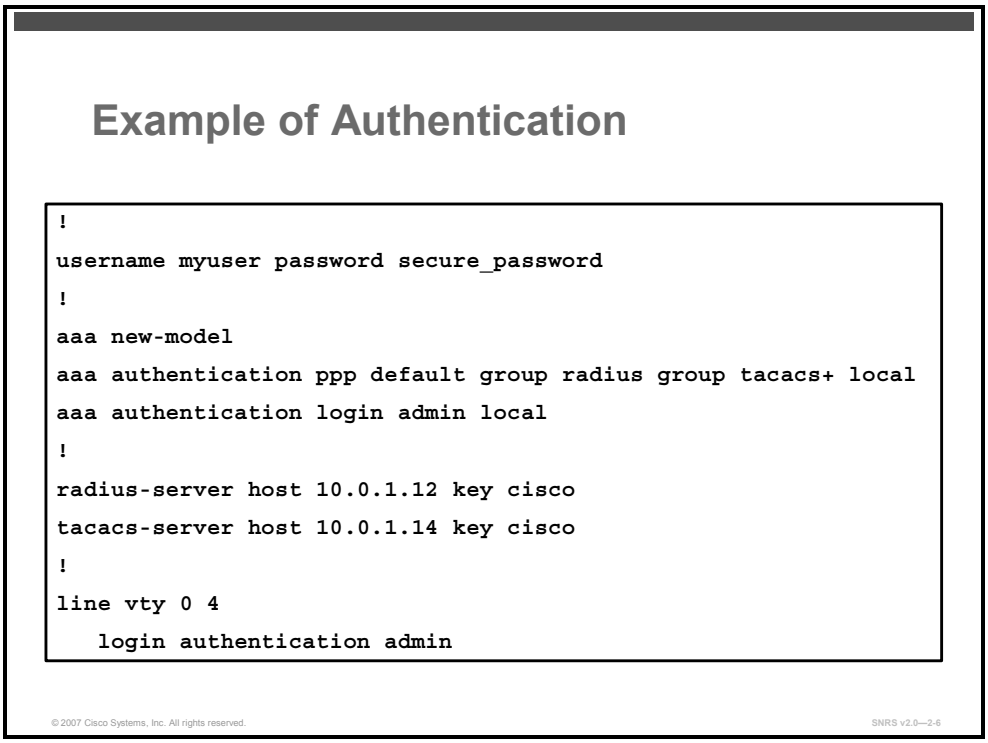
enable	Uses the enable password for authentication.
krb5	Uses Kerberos 5 for authentication.
krb5-telnet	Uses Kerberos 5 telnet authentication protocol when using Telnet to connect to the router.
line	Uses the line password for authentication.
local	Uses the local username database for authentication.
local-case	Uses case-sensitive local username authentication.
none	Uses no authentication.
cache *group-name*	Uses a cache server group for authentication.
group radius	Uses the list of all RADIUS servers for authentication.
group tacacs+	Uses the list of all TACACS+ servers for authentication.
group *group-name*	Uses a subset of RADIUS or TACACS+ servers for authentication as defined by the **aaa group server radius** or **aaa group server tacacs+** command.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Authentication Example

Here is an example of AAA authentication.



This example shows a security solution where some interfaces will use the same authentication methods to authenticate PPP connections but the vty will use a named method list.

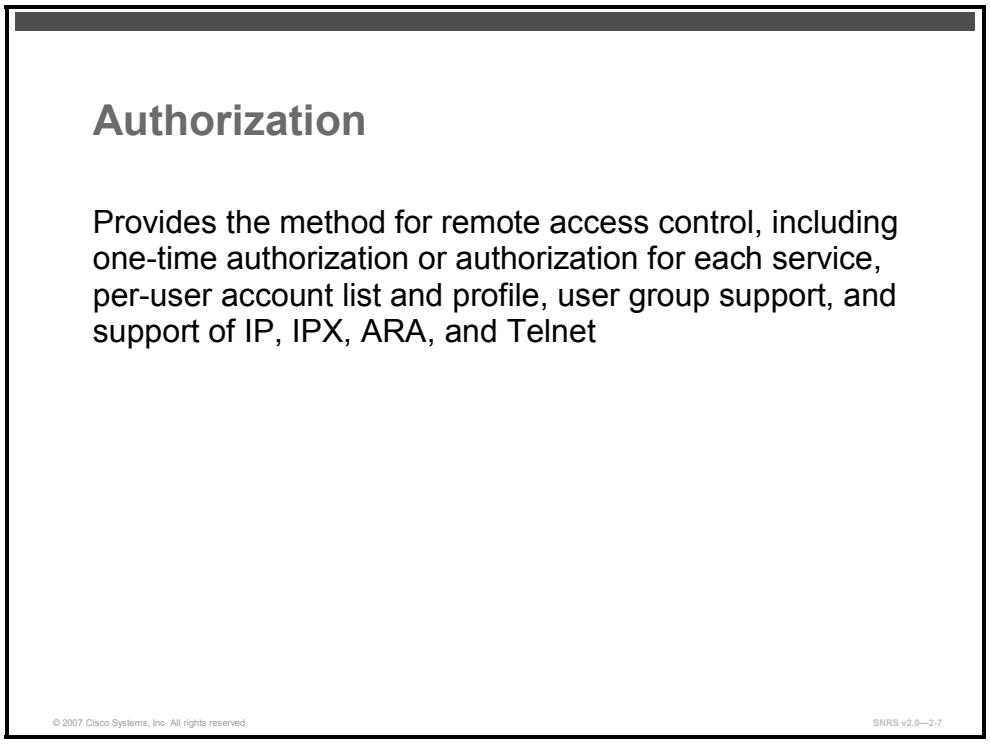
For PPP connections, the RADIUS servers are contacted first for authentication information, then if there is no response, the TACACS+ group is contacted. If all designated servers fail to respond, authentication falls to the local username database on the access server itself.

For vty connections, the “admin” list is used, which specifies the local username database on the access server.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Authorization

This topic describes authorization.



To set parameters that restrict user access to a network, use the **aaa authorization** command.

Authorization provides the method for remote access control, including one-time authorization or authorization for each service, per-user account list and profile, user group support, and support of IP, IPX, AppleTalk Remote Access (ARA), and Telnet.

Authorization defines a set of attributes that describe what the user is authorized to perform. These attributes are compared to the information contained in a database for a given user and the result is returned to AAA to determine the actual capabilities and restrictions of the user. The database can be located locally on the access server or router, or it can be hosted remotely on a RADIUS or TACACS+ security server. Remote security servers, such as RADIUS and TACACS+, authorize users by associating attribute-value pairs, which define those rights with the appropriate user.

All authorization methods must be defined through AAA.

Authorization is configured by defining a named list of authorization methods, and then applying that list to an interface.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Common Cisco IOS AAA Authorization Configuration

To enable AAA authorization and create an authorization method list for a particular authorization type, use the **aaa authorization** command.

```
dev(config)# aaa authorization {auth-proxy | network | exec |
commands level | reverse-access | configuration} {default | list-
name} [method1 [method2...]]
```

Syntax Description

auth-proxy	Applies specific security policies on a per-user basis.
network	Runs authorization for all network-related service requests, including Serial Line Internet Protocol (SLIP), PPP, PPP Network Control Programs (NCPs), and AppleTalk Remote Access (ARA).
exec	Runs authorization to determine if the user is allowed to run an EXEC shell. This facility might return user profile information such as autocommand information.
commands	Runs authorization for all commands at the specified privilege level.
level	Specific command level that should be authorized. Valid entries are 0 through 15.
reverse-access	Runs authorization for reverse access connections, such as reverse Telnet.
configuration	Downloads the configuration from the authentication, authorization, and accounting (AAA) server.
default	Uses the listed authorization methods that follow this keyword as the default list of methods for authorization.
list-name	Character string used to name the list of authorization methods.
method1 **[***method2...***]**	Identifies an authorization method or multiple authorization methods to be used for authorization. A method may be any one of the keywords listed in the table "AAA Authorization Methods".

Use the **aaa authorization** command to enable authorization and to create named methods lists, defining authorization methods that can be used when a user accesses the specified function. Method lists for authorization define the ways authorization will be performed and the sequence in which these methods will be performed. A method list is a named list describing the authorization methods to be used (such as RADIUS or TACACS+), in sequence. Method lists enable you to designate one or more security protocols to be used for authorization, thus ensuring a backup system in case the initial method fails. Cisco IOS software uses the first method listed to authorize users for specific network services; if that method fails to respond, the Cisco IOS software selects the next method listed in the method list. This process continues until there is successful communication with a listed authorization method, or all methods defined are exhausted.

Note The Cisco IOS software attempts authorization with the next listed method only when there is no response from the previous method. If authorization fails at any point in this cycle—meaning that the security server or local username database responds by denying the user services—the authorization process stops and no other authorization methods are attempted.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

If the **aaa authorization** command for a particular authorization type is issued without a named method list specified, the default method list is automatically applied to all interfaces or lines (where this authorization type applies) except those that have a named method list explicitly defined. (A defined method list overrides the default method list.) If no default method list is defined, then no authorization takes place.

Use the **aaa authorization** command to create a list by entering values for the *list-name* and the *method* arguments, where *list-name* is any character string used to name this list (excluding all method names) and *method* identifies the list of authorization methods tried in the given sequence.

Note In the following table the **group radius**, **group tacacs+**, and **group** *group-name* methods refer to a set of previously defined RADIUS or TACACS+ servers. Use the **radius-server host** and **tacacs-server host** commands to configure the host servers. Use the **aaa group server radius** and **aaa group server tacacs+** commands to create a named group of servers.

This table lists AAA authorization methods.

AAA Authorization Methods

cache *group-name*	Uses a cache server group for authorization
group radius	Uses the list of all RADIUS servers for authentication as defined by the **aaa group server radius** command.
group tacacs+	Uses the list of all TACACS+ servers for authentication as defined by the **aaa group server tacacs+** command.
group *group-name*	Uses a subset of RADIUS or TACACS+ servers for accounting as defined by the **server group** *group-name* command.
if-authenticated	Allows the user to access the requested function if the user is authenticated.
krb5-instance	Uses the instance defined by the **kerberos instance map** command.
local	Uses the local database for authorization.
none	No authorization is performed.

Cisco IOS software supports the following methods for authorization:

- Cache Sever Groups—The router consults its cache server groups to authorize specific rights for users.
- RADIUS—The network access server requests authorization information from the RADIUS security server group. RADIUS authorization defines specific rights for users by associating attributes, which are stored in a database on the RADIUS server, with the appropriate user.
- TACACS+—The network access server exchanges authorization information with the TACACS+ security daemon. TACACS+ authorization defines specific rights for users by associating attribute-value (AV) pairs, which are stored in a database on the TACACS+ security server, with the appropriate user.
- If-Authenticated—The user is allowed to access the requested function provided the user has been authenticated successfully.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

- None—The network access server does not request authorization information; authorization is not performed over this line or interface.
- Local—The router or access server consults its local database, as defined by the **username** command, to authorize specific rights for users. Only a limited set of functions can be controlled via the local database.
- Kerberos Instance Map—The network access server uses the instance defined by the **kerberos instance map** command for authorization.

Method lists are specific to the type of authorization being requested. AAA supports five different types of authorization:

- Network—Applies to network connections. This can include a PPP, SLIP, or ARA connection.
- EXEC—Applies to the attributes associated with a user EXEC terminal session.
- Commands—Applies to the EXEC mode commands a user issues. Command authorization attempts authorization for all EXEC mode commands, including global configuration commands, associated with a specific privilege level.
- Reverse Access—Applies to reverse Telnet sessions.
- Configuration—Applies to the configuration downloaded from the AAA server.

When you create a named method list, you are defining a particular list of authorization methods for the indicated authorization type.

Once defined, method lists must be applied to specific lines or interfaces before any of the defined methods will be performed.

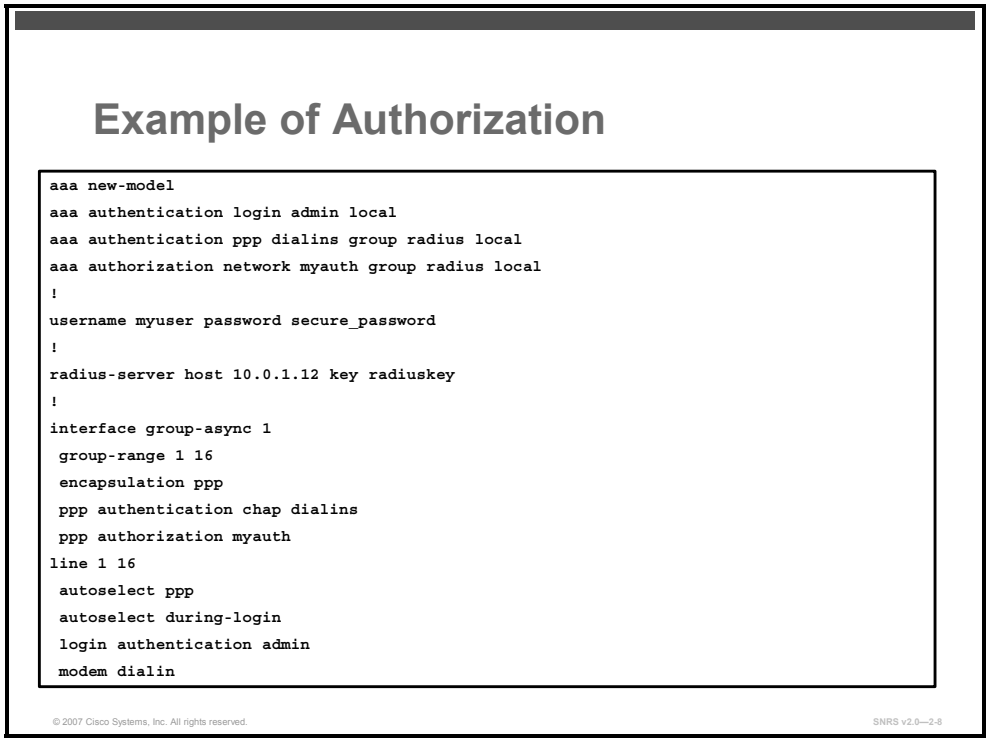
The authorization command causes a request packet containing a series of AV pairs to be sent to the RADIUS or TACACS daemon as part of the authorization process. The daemon can do one of the following:

- Accept the request as is.
- Make changes to the request.
- Refuse the request and refuse authorization.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Authorization Example

Here is an example of authorization.

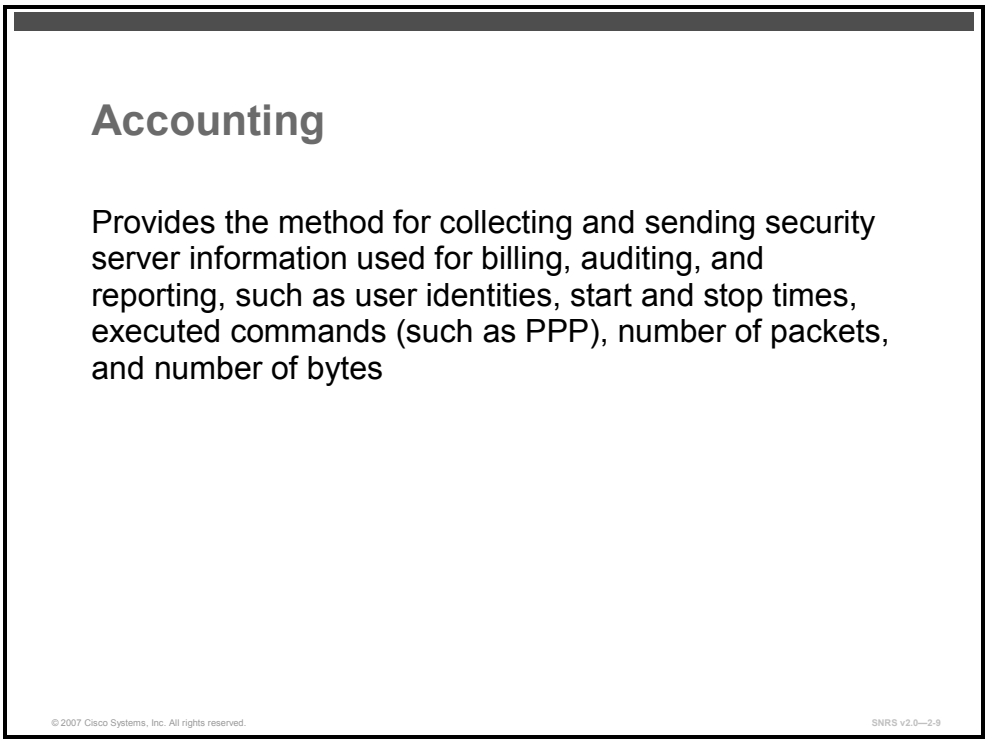


This example shows how to configure a Cisco AS5300 Series Universal Access Servers (enabled for AAA and communication with a RADIUS security server) for AAA services to be provided by the RADIUS server. If the RADIUS server fails to respond, the local database will be queried for authentication and authorization information.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Accounting

This topic describes accounting.



Accounting provides the ability to collect and send security server information to be used for billing, auditing, and reporting purposes. Accounting collects information such as user identities, start and stop times, executed commands (such as PPP), number of packets, and number of bytes.

Accounting enables you to track the services that users are accessing and the amount of network resources that they are consuming. When you activate AAA accounting, the NAS reports user activity to a security server in the form of accounting records. Each accounting record comprises accounting attribute-value pairs and is stored on the Cisco Secure ACS. This data can then be analyzed for network management, client billing, or auditing.

All accounting methods must be defined through AAA.

Accounting is configured by defining a named list of accounting methods, and then applying that list to an interface.

Common Cisco IOS AAA Accounting Configuration

Use the **aaa accounting** command to enable accounting and to create named method lists that define specific accounting methods on a per-line or per-interface basis.

```
dev(config)# aaa accounting {auth-proxy | system | network | exec |
connection | commands level} {default | list-name} [vrf vrf-name]
{start-stop | stop-only | none} [broadcast] group group-name
```


The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Syntax Description

auth-proxy	Provides information about all authenticated-proxy user events.
system	Performs accounting for all system-level events not associated with users, such as reloads. **Note** When system accounting is used and the accounting server is unreachable at system startup time, the system will not be accessible for approximately two minutes.
network	Runs accounting for all network-related service requests, including Serial Line Internet Protocol (SLIP), PPP, PPP Network Control Protocols (NCPs), and AppleTalk Remote Access Protocol (ARAP).
exec	Runs accounting for EXEC shell session. This keyword might return user profile information such as what is generated by the **autocommand** command.
connection	Provides information about all outbound connections made from the network access server, such as Telnet, local-area transport (LAT), TN3270, packet assembler and disassembler (PAD), and rlogin.
commands *level*	Runs accounting for all commands at the specified privilege level. Valid privilege level entries are integers from 0 through 15.
default	Uses the listed accounting methods that follow this argument as the default list of methods for accounting services.
list-name	Character string used to name the list of at least one of the following accounting methods: ■ **group radius**—Uses the list of all RADIUS servers for authentication as defined by the **aaa group server radius** command. ■ **group tacas+**—Uses the list of all TACACS+ servers for authentication as defined by the **aaa group server tacacs+** command. ■ **group** *group-name*—Uses a subset of RADIUS or TACACS+ servers for accounting as defined by the server group *group-name.*
vrf *vrf-name*	(Optional) Specifies a virtual route forwarding (VRF) configuration. VRF is used *only* with system accounting.
start-stop	Sends a "start" accounting notice at the beginning of a process and a "stop" accounting notice at the end of a process. The "start" accounting record is sent in the background. The requested user process begins regardless of whether the "start" accounting notice was received by the accounting server.
stop-only	Sends a "stop" accounting notice at the end of the requested user process.
none	Disables accounting services on this line or interface.
broadcast	(Optional) Enables sending accounting records to multiple AAA servers. Simultaneously sends accounting records to the first server in each group. If the first server is unavailable, failover occurs using the backup servers defined within that group.
group *group-name*	Specifies the accounting method list. Enter at least one of the following keywords: ■ **auth-proxy**—Uses the list of all RADIUS servers for authentication as defined by the **aaa group server radius** command. ■ **commands**—Creates a method list to provide accounting information about specific, individual EXEC commands associated with a specific privilege level. ■ **connection**—Creates a method list to provide accounting information about all outbound connections made from the network access server.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

	■ **exec**—Creates a method list to provide accounting records about user EXEC terminal sessions on the network access server, including username, date, and start and stop times. ■ **network**—Creates a method list to provide accounting information for SLIP, PPP, NCPs, and ARAP sessions. ■ **resource**—Creates a method list to provide accounting records for calls that have passed user authentication or calls that failed to be authenticated. ■ **tunnel**—Creates a method list to provide accounting records (Tunnel-Start, Tunnel-Stop, and Tunnel-Reject) for virtual private dialup network (VPDN) tunnel status changes. ■ **tunnel-link**—Creates a method list to provide accounting records (Tunnel-Link-Start, Tunnel-Link-Stop, and Tunnel-Link-Reject) for VPDN tunnel-link status changes.

The following table contains descriptions of keywords for aaa accounting methods.

AAA Accounting Methods

group radius	Uses the list of all RADIUS servers for authentication as defined by the **aaa group server radius** command.
group tacacs+	Uses the list of all TACACS+ servers for authentication as defined by the **aaa group server tacacs+** command.
group *group-name*	Uses a subset of RADIUS or TACACS+ servers for accounting as defined by the server group *group-name*.

In the table above, the **group radius** and **group tacacs+** methods refer to a set of previously defined RADIUS or TACACS+ servers. Use the **radius-server host** and **tacacs-server host** commands to configure the host servers. Use the **aaa group server radius** and **aaa group server tacacs+** commands to create a named group of servers.

Cisco IOS software supports the following two methods of accounting:

- RADIUS—The network access server reports user activity to the RADIUS security server in the form of accounting records. Each accounting record contains accounting attribute-value (AV) pairs and is stored on the security server.
- TACACS+—The network access server reports user activity to the TACACS+ security server in the form of accounting records. Each accounting record contains accounting AV pairs and is stored on the security server.

Method lists for accounting define the way accounting will be performed. Named accounting method lists enable you to designate a particular security protocol to be used on specific lines or interfaces for particular types of accounting services. Create a list by entering the *list-name* and the *method*, where *list-name* is any character string used to name this list (excluding the names of methods, such as RADIUS or TACACS+) and *method* identifies the methods to be tried in sequence as given.

If the **aaa accounting** command for a particular accounting type is issued without a named method list specified, the default method list is automatically applied to all interfaces or lines (where this accounting type applies) except those that have a named method list explicitly defined. (A defined method list overrides the default method list.) If no default method list is defined, then no accounting takes place.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Named accounting method lists are specific to the indicated type of accounting. Method list keywords are described in the following table.

AAA Accounting Method List Keywords

auth-proxy	Creates a method list to provide accounting information about all authenticated hosts that use the authentication proxy service.
commands	Creates a method list to provide accounting information about specific, individual EXEC commands associated with a specific privilege level.
connection	Creates a method list to provide accounting information about all outbound connections made from the network access server.
exec	Creates a method list to provide accounting records about user EXEC terminal sessions on the network access server, including username, date, and start and stop times.
network	Creates a method list to provide accounting information for SLIP, PPP, NCPs, and ARAP sessions.
resource	Creates a method list to provide accounting records for calls that have passed user authentication or calls that failed to be authenticated.
tunnel	Creates a method list to provide accounting records (Tunnel-Start, Tunnel-Stop, and Tunnel-Reject) for virtual private dialup network (VPDN) tunnel status changes.
tunnel-link	Creates a method list to provide accounting records (Tunnel-Link-Start, Tunnel-Link-Stop, and Tunnel-Link-Reject) for VPDN tunnel-link status changes.

For minimal accounting, include the **stop-only** keyword to send a "stop" record accounting notice at the end of the requested user process. For more accounting, you can include the **start-stop** keyword, so that RADIUS or TACACS+ sends a "start" accounting notice at the beginning of the requested process and a "stop" accounting notice at the end of the process. Accounting is stored only on the RADIUS or TACACS+ server. The **none** keyword disables accounting services for the specified line or interface.

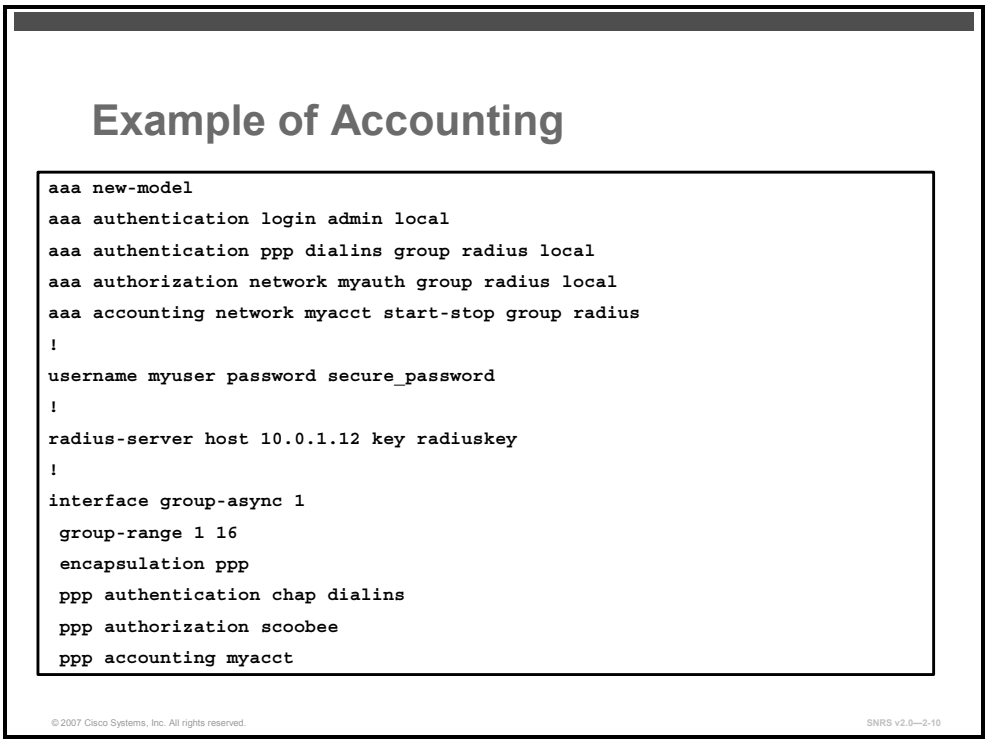
When AAA accounting is activated, the network access server monitors either RADIUS accounting attributes or TACACS+ AV pairs pertinent to the connection, depending on the security method you have implemented. The network access server reports these attributes as accounting records, which are then stored in an accounting log on the security server.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Accounting Example

Here is an example of accounting.

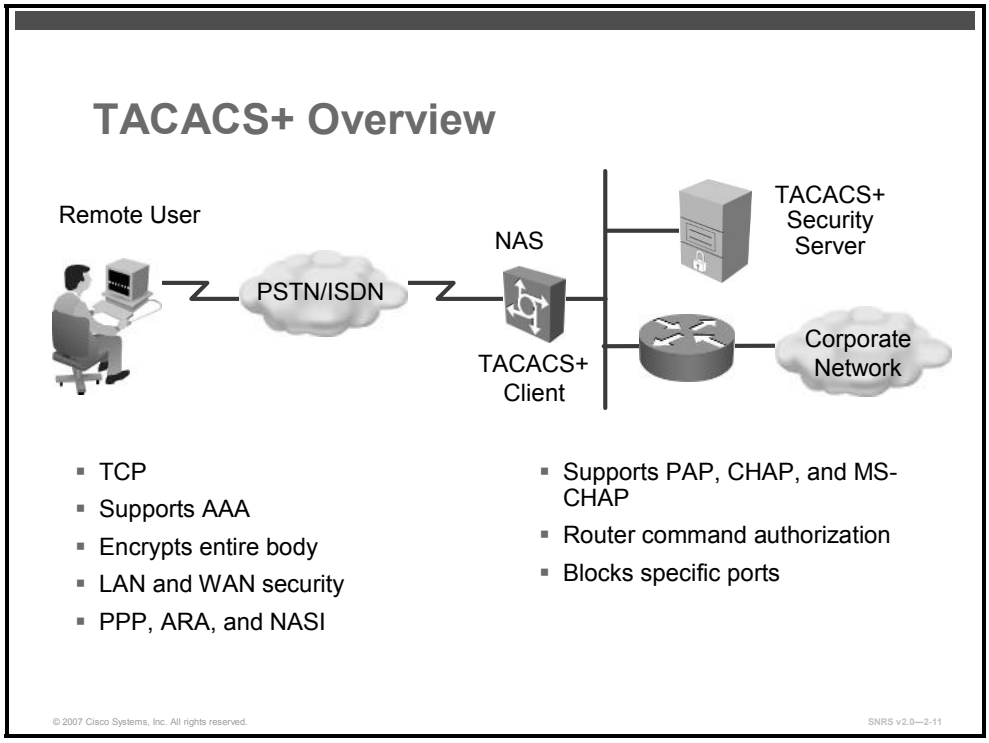


In this example, the accounting services will be handled by the RADIUS server.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

TACACS+

This topic describes TACACS+.



TACACS+ is a security application that provides centralized validation of users attempting to gain access to a router or NAS. TACACS+ services are maintained in a database on a TACACS+ daemon running, typically, on a UNIX or Microsoft Windows NT workstation. You must have access to and must configure a TACACS+ server before the configured TACACS+ features on your NAS are available.

TACACS+ provides for separate and modular AAA facilities. TACACS+ allows for a single Cisco Secure ACS (the TACACS+ daemon) to provide each service—authentication, authorization, and accounting—independently. Each service can be tied into its own database to take advantage of other services available on that server or on the network, depending on the capabilities of the daemon.

NADs enable traditional "dumb" terminals, terminal emulators, workstations, PCs, and routers in conjunction with suitable adapters (for example, modems or ISDN adapters) to communicate using protocols such as PPP, SLIP, compressed SLIP (CSLIP), or ARAP. In other words, an NAS provides connections to a single user, to a network or subnetwork, and to interconnected networks. The entities connected to the network through an NAS are called "network access clients."

TACACS+, administered through the AAA security services, can provide these services:

- **Authentication:** Provides complete control of authentication through login and password dialog, challenge and response, and messaging support

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

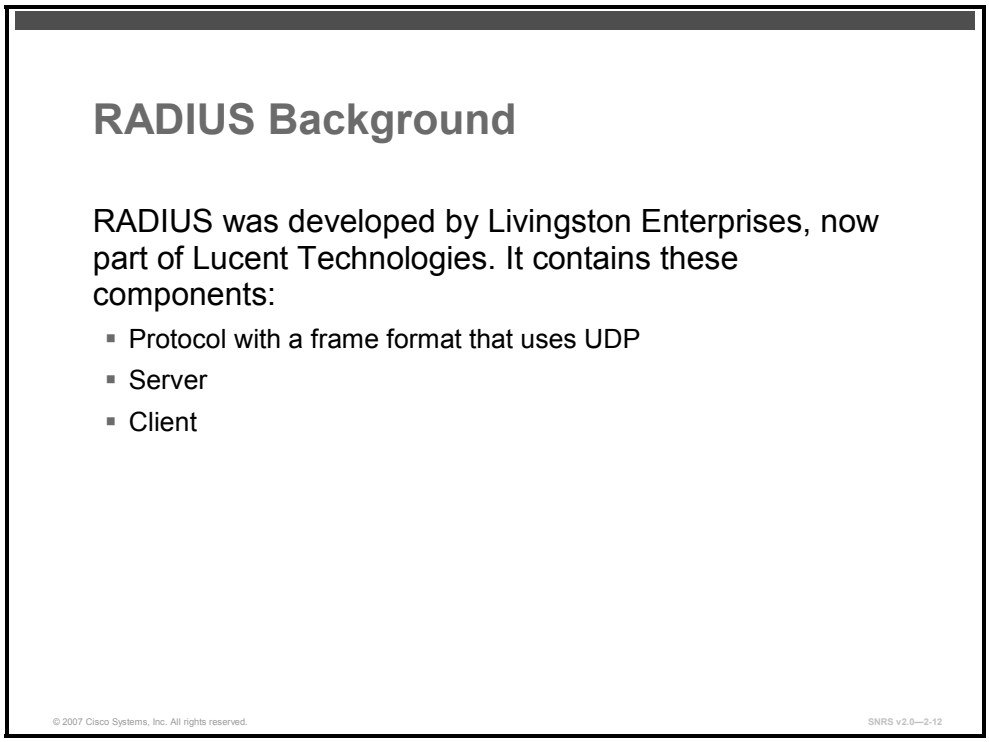
- **Authorization:** Provides fine-grained control over user capabilities for the duration of the user session, including but not limited to setting auto commands, access control, session duration, or protocol support, and also what commands a user may execute
- **Accounting:** Collects and sends information used for billing, auditing, and reporting to the TACACS+ daemon

The TACACS+ protocol provides authentication between the NAS and the TACACS+ daemon, and it ensures confidentiality because all protocol exchanges between a NAS and a TACACS+ daemon are encrypted.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

RADIUS

This topic describes RADIUS.



RADIUS is an access server AAA protocol developed by Livingston Enterprises (now part of Lucent Technologies). It is a system of distributed security that secures remote access to networks and network services against unauthorized access. RADIUS comprises these three components:

- Protocol with a frame format that uses UDP
- Server
- Client

The server runs on a central computer, typically at the customer site, while the clients reside in the dial-in access servers and can be distributed throughout the network. Cisco incorporated the RADIUS client into Cisco IOS Software, starting with Cisco IOS Software Release 11.1.

Client-Server Model

A router operates as a client of RADIUS. The client is responsible for passing user information to designated RADIUS servers, and then acting on the response that is returned. RADIUS servers are responsible for receiving user connection requests, authenticating the user, and then returning all configuration information necessary for the client to deliver service to the user. The RADIUS servers can act as proxy clients to other kinds of authentication servers.

Network Security

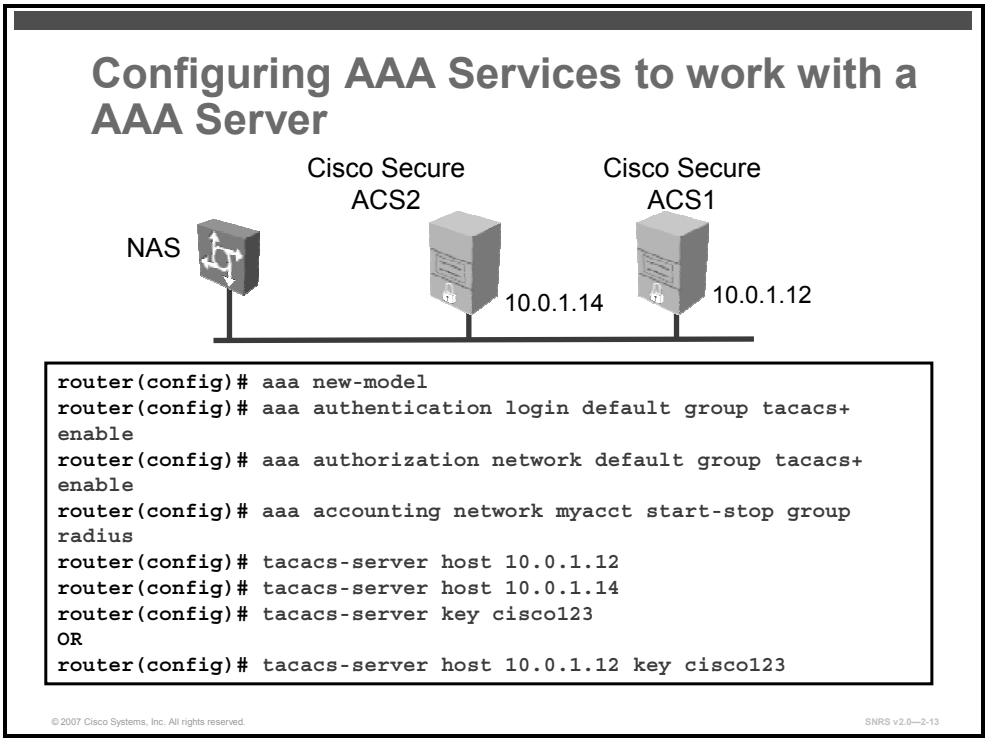
Transactions between the client and RADIUS server are authenticated using a shared secret, which is never sent over the network. In addition, any user passwords are sent encrypted between the client and RADIUS server, to eliminate the possibility that someone snooping on an unsecured network could determine a user password.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Configuring AAA to Work with External AAA Servers

This topic describes how to configure AAA services to work with external TACACS and RADIUS servers.



Several steps are required to configure AAA services to work with external AAA servers using TACACS or RADIUS.

Follow these steps to enable the router to use an external AAA server. At a minimum, the following commands should be entered.

Step 1 Globally enable AAA

```
router(config)# aaa new-model
```

Step 2 Specify authentication, authorization, and accounting lists and methods.

```
router(config)# aaa authentication
router(config)# aaa authorization
router(config)# aaa accounting
```

Step 3 Specify AAA server hosts addresses.

```
router(config)# tacacs-server host ip-address
or
router(config)# radius-server host ip-address
```


The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Step 4 Specify encryption keys used to encrypt data between the NAS and the AAA server.

```
router(config)# tacacs-server key keyword
or
router(config)# radius-server key keyword
```

The **tacacs-server key** and **radius-server key** commands are used when two or more servers share the same key. If you need to configure multiple TACACS+ or RADIUS servers, each with its own specific key, you need to use the method shown next..

You can specify multiple Cisco Secure ACS servers, each with its own key, by repeating the **tacacs-server host** or **radius-server host** commands (one for each TACACS+ or RADIUS host and its specific key) as follows:

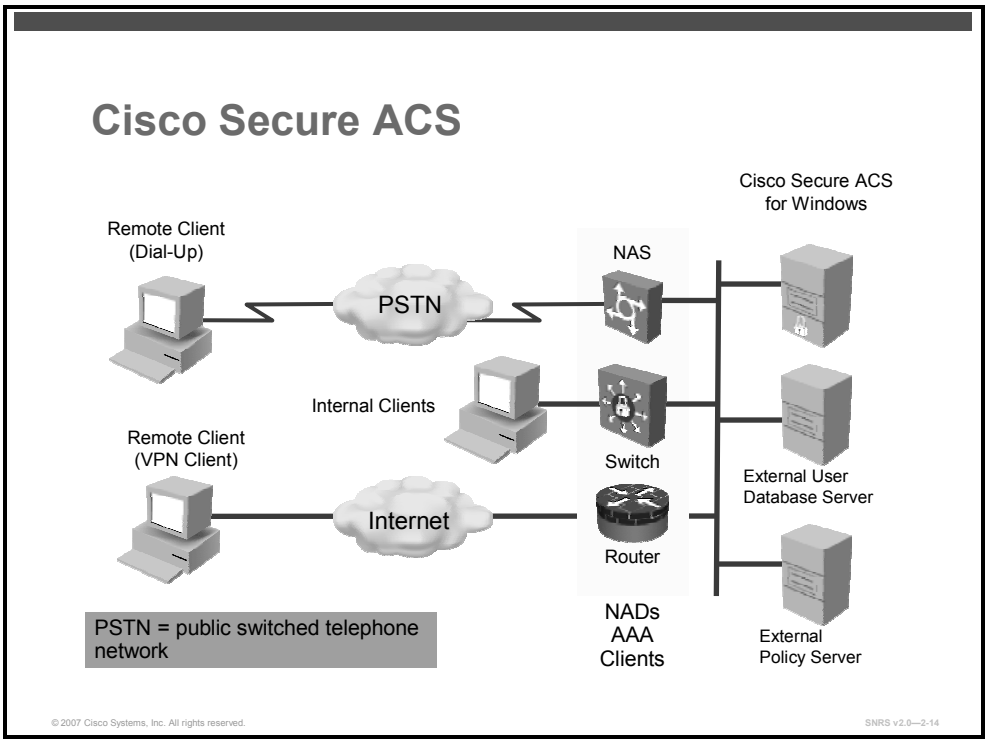
```
router(config)# tacacs-server host ip-address key keyword
or
router(config)# radius-server host ip-address key keyword
```

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Cisco Secure ACS as a AAA Server

This topic describes how Cisco Secure ACS functions as a AAA server for NADs.



Cisco Secure ACS functions as the AAA server from the perspective of the NAD. You must configure the device, which functions as a AAA client from the Cisco Secure ACS perspective, to direct all end-user host access requests to Cisco Secure ACS, via the TACACS+ or RADIUS protocols.

Basically, the NAD serves as the network gatekeeper and sends an access request to Cisco Secure ACS on behalf of the user. Cisco Secure ACS verifies the username, password, and possibly other data by using its internal database or one of the configured external identity directories. Cisco Secure ACS ultimately responds to the NAD with an access denied message or an Access-Accept message with a set of authorization attributes. When Cisco Secure ACS is used in the context of the NAC architecture, additional machine data, known as "posture," is validated as well, before the user is granted access to the network.

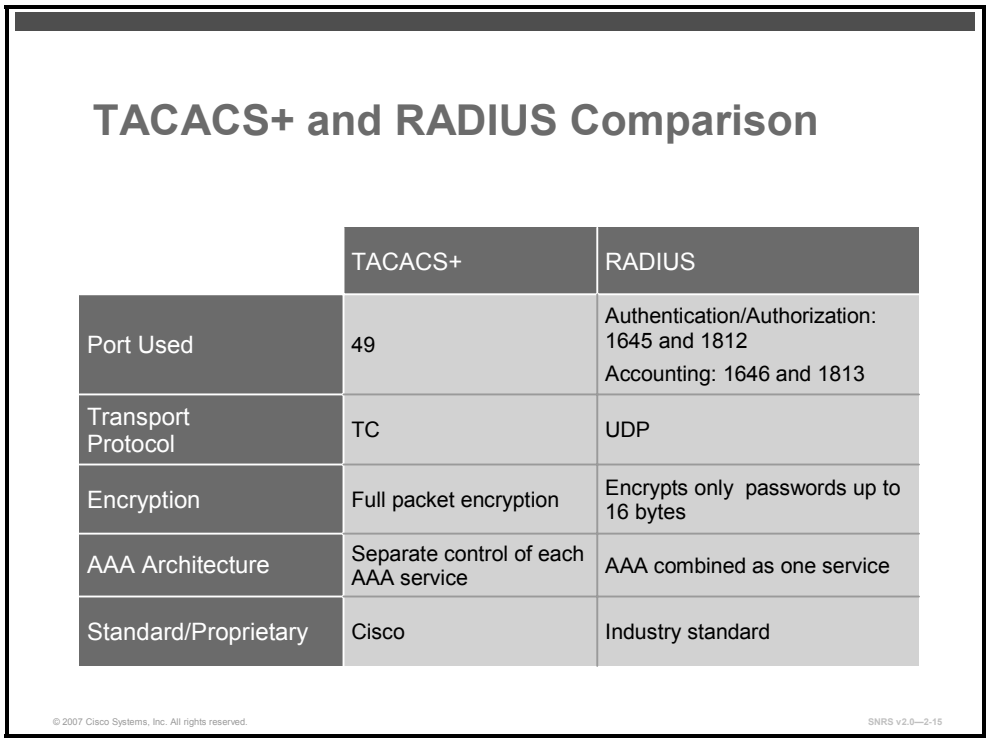
Note	TACACS+ is traditionally used to provide authorization for network administrative operations on the network infrastructure itself; RADIUS is universally used to secure the access of end users to network resources.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

TACACS+ and RADIUS

Cisco Secure ACS can use both the TACACS+ and RADIUS AAA protocols.

The diagram gives a comparison of TACACS+ and RADIUS.



TACACS+

Cisco Secure ACS conforms to the TACACS+ protocol as defined by Cisco Systems.

RADIUS

Cisco Secure ACS conforms to the RADIUS protocol as defined in these RFCs:

- RFC 2138, *Remote Authentication Dial In User Service (RADIUS)*
- RFC 2139, *RADIUS Accounting*
- RFC 2284, *PPP Extensible Authentication Protocol (EAP)*
- RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*, Obsoletes 2138
- RFC 2866, *RADIUS Accounting*
- RFC 2867, *RADIUS Accounting Modifications for Tunnel Protocol Support*, Updates 2866
- RFC 2868, *RADIUS Attributes for Tunnel Protocol Support*, Updates 2865
- RFC 2869, *RADIUS Extensions*

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

The ports used for authentication and accounting have changed in RADIUS RFC documents. To support the older and newer RFCs, Cisco Secure ACS accepts authentication requests on port 1645 and port 1812. For accounting, Cisco Secure ACS accepts accounting packets on ports 1646 and 1813.

In addition to support for standard Internet Engineering Task Force (IETF) RADIUS attributes, Cisco Secure ACS includes support for RADIUS vendor-specific attributes (VSAs). The following VSAs have been predefined in Cisco Secure ACS:

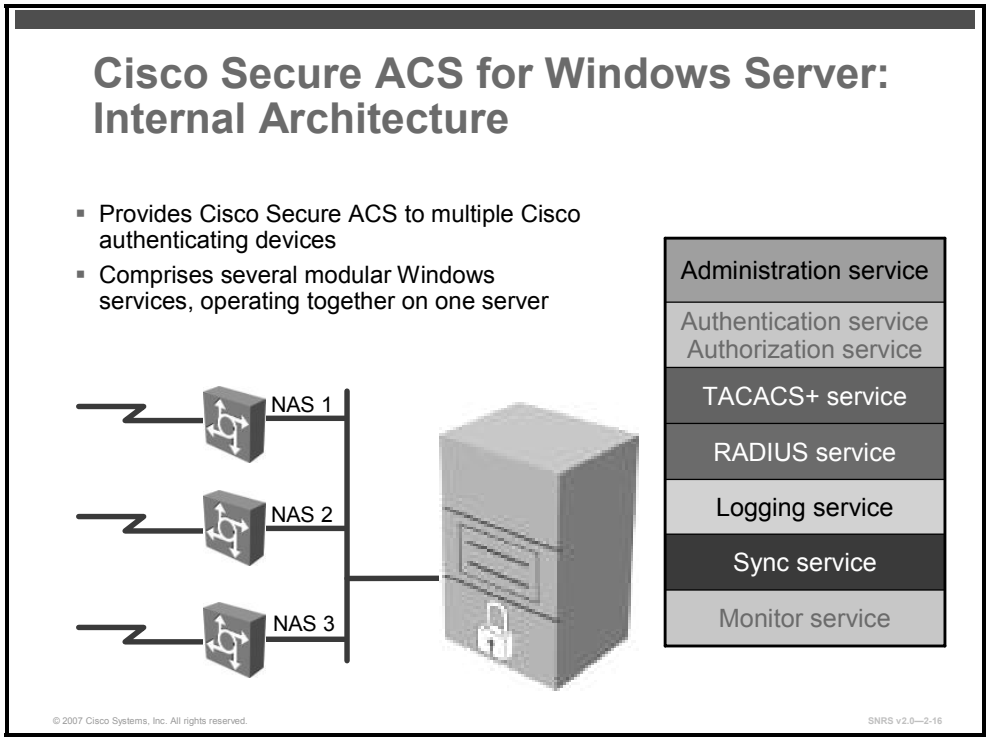
- Cisco Building Broadband Service Manager (BBSM)
- Cisco IOS/PIX 6.0
- Cisco VPN 3000/ASA/PIX 7.x+
- Cisco VPN 5000
- Cisco Aironet
- Ascend
- Juniper
- Microsoft
- Nortel

Cisco Secure ACS also supports up to 10 RADIUS VSAs that an administrator may define. After defining a new RADIUS VSA, you may then use it as you would one of the RADIUS VSAs that come predefined in Cisco Secure ACS. In the Network Configuration section of the Cisco Secure ACS web interface, you can configure AAA clients to use a user-defined RADIUS VSA as the AAA protocol. In Interface Configuration, you can enable user-level and group-level attributes for user-defined RADIUS VSAs. In User Setup and Group Setup, you can configure the values for enabled attributes of a user-defined RADIUS VSA.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Cisco Secure ACS for Microsoft Windows Architecture

This topic describes the Cisco Secure ACS architectural components.



When you install Cisco Secure ACS, the installation adds several Microsoft Windows services. The services provide the core of Cisco Secure ACS functionality. The Cisco Secure ACS services on the computer running Cisco Secure ACS include those listed here.

CSAdmin

CSAdmin is the service that provides the web server for the Cisco Secure ACS web interface. After Cisco Secure ACS is installed, you must configure it from its web interface; therefore, CSAdmin must be running when you configure Cisco Secure ACS.

CSAuth

CSAuth is the authentication and authorization service. It permits or denies access to users by processing authentication and authorization requests. This is the Cisco Secure ACS database manager.

CSDBSync

CSDBSync is the service used to synchronize the Cisco Secure ACS database with third-party RDBMS systems.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

CSLog

CSLog is the service used to capture and store logging information.

CSMonitor

CSMonitor works for TACACS+ and RADIUS and automatically detects which protocols are in use.

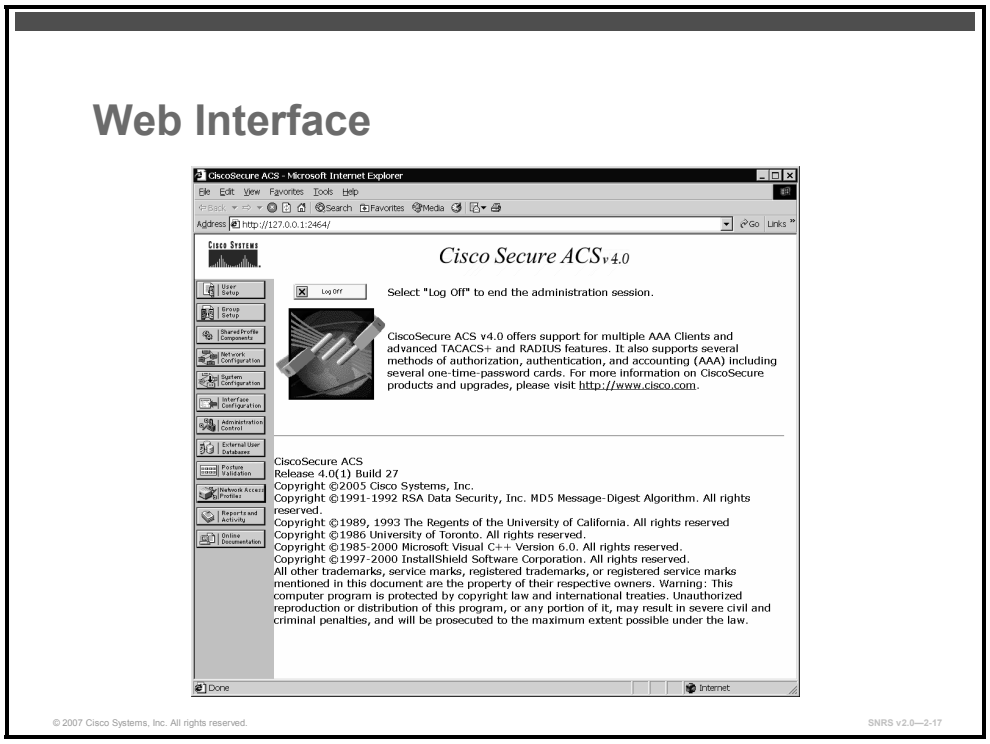
CSTacacs and CSRadius

The CSTacacs and CSRadius services communicate between the CSAuth module and the access device that is requesting authentication and authorization services.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Administering Cisco Secure ACS

This topic describes the Cisco Secure ACS web interface.



Nearly all Cisco Secure ACS administration tasks can be performed through the Cisco Secure ACS web interface. You use the web interface to easily modify and view the Cisco Secure ACS configuration from any connection on the LAN or WAN, by using a web browser.

Note	Accessing the web interface requires a valid administrator name and password.

HTTP port 2002 is used for administrative access through a web browser.

Layout

The web interface has three vertical frames:

- **Navigation bar:** The navigation bar is the gray frame on the left side of the browser window that contains the task buttons. Each button changes the configuration area to a unique section of the Cisco Secure ACS application, such as the User Setup section or the Interface Configuration section. This frame does not change; it always contains the following buttons:
  - **User Setup:** Add and edit user profiles
  - **Group Setup:** Configure network services and protocols for groups of users
  - **Shared Profile Components**: Add and edit network access restriction and command authorization sets, to be applied to users and groups
  - **Network Configuration:** Add and edit NADs and configure distributed systems

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

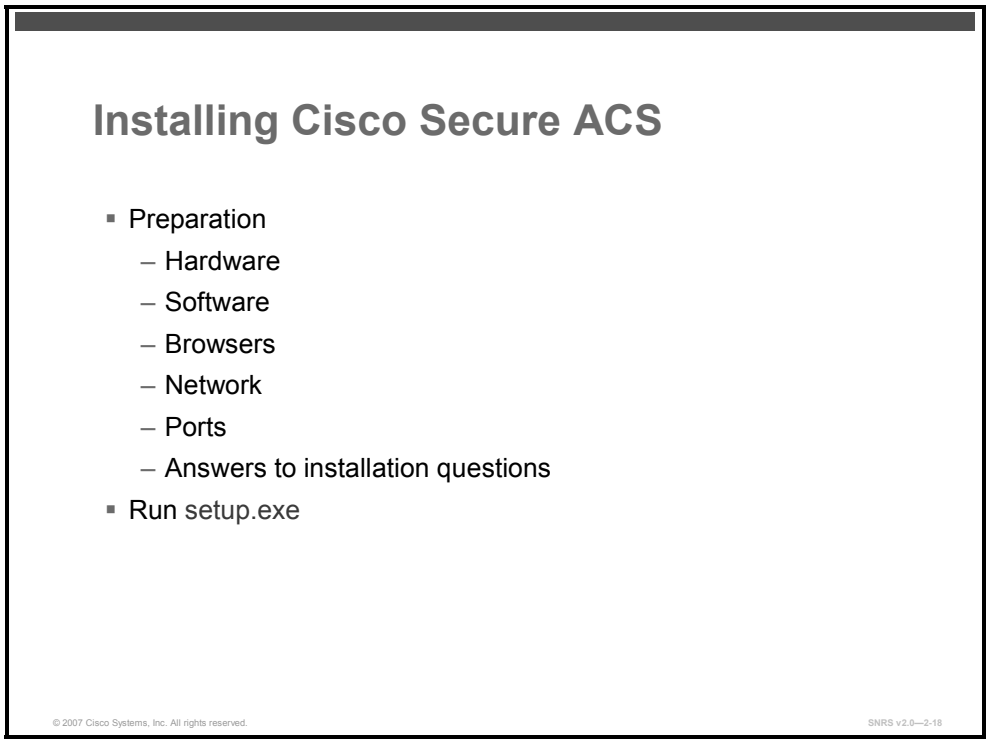
- **System Configuration:** Configure system-level features
- **Interface Configuration:** Display or hide product features and options to configure
- **Administration Control:** Define and configure access policies
- **External User Databases:** Configure databases, the unknown user policy, and user group mapping
- **Posture Validation:** Control the degree of access that is permitted from computers that access your network through AAA clients that are configured to enforce NAC
- **Network Access Profiles:** Set up the conditions that allow a user to connect to the network and identify the way that requests to access the network are applied
- **Reports and Activity:** Display accounting and logging information
- **Online Documentation:** View the user guide

- **Configuration area:** The frame in the middle of the browser window, the configuration area displays web pages that belong to one of the sections represented by the buttons in the navigation bar. The configuration area is where you add, edit, or delete information.
- **Display area:** Shows one of the following options:
  - **Online Help**
  - **Reports or Lists**
  - **System Messages**

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Installing Cisco Secure ACS

This topic describes how to install Cisco Secure ACS on a Microsoft Windows server.



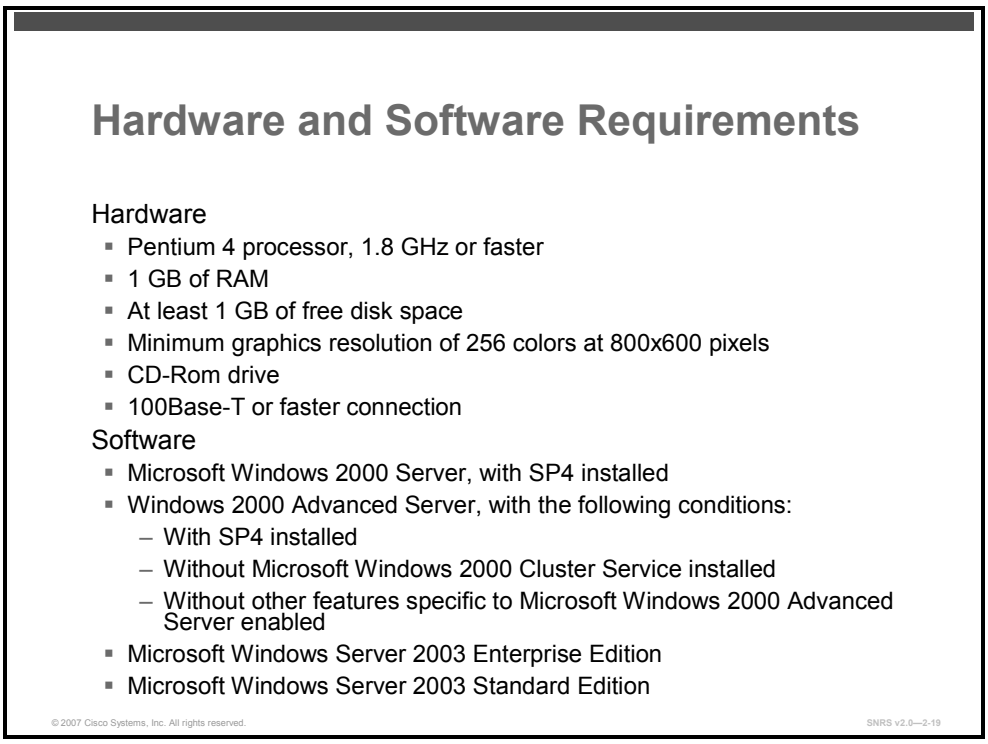
Cisco Secure ACS operates on Microsoft Windows 2000 Server and Microsoft Windows Server 2003.

Preparation

When you are installing Cisco Secure ACS, certain minimum hardware, OS, and third-party software requirements must be met previous to installation. Some of these requirements for preparation include:

- Hardware configuration of the Microsoft Windows server
- Software on the Microsoft Windows server
- Types of browsers used to administer Cisco Secure ACS
- Network requirements
- Ports that are used to communicate with Cisco Secure ACS
- Answers to installation questions such as administrator and database passwords

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.



Hardware Requirements

Any server running Cisco Secure ACS must meet the following minimum hardware requirements:

- IBM PC-compatible with a Pentium 4 processor, 1.8 GHz or faster
- 1 GB of RAM
- At least 1 GB of free disk space (If you are running the database on the same computer, more disk space is required.)
- Minimum graphics resolution of 256 colors at 800x600 pixels
- CD-ROM drive
- 100Base-T or faster connection

OS Requirements

Cisco Secure ACS 4.0 for Windows supports the Microsoft Windows operating systems listed. Both the OS and the service pack must be English-language versions.

- Microsoft Windows 2000 Server (Service Pack 4 [SP4])
- Microsoft Windows 2000 Advanced Server, with the following conditions:
  - SP4 installed
  - Without Microsoft Windows 2000 Cluster Service installed
  - Without other features specific to Microsoft Windows 2000 Advanced Server enabled, such as Microsoft Windows 2000 Terminal Services

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Note The multiprocessor feature of Microsoft Windows 2000 Advanced Server has not yet been tested, and cannot be considered supported. Microsoft Windows 2000 Data Center Server is not a supported OS.

- — Microsoft Windows Server 2003 Enterprise Edition (Service Pack 1 [SP1])
- — Microsoft Windows Server 2003 Standard Edition (SP1)

Browsers

Cisco Secure ACS is compatible with the following browsers:

- Microsoft Internet Explorer 6 SP1 and Microsoft Internet Explorer 5.5 for Microsoft Windows—English and Japanese version
- Netscape 7.0, 7.1, and 7.2 for Microsoft Windows—English and Japanese version

Caution Several known problems are related to using Netscape Communicator with Cisco Secure ACS. For more information, see the *Release Notes for Cisco Secure ACS for Windows* on Cisco.com.

- Sun Java Runtime Environment (JRE) 1.4.2_04 or Microsoft Java Virtual Machine (JVM)

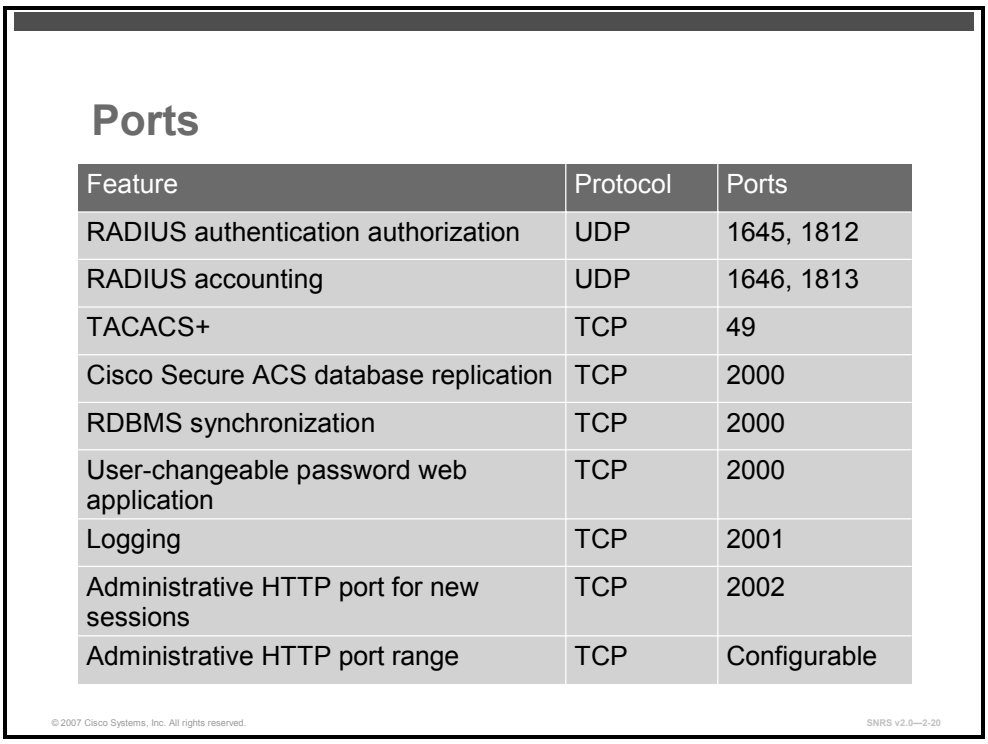
Network Requirements

Your network should meet the following requirements before you begin deploying Cisco Secure ACS:

- The computer that is running Cisco Secure ACS must be able to ping all AAA clients.
- Dial-in, VPN, or wireless clients must be able to connect to the applicable AAA clients.
- For full TACACS+ and RADIUS support on Cisco IOS devices, AAA clients must run Cisco IOS Release 11.1 or later.
- AAA clients from other vendors must be configured with TACACS+, RADIUS, or both.
- Gateway devices (such as firewalls) between Cisco Secure ACS and the AAA clients must permit communication over the ports needed to support the applicable feature or protocol.
- A supported web browser must be installed on the computer that is running Cisco Secure ACS.
- All network cards in the computer that is running Cisco Secure ACS must be enabled.
- If you want Cisco Secure ACS to use the Grant Dial-in Permission to User feature in Microsoft Windows when authorizing network users, you must select this option in the Windows User Manager or Active Directory Users and Computers for the applicable user accounts.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.



Ports

The table lists the ports that are used as Cisco Secure ACS listens for communications with AAA clients, other Cisco Secure ACS machines and applications, and web browsers.

Cisco Secure ACS uses other ports to communicate with external user databases; however, it initiates those communications rather than listening to specific ports. For example, if Cisco Secure ACS initiates communications with LDAP or RADIUS token server databases, you can configure these destination ports in Cisco Secure ACS.

Questions for Installation

At the beginning of the installation process, you will be asked to check the following:

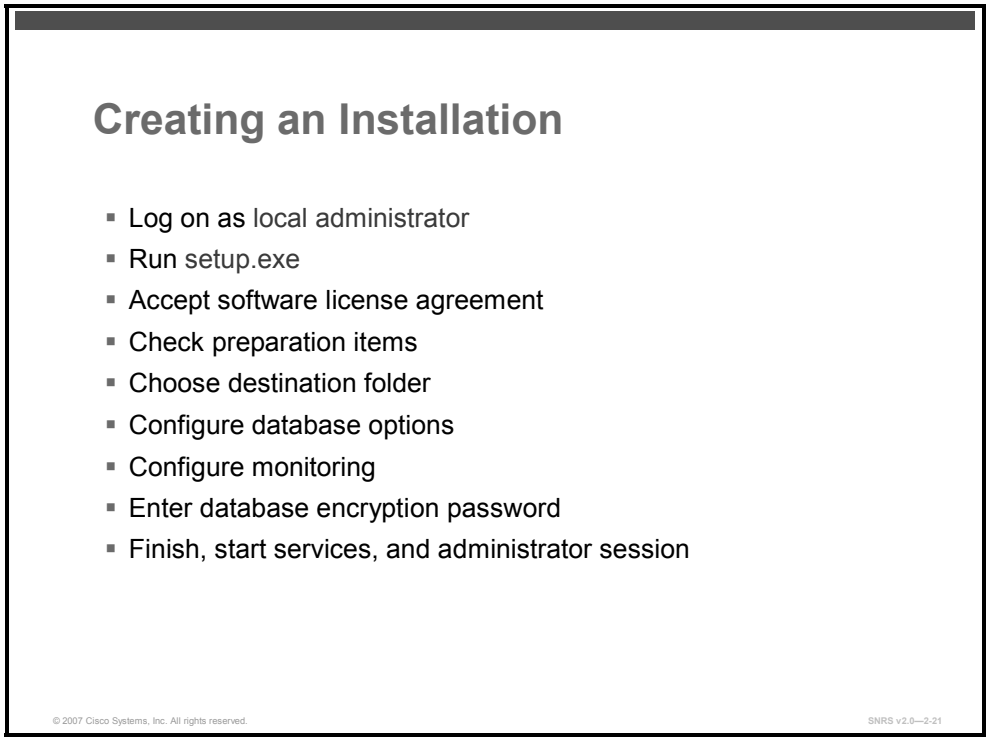
- End-user clients can successfully connect to AAA clients.
- This Microsoft Windows server can ping the AAA clients.
- Cisco IOS clients are running Cisco IOS Release 11.1 or later.
- Microsoft Internet Explorer 6 SP1 or Netscape 7.02 is installed.

You will also need to create an administrator password and a database password.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Creating an Installation

This topic describes how to perform a Cisco Secure ACS installation using setup.exe on the Cisco Secure ACS CD-ROM.



Complete the following steps to install Cisco Secure ACS for the first time.

Step 1 Log onto the computer using a local administrator account.

Step 5 Click **setup.exe**, located in the root directory of the CD-ROM.

Step 6 Read the software license agreement and click **Accept**.

Step 7 Click **Next** after reading the Welcome screen. The Before You Begin dialog box appears.

Step 8 Click **Next** after you have completed all items in the Before You Begin dialog box and checked the corresponding check box for each item. The Choose Destination Location dialog box appears.

Step 9 Click **Next**, unless you want to change the destination folder. The Authentication Database Configuration dialog box appears.

Step 10 There are two options:

- Click **Check the Cisco Secure ACS Database Only** if you want to authenticate users with the Cisco Secure ACS internal database only.
- If you want to authenticate users with a Windows Security Access Manager (SAM) user database or Active Directory user database in addition to the Cisco Secure ACS internal database, you can do the following:

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

— Click Also Check the Windows User Database. The Yes, Refer to "Grant Dial-in Permission to User" Setting check box becomes available.

— Click Yes, Refer to "Grant Dial-in Permission to User" Setting if you want to allow access by users who are authenticated by a Microsoft Windows domain user database only when they have dial-in permission in their Microsoft Windows account.

Note The Yes, Refer to "Grant Dial-in Permission to User" Setting check box applies to all forms of access that Cisco Secure ACS controls, not just dial-in access. For example, a user accessing your network through a VPN tunnel is not dialing into an NAS; however, if the Yes, Refer to "Grant Dial-in Permission to User" setting check box is selected, Cisco Secure ACS applies the Microsoft Windows user dial-in permissions to determine whether to grant the user access to your network.

Step 11 Click **Next**. The setup program installs Cisco Secure ACS and updates its configuration.

Step 12 The Advanced Options dialog box displays several features of Cisco Secure ACS that are not enabled by default. Check the corresponding check box for each feature that you want to enable. Click **Next**. The Active Service Monitoring dialog box appears.

Step 13 Click **Next**. The Database Encryption Password dialog box appears.

Step 14 Enter a password for database encryption and click **Next**. The setup program ends and the Cisco Secure ACS Service Initiation dialog box appears.

Step 15 For each option that you require, check the corresponding check box; then click **Next**.

- Yes, I Want to Start the Cisco Secure ACS Service Now
- Yes, I Want Setup to Launch the Cisco Secure ACS Administrator from My Browser Following Installation
- Yes, I Want to View the Readme File

Step 16 Click **Finish**. The setup program exits.

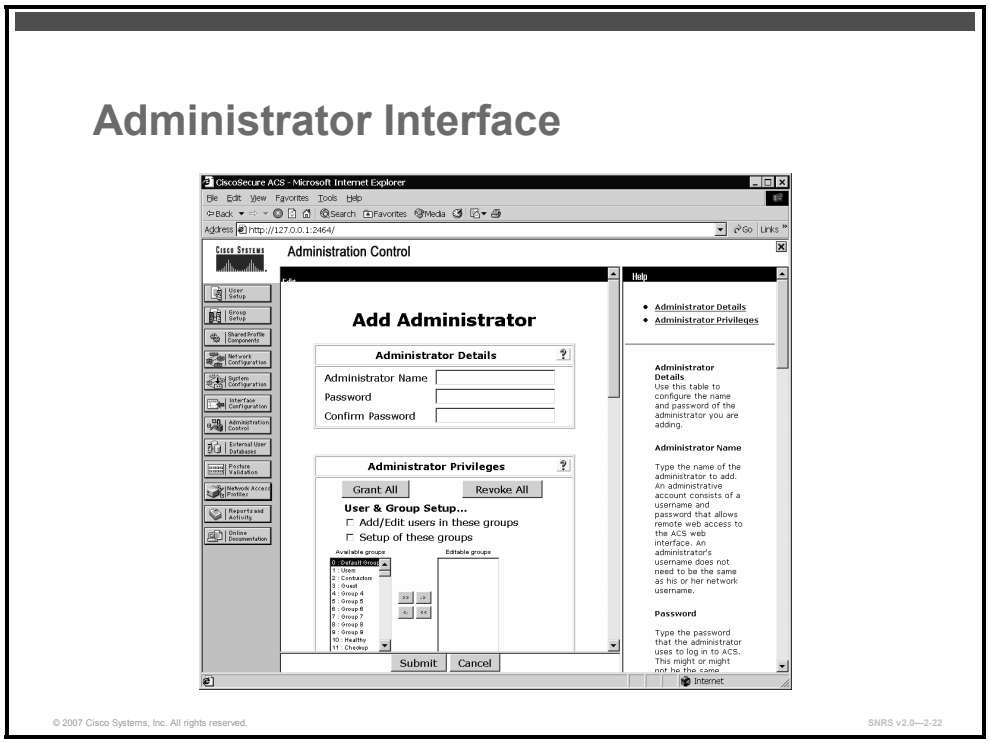
On the computer that is running Cisco Secure ACS, you can access the Cisco Secure ACS HTML interface by using the Cisco Secure ACS Admin desktop icon, or you can use one of the following URLs:

```
http://127.0.0.1:2002
        or
http://localhost:2002
```


The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Adding an Administrator

This topic describes how to add an administrator to Cisco Secure ACS.



If you plan to administer the Cisco Secure ACS from the network, you will have to create and enable an administrator first. An administrative account is not created by default.

To create an administrative account, follow these steps:

Step 1 Click **Administration Control**.

Step 17 Click **Add Administrator**.

Step 18 Complete the boxes in the Administrator Details table:

- Type the login name (up to 32 characters) in the Administrator Name box.
- Type the password (up to 32 characters) in the Password box.
- Type the password a second time in the Confirm Password box.

Step 19 Click **Grant All** to choose all privileges, including user group editing privileges for all user groups.

All privilege options are selected. All user groups move to the Editable groups list.

Note To clear all privileges, including user group editing privileges for all user groups, click **Revoke All**.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Access Policy

The Cisco Secure ACS Access Policy feature affects access to the Cisco Secure ACS HTML interface. You can limit access by IP address and by the TCP port range used for administrative sessions. You can also enable Secure Sockets Layer (SSL) for access to the HTML interface.

The IP address options include the following:

- Allow All IP Addresses to Connect
- Allow Only Listed IP Addresses to Connect
- Reject Connections from Listed IP Addresses

Note The IP addresses entered to define a range must differ only in the last octet.

The port options include the following:

- Allow Any TCP Ports to Be Used for Administration HTTP Access
- Restrict Administration Sessions to the Following Port Range
- Use HTTPS Transport for Administration Access (HTTPS stands for “secure HTTP.”)

Session Policy

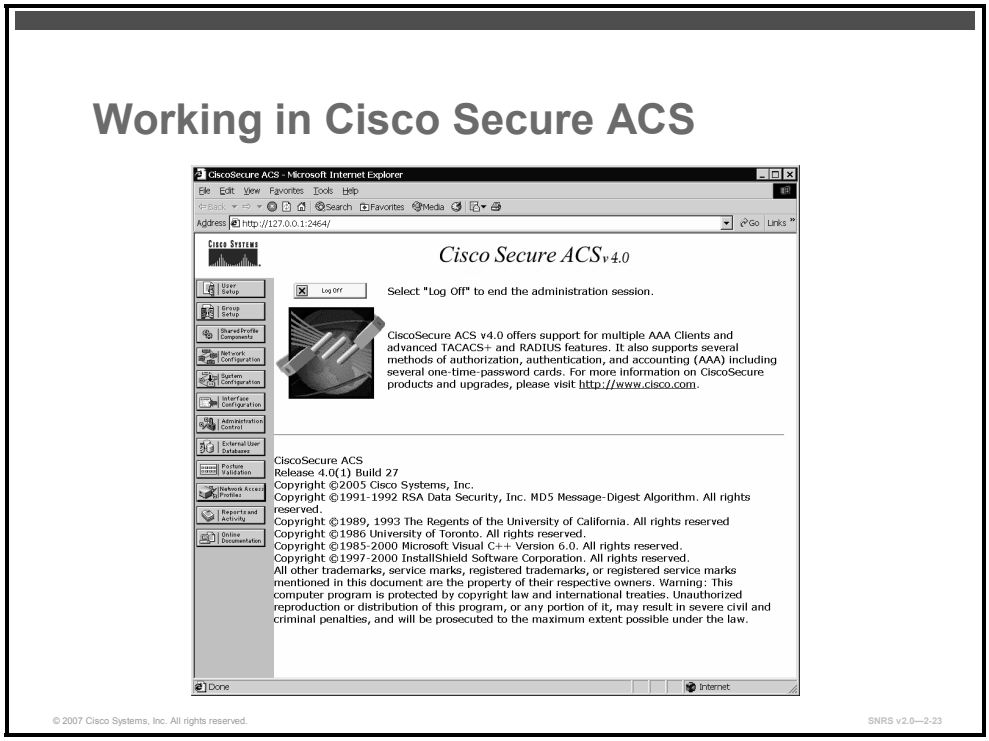
The Session Policy feature controls various aspects of Cisco Secure ACS administrative sessions such as the following:

- Session idle timeout
- Allow Automatic Local Login
- Invalid IP address response policy
- Failed administrative login attempts policy

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Working in Cisco Secure ACS

This topic describes how to navigate the HTML interface and provides a brief overview of each section of the interface.



The next sections will look at each menu item on the navigation bar and what types of configuration can be performed at each level.

User Setup

This button allows an administrator to add a new user, search for an existing user, find users alphabetically or numerically, or list all users at once.

Group Setup

This button allows an administrator to apply configurations from Shared Profile Components, as well as specific TACACS+ and RADIUS attributes. Group Setup is where an administrator can configure any parameters that are common to a group of users.

Shared Profile Components

This button allows an administrator to specify shell command authorization sets. By creating these command authorization sets, an administrator can control the commands that a user can execute on a device by applying the command authorization set to the user profile in the TACACS+ settings or at the group level.

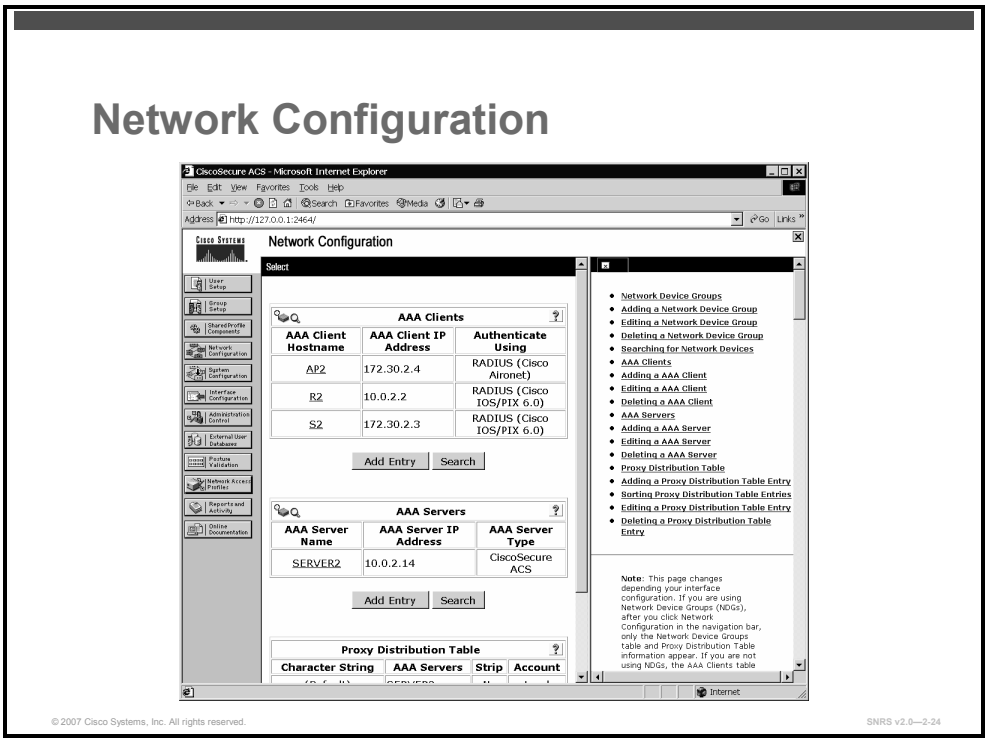
This is where you also configure downloadable ACLs and RADIUS Authorization Components. For these options to be visible, you must choose them in the Interface

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Configuration page. Another benefit is the ability to configure shared network access restrictions.

Network Configuration

This slide shows the Network Configuration screen..



This button is where an administrator can add, delete, or modify settings for AAA clients (NADs). The layout of this page changes depending on the settings for interface configuration. If you are using NDGs, after you click Network Configuration in the navigation bar, only the Network Device Groups table and Proxy Distribution Table information appears. If you are not using NDGs, the AAA Clients table and the AAA Servers table appear in place of the Network Device Groups table.

System Configuration

This button allows the administrator to configure several of the basic system settings and to restart the Cisco Secure ACS service. The following subconfiguration links may be found under System Configuration:

- **Service Control:** From Service Control, you can start and stop the Cisco Secure ACS services. You can also start and stop the Cisco Secure ACS services in the Control Panel of Microsoft Windows. By stopping the Cisco Secure ACS service from within Cisco Secure ACS, you do not stop the Cisco Secure ACS web server. If you want to stop the Cisco Secure ACS web server, you need to do so in the Control Panel of Microsoft Windows. This service is called CSAdmin.
- **Logging:** From the Logging task, you can configure local logging, such as failed attempts as well as TACACS+ and RADIUS accounting. You can also configure Open Database Connectivity (ODBC) and remote logging here, as well as other Cisco Secure ACS devices.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

- **Date Format Control:** The Date Format Control option allows you to change the format of the date displayed on reports. After you change the format, you must log out of the server to actually see the changes take effect.
- **Local Password Management:** From this task, you can set password length and password options. You can also configure options for Remote Password Change and logging of password changes.
- **ACS Backup:** This option allows you to schedule backups to be performed manually or automatically at specific times. You can also specify a location for the backup files to be stored and manage the backup files. When Cisco Secure ACS is backed up, it creates a file with the extension of .dmp. This file will be present when you enter the ACS Restore link. Here you have the ability to select from numerous backup files and to determine if you want to restore the users and groups, system configuration, or both.
- **ACS Service Management:** ACS Service Management enables the administrator to determine how often to test the availability of Cisco Secure ACS authentication services. This is the CSMonitor service configuration. CSMonitor allows Cisco Secure ACS to test itself and take action when the test is unsuccessful. If no authentications are recorded, the available actions are as follows:
    - Restart all
    - Restart RADIUS/TACACS
    - Reboot
    - Take no action

Note	If the reboot option is chosen, this causes the server running Cisco Secure ACS to reboot.

Interface Configuration

In the Interface Configuration section, you will find a selection from the following subconfiguration links, depending on whether you have chosen TACACS+ or a form of RADIUS when you entered your AAA client:

- User Data Configuration
- TACACS+ (Cisco IOS Software)
- RADIUS (Microsoft)
- RADIUS (Ascend)
- RADIUS (IETF)
- RADIUS (IOS/PIX)
- Advanced Options

Caution	If you do not see RADIUS options here, you need to add a AAA client that uses the RADIUS protocol. Interface Configuration is directly affected by the settings in Network Configuration.

The User Data Configuration link enables you to customize the fields that appear in the User Setup and Configuration page. Here you can add fields such as Phone Number, Work Location, Supervisor Name, or any other information that you feel is pertinent.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

The TACACS+ (Cisco IOS) link enables the administrator to configure TACACS+ settings and to add new TACACS+ services. You can also configure advanced options that affect what you see in your interface. It is important that you understand how this works. Depending on the current configuration of your server, if you go to the TACACS+ link, you may or may not see two columns. If you *do* see two columns, you are able to configure user-level settings and group-level settings.

Administration Control

The Administration Control section is where you configure all aspects of Cisco Secure ACS for administrative access. Here you have the ability to add administrators and configure access policy. Information, such as which IP addresses are allowed or not allowed to access Cisco Secure ACS, and HTTP port allocation, can be configured here.

Remember that Cisco Secure ACS uses port 2002 as the listening port, but after a connection is made to that port, you are redirected to a random port number. When Cisco Secure ACS is positioned behind a firewall, this random port assignment becomes a security issue. You have the ability to specify a range of ports used so that you can configure access restrictions within your firewall. This is especially helpful when using a Cisco PIX Firewall.

External User Databases

The External User Databases section consists of three subsections. In addition to configuring the parameters to communicate with the external databases, you can configure how Cisco Secure ACS handles requests from users that are not in the local Cisco Secure ACS database (Unknown User Policy), and a mapping from the external database group to the local Cisco Secure ACS database group.

In this section, you configure an unknown user policy. You also configure database group mappings to external user databases and perform the actual database configuration. Further, you are given a list of compatible databases, and you can choose which one you will configure to be used with Cisco Secure ACS.

The following servers are available for use as an external database:

- Microsoft Windows Server
- Novell NDS
- Generic LDAP
- External ODBC database
- LEAP proxy RADIUS server
- RADIUS token server
- VASCO token server
- ActivCard token server
- PassGo Defender token server
- CRYPTOCard token server
- SafeWord token server
- RSA SecurID token server

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Each version of Cisco Secure ACS includes more support for external databases, while greatly improving the functionality of the Cisco Secure ACS database.

Posture Validation

Cisco Secure ACS supports the NAC initiative. NAC is used to ensure that every endpoint conforms to the security policy before being granted access to the network. Posture validation is the mechanism used to determine the state of the endpoint that is requesting network access. The Posture Validation section is used to configure the different ways you can set up posture validation policies.

The administrator can choose to create and modify internal posture validation policies, or configure Cisco Secure ACS to forward selected credentials to external policy validation servers. Audit servers can also be configured to determine the posture of endpoints without the presence of a posture agent (used to forward endpoint credentials).

Network Access Profiles

A NAP is a means to classify access requests, according to the IP addresses of AAA clients, membership in an NDG, protocol types, or other specific RADIUS attribute values sent by the network device through which the user connects. The use of NAPs allows the administrator to configure different authentication mechanisms and authorizations depending on the characteristics of the access request, resulting in increased flexibility.

The Network Access Profile section is used to create profiles, and associate a set of rules and policies with them.

Reports and Activity

The Reports and Activity section provides a wealth of tools for both troubleshooting and monitoring the network.

Caution	Logging consumes resources, and the log files should be checked periodically for content and size.

The available logs that Cisco Secure ACS keeps are as follows:

- **TACACS+ Accounting:** The information that is included in these reports is configurable by the administrator in the System Configuration section under Logging.
- **TACACS+ Administration:** These reports include all of the command requests from AAA clients, such as routers or firewalls where command authorization is configured.
- **RADIUS Accounting:** These reports include the same type of information as mentioned under the TACACS+ Accounting log. The information included in this report is also configurable by the administrator under System Configuration.
- **VoIP Accounting:** These reports include information on VoIP sessions, including session duration and AAA usernames.
- **Passed Authentications:** This report lists successful authentication requests. The information contained in these reports assists with user administration and in troubleshooting users that are failing authentication.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

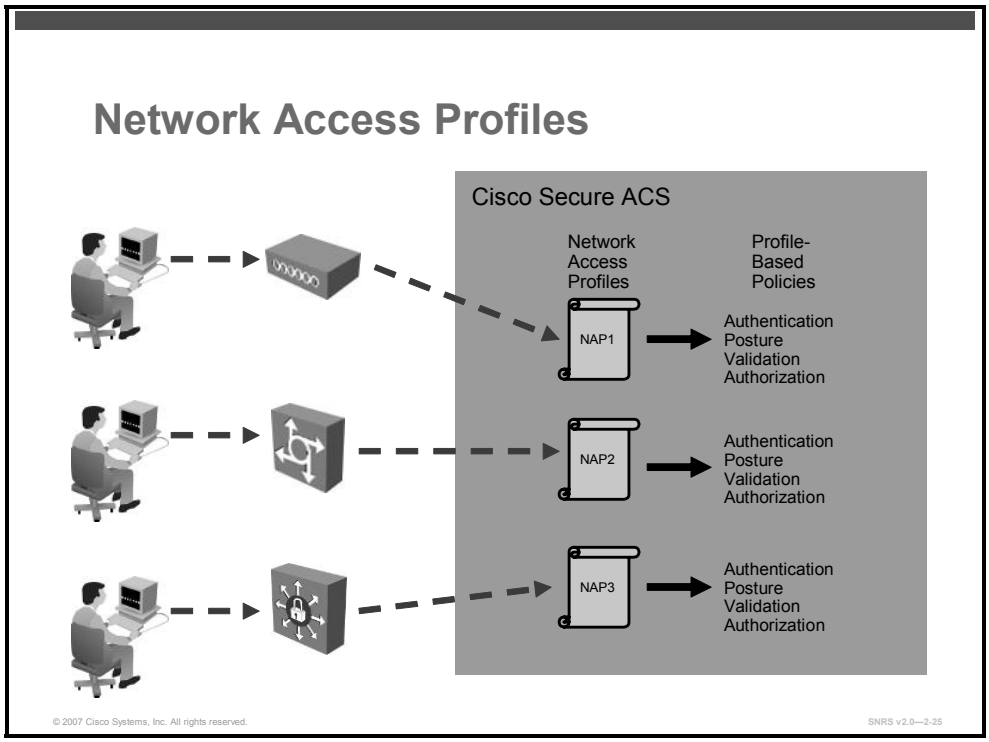
- **Failed Attempts:** This report lists authentication and authorization failures with an indication of the cause. This report also assists with user administration and in troubleshooting users that are failing authentication.
- **Logged-in Users:** The Logged-in Users file is rather unique. Most of the logging files in Cisco Secure ACS are stored as comma-separated values (CSV) files and are stored for a period of time, usually one day, on the hard drive of the server. The Logged-in Users file is not saved as a CSV file. As users log in, they are maintained in this file, organized by the name of the AAA client. This report also assists with user administration and in troubleshooting users that are failing authentication.
- **Disabled Accounts:** This report enables you to view accounts that have been disabled.
- **ACS Backup and Restore:** This log maintains a history of the dates and times that Cisco Secure ACS was backed up or restored.
- **Remote Database Management Source (RDBMS) Synchronization:** This log allows Cisco Secure ACS to keep report information on RDBMS synchronization. It logs the time of and reason for RDBMS synchronization.
- **Database Replication:** This report logs the time that the ACS database was replicated to the backup server.
- **Administration Audit:** This log keeps track of who logged in, what users and groups they made changes to, and what time they logged out. It logs all of the activity in Cisco Secure ACS that is performed by administrators.
- **User Password Changes:** This report tracks changes that were made to user passwords using the User Changeable Password Module.
- **ACS Service Monitoring:** This report keeps track of all of the events that occur to Cisco Secure ACS services that are monitored. An example of a service that might be monitored is CSAdmin or CSTacacs. By default, CSMonitor is enabled; however, this is configurable. To configure this, choose **System Configuration > ACS Service Management**. Here, you can choose to monitor the login process, generate events when someone tries to log into disabled accounts, and so on.

It is possible to view these reports in the Cisco Secure ACS HTML interface or from the hard drive of the Cisco Secure ACS server. The logs are stored as CSV files, except where noted; therefore, you can view them in a spreadsheet program such as Microsoft Excel. You can also import the files into third-party software, such as Crystal Reports. If you do not have access to the hard drive of the Cisco Secure ACS server, you have the ability to download the logs using the Cisco Secure ACS HTML interface.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Network Access Profiles

This topic describes the use of NAPs.



Cisco Secure ACS 4.0 for Windows introduces the concept of NAPs. NAPs provide the ability to process network access requests differently depending on the characteristics of the request.

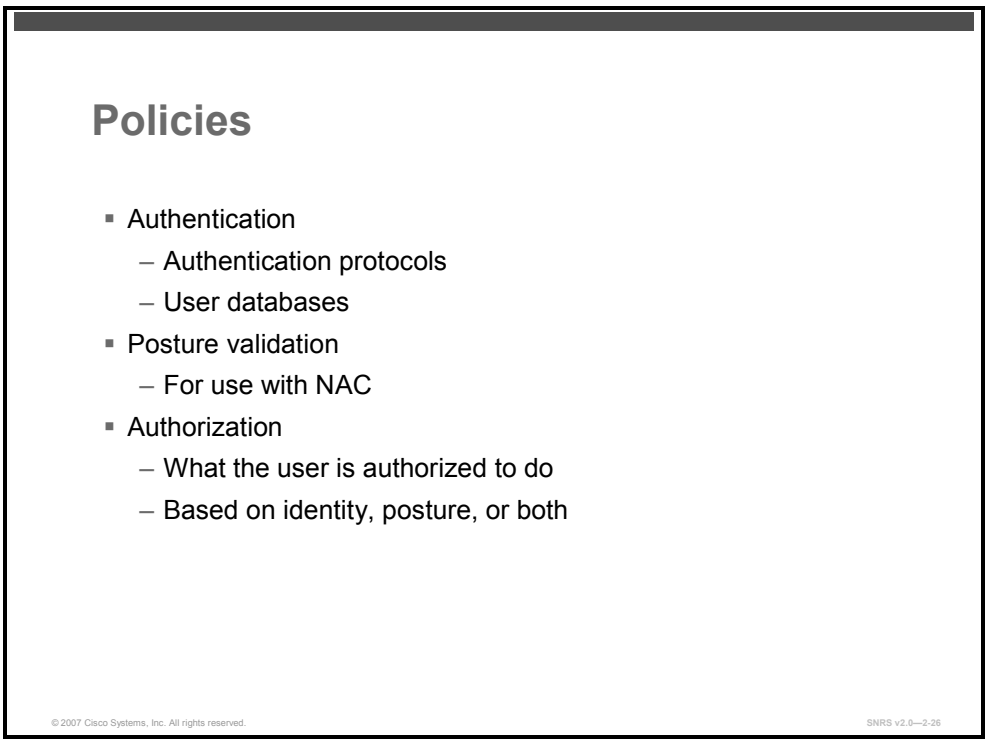
Most organizations have various kinds of users who access the network in different ways and for different purposes. Correspondingly, you must apply different security policies to the different use cases. For example, you might have to apply a tighter and more limiting security policy to the wireless access points of the lobby area of your building, versus the physically secured production plant. Or, you might have to treat remote access users who use a VPN differently from users that log in from behind a firewall. Users who connect through certain subnetworks might be authenticated differently from other users. Wireless access is often treated more strictly than wired access, as is any form of remote access (dialup, VPN, home wireless).

The configuration of a NAP consists of defining the characteristics of an access request that will determine the membership of an access request in the profile, and the authentication and authorization policies that are to be applied to an access request matching the profile.

There are three characteristics of an access request that are used to determine how it is classified and mapped to a profile. They are as follows:

1. The NAF, which are groupings of AAA client configurations (which may represent multiple network devices), NDGs, or IP addresses of specific AAA client devices.
2. The protocol type, which are AAA client vendor types from which access requests are allowed.
3. The specific RADIUS attribute value (or values) sent by the NAD that an endpoint attaches to when requesting network access.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.



Policies

Each NAP created must have a set of policies to reflect the security policy for the access type represented by the profile (that is, access via VPN or wireless). These associations are called profile-based policies. Profile-based policies include rules for authentication, authorization, and posture validation.

Configuring a profile-based policy includes creating rules for the following:

- **Authentication rules:** Allowed protocols and credential validation databases
- **Posture validation rules:** Used when implementing NAC
  - A posture is the term used to describe the set of attributes sent by the posture agent of the endpoint requesting network access defining its state and health.
  - A posture validation rule is made up of a condition and actions. The condition is a set of required credential types. The actions are to determine which internal policies or external servers should be used for posture validation.
- **Authorization rules:** What the user is allowed to do on the network. Rules may be based on group membership, the posture of the machine that is used to access the network, or both.

How Policies Are Applied

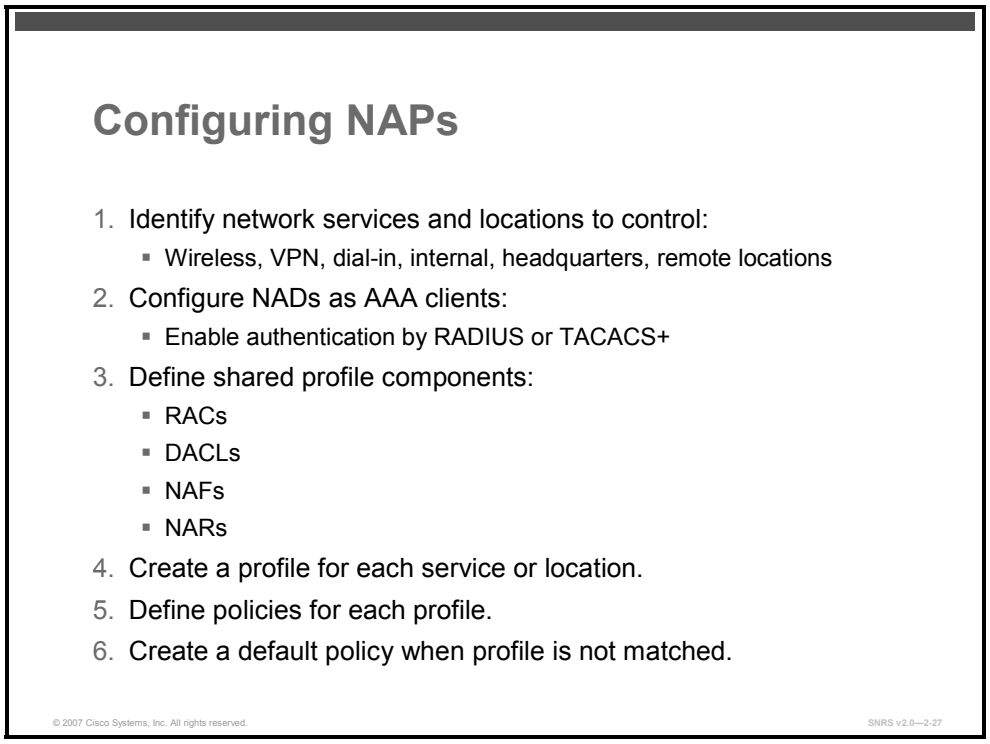
Policies are applied by Cisco Secure ACS going down the list of active NAPs. It goes down the list until the first match is made. Actions defined in the policies associated with this profile are executed.

When a policy is not matched, you can configure Cisco Secure ACS to either deny access or grant access using global configuration for authentication and authorization.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Configuring Cisco Secure ACS NAPs

This topic describes how to configure NAPs.



There are several steps required to configure NAPs.

One of the first things you want to do is determine how you want to group your profiles into services and into locations.

Common Configurations in Cisco Secure ACS

Several Cisco Secure ACS elements must be configured first to configure a NAP and its policies.

- AAA clients
  - Addresses of NADs
  - Encryption key
  - Authentication using TACACS+ or RADIUS <*vendor*>
- Interface configuration
  - In Advanced Options, allow the following:
    - Per-user TACACS+ or RADIUS attributes
    - Group-level shared network access restrictions
    - Group-level downloadable ACLs
    - Network access filtering
    - NDGs

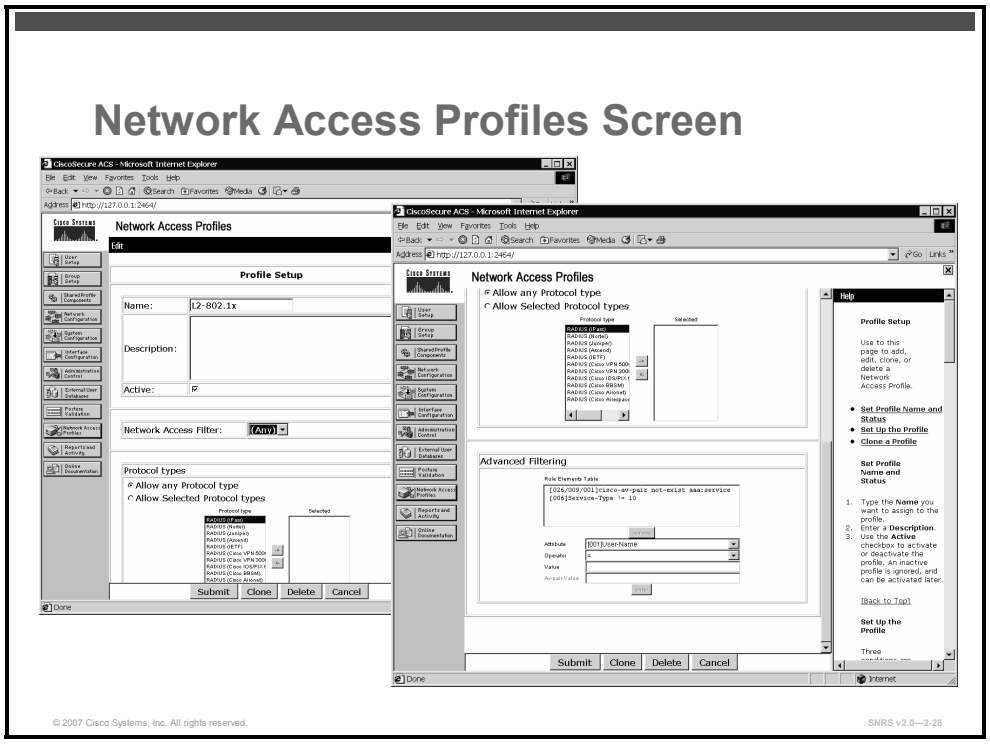
The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

- NAFs
  - Group NADs into locations or by other criteria
- RADIUS authorization components
  - Configure VSAs
  - Assign VLAN with 802.1x
- Downloadable ACLs
  - Create ACLs to manage access
- NARs
  - Allow access by IP address, MAC address
- Reports and logs
  - Failed authentication report
  - Passed authentication report

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Creating a NAP

This topic describes how to create a NAP.



When you edit the Profile Setup page, you are actually defining how Cisco Secure ACS classifies access requests. Each profile has a name and is either active or inactive. Profiles are associated with RADIUS attributes. Three conditions determine how Cisco Secure ACS classifies an access request and maps it to a profile. To create a NAP, you must configure the following conditions:

- NAF
- Allowed protocols
- Advanced filtering

The profile is selected when all of the selected conditions match.

You can manually build your profile or use one or more of the built-in templates.

Templates

You can also use the built-in templates supplied with Cisco Secure ACS 4.0 for Windows. These templates are very useful, especially when you are implementing NAC.

These templates are included in Cisco Secure ACS 4.0 for Windows:

- IEEE 802.1x
- NAC Layer 2 802.1x
- NAC Layer 2 IP

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

- NAC Layer 3 IP
- Wireless (NAC Layer 2 802.1x)
- NAC agentless host (NAH)
- MAC authentication bypass

NAFs

NAFs are used to define flexible network device restriction policies. This provides a flexible mechanism for applying network access restrictions and downloadable ACLs by network device names, NDGs, or IP addresses. NAFs can use IP address ranges and wildcards.

NAFs can be used to group (and name) sets of devices or to differentiate user requests on the same type of device.

Protocol Types

Protocol types are used to classify a user request based on the type of protocol that is used to request access to the network. Protocol types are a subset of the VSAs that RADIUS supports.

Note The TACACS+ protocol for NAPs is not supported in Cisco Secure ACS 4.0 for Windows.

Advanced Filtering

Advanced filtering is based on a Boolean AND expression of RADIUS attributes. Advanced filtering is used to create rules based on specific RADIUS attributes and values (including Cisco attribute-value pairs).

Multiple-rule elements are allowed, which are treated as a Boolean AND expression. The operators “contains,” “start with,” and “regular expression” apply only to string-type attribute values.

The rule elements table is used to dictate the rule elements that make up a rule based on a RADIUS attribute.

The components of the rule elements table are as follows:

- Attribute
  - — Lists all RADIUS attributes that you can use to specify rules
- Operator
  - — The comparison method by which Cisco Secure ACS evaluates whether the rule element is true or false
- Value
  - — The value to which Cisco Secure ACS compares the contents of the attribute
- Attribute-value pair key
  - — Enabled when you select the Cisco attribute-value pair attribute and Cisco Secure ACS compares the contents of the attribute-value pair key attribute from the access request

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

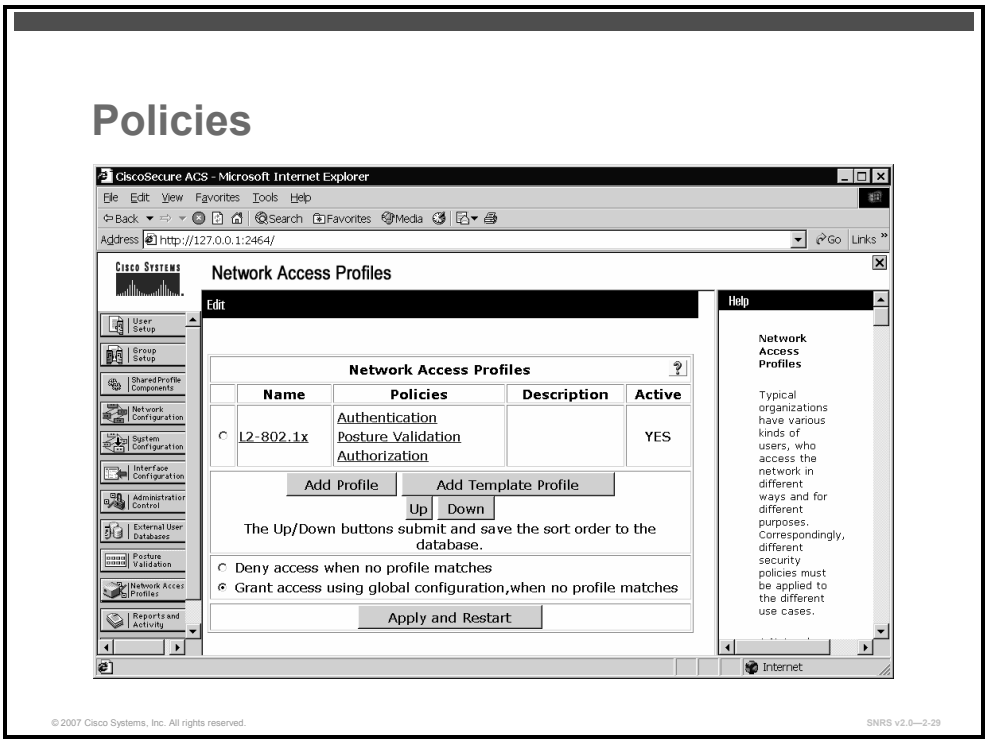
- Attribute-value pair value
  - Enabled when you select the Cisco attribute-value pair attribute and Cisco Secure ACS compares the contents of the attribute-value pair value attribute from the access request
- Enter button
  - Adds the rule element that is defined in the rule elements table

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Configuring Profile-Based Policies

This topic describes how to configure profile-based policies.



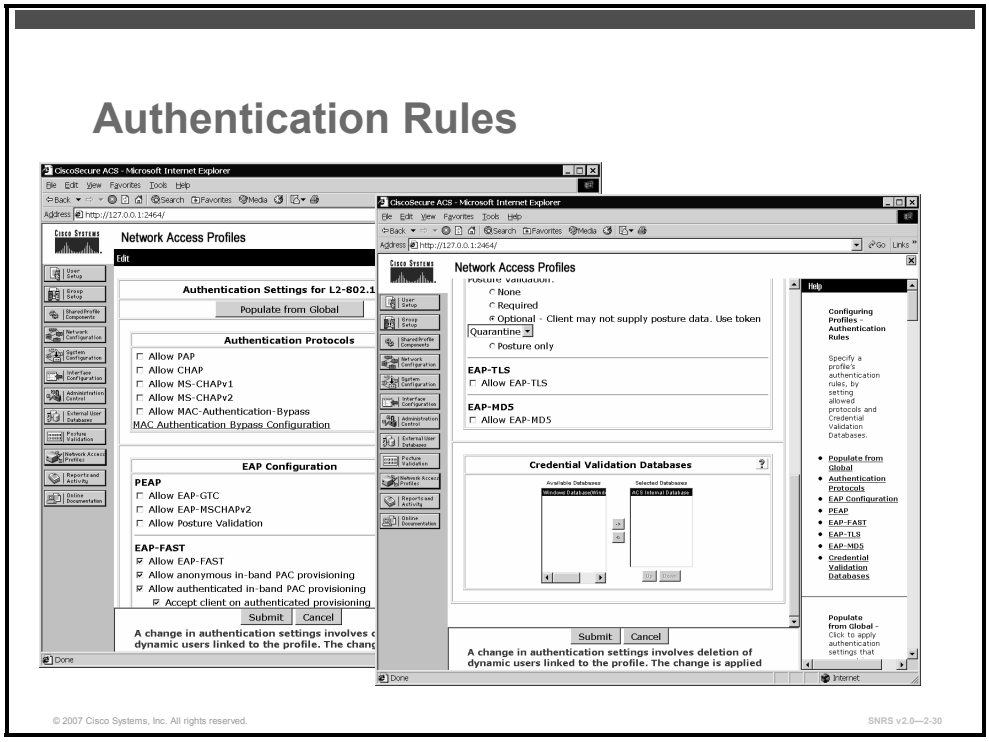
Once you create a profile, you must associate some rules or policies with this profile. It is important that the rules that comprise the profile-based policies reflect organizational security policies.

Before You Begin

To set up authorization rules for a profile, it is assumed that some other elements of Cisco Secure ACS have been set up, including the following:

- Users
- User groups
- Shared profile components
  - RADIUS authorization components (RACs)
  - Downloadable ACLs
  - NAFs
  - NARs
- Global authentication
  - Allow required protocols

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.



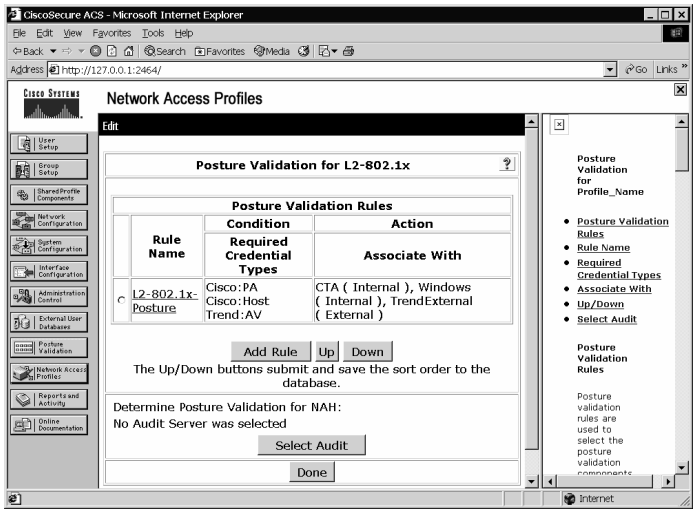
Authentication Rules

From here, you can do the following:

- Specify authentication protocols
  - Password Authentication Protocol (PAP)
  - Challenge Handshake Authentication Protocol (CHAP)
  - Microsoft CHAP version 1 (MS-CHAPv1)
  - Microsoft CHAP version 2 (MS-CHAPv2)
  - MAC authentication bypass
- Configure EAP types
  - PEAP
  - EAP-FAST (using Protected Access Credentials [PACs])
  - EAP-TLS (TLS/certificates)
  - EAP-MD5 (MD5/CHAP over EAP)
  - EAP-generic token card (EAP-GTC) (one-time password [OTP] tokens)
- For NAC, you can do the following:
  - Allow posture validation
  - EAP- Type, Length, Value (TLV) (posture credentials, attribute-value pairs, posture notifications)
  - Status query
  - EAP over UDP (EAPoUDP) (for Layer 2 and Layer 3 transport)
- Specify internal and external user databases

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Posture Validation Rules



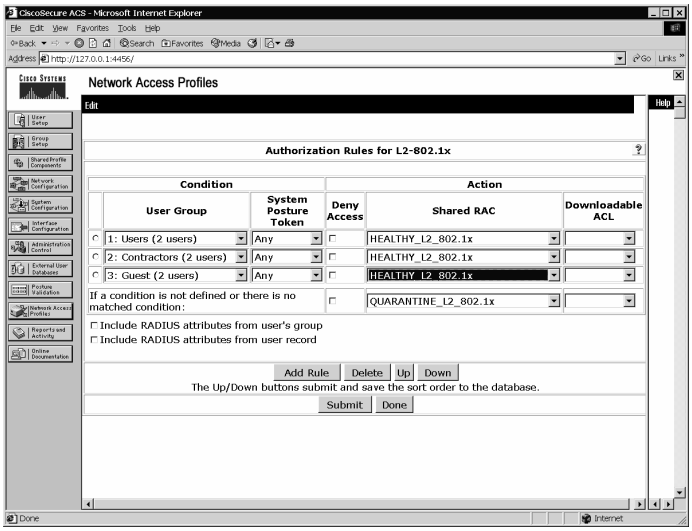
© 2007 Cisco Systems, Inc. All rights reserved.

Posture Validation Rules

Posture validation rules are used with NAC and is beyond the scope of this course. The NAC framework course is called *Implementing Cisco Network Admission Control* v3.0.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Authorization Rules



© 2007 Cisco Systems, Inc. All rights reserved.

Authorization Rules

Authorization rules are used for authorizing authenticated users. These rules can be based on group membership, posture validation (NAC), or both. Authorization actions contain downloadable ACLs and RADIUS authorization components.

Authorization rules consist of the following:

- Conditions
- Actions
- Action for nonmatched condition

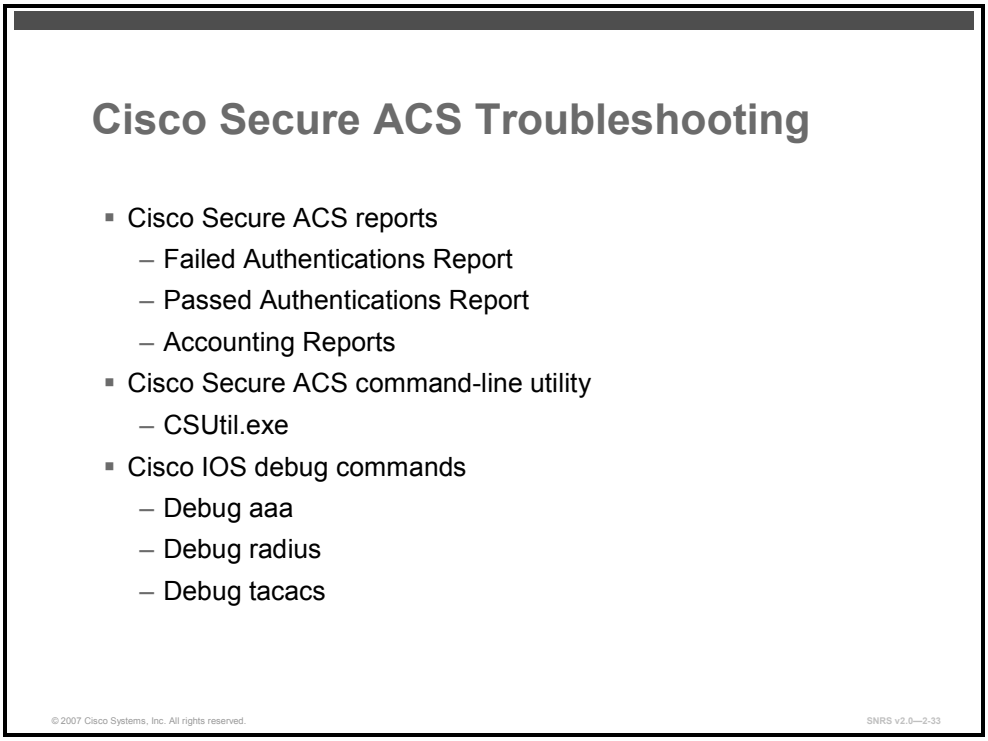
Authorization rules include the following:

- RADIUS attributes from user record
- RADIUS attributes from user group

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Troubleshooting Cisco Secure ACS

This topic describes three techniques for troubleshooting network and device access in a Cisco Secure ACS environment.



Basically, there are three tools to help with troubleshooting a Cisco Secure ACS environment. These tools can be used to help determine where the problem exists, including a third-party database back-end authentication problem.

Reports

Reports are the tool that you will use first when troubleshooting Cisco Secure ACS. Specifically, you will use the following reports:

- **Failed Authentications Report:** Use this report as the starting point to determine the source of the problem—endpoint, the AAA client, the Cisco Secure ACS server, or the database back end.
- **Passed Authentications Report:** Use this report to verify that the appropriate authorization mappings were made.
- **Accounting Reports:** Use these reports also to verify that the appropriate authorization mappings were made.

Cisco Secure ACS Command-Line Utility

Cisco Secure ACS has a few command-line utilities that can be useful when debugging access issues. One of these utilities is CSUtil.exe.

CSUtil.exe can be used with the "-e" option to decode error numbers found in Cisco Secure ACS service logs. This is useful in determining the cause of the error.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Cisco IOS Debug

Cisco IOS Software includes several debugging commands that can be used to provide detailed information about the processing of AAA requests by the AAA client.

For general information about AAA processing, including which protocol is being used, use one of these commands:

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

For details about TACACS or RADIUS in particular, use one of these commands:

- **debug radius**
- **debug tacacs**

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Summary

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

Summary

- Cisco Secure ACS for Windows provides a centralized identity networking solution and simplified user management.
- AAA provides the framework through which you set up network access control.
- Authentication is used to identify users before they gain access to the network and network services.
- Authorization provides the method for remote access control.
- Accounting provides the ability to collect and send security server information to be used for billing, auditing, and reporting purposes.

© 2007 Cisco Systems, Inc. All rights reserved.

Summary (Cont.)

- TACACS+ is a security application that provides centralized validation of users.
- RADIUS is an access server AAA protocol that secures remote access to networks and network services.
- Several commands are required to enable TACACS+ and RADIUS on a Cisco router or switch.
- Cisco Secure ACS may serve as AAA server for network access devices.
- Several Microsoft Windows services provide the core of Cisco Secure ACS functionality.

© 2007 Cisco Systems, Inc. All rights reserved.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Summary (Cont.)

- Nearly all Cisco Secure ACS administration tasks can be performed through the Cisco Secure ACS web interface.
- Check hardware and software requirements for your system and then use the setup.exe to install Cisco Secure ACS.
- An administrator must be added if you want to manage the Cisco Secure ACS server over the network.
- The navigation bar has several buttons to administer Cisco Secure ACS.
- Cisco Secure ACS 4.0 for Windows introduces NAPs.

© 2007 Cisco Systems, Inc. All rights reserved.

Summary (Cont.)

- There are several steps required to configure NAPs.
- Three conditions determine how Cisco Secure ACS classifies an access request and maps it to a profile.
- Once you create a profile, you must associate some rules or policies with the profile.
- There are tools to help with troubleshooting a Cisco Secure ACS environment.

© 2007 Cisco Systems, Inc. All rights reserved.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

References

For additional information, refer to these resources:

- *ACS 4.0 Primer:* http://www.cisco.com/application/pdf/en/us/guest/products/ps2086/c1161/cdccont_0900aecd8040daa7.pdf.
- *Installation Guide for Cisco Secure ACS for Windows 4.0:* http://www.cisco.com/en/US/partner/products/sw/secursw/ps2086/products_installation_guide_chapter09186a008060f73e.html.
- *User Guide for Cisco Secure ACS for Windows 4.0:* http://www.cisco.com/en/US/partner/products/sw/secursw/ps2086/products_user_guide_book09186a0080533dd8.html.
- *Cisco IOS Security Configuration Guide, Release 12.4:* http://www.cisco.com/en/US/partner/products/ps6350/products_configuration_guide_book09186a008043360a.html.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Lesson 2

Implementing Cisco IBNS

Overview

Cisco Identity-Based Networking Services (IBNS) is an integrated solution combining several Cisco products that offer authentication, access control, and user policies to secure network connectivity and resources. The Cisco IBNS solution enables greater security while simultaneously offering cost-effective management of changes throughout the organization. This lesson introduces you to Cisco IBNS. You will be introduced to 802.1x, RADIUS, and Extensible Authentication Protocol (EAP) as they relate to Cisco IBNS. You will then configure the Cisco Secure Access Control Server (ACS) to authenticate using EAP and RADIUS.

Objectives

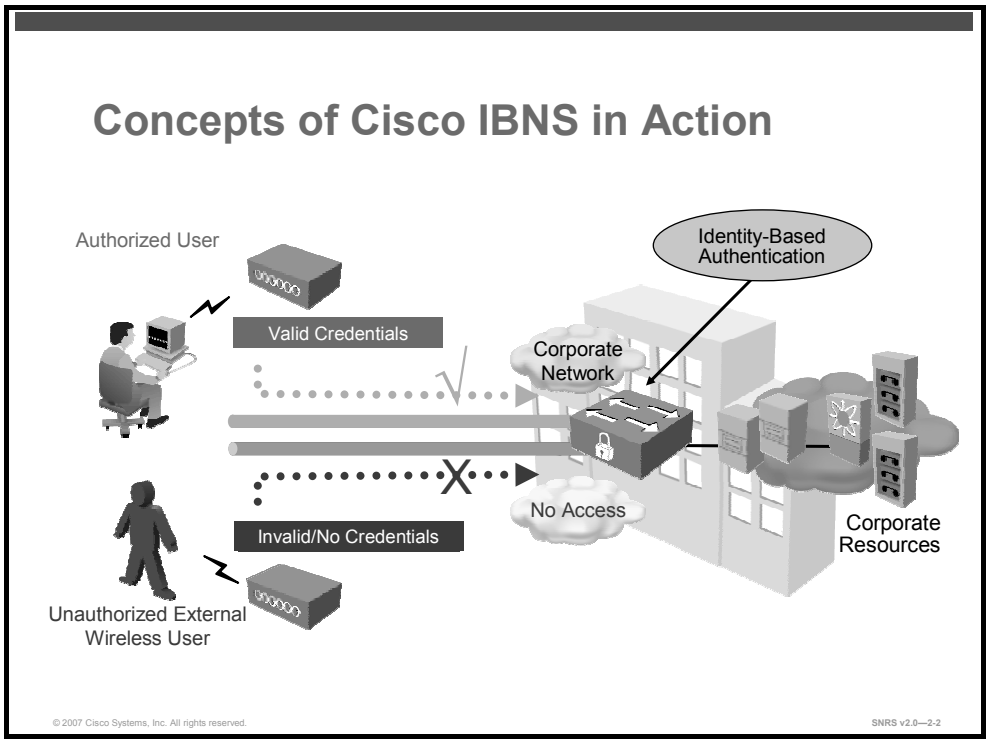
Upon completing this lesson, you will be able to describe the Cisco IBNS model and how to configure 802.1x on a Cisco Catalyst switch. This ability includes being able to meet these objectives:

- Explain how Cisco IBNS improves the security of LANs
- Describe how port-based access control
- Describe the IEEE 802.1x standard and 802.1X components
- Describe the various EAP types used in 802.1x
- Describe the use of 802.1x with port security
- Describe the use of 802.1x with VLAN assignment
- Describe the use of 802.1x with guest VLANs
- Describe the use of 802.1x with restricted VLANs
- Describe the commands used to configure and verify 802.1x operation on a Cisco Catalyst switch

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Cisco IBNS Overview

This topic describes the Cisco IBNS.



Cisco IBNS is an integrated solution combining several Cisco products that offer authentication, access control, and user policies to secure network connectivity and resources. The Cisco IBNS solution enables greater security while simultaneously offering cost-effective management of changes throughout the organization.

Cisco IBNS provides the network with the following services and capabilities:

- User or device authentication, or both
- Mapping the identity of a network entity to a defined set of policies configured by management
- Granting or denying network access, at the port level, based on configured authorization policies
- Enforcement of additional policies, such as resource access, when access is granted

These capabilities are introduced when a Cisco end-to-end system is implemented with the Cisco Catalyst Family of switches, wireless LAN access points and controllers, and Cisco ACS. Additional components of the system include an IEEE 802.1x compliant client OS, such as Microsoft Windows XP, and an optional X.509 public key infrastructure (PKI) certificate architecture. Cisco IP phones also interoperate with an identity-based networking system based on IEEE 802.1x when deployed on a Cisco end-to-end infrastructure.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

With 8021x, you can set up two different DHCP pools, you can assign addresses in one address range (for example the 10.0.0.0 network) to devices that can authenticate properly, and for a device that doesn't offer the right credentials, you can assign them an address in another address range (for example the 192.168.0.0 network) dynamically.

When a teleworker starts up or connects the PC on the home LAN, the PC usually first requests its network identity (IP address) and other needed information from a DHCP server; for PCs enabled for 802.1X, the first request is an Extensible Authentication Protocol over LAN (EAPOL) request.

When the access device (such as a Cisco switch) sees this request, it challenges the PC, which responds with the appropriate credentials (userid and password for example). The switch checks with the AAA server across the VPN to authenticate the user's credentials via RADIUS; if the teleworker logs in successfully, the PC is provided a network address and other information via DHCP on a subnet which allow access to the enterprise via the switch.

If a PC is not 802.1X capable, or the user does not log in successfully, the PC will be provided a network address on a subnet that only allows Internet access.

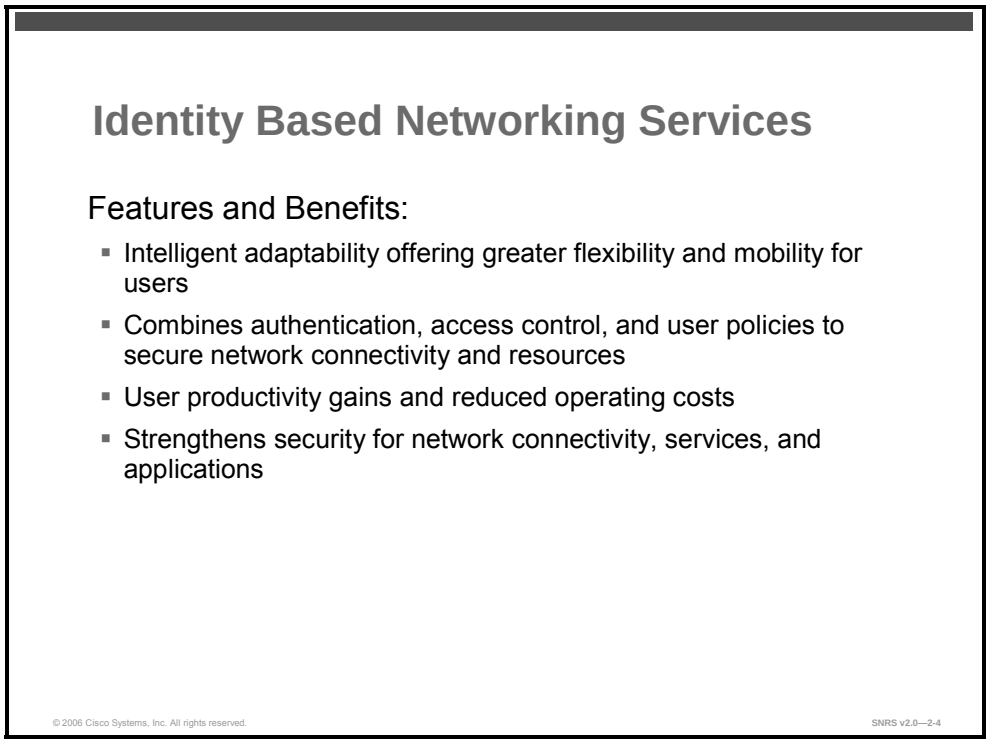
The following are access control using 802.1X authentication feature advantages:

- User is prompted upon PC start up or plug-in to the LAN; web access to a protected site to initiate challenge is not required (as in Authentication Proxy)
- The IP phone can be automatically allowed through the VPN; CDP is used for IP phone discovery
- A separate address range for spouse-and-child PCs allows for standardized addressing and access control, and a smaller enterprise addressable subnet for each teleworker home
- The teleworker can still communicate with non-enterprise PCs, print servers, and the like, if permitted—allowing for sharing between all home workstations
- Multiple authentication types are supported, including two-way authentication and the use of certificates, as permitted in the 802.1X standard; EAP-MD5, PEAP, and EAP-TLS are among the supported authentication methods
- PCs with static IP addresses in the enterprise addressable subnet cannot access the VPN until 802.1X authentication occurs; this reduces rogue access
- Any wireless PC (teleworker, spouse, child, or rogue) by default cannot gain enterprise access; this reduces rogue access

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Features and Benefits

The figure lists some features and benefits of Cisco IBNS.



The Cisco IBNS solution provides the following benefits:

- **Intelligent adaptability for offering greater flexibility and mobility to stratified users:** Creating user or group profiles with policies that define trust relationships between users and network resources allows organizations to easily authenticate, authorize, and account for all users of wired or wireless networks. This architecture-secure flexibility is a primary enabler for the networked virtual organization (NVO).
- **Combination of authentication, access control, and user policies to secure network connectivity and resources**: Because policies are associated with users and not physical ports, users obtain more mobility and freedom, and IT administration is simplified. Greater scalability and ease of management is achieved through policy enforcement and dynamic provisioning.
- **User productivity gains and reduced operating costs:** Providing security and greater flexibility for wired or wireless network access provides enterprises with the capability to have cross-functional or new project teams form more quickly, enables secure access for trusted partners and vendors, and facilitates secure conference room connectivity. Enabling flexibility with secure network access through centralized policy-based administration decreases the time, complexity, and effort associated with port security techniques at the MAC level.
- **Strengthens security for network connectivity, services, and applications**

Cisco IBNS and 802.1x are supported on all Cisco Catalyst switches, including Cisco Catalyst 6500, 4500, 3550, and 2950 Series Switches, Cisco Secure ACS and Cisco Aironet Access Points.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Cisco IBNS is a solution for increasing the security of physical and logical access to an enterprise network that is built on the IEEE 802.1x standard.

Cisco IBNS allows the network administrator to implement true identity-based network access control and policy enforcement at the user and port levels. It provides user and device identification using secure and reliable strong authentication technologies. This solution associates identified entities with policies. The policies are created and administered by management and provide increased granularity of control.

Cisco IBNS is a standards-based implementation of port security that is centrally managed by a RADIUS server (Cisco Secure ACS). Additionally, Cisco IBNS offers greater flexibility and mobility to users by combining access control and user profiles to secure network connectivity, services, and applications. This allows enterprises to increase user productivity and reduce operating costs.

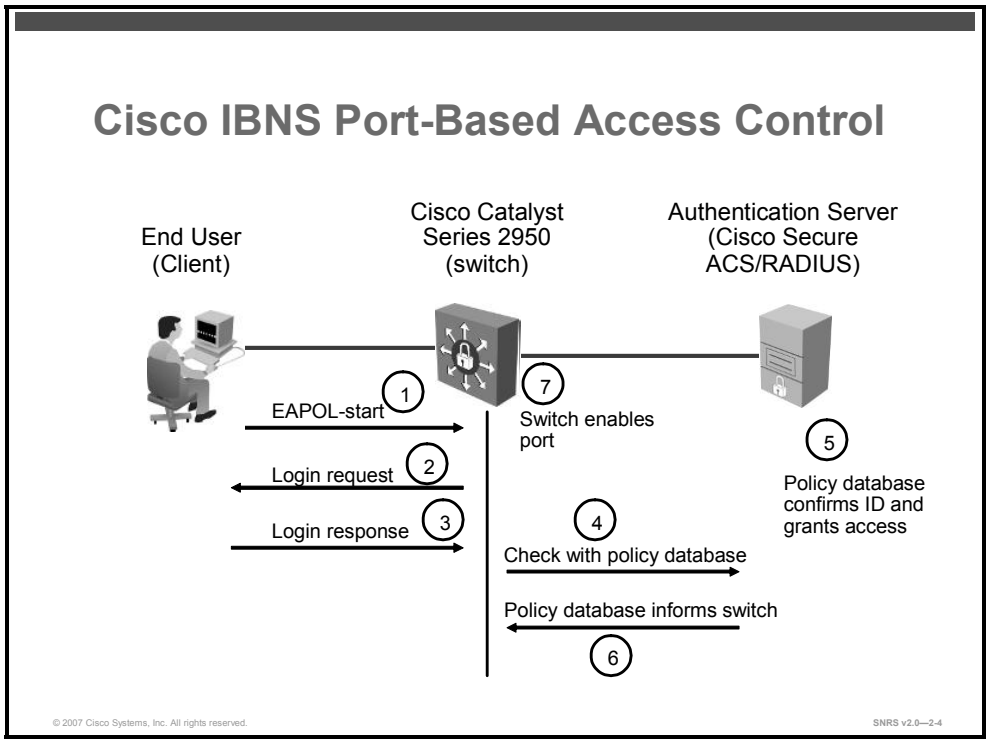
Cisco Catalyst switches support Microsoft Windows XP, Linux, and HP UNIX, with additional 802.1x client support anticipated in the future. Cisco Aironet products support all current versions of Microsoft Windows, Microsoft Windows CE, MAC OS, Linux, and MS-DOS.

The Cisco IBNS solution is based on standard RADIUS and 802.1x implementations. It interoperates with all Internet Engineering Task Force (IETF) authentication servers that comply with these two standards. Cisco has particularly enhanced its Cisco Secure ACS to provide a tight integration across all Cisco switches.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Port-Based Access Control

This topic describes port-based access control.



In compliance with the IEEE 802.1x standard, Cisco Catalyst switches can perform basic port-based authentication and Network Access Control (NAC). Once the IEEE 802.1x-compliant client software is configured on the end device (client), the Cisco Catalyst switches running IEEE 802.1x features authenticate the requesting user or system in conjunction with a back-end Cisco Secure ACS or other RADIUS server.

The diagram illustrates how IEEE 802.1X port-based access control works within Cisco IBNS.

The process follows these steps:

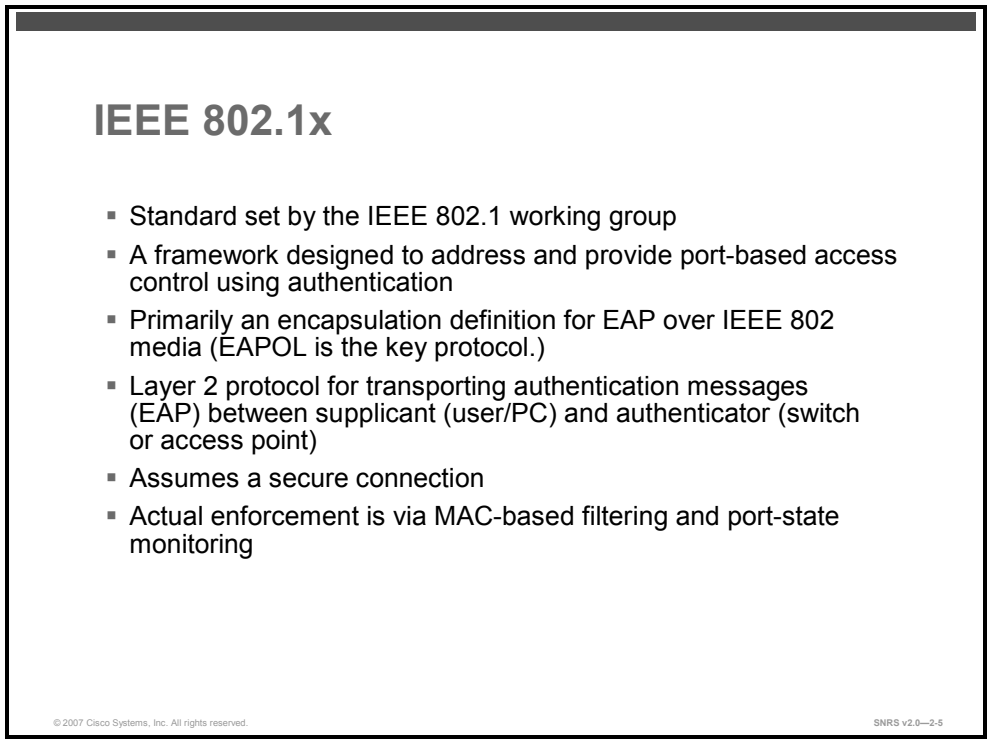
1. A client connects to an IEEE 802.1x-enabled network and sends a start message to the LAN switch.
2. The LAN switch sends a login request to the client.
3. The client replies with a login response.
4. The switch forwards the response to the policy database.
5. The policy database authenticates the user.
6. The policy database authorizes network access for the user and informs the LAN switch.
7. The LAN switch then enables the port connected to the client.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

IEEE 802.1x

This topic describes the IEEE 802.1x standard and 802.1X components..



802.1x is a protocol standard defined by the IEEE, designed to provide port-based network access. IEEE 802.1x authenticates network clients using information unique to the client and with credentials known only to the client. This service is called port-level authentication because, for security reasons, it is offered to a single endpoint for a given physical port.

The IEEE 802.1x standard defines a client-server-based access control and authentication protocol that prevents unauthorized clients from connecting to a LAN through publicly accessible ports unless they are properly authenticated. The authentication server authenticates each client connected to a switch port before making available any services offered by the switch or the LAN.

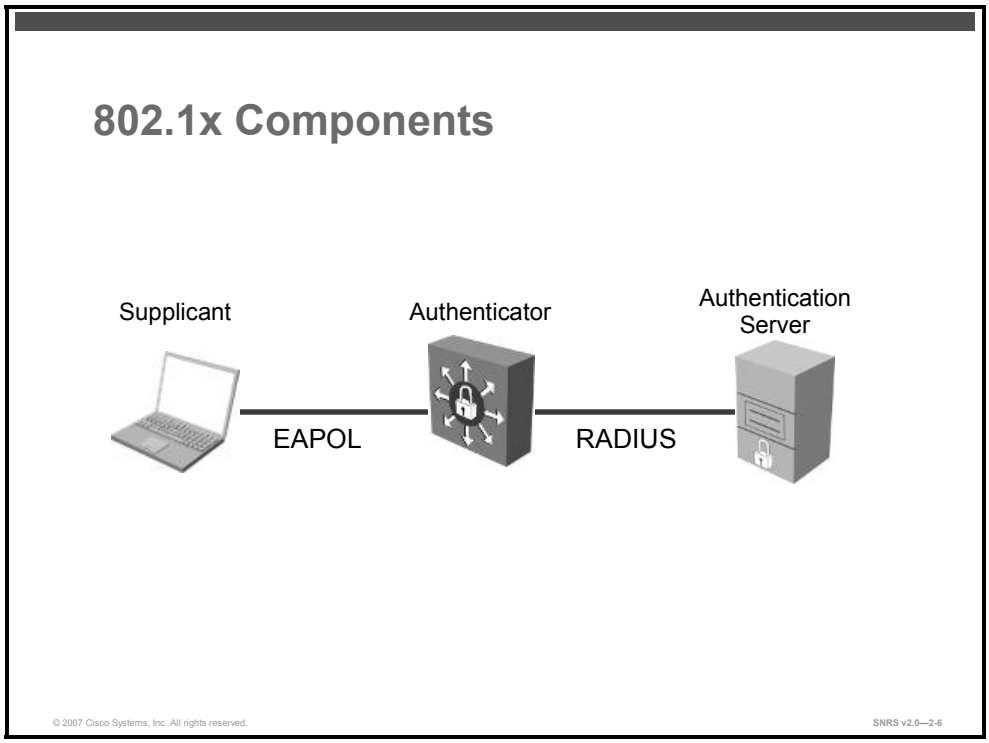
Until the client is authenticated, IEEE 802.1x access control allows only Extensible Authentication Protocol over LAN (EAPOL), Cisco Discovery Protocol (CDP), and Spanning Tree Protocol (STP) traffic through the port to which the client is connected. After authentication is successful, normal traffic can pass through the port.

IEEE 802.1x provides an encapsulation definition for the transport of EAP at the MAC layer over any PPP or IEEE 802 media. IEEE 802.1x enables the implementation of port-based NAC to a network device. IEEE 802.1x transports EAP messages between a supplicant and an authenticator. The authenticator then typically relays the EAP information to an authentication server via the RADIUS protocol. IEEE 802.1x not only provides the capability to permit or deny network connectivity based on user or machine identity, but also works in conjunction with higher layer protocols to enforce network policy.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

802.1x Components

This section describes the 802.1x components.



With 802.1x port-based authentication, the devices in the network have specific roles as shown in the diagram.

Supplicant

The supplicant is a device (workstation, laptop, and so on) that requests access to the LAN and switch services and responds to requests from the authenticator (switch). The device must be running IEEE 802.1x-compliant client software, such as that offered in the Microsoft Windows XP OS. The client is the supplicant in the IEEE 802.1x specification.

Authenticator

The authenticator is a device (such as a Cisco Catalyst switch) that controls physical access to the network based on the authentication status of the client. The authenticator usually acts as an intermediary (proxy) between the client and the authentication server. The authenticator requests identity information from the client via EAP, verifies that information with the authentication server via RADIUS, and then relays a response to the client based on the response from the authentication server.

When the switch receives EAP over LAN (EAPOL) frames and relays them to the authentication server, the Ethernet header and EAP frame are re-encapsulated into the RADIUS format. The EAP frames are not modified or examined during encapsulation, and the authentication server must support EAP within the native frame format. When the switch receives frames from the authentication server, the RADIUS header is removed, leaving the EAP frame, which is then encapsulated in the IEEE 802.1x format and sent to the client.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Authentication Server

The authentication server performs the actual authentication of the client. The authentication server validates the identity of the client and notifies the switch whether the client is authorized to access the LAN and switch services. Because the switch acts as the proxy, the authentication server is transparent to the client. The RADIUS security system with EAP extensions is the only supported authentication server. RADIUS uses a client-server model in which secure authentication information is exchanged between the RADIUS server and one or more RADIUS clients.

Terminology

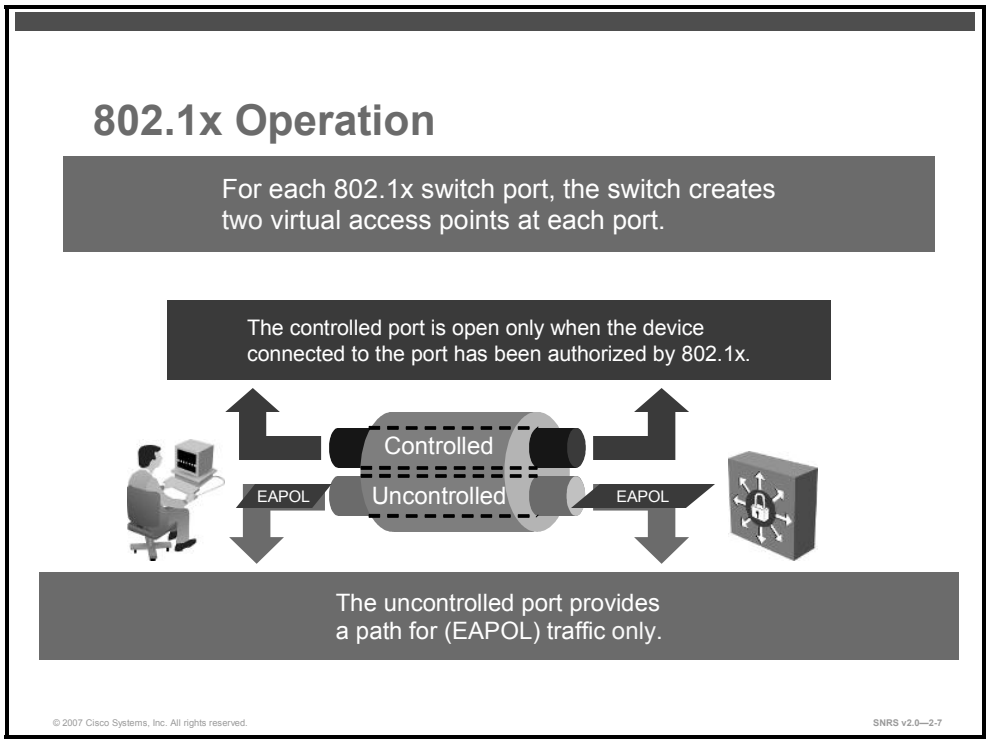
IEEE Term	Also Know As
Supplicant	Client
Authenticator	Network access device (NAD)
Authentication server	Authentication, authorization, and accounting (AAA) and RADIUS server

The table lists some familiar terms and their IEEE counterparts.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

How 802.1x Works

This section describes 802.1x operation.



The default security mode for 802.1x is to create two virtual access points at each port.

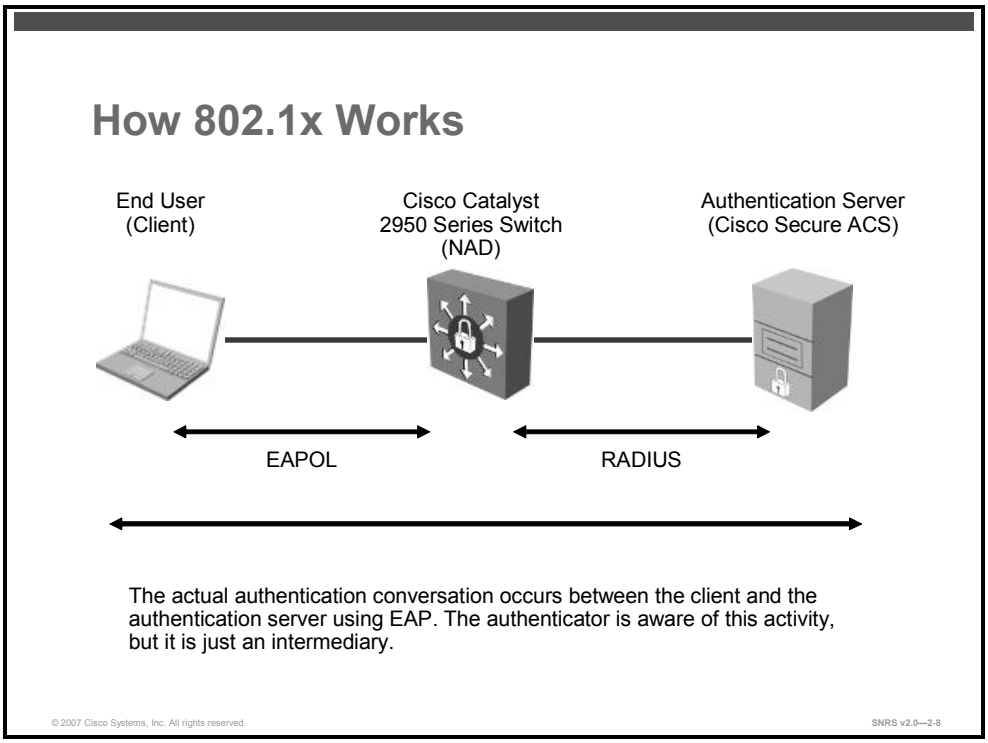
- **Controlled port:** Opened only after authorization by 802.1x
- **Uncontrolled port:** Path for EAPOL traffic only

Only EAPOL traffic is allowed to flow on the uncontrolled port until the device has been authorized.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Authentication Initiation and Message Exchange

This topic describes the authentication initiation, message exchanges and port states.



The switch or the client can initiate authentication. If you enable authentication on a port by using the **dot1x port-control auto** interface configuration command, the switch must initiate authentication when it determines that the port link-state transitions from down to up. It then sends an EAP Request/Identity frame to the client to request its identity (typically, the switch sends an initial Request/Identity frame followed by one or more requests for authentication information). Upon receipt of the frame, the client responds with an EAP Response/Identity frame.

However, if during bootup the client does not receive an EAP Request/Identity frame from the switch, the client can initiate authentication by sending an EAPOL-start frame, which prompts the switch to request the identity of the client.

Note	If 802.1x is not enabled or supported on the NAD, any EAPOL frames from the client are dropped. If the client does not receive an EAP Request/Identity frame after three attempts to start authentication, the client transmits frames as if the port is in the authorized state. A port in the authorized state effectively means that the client has been successfully authenticated.

When the client supplies its identity, the switch begins its role as the intermediary, passing EAP frames between the client and the authentication server until authentication succeeds or fails. If the authentication succeeds, the switch port becomes authorized.

The specific exchange of EAP frames depends on the authentication method being used. The diagram shows a message exchange initiated by the client using the one-time password (OTP) authentication method with a RADIUS server.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Ports in Authorized and Unauthorized States

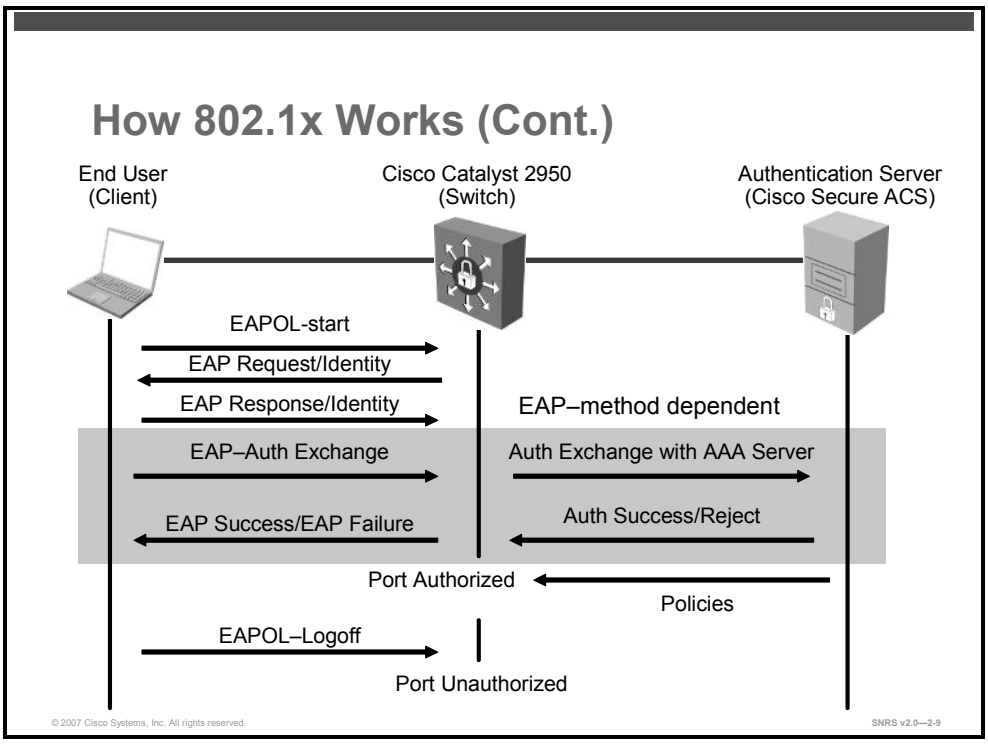
The switch port state determines whether the client is granted access to the network. The port starts in the unauthorized state. While in this state, the port disallows all ingress and egress traffic except for 802.1x packets. When a client is successfully authenticated, the port transitions to the authorized state, allowing all traffic for the client to flow normally.

If a client that does not support 802.1x is connected to an unauthorized 802.1x port, the switch requests the identity of the client. In this situation, the client does not respond to the request, the port remains in the unauthorized state, and the client is not granted access to the network.

In contrast, when an 802.1x-enabled client connects to a port that is not running 802.1x, the client initiates the authentication process by sending the EAPOL-start frame. When no response is received, the client sends the request for a fixed number of times. Because no response is received, the client begins sending frames as if the port is in the authorized state.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.



You control the port authorization state by using the **dot1x port-control** interface configuration command and these keywords:

- **force-authorized:** This keyword disables 802.1x and causes the port to transition to the authorized state without any authentication exchange required. The port transmits and receives normal traffic without 802.1x-based authentication of the client. This is the default setting.
- **force-unauthorized:** This keyword causes the port to remain in the unauthorized state, ignoring all attempts by the client to authenticate. The switch cannot provide authentication services to the client through the interface.
- **auto:** This keyword enables 802.1x authentication and causes the port to begin in the unauthorized state, allowing only EAPOL frames to be sent and received through the port. The authentication process begins when the link state of the port transitions from down to up, or when an EAPOL-start frame is received. The switch requests the identity of the client and begins relaying authentication messages between the client and the authentication server. Each client attempting to access the network is uniquely identified by the switch by using the MAC address of the client.

If the client is successfully authenticated (receives an Accept frame from the authentication server), the port state changes to authorized, and all frames from the authenticated client are allowed through the port. If the authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the switch can retransmit the request. If no response is received from the server after the specified number of attempts, authentication fails and network access is not granted.

When a client logs off, it sends an EAPOL-logoff message, causing the switch port to transition to the unauthorized state.

If the link state of a port transitions from up to down, or if an EAPOL-logoff frame is received, the port returns to the unauthorized state.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

IEEE 802.1x Host Mode

IEEE 802.1x ports can be configured for single-host or multiple-host mode.

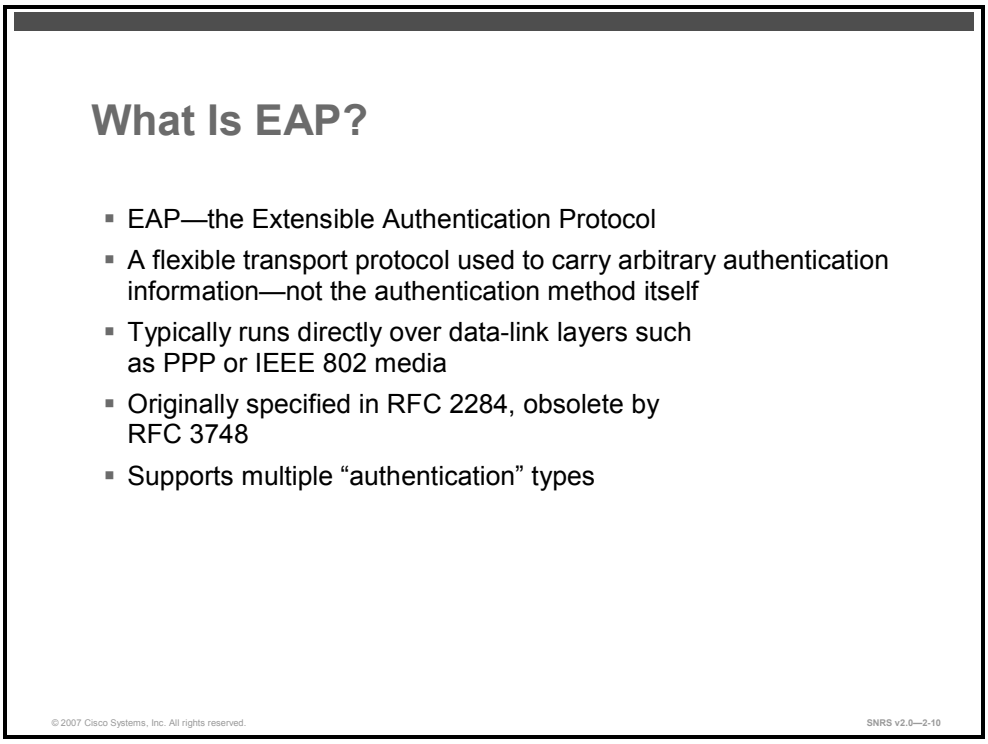
- Single-host mode
  - Only one client can be connected to the IEEE 802.1x-enabled switch port.
  - The switch detects the client by sending an EAPOL frame when the port link state changes to the up state. If a client leaves or is replaced with another client, the switch changes the port link state to down, and the port returns to the unauthorized state.
- Multiple-host mode
  - Multiple hosts may be attached to a single IEEE 802.1x-enabled port.
  - Only one of the attached clients must be authorized for all clients to be granted network access.
  - If the port becomes unauthorized (reauthentication fails or an EAPOL-logoff message is received), the switch denies network access to all of the attached clients.
  - You can use IEEE 802.1x authentication to authenticate the port and port security to manage network access for all MAC addresses, including that of the client.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Selecting the Correct EAP

This topic describes the various EAP types used in 802.1x.



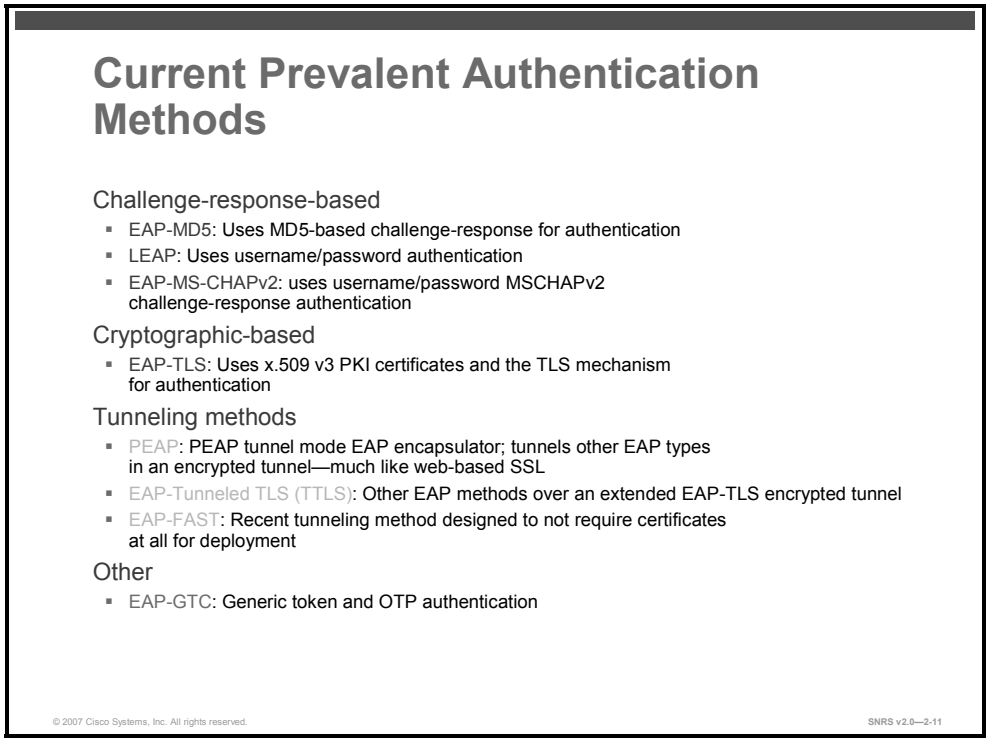
EAP, based on IETF 802.1x, is an end-to-end framework that allows the creation of authentication types without changing AAA client configurations.

EAP has the following characteristics:

- An extension of PPP to provide additional authentication features
- A flexible protocol used to carry arbitrary authentication information
- Typically rides on top of another protocol, such as 802.1x or RADIUS
- Specified in RFC 2284, but made obsolete by RFC 3748
- Support multiple authentication types, such as the following:
  - EAP- Message Digest 5 (MD5): Plain password hash (Challenge Handshake Authentication Protocol [CHAP] over EAP)
  - EAP-Transport Layer Security (TLS) (based on X.509 certificates)
  - Lightweight EAP (LEAP) (also called EAP-Cisco Wireless)
  - Protected EAP (PEAP)
  - EAP-Flexible Authentication via Secure Tunneling (FAST)

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

This section describes some common authentication methods.



Cisco Secure ACS supports the following varieties of EAP and others:

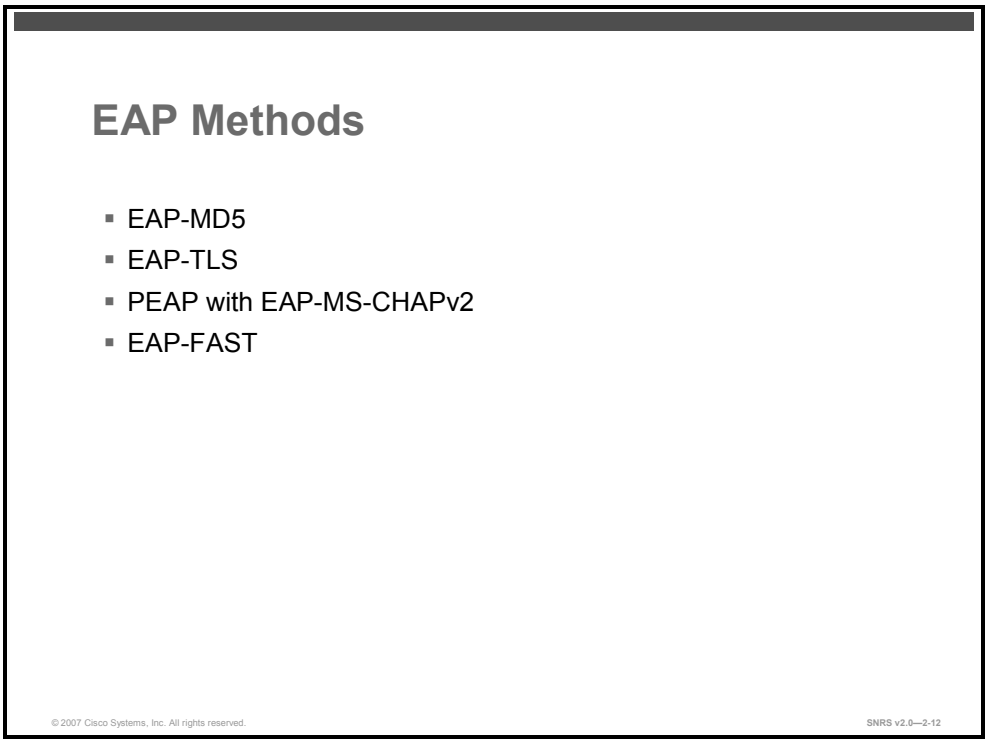
- **EAP-MD5:** Uses MD5-based challenge and response for authentication
- **EAP-TLS:** EAP incorporating TLS using X.509 PKI certificates
- **LEAP:** -An EAP protocol used by Cisco Aironet wireless equipment using username and password authentication
- **PEAP:** PEAP tunnel mode, which is implemented with EAP-generic token card (GTC) and EAP-Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2)
- **EAP-TTLS:** Uses other EAP methods over an extended EAP-TLS encrypted tunnel
- **EAP-FAST:** A faster means of encrypting EAP authentication, supports EAP-GTC authentication; designed to not require certificates

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

EAP Methods

This section describes various EAP methods.



IEEE 802.1x supports several different EAP methods for providing identity-based NAC. This section covers these four types of EAP methods:

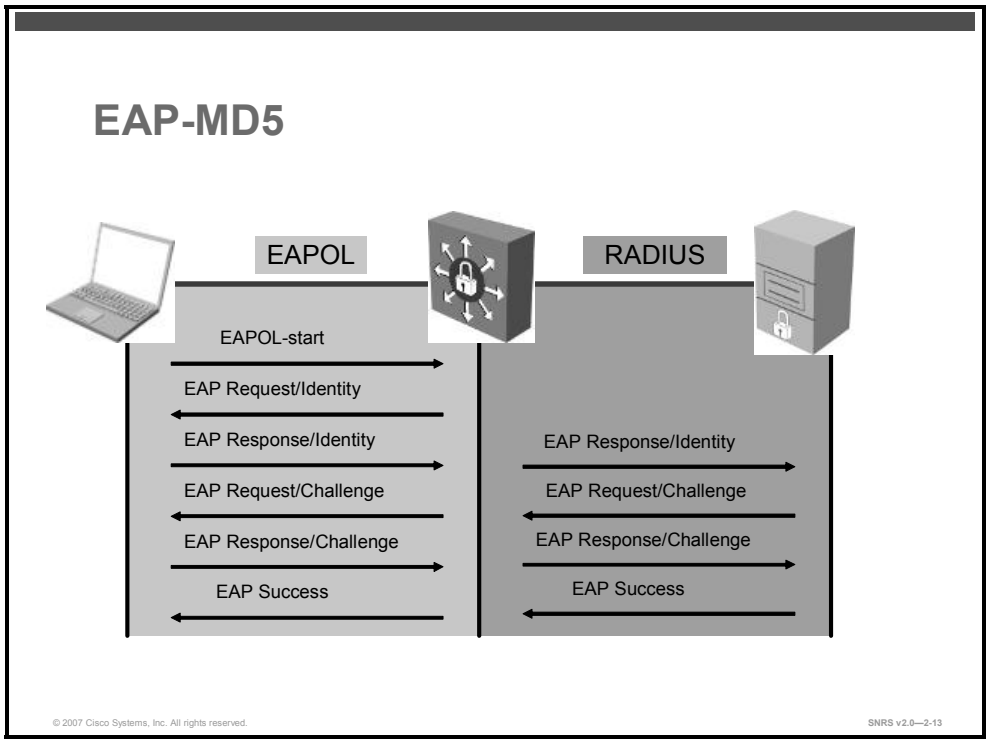
- EAP-MD5
- EAP-TLS
- PEAP with EAP-MS-CHAPv2
- EAP-FAST

The next sections examine each method further.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

EAP-MD5

This section describes EAP-MD5.



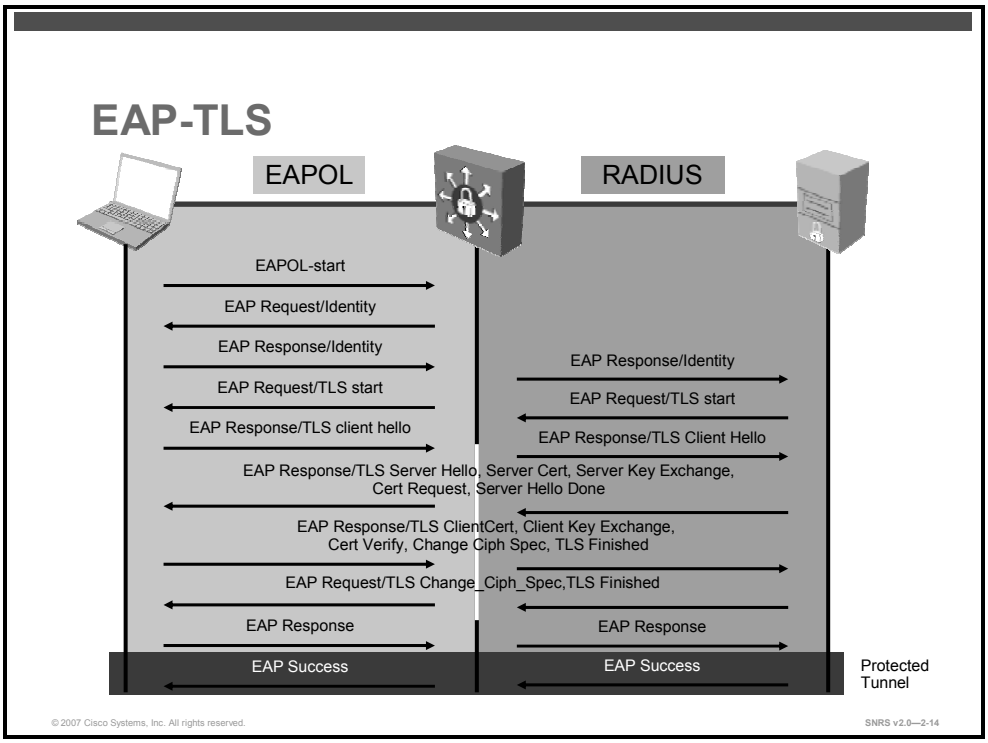
EAP-MD5 is a standard, nonproprietary EAP type. It is based on RFC 1994 (CHAP) and RFC 2284 (EAP). An MD5-Challenge within an EAP message is analogous to the PPP CHAP protocol, with MD5 specified as the hash algorithm. Because MD5 support is included in RFC 3748, all EAP deployments should support the MD5-Challenge mechanism.

The diagram illustrates the EAP-MD5 message exchange between the supplicant, authenticator, and authentication server. First, a client running the IEEE 802.1x supplicant connects to the network and sends an EAPOL-start message to the authenticator. The authenticator sends an EAP Request/Identity to the supplicant and the supplicant replies with an EAP Response/Identity. The authenticator forwards the response to the authentication server via RADIUS. The authentication server sends an EAP-MD5 Challenge to the supplicant and the supplicant replies with a response. The authentication server confirms the user identity and instructs the authenticator to authorize network access for the user. The authenticator then enables the port connected to the supplicant.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

EAP-TLS

This section describes EAP-TLS.



EAP-TLS was developed by Microsoft Corporation to enable the use of EAP as an extension of PPP to provide authentication within PPP and TLS to provide integrity-protected cipher suite negotiation and key exchange. EAP-TLS, which is defined in RFC 2716, uses X.509 PKI certificate-authenticated IEEE 802.1x port-based access control and is specifically targeted to address a number of weaknesses in other EAP protocols such as EAP-MD5. In addressing these weaknesses, however, the complexity of deployment increases because not only servers, but also clients, require certificates for mutual authentication.

Here are some of the benefits of EAP-TLS:

- The ability to provide per-packet confidentiality and integrity protection, which protects user identity
- A standardized mechanism for key exchange
- Built-in support for fragmentation and reassembly
- Support for acknowledged success and failure indications

Within IEEE 802.1x, the EAP-TLS exchange of messages provides mutual authentication, negotiation of the encryption method, and encrypted key determination between a supplicant and an authentication server.

The diagram illustrates the EAP-TLS message exchange between the supplicant, authenticator, and authentication server. First, a client running the IEEE 802.1x supplicant connects to the network and sends an EAPOL-start message to the authenticator. The authenticator sends an EAP Request/Identity to the supplicant and the supplicant replies with an EAP Response/Identity. The authenticator forwards the response to the authentication server via

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

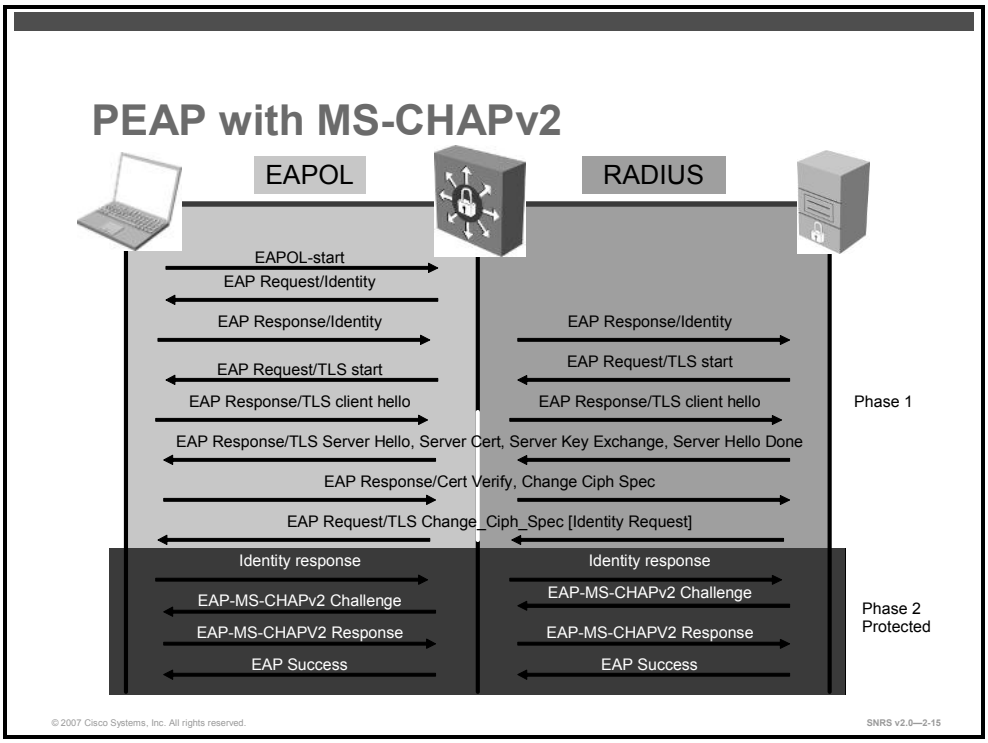
RADIUS. The authentication server sends an EAP-TLS start message to the supplicant and the supplicant replies with an EAP-TLS client hello packet. The authentication server sends its X.509 PKI certificate to the supplicant and requests that the supplicant send its certificate. The supplicant verifies the certificate with the authentication public key of the server and sends its certificate to the authentication server along with an updated cipher suite. The authentication server verifies the certificate of the supplicant, thus authenticating the identity of the user, and confirms the cipher suite. With the TLS tunnel now established, the authentication server instructs the authenticator to authorize network access for the user. The authenticator then enables the port connected to the supplicant.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

PEAP with MS-CHAPv2

This section describes PEAP with MS-CHAPv2.



This diagram illustrates PEAP with MS-CHAPv2 message exchange between the supplicant, authenticator, and authentication server. First, a client running the IEEE 802.1x supplicant connects to the network and sends an EAPOL-start message to the authenticator. The authenticator sends an EAP Request/Identity to the supplicant and the supplicant replies with an EAP Response/Identity. The authenticator forwards the response to the authentication server via RADIUS. The authentication server sends an EAP-TLS start message to the supplicant, and the supplicant replies with an EAP-TLS client hello packet. The authentication server sends its X.509 PKI certificate to the supplicant. The supplicant verifies the certificate with the public key of the authentication server and sends an updated cipher suite. The authentication server agrees to the cipher suite. With the TLS tunnel now established, the authentication server sends an EAP-MS-CHAPv2 challenge to the supplicant and the supplicant replies with a response. The authentication server confirms the identity of the user and instructs the authenticator to authorize network access for the user. The authenticator then enables the port connected to the supplicant.

PEAP was developed by Cisco Systems, Microsoft Corporation, and RSA Security. PEAP is an EAP type that addresses security issues by first creating a secure channel that is both encrypted and integrity-protected with TLS. Then, a new EAP negotiation with virtually any EAP type (for example, EAP-MS-CHAPv2) occurs, authenticating the network access attempt of the client. Because the TLS channel protects EAP negotiation and authentication for the network access attempt, password-based authentication protocols that are normally susceptible to an offline dictionary attack can be used for authentication. By wrapping the EAP messages within TLS, any EAP method running within PEAP is provided with built-in support for key exchange, session resumption, fragmentation, and reassembly. Furthermore, PEAP makes it possible to authenticate LAN clients without requiring them to have certificates, simplifying the architecture of secure wired LANs and wireless LANs (WLANs).

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

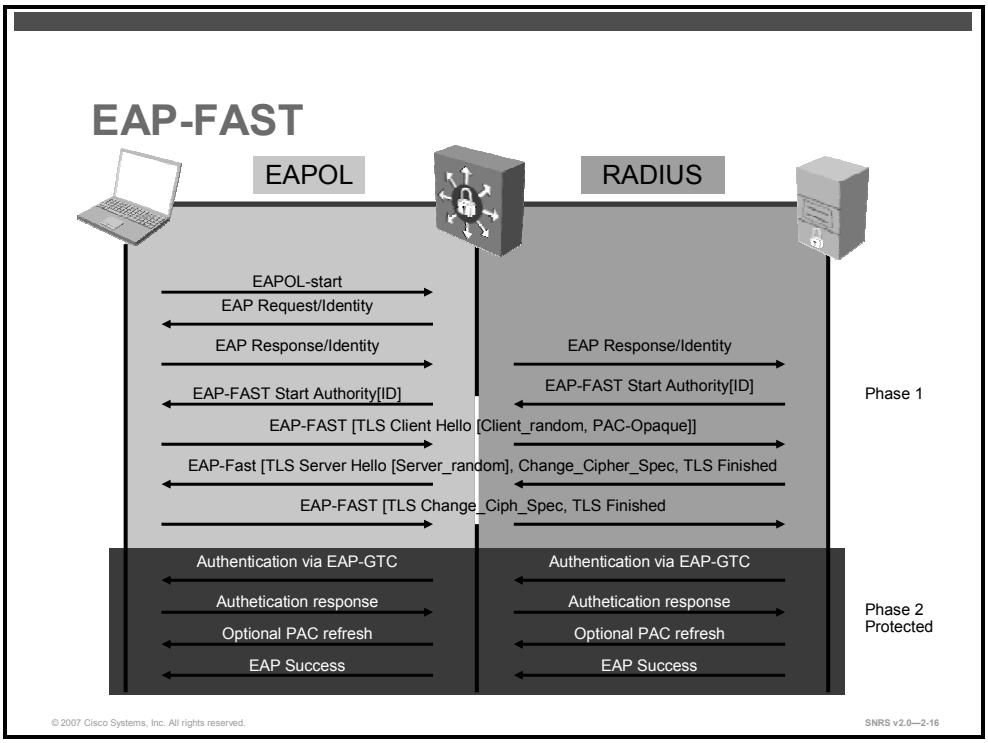
MS-CHAPv2 is a password-based, challenge-response, mutual authentication protocol that uses Message Digest 4 (MD4) and Data Encryption Standard (DES) to encrypt responses. The authenticator challenges a supplicant and the supplicant can challenge the authentication server. If either challenge is not correctly answered, the connection can be rejected. MS-CHAPv2 was originally designed by Microsoft as a PPP authentication protocol to provide better protection for dial-up and VPN connections, although it is now an EAP type as well. Although MS-CHAPv2 provides better protection than previous challenge-response authentication protocols, it is still susceptible to an offline dictionary attack. A malicious user can capture a successful MS-CHAPv2 exchange and guess passwords until the correct one is determined. Used in combination with PEAP, however, the MS-CHAPv2 exchange is protected with the strong security of the TLS channel.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

EAP-FAST

This section describes EAP-FAST.



This diagram illustrates the EAP-FAST message exchange between the supplicant, authenticator, and authentication server using EAP-GTC as the inner method. First, a client running the IEEE 802.1x supplicant connects to the network and sends an EAPOL-start message to the authenticator. The authenticator sends an EAP Request/Identity to the supplicant and the supplicant replies with an EAP Response/Identity. The authenticator forwards the response to the authentication server via RADIUS. The authentication server sends an EAP-FAST start message, which includes the authority ID, to the supplicant. Based on the authority ID sent by the authentication server, the supplicant selects a stored Protected Access Credential (PAC), which is a unique shared key used to mutually authenticate the supplicant and server. The supplicant then replies to the authentication server with a PAC opaque (based on the PAC key). The authentication server decrypts the PAC opaque using a master key to derive the PAC key. At this point, both the supplicant and server possess the same PAC key and create a TLS tunnel.

The authentication server sends an EAP-GTC request to the supplicant and the supplicant replies with a response. The authentication server confirms the user identity and instructs the authenticator to authorize network access for the user. The authenticator then enables the port connected to the supplicant.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

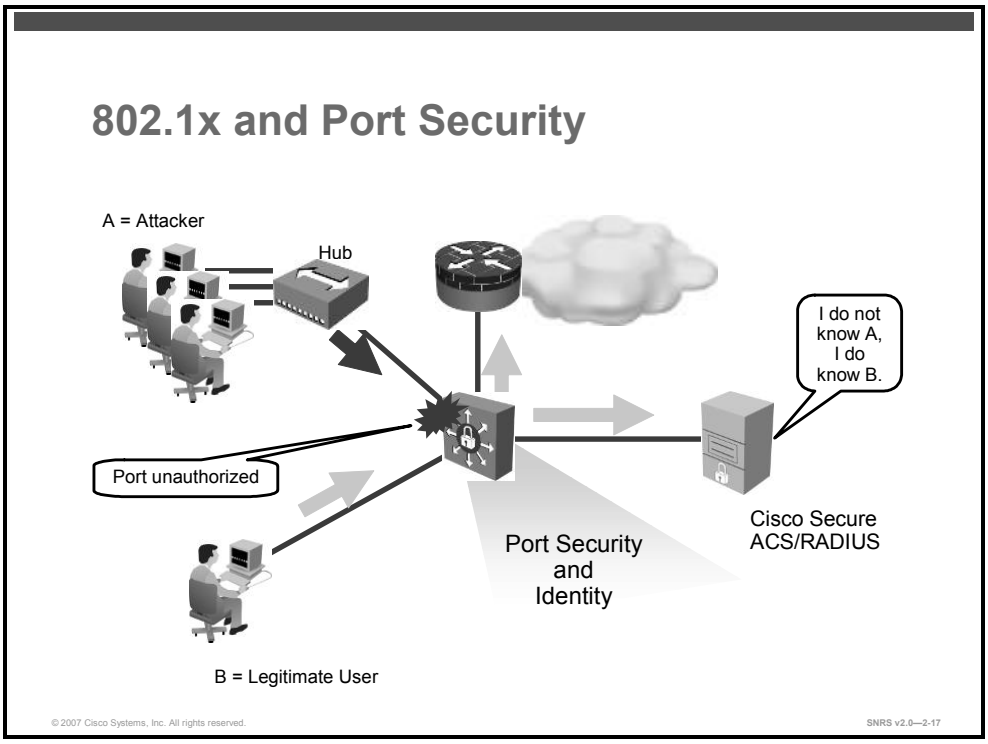
EAP-FAST was developed by Cisco Systems and submitted to the IETF as an Internet draft in February 2004. The Internet draft was revised and submitted in April 2005. The EAP-FAST protocol is a client-server security architecture that encrypts EAP transactions within a TLS tunnel. While similar to PEAP in this respect, EAP-FAST differs significantly in that the EAP-FAST tunnel establishment is based upon strong shared secret keys that are unique to users. These secrets are called PACs and may be distributed automatically (automatic or in-band provisioning) or manually (manual or out-of-band provisioning) to client devices. Because handshakes based upon shared secrets are intrinsically faster than handshakes based upon a PKI infrastructure, EAP-FAST is the significantly faster of the two solutions that provide encrypted EAP transactions.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

802.1x and Port Security

This topic describes the use of 802.1x with port security.



You can configure an IEEE 802.1x port with port security in either single-host or multiple-hosts mode. (You must also configure port security on the port by using the **switchport port-security** interface configuration command.) When you enable port security and IEEE 802.1x on a port, IEEE 802.1x authenticates the port, and port security manages network access for all MAC addresses, including that of the client. You can then limit the number or group of clients that can access the network through an IEEE 802.1x port.

Here are some examples of the interaction between IEEE 802.1x and port security on the switch:

- When a client is authenticated, and the port security table is not full, the client MAC address is added to the port security list of secure hosts. The port then proceeds to come up normally.

  When a client is authenticated and manually configured for port security, it is guaranteed an entry in the secure host table (unless port security static aging has been enabled).

  A security violation occurs if the client is authenticated, but the port security table is full. This can happen if the maximum number of secure hosts has been statically configured or if the client ages out of the secure host table. If the client address is aged out, its place in the secure host table can be taken by another host.
- When an IEEE 802.1x client logs off, the port transitions back to an unauthenticated state, and all dynamic entries in the secure host table are cleared, including the entry for the client. Normal authentication then takes place.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

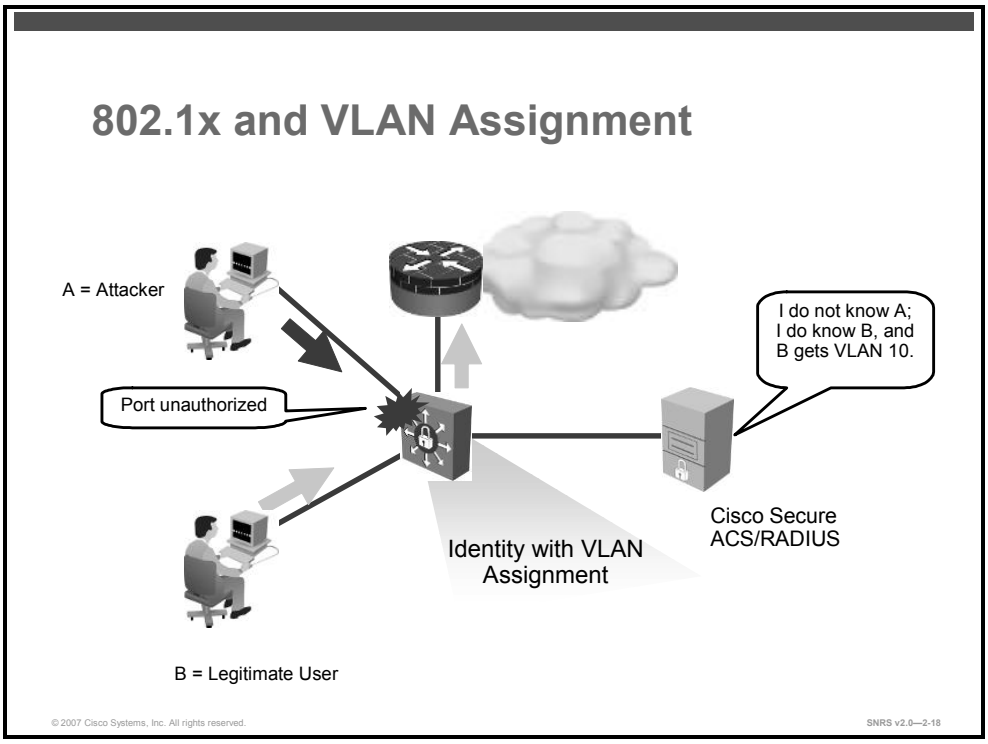
- You should reauthenticate the IEEE 802.1x client by using the **dot1x re-authenticate interface** *interface-id* command when you manually remove an IEEE 802.1x client address from the port security table by using the **no switchport port-security mac-address** *mac-address* command.
- If the port is administratively shut down, the port becomes unauthenticated, and all dynamic entries are removed from the secure host table.
- Port security and a voice VLAN can be configured simultaneously on an IEEE 802.1x port that is in either single-host or multiple-hosts mode. Port security applies to both the voice VLAN ID (VVID) and the port VLAN ID (PVID).

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

802.1x and VLAN Assignment

This topic describes the use of 802.1x with VLAN assignment.



A common security policy is to limit network access for certain users by using VLAN assignment.

You will accomplish this using the **aaa authorization network {default} group radius** command.

After successful IEEE 802.1x authentication of a port, the RADIUS server sends the VLAN assignment to configure the switch port. The RADIUS server database maintains the username-to-VLAN mappings, which assigns the VLAN based on the username of the client connected to the switch port.

When configured on the switch and the RADIUS server, IEEE 802.1x with VLAN assignment has these characteristics:

- If no VLAN is supplied by the RADIUS server or if IEEE 802.1x authorization is disabled, the port is configured in its access VLAN after successful authentication.
- If IEEE 802.1x authorization is enabled but the VLAN information from the RADIUS server is not valid, the port returns to the unauthorized state and remains in the configured access VLAN. This prevents ports from appearing unexpectedly in an inappropriate VLAN because of a configuration error.
- If IEEE 802.1x authorization is enabled and all information from the RADIUS server is valid, the port is placed in the specified VLAN after authentication.
- If the multiple-host mode is enabled on an IEEE 802.1x port, all hosts are placed in the same VLAN (specified by the RADIUS server) as the first authenticated host.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

- If IEEE 802.1x and port security are enabled on a port, the port is placed in the RADIUS server-assigned VLAN.
- If IEEE 802.1x is disabled on the port, it is returned to the configured access VLAN.

When the port is in the force authorized, force unauthorized, unauthorized, or shutdown state, it is placed in the configured access VLAN.

If an IEEE 802.1x port is authenticated and put in the RADIUS server-assigned VLAN, any change to the port access VLAN configuration does not take effect.

Note The IEEE 802.1x with VLAN assignment feature is not supported on trunk ports, dynamic ports, or with dynamic access port assignment through a VLAN Membership Policy Server (VMPS).

Configuring VLAN Assignment

Perform these tasks to configure VLAN assignment:

1. Enable AAA authorization.
2. Enable IEEE 802.1x
3. Assign vendor-specific tunnel attributes in the RADIUS (Cisco Secure ACS) server. The RADIUS server must return these attributes to the switch:
   - [64] Tunnel-Type = VLAN
   - [65] Tunnel-Medium-Type = IEEE 802
   - [81] Tunnel-Private-Group-ID = VLAN name or VLAN ID

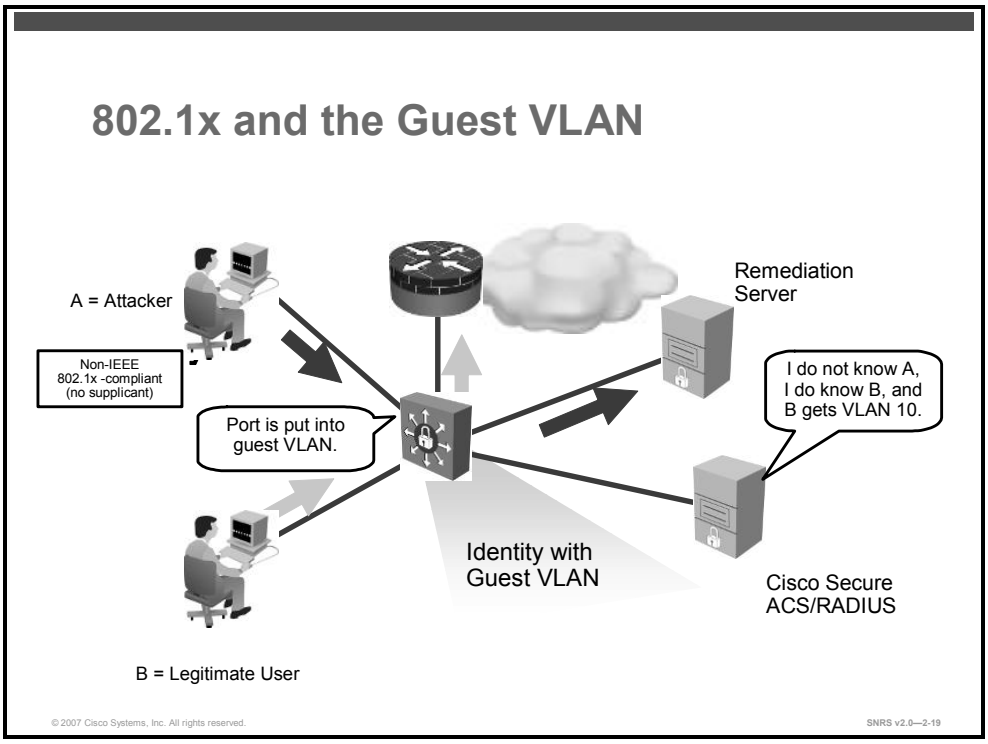
Note Attribute [64] must contain the value **VLAN** (type 13). Attribute [65] must contain the value **IEEE 802** (type 6). Attribute [81] specifies the **VLAN name** or **VLAN ID** assigned to the IEEE 802.1x-authenticated user.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

802.1x and Guest VLANs

This topic describes the use of 802.1x with guest VLANs.



It is possible to configure a guest VLAN for each IEEE 802.1x port on the switch to provide limited services to clients, such as Internet access or downloading the IEEE 802.1x client. Some clients might be upgrading their system for IEEE 802.1x authentication, while others, such as Microsoft Windows 98 systems, might not be IEEE 802.1x-capable.

With a guest VLAN enabled on an IEEE 802.1x port, the switch assigns clients to a guest VLAN when EAPOL packets are not sent by the client or when the authentication server does not receive a response to its EAPOL Request/Identity frame.

Before Cisco IOS Release 12.1(22)EA2, the switch did not maintain the EAPOL packet history and allowed clients that failed authentication access to the guest VLAN, regardless of whether EAPOL packets had been detected on the interface. You can enable this optional behavior by using the **dot1x guest-vlan supplicant** global configuration command.

Note	With Cisco IOS Release 12.1(22)EA2 and later, the switch maintains the EAPOL packet history. If another EAPOL packet is detected on the interface during the lifetime of the link, network access is denied. The EAPOL history is reset upon loss of the link.

Any number of IEEE 802.1x-incapable clients are allowed access when the switch port is moved to the guest VLAN.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Caution If an IEEE 802.1x-capable client joins the same port on which the guest VLAN is configured, the port is put into the unauthorized state in the user-configured access VLAN, and authentication is restarted.

Guest VLANs are supported on 802.1x ports in single-host or multiple-hosts mode.

Any active VLAN, except an RSPAN VLAN or a voice VLAN, can be configured as an IEEE 802.1x guest VLAN. The guest VLAN feature is supported only on access ports.

Configuring a Guest VLAN on a Port

When you configure a guest VLAN, clients that are not IEEE 802.1x-capable are put into the guest VLAN when the server does not receive a response to its EAPOL Request/Identity frame. Clients that are IEEE 802.1x-capable but fail authentication are not granted access to the network. The switch supports guest VLANs in single-host or multiple-hosts mode.

Perform these tasks to configure a guest VLAN on a switch port:

Step 1 Enable AAA.

Step 2 Enable 802.1x guest VLAN behavior globally.

Step 3 Configure the switch port as an access port.

Step 4 Configure dot1x port control as **auto.**

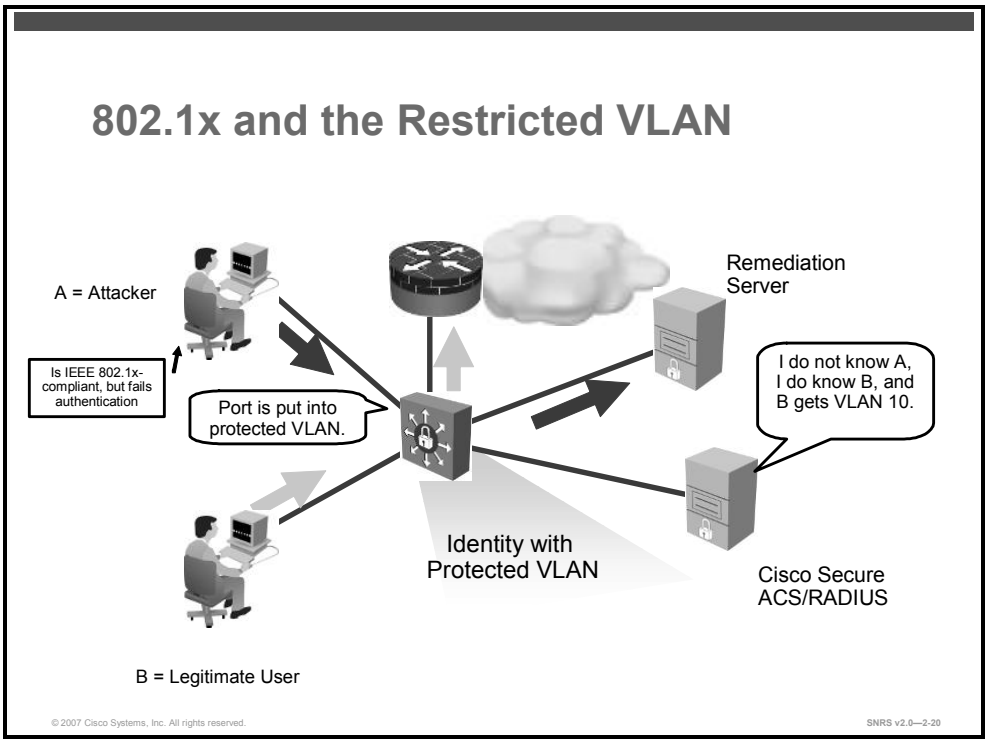
Step 5 Specify an active VLAN as a guest VLAN.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

802.1x and Restricted VLANs

This topic describes the use of 802.1x with restricted VLANs.



Another security feature allows you to configure a restricted VLAN for each IEEE 802.1x port to provide limited services to clients that cannot access the guest VLAN. Clients that are IEEE 802.1x-compliant and cannot access another VLAN because they fail the authentication process will be put in this VLAN. A restricted VLAN allows users without valid credentials in an authentication server (typically, visitors to an enterprise) to access a limited set of services. The administrator can control the services available to the restricted VLAN.

Note	It is possible to configure a VLAN to be both the guest VLAN and the restricted VLAN if you want to provide the same services to both types of users.

The switch (authenticator) keeps a count of failed authentication attempts for that particular client. When the count exceeds the maximum number of authentication attempts configured, the port moves to the restricted VLAN. When RADIUS replies with either an EAP Failure or an empty response that contains no EAP packet, the failed attempt count is incremented. The failed attempt counter resets when the port moves into the restricted VLAN.

Users failing authentication will remain in the restricted VLAN until the next reauthentication attempt. A port in the restricted VLAN tries to reauthenticate at configured intervals (the default is 60 seconds). If reauthentication fails, the port remains in the restricted VLAN. If reauthentication is successful, the port moves to either the configured VLAN or to a VLAN sent by the RADIUS server. Reauthentication may be disabled. If you do this, the only way to start the authentication process again is for the port to receive a link down event or EAP logoff event. It is recommended that reauthentication be enabled if a client might connect through a

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

hub. When a client disconnects from the hub, the port might not receive the link down event or EAP logoff event.

After a port moves to the restricted VLAN, it sends a simulated EAP Success message to the client. This prevents clients from attempting authentication indefinitely. Some clients (for example, devices running Microsoft Windows XP) cannot implement DHCP without EAP success.

802.1X restrictions are as follows:

- Supported only on access ports
- Not supported on trunk ports
- Supported only on IEEE 802.1x ports in single-host mode on a Layer 2 port
- Cannot configure on a Remote Switched Port Analyzer (RSPAN) VLAN or a voice VLAN

Compatibility with Other Security Features

Protected VLANs work with port security. As soon as the port is authorized, a MAC address is provided to port security. If port security does not permit the MAC address, or if the maximum secure address count is reached, the port becomes unauthorized and error-disabled.

Dynamic Address Resolution Protocol (ARP) inspection, DHCP snooping, and IP source guard can be configured independently.

Configuring a Restricted VLAN

IEEE 802.1x-compliant clients are moved into the restricted VLAN when the authentication server does not receive a valid username and password from the client. Restricted VLANs are supported only in single-host mode.

Perform these tasks to configure a restricted VLAN on a switch port:

Step 1 Enable AAA.

Step 2 Configure the switch port as access.

Step 3 Configure dot1x port control as **auto.**

Step 4 Specify an active VLAN as a restricted VLAN.

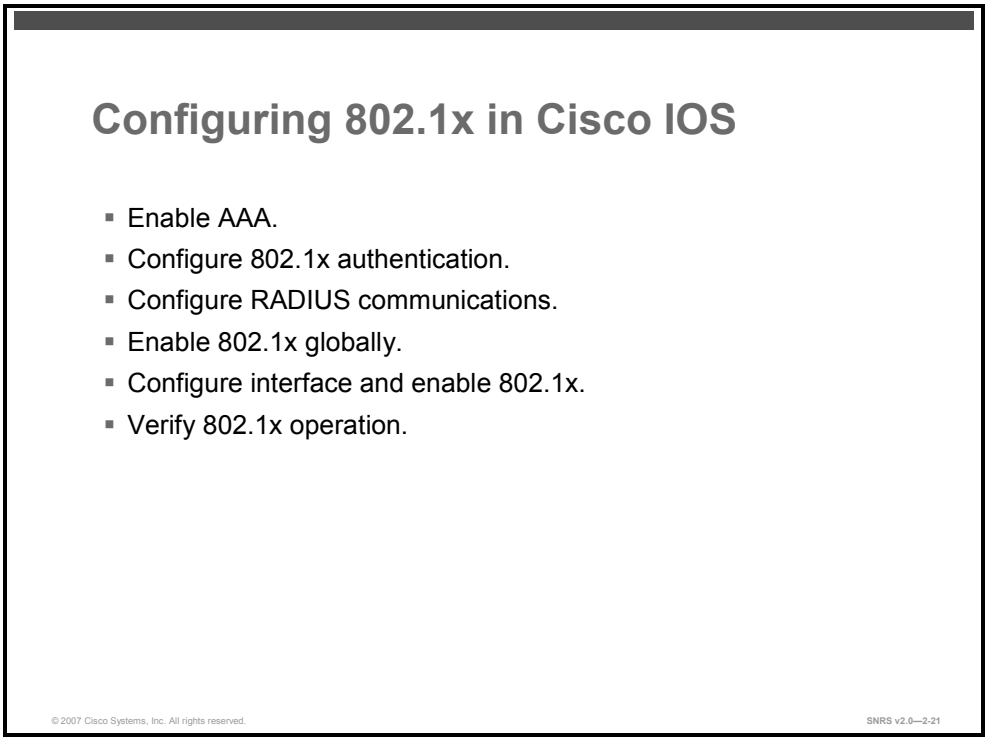
You can configure the maximum number of authentication attempts allowed before a user is assigned to the restricted VLAN by using the **dot1x auth-fail max-attempts** interface configuration command. The range of allowable authentication attempts is one to three. The default is three attempts.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Configuring 802.1x

This topic describes the commands used to configure and verify 802.1x operation on a Cisco Catalyst switch.



The basic configuration of the Cisco Catalyst switch or Cisco Aironet wireless LAN access point remains constant within any IEEE 802.1x deployment regardless of the EAP method chosen for authentication. The EAP method is agreed upon by the client and authentication server, and the authenticator simply proxies the information between the two.

The switch, as the authenticator, controls the physical access to the network based on the authentication status of the client. The authenticator acts as an intermediary between the client and the authentication server, requesting identity information from the client, verifying that information with the authentication server and relaying a response to the client. The authenticator communicates with the client via EAPOL and with the authentication server via RADIUS.

The following steps are required to enable 802.1x on the switch:

Step 1 Enable AAA.

Step 2 Configure 802.1x authentication.

Step 3 (Optional) Configure 802.1x authorization.

Step 4 Configure RADIUS communications.

Step 5 Enable 802.1x globally on the switch.

Step 6 Verify 802.1x operation

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Guidelines

The next figures will cover the commands required to complete these steps. But first, these are some configuration guidelines:

- When IEEE 802.1x authentication is enabled, ports are authenticated before any other Layer 2 features are enabled.
- The IEEE 802.1x protocol is supported on Layer 2 static access ports and voice VLAN ports, but it is not supported on the following port types:
  - Trunk port
  - Dynamic ports
  - Dynamic access ports
  - EtherChannel ports
  - Switched Port Analyzer (SPAN) and RSPAN destination ports
  - Cisco Long-Reach Ethernet (LRE) switch ports
- Remove the EtherChannel configuration from the interfaces on which IEEE 802.1x authentication will be enabled and before globally enabling IEEE 802.1x authentication on a switch.

VLAN Assignment, Guest VLANs, and Restricted VLANs

Here are some points to keep in mind when configuring IEEE 802.1x with VLAN assignment and with guest and restricted VLANs:

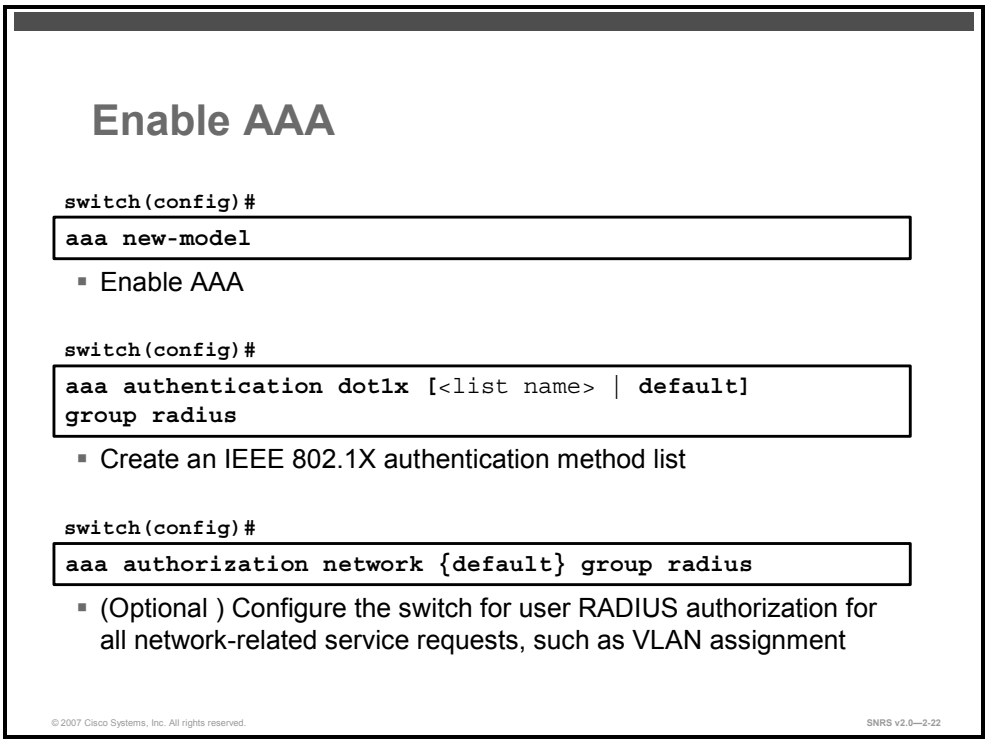
- Authentication with the VLAN assignment feature is not supported on trunk ports, dynamic ports, or with dynamic access port assignment through a VMPS.
- Guest VLANs are supported only on access ports.
- You can configure any VLAN as a guest VLAN except an RSPAN VLAN or a voice VLAN.
- You can configure any VLAN as a restricted VLAN except an RSPAN VLAN or a voice VLAN.
- Restricted VLANs are supported only on access ports.
- You can change the settings for restarting the authentication process on the switch before the DHCP process on the client times out and tries to get a host IP address from the DHCP server.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Configuring AAA

This section describes how to configure AAA to work with IBNS.



Complete these steps to enable AAA services on the switch.

- Enable AAA.

```
switch(config)# aaa new-model
```

- Specify one or more authentication, authorization, and accounting (AAA) methods for use on interfaces running IEEE 802.1X.

```
switch(config)# aaa authentication dot1x {default | listname}
method1 [method2...]
```

Syntax Description

default	Uses the listed authentication methods that follow this argument as the default list of methods when a user logs in.
listname	Character string used to name the list of authentication methods tried when a user logs in.
method1 [method2...]	At least one of these keywords: ■ **enable**—Uses the enable password for authentication. ■ **group radius**—Uses the list of all RADIUS servers for authentication. ■ **line**—Uses the line password for authentication. ■ **local**—Uses the local username database for authentication. ■ **local-case**—Uses the case-sensitive local username database for authentication. ■ **none**—Uses no authentication. The client is automatically authenticated by the switch without using the information supplied by the client.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

The *method* argument identifies the list of methods that the authentication algorithm tries in the given sequence to validate the password provided by the client. The only method that is truly 802.1X-compliant is the **group radius** method, in which the client data is validated against a RADIUS authentication server. The remaining methods enable AAA to authenticate the client by using locally configured data. For example, the **local** and **local-case** methods use the username and password that are saved in the Cisco IOS configuration file. The **enable** and **line** methods use the **enable** and **line** passwords for authentication.

If you specify **group radius**, you must configure the RADIUS server by entering the **radius-server host** global configuration command. If you are not using a RADIUS server, you can use the **local** or **local-case** methods, which access the local username database to perform authentication. By specifying the **enable** or **line** methods, you can supply the clients with a password to provide access to the switch.

A named method list can be defined or the keyword **default** can be used and applied to all ports. Though other methods appear as configuration options, only "group radius" is supported.

- Set parameters that restrict user access to a network.

```
Switch(config)# aaa authorization {network | exec | commands
level | reverse-access | configuration} {default | list-name}
[method1 [method2...]]
```

Syntax Description

network	Runs authorization for all network-related service requests, including Serial Line Internet Protocol (SLIP), PPP, PPP Network Control Programs (NCPs), and AppleTalk Remote Access (ARA).
exec	Runs authorization to determine if the user is allowed to run an EXEC shell. This facility might return user profile information such as autocommand information.
commands	Runs authorization for all commands at the specified privilege level.
level	Specific command level that should be authorized. Valid entries are 0 through 15.
reverse-access	Runs authorization for reverse access connections, such as reverse Telnet.
configuration	Downloads the configuration from the authentication, authorization, and accounting (AAA) server.
default	Uses the listed authorization methods that follow this keyword as the default list of methods for authorization.
list-name	Character string used to name the list of authorization methods.
method1 [*method2...*]	Identifies an authorization method or multiple authorization methods to be used for authorization. A method may be any one of the keywords listed in the table "AAA Authorization Methods".

Use the **aaa authorization** command to enable authorization and to create named methods lists, defining authorization methods that can be used when a user accesses the specified function. Method lists for authorization define the ways authorization will be performed and the sequence in which these methods will be performed. A method list is a named list describing the authorization methods to be used (such as RADIUS or TACACS+), in sequence. Method lists enable you to designate one or more security protocols to be used for authorization, thus ensuring a backup system in case the initial method fails. Cisco IOS software uses the first method listed to authorize users for specific network services; if that method fails to respond,

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

the Cisco IOS software selects the next method listed in the method list. This process continues until there is successful communication with a listed authorization method, or all methods defined are exhausted.

Note The Cisco IOS software attempts authorization with the next listed method only when there is no response from the previous method. If authorization fails at any point in this cycle—meaning that the security server or local username database responds by denying the user services—the authorization process stops and no other authorization methods are attempted.

If the **aaa authorization** command for a particular authorization type is issued without a named method list specified, the default method list is automatically applied to all interfaces or lines (where this authorization type applies) except those that have a named method list explicitly defined. (A defined method list overrides the default method list.) If no default method list is defined, then no authorization takes place.

Use the **aaa authorization** command to create a list by entering values for the *list-name* and the *method* arguments, where *list-name* is any character string used to name this list (excluding all method names) and *method* identifies the list of authorization methods tried in the given sequence.

Following is a table of AAA authorization methods that may be used with the **aaa authorization** command.

AAA Authorization Methods

cache *group-name*	Uses a cache server group for authorization.
group radius	Uses the list of all RADIUS servers for authentication as defined by the **aaa group server radius** command.
group tacacs+	Uses the list of all TACACS+ servers for authentication as defined by the **aaa group server tacacs+** command.
group	Uses a subset of RADIUS or TACACS+ servers for accounting as defined by the **server group** *group-name* command.
if-authenticated	Allows the user to access the requested function if the user is authenticated.
krb5-instance	Uses the instance defined by the **kerberos instance map** command.
local	Uses the local database for authorization.

Cisco IOS software supports the following methods for authorization:

- Cache Sever Groups—The router consults its cache server groups to authorize specific rights for users.
- RADIUS—The network access server requests authorization information from the RADIUS security server group. RADIUS authorization defines specific rights for users by associating attributes, which are stored in a database on the RADIUS server, with the appropriate user.
- TACACS+—The network access server exchanges authorization information with the TACACS+ security daemon. TACACS+ authorization defines specific rights for users by

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

associating attribute-value (AV) pairs, which are stored in a database on the TACACS+ security server, with the appropriate user.

- If-Authenticated—The user is allowed to access the requested function provided the user has been authenticated successfully.
- None—The network access server does not request authorization information; authorization is not performed over this line or interface.
- Local—The router or access server consults its local database, as defined by the **username** command, to authorize specific rights for users. Only a limited set of functions can be controlled via the local database.
- Kerberos Instance Map—The network access server uses the instance defined by the **kerberos instance map** command for authorization.

Method lists are specific to the type of authorization being requested. AAA supports five different types of authorization:

- Network—Applies to network connections. This can include a PPP, SLIP, or ARA connection.
- EXEC—Applies to the attributes associated with a user EXEC terminal session.
- Commands—Applies to the EXEC mode commands a user issues. Command authorization attempts authorization for all EXEC mode commands, including global configuration commands, associated with a specific privilege level.
- Reverse Access—Applies to reverse Telnet sessions.
- Configuration—Applies to the configuration downloaded from the AAA server.

When you create a named method list, you are defining a particular list of authorization methods for the indicated authorization type.

Once defined, method lists must be applied to specific lines or interfaces before any of the defined methods will be performed.

The authorization command causes a request packet containing a series of AV pairs to be sent to the RADIUS or TACACS daemon as part of the authorization process. The daemon can do one of the following:

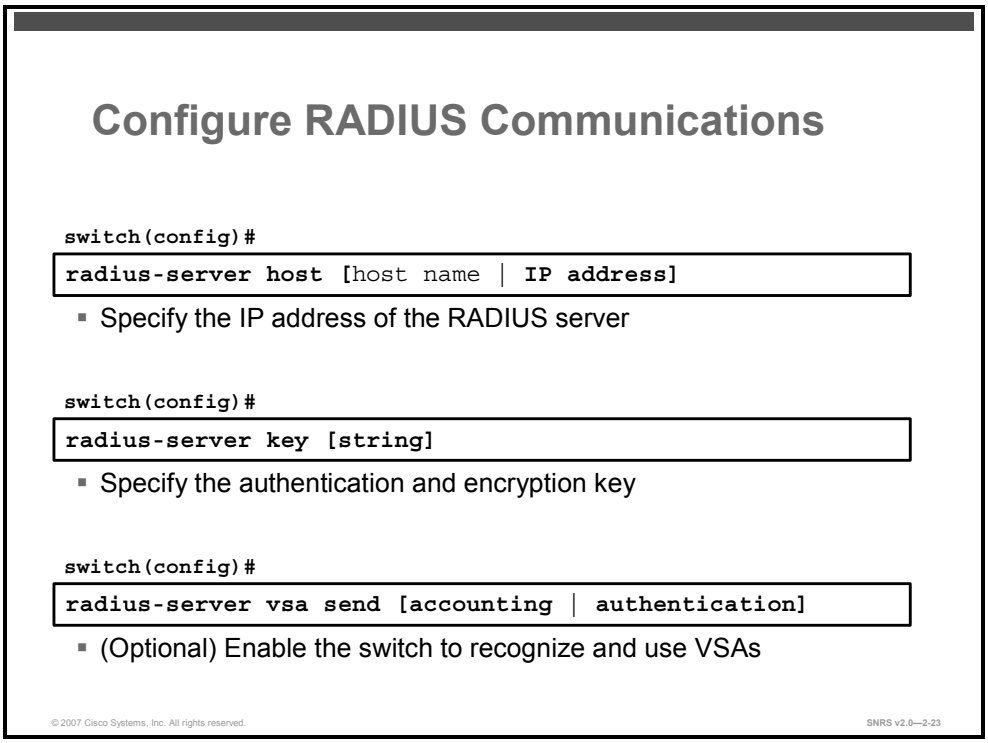
- Accept the request as is.
- Make changes to the request.
- Refuse the request and refuse authorization.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

RADIUS Communications

This section describes how to configure RADIUS to work with IBNS.



Configure RADIUS communications using the following commands:

- **radius-server host** [*host name* | *IP address*] auth-port [port] acct-port [port]

This command specifies the IP address of the RADIUS server. Additionally, the authentication and accounting port numbers can be changed from the default values of 1645 and 1646.

- **radius-server key** [**string**]

This command specifies the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server.

- **radius-server vsa send** [**accounting** | **authentication**]

This command enables the switch to recognize and use vendor-specific attributes (VSAs) as defined by RADIUS IETF attribute 26.

  — Use the **accounting** keyword to limit the set of recognized VSAs to only accounting attributes.

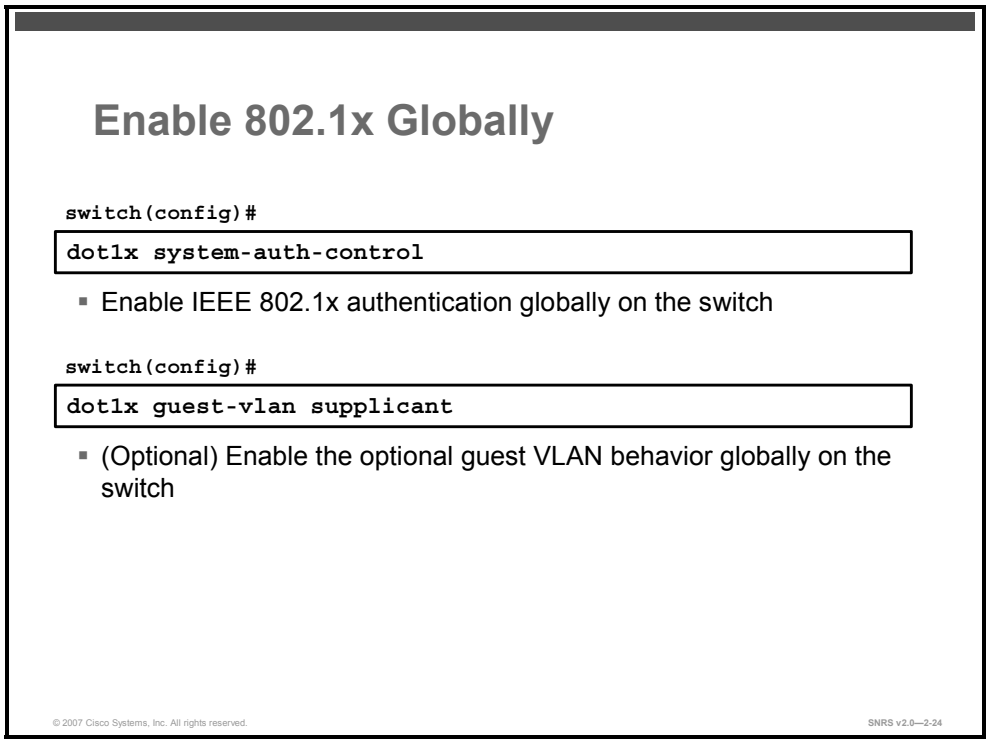
  — Use the **authentication** keyword to limit the set of recognized VSAs to only authentication attributes.

Note	If you enter this command without keywords, both accounting and authentication VSAs are used.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Enabling 802.1X

This section describes how to globally enable 802.1X



Enable 802.1x globally on the switch using the following commands:

- **dot1x system-auth-control**

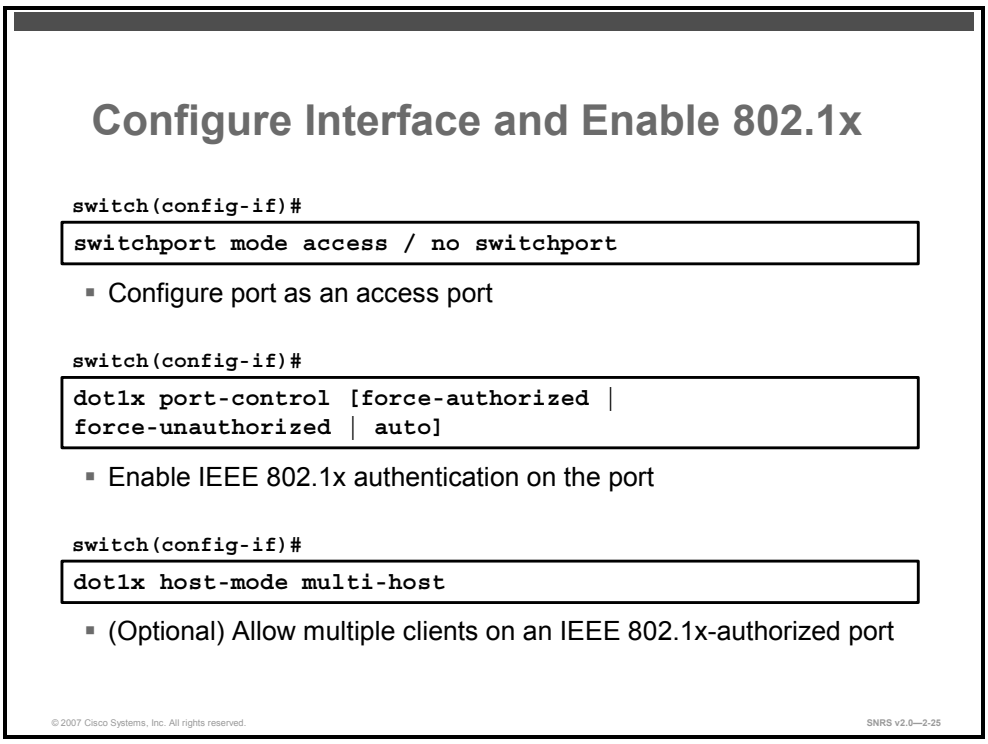
This command globally enables IEEE 802.1x authentication on the switch.

- (Optional) **dot1x guest-vlan supplicant**

Before Cisco IOS Release 12.1(22)EA2, the switch did not maintain the EAPOL packet history and allowed clients that failed authentication access to the guest VLAN, regardless of whether EAPOL packets had been detected on the interface. Use this command to enable this optional behavior.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.



Configure the interface using the following commands:

- **switchport mode access**

IEEE 802.1x can only be configured on static Layer 2 access ports.

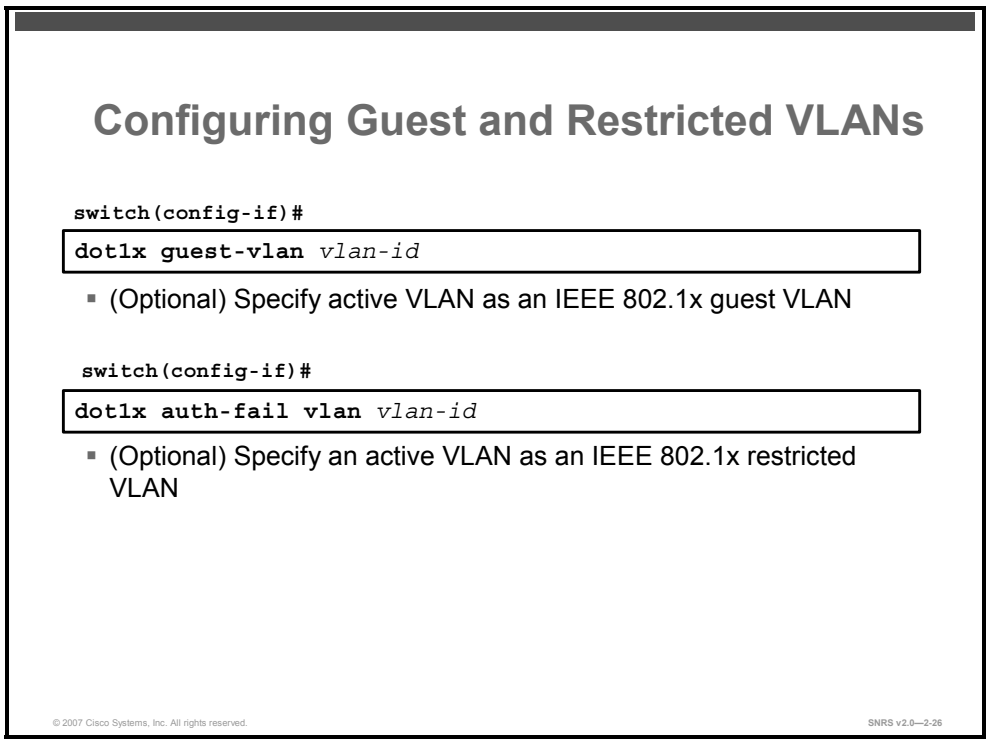
- **dot1x port-control** [*force-authorized* | *force-unauthorized* | **auto**]

This command enables IEEE 802.1x authentication on the port. The default is *force-authorized*.

- (Optional) **dot1x host-mode multi-host**

This command allows multiple hosts on an IEEE 802.1x-authorized port.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.



- **dot1x guest-vlan** *vlan-id*
  - This command specifies an active VLAN as an IEEE 802.1x guest VLAN. The range is 1 to 4094.
  - You can configure any active VLAN, except an RSPAN VLAN or a voice VLAN, as an IEEE 802.1x guest VLAN.
- **dot1x auth-fail vlan** *vlan-id*
  - This command specifies an active VLAN as an IEEE 802.1x restricted VLAN. The range is 1 to 4094.
  - You can configure any active VLAN, except an RSPAN VLAN or a voice VLAN, as an IEEE 802.1x restricted VLAN.

Periodic Reauthentication

You can enable periodic IEEE 802.1x reauthentication and specify how often it occurs. The default number of seconds between reauthentication attempts is 3600.

```
switch(config)# dot1x reauthentication
switch(config)# dot1x timeout reauth-period {seconds | server}
```

- *seconds*: The number of seconds from 1 to 65535
- **server:** Sets the number of seconds based on the value of the Session-Timeout RADIUS attribute (Attribute[27]) and Termination-Action RADIUS attribute (Attribute [29]). These attributes are set in the Radius Authorization Components section of the Shared Profile Components of Cisco Secure ACS.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Manually Reauthenticating a Client

You can manually reauthenticate the client connected to a specific port at any time by entering the **dot1x re-authenticate interface** *interface-id* privileged EXEC command.

```
switch# dot1x re-authenticate interface fastethernet0/1
```

Adjusting the Quiet Period

The quiet period is the time that the switch remains idle before attempting to authenticate a client again after a failed authentication exchange. This time is determined by the quiet-period value. Faster response time to the user may be provided by entering a smaller number than the default of 60 seconds.

Use the **dot1x timeout quiet-period** command to configure the quiet period.

```
switch(config)# dot1x timeout quiet-period seconds
```

Adjusting the Switch-to-Client Retransmission Time

When the switch does not receive an EAP Response/Identity from the client, it waits a specific amount of time and then resends the request. The default time is 30 seconds.

Use the **dot1x timeout tx-period** command to adjust the retransmission time.

```
switch(config)# dot1x timeout tx-period seconds
```

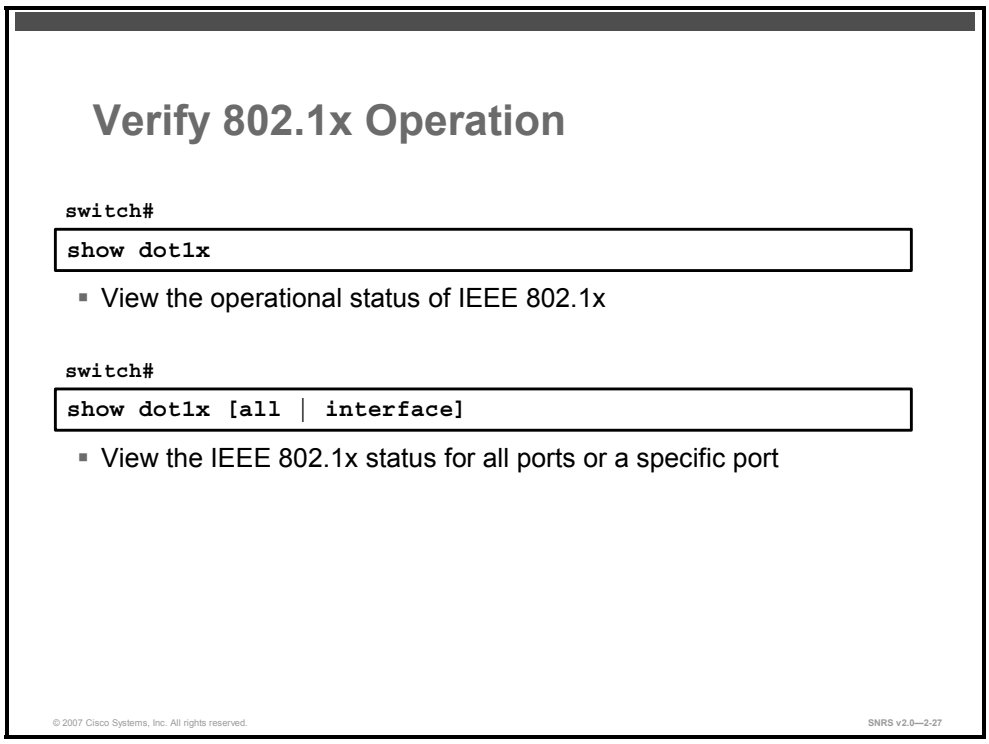
Note	You should change the default value of this command only to adjust for unusual circumstances, such as unreliable links or specific behavioral problems with certain clients and authentication servers.

Adjusting Timers for DHCP

The following example shows how to enable a VLAN as an IEEE 802.1x guest VLAN when an IEEE 802.1x port is connected to a DHCP client. You will set 3 as the quiet time on the switch, and set 15 as the number of seconds that the switch waits for a response to an EAP Request/Identity frame from the client before resending the request.

```
Switch(config-if)# dot1x timeout quiet-period 3
Switch(config-if)# dot1x timeout tx-period 15
Switch(config-if)# dot1x guest-vlan 20
```


The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.



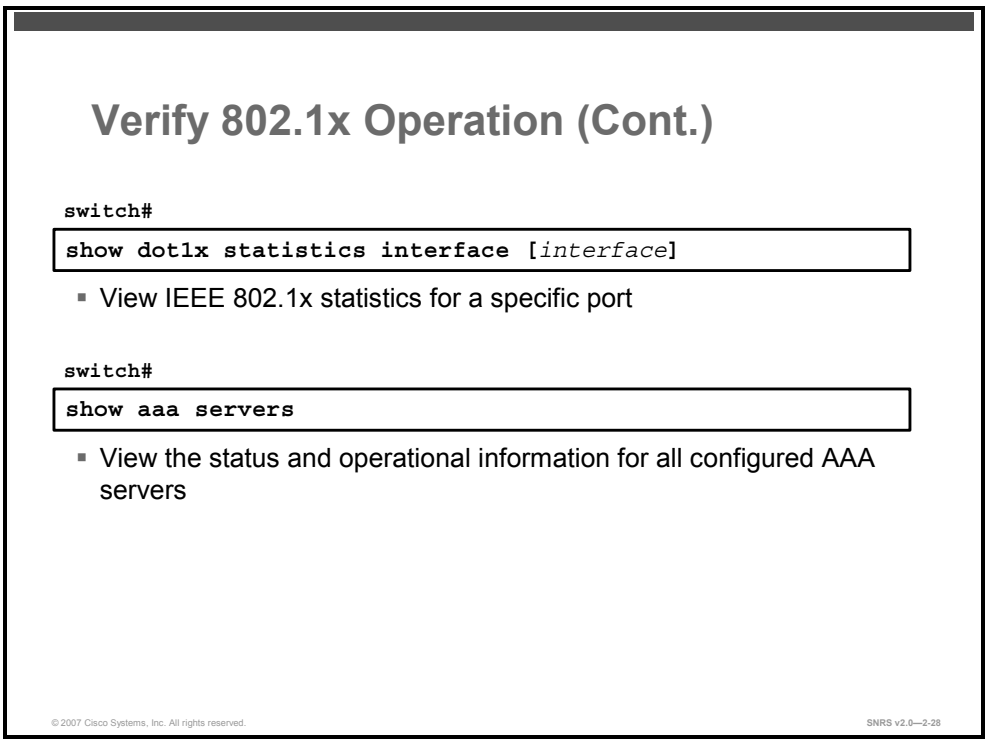
The following commands are used to verify 802.1x operation on the switch:

- **show dot1x**
  - This command displays the operational status of IEEE 802.1x.
  - Check the Status column in the IEEE 802.1x Port Summary section of the display. An "Enabled" status means that the port-control value is set to either **auto** or to **force-unauthorized**.
- **show dot1x** [*all* | *interface]*

This command displays the IEEE 802.1x status for all ports or a specific port.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.



The following commands are also used to verify 802.1x operation on the switch:

- **show dot1x statistics interface** [*interface]*

This command displays IEEE 802.1x statistics for a specific port.

- **show aaa servers**

This command displays the status and operational information for all configured AAA servers.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Summary

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

Summary

- Cisco IBNS combines several Cisco products that offer authentication, access control, and user policies to secure network connectivity and resources.
- 802.1x is a standardized framework defined by the IEEE, designed to provide port-based network access.
- 802.1x roles include the supplicant, authenticator, and authentication server.
- 802.1x uses EAP and RADIUS for authentication.
- Various types of EAP methods are available for use with 802.1x.

© 2007 Cisco Systems, Inc. All rights reserved.

Summary (Cont.)

- 802.1x works with port security.
- 802.1x works with VLAN assignment.
- 802.1x works with guest VLANs.
- 802.1x works with restricted VLANs.
- Various commands are used to configure and verify operation of 802.1x on a Cisco Catalyst switch.

© 2007 Cisco Systems, Inc. All rights reserved.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

References

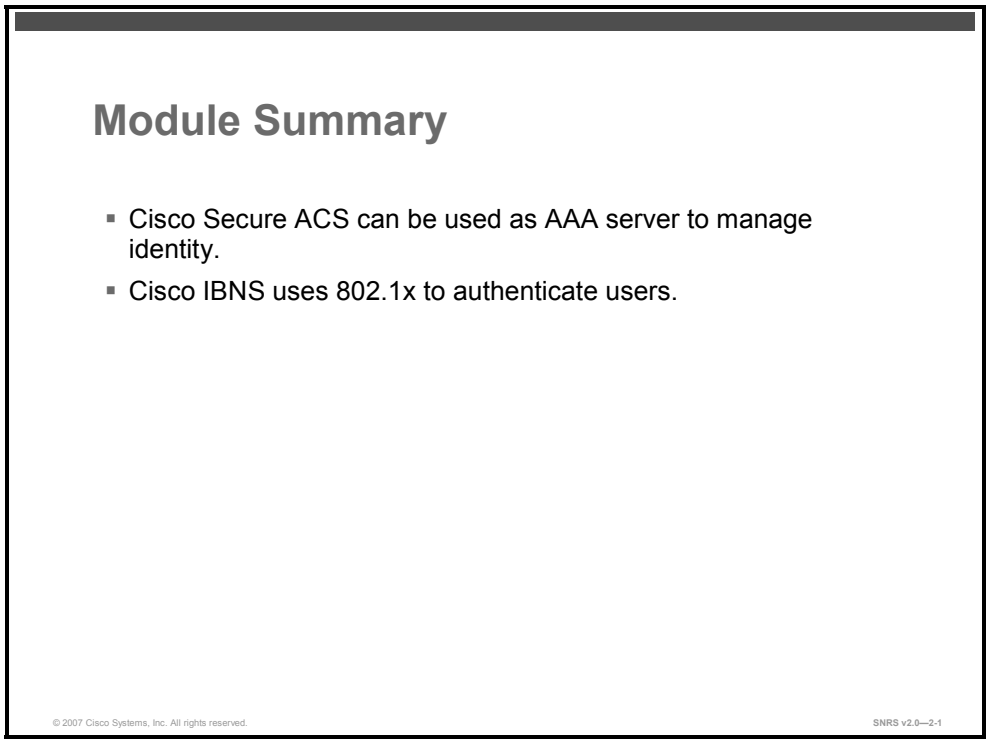
For additional information, refer to these resources:

- Cisco Systems, Inc. *Identity-Based Networking Systems Configuration Guide*, Version 1.0. San Jose, California, December 2005. http://www.cisco.com/application/pdf/en/us/guest/netsol/ns75/c654/cdccont_0900aecd803fab62.pdf.
- *Configuring IEEE 802.1x Port-Based Authentication*: http://www.cisco.com/en/US/products/hw/switches/ps628/products_configuration_guide_chapter09186a0080648d7a.html
- Cisco Systems, Inc. *Catalyst 2950 and Catalyst 2955 Switch Software Configuration Guide, 12.1(22)EA7*. http://www.cisco.com/en/US/products/hw/switches/ps628/products_configuration_guide_book09186a008064737d.html.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Module Summary

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



In this module you were introduced to IBNS and Cisco Secure ACS. Cisco Identity-Based Network Services (IBNS) use 802.1x as a framework. The Cisco Secure Access Control Server (ACS) is used as an authentication, authorization, and accounting (AAA) server in conjunction with Cisco IBNS to provide an identity management solution. Both of these technologies are combined to implement a port-based authentication and authorization solution to help secure the network.

References

For additional information, refer to these resources:

- Cisco Systems, Inc. *Identity-Based Networking Systems Configuration Guide*, Version 1.0. San Jose, California, December 2005. http://www.cisco.com/application/pdf/en/us/guest/netsol/ns75/c654/cdccont_0900aecd803fab62.pdf
- Cisco Systems, Inc. *Catalyst 2950 and Catalyst 2955 Switch Software Configuration Guide, 12.1(22)EA7.* http://www.cisco.com/en/US/products/hw/switches/ps628/products_configuration_guide_book09186a008064737d.htm
- *Configuring IEEE 802.1x Port-Based Authentication*: http://www.cisco.com/en/US/products/hw/switches/ps628/products_configuration_guide_chapter09186a0080648d7a.html

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

- *ACS 4.0 Primer: http://www.cisco.com/application/pdf/en/us/guest/products/ps2086/c1161/cdccont_0900aecd8040daa7.pdf*
- *Installation Guide for Cisco Secure ACS for Windows 4.0: http://www.cisco.com/en/US/partner/products/sw/secursw/ps2086/products_installation_guide_chapter09186a008060f73e.html*
- *User Guide for Cisco Secure ACS for Windows 4.0: http://www.cisco.com/en/US/partner/products/sw/secursw/ps2086/products_user_guide_book09186a0080533dd8.html*
- *Cisco IOS Security Configuration Guide, Release 12.4: http://www.cisco.com/en/US/partner/products/ps6350/products_configuration_guide_book09186a008043360a.html*

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Module 3

Cisco Network Foundation Protection

Overview

In the modern business environment, connecting to the Internet is crucial to business operations. As a result, network elements and infrastructure devices are exposed to the risks and threats of the modern hacker or anyone else with malicious intent. To meet the business needs of IP services such as network availability and rapid deployment, it is critical to utilize network security features and services. There are security best practices that help secure the network foundation by protecting network elements and their interactions. Furthermore, to address the increasing complexity of the attacks in a heightened security environment, Cisco has enhanced Cisco IOS security features and services for both network elements and infrastructure, thus ensuring the availability of the network elements under all circumstances. Cisco Network Foundation Protection (NFP) provides an umbrella strategy for infrastructure protection by encompassing Cisco IOS security features.

Module Objectives

Upon completing this module, you will be able to implement command-line Cisco NFP to protect infrastructure devices. This ability includes being able to meet these objectives:

- Describe the Cisco NFP strategy
- Describe some tools used to secure the control plane of a router and configure the Cisco IOS CPPr feature
- Configure the Cisco IOS MPP feature
- Describe some tools used to protect the data plane and configure the Cisco IOS FPM feature

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Lesson 1

Introducing Cisco NFP

Overview

Cisco Network Foundation Protection (NFP) provides a strategy for network infrastructure protection by utilizing Cisco IOS security features. This lesson provides an overview of Cisco NFP and some of its features.

Objectives

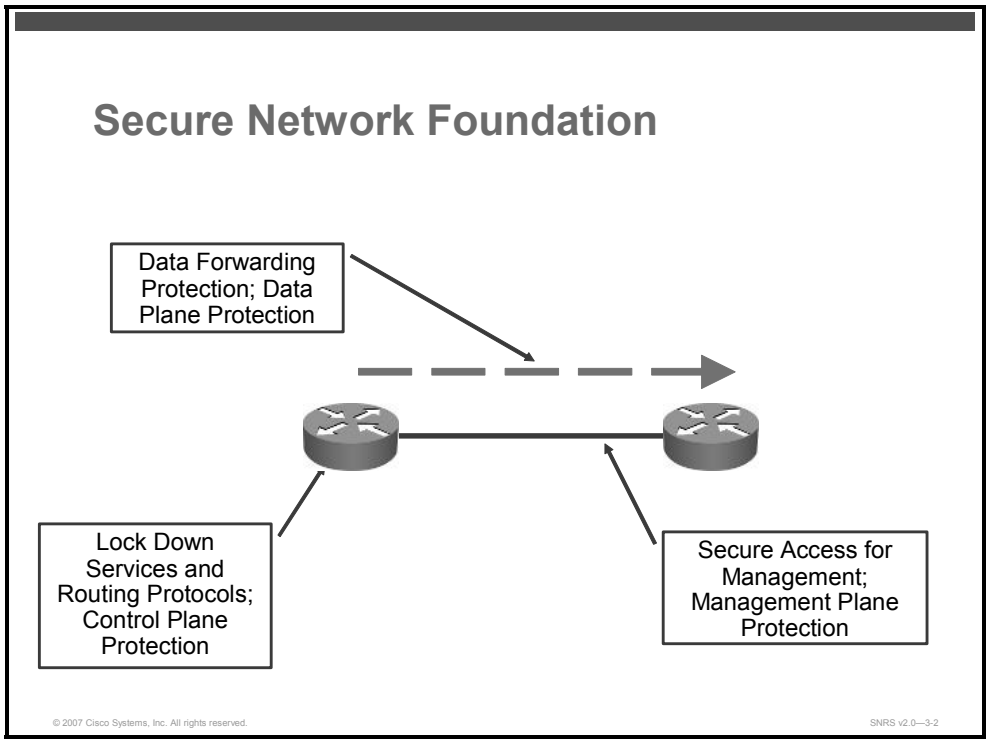
Upon completing this lesson, you will be able to describe the functions of Cisco NFP. This ability includes being able to meet these objectives:

- Describe Cisco NFP
- Describe some tools used to secure the network infrastructure

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Cisco NFP Overview

This topic gives an overview of Cisco NFP.



The network environment of today is complex, while networking devices offer a feature-rich set of services to cater to different business needs. Because connecting to the Internet is imperative, network devices and infrastructure are exposed to many risks and threats. To meet the business needs of IP services such as network availability and rapid deployment, it is critical to utilize security features and services. Deploying security best practices help secure the network foundation by protecting network elements and their interactions. Furthermore, to address the increasing complexity of the attacks in a heightened security environment, Cisco has enhanced Cisco IOS security features and services for both network devices and as infrastructure, thus ensuring the availability of the network devices under all circumstances.

The features of Cisco NFP provide a strategy for infrastructure protection by utilizing Cisco IOS security features.

Network Device Planes

Cisco NFP divides the device into three planes:

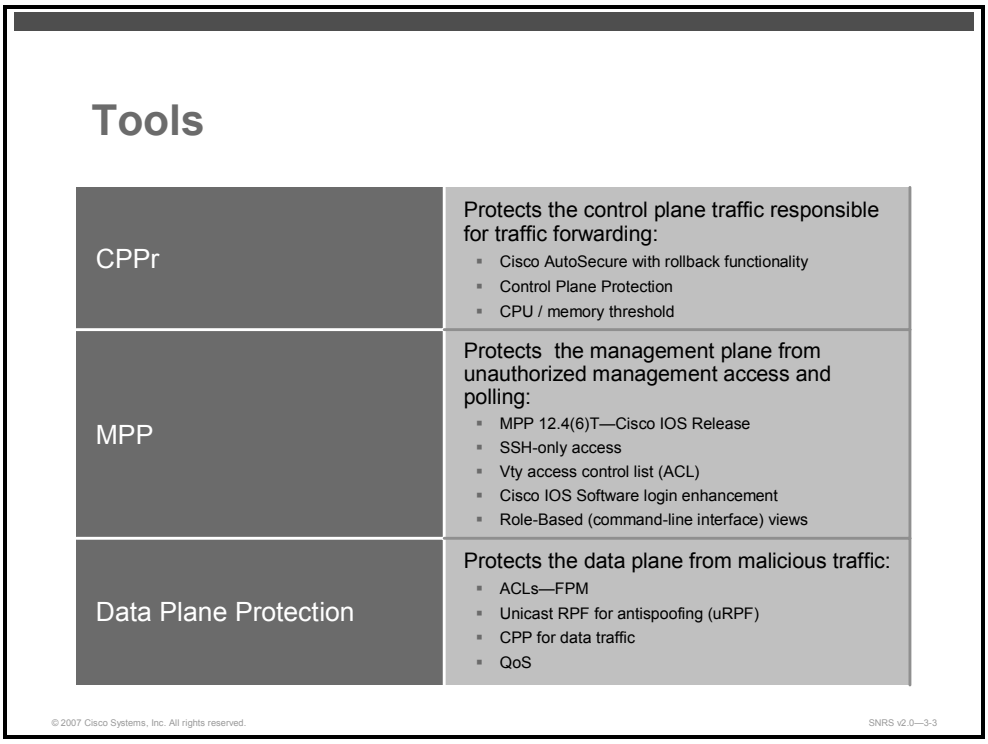
- **Control plane:** Represents the ability to route traffic
- **Management plane:** Represents the ability to manage the device
- **Data plane:** Represents the ability to forward data

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Cisco IOS Tools for a Secure Infrastructure

This topic describes some tools used to secure the infrastructure.



There are several tools that make up the Cisco NFP strategy. In the lessons in this module, the three planes and some tools used to secure each plane are examined.

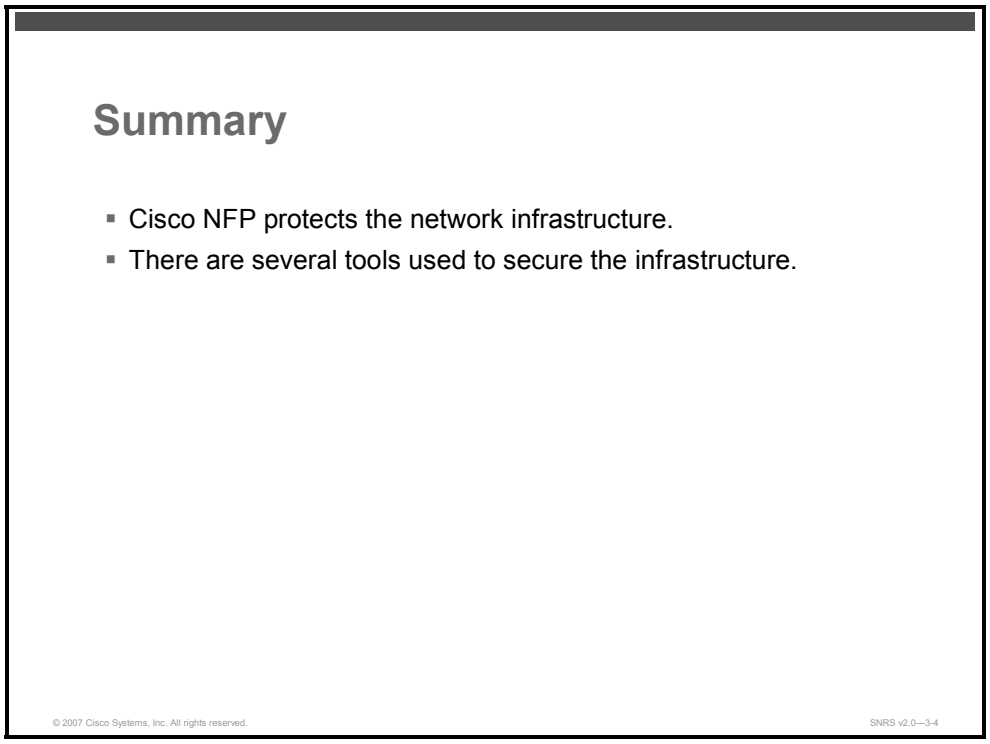
Listed here are some of the Cisco NFP tools available in Cisco IOS software:

- Cisco AutoSecure
- Control Plane Policing (CoPP)
- Control Plane Protection (CPPr)
- Flexible Packet Matching (FPM)
- Management Plane Protection (MPP)
- Quality of service (QoS) tools
- Unicast Reverse Path Forwarding (uRPF)

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Summary

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



References

For additional information, refer to this resource:

- Cisco Systems, Inc. Cisco Network Foundation Protection (NFP): Introduction. http://www.cisco.com/en/US/partner/products/ps6642/products_ios_protocol_group_home.html.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Lesson 2

Securing the Control Plane

Overview

Securing the control plane of a router is essential to a secure infrastructure. Control Plane Policing (CoPP) allows administrators to configure a quality of service (QoS) filter that will manage the traffic flow of control plane packets to protect the control plane of Cisco IOS routers and switches against reconnaissance and denial of service (DoS) attacks. Control Plane Protection (CPPr) extends the capabilities of CoPP. This lesson will teach the learner to configure both CoPP and CPPr.

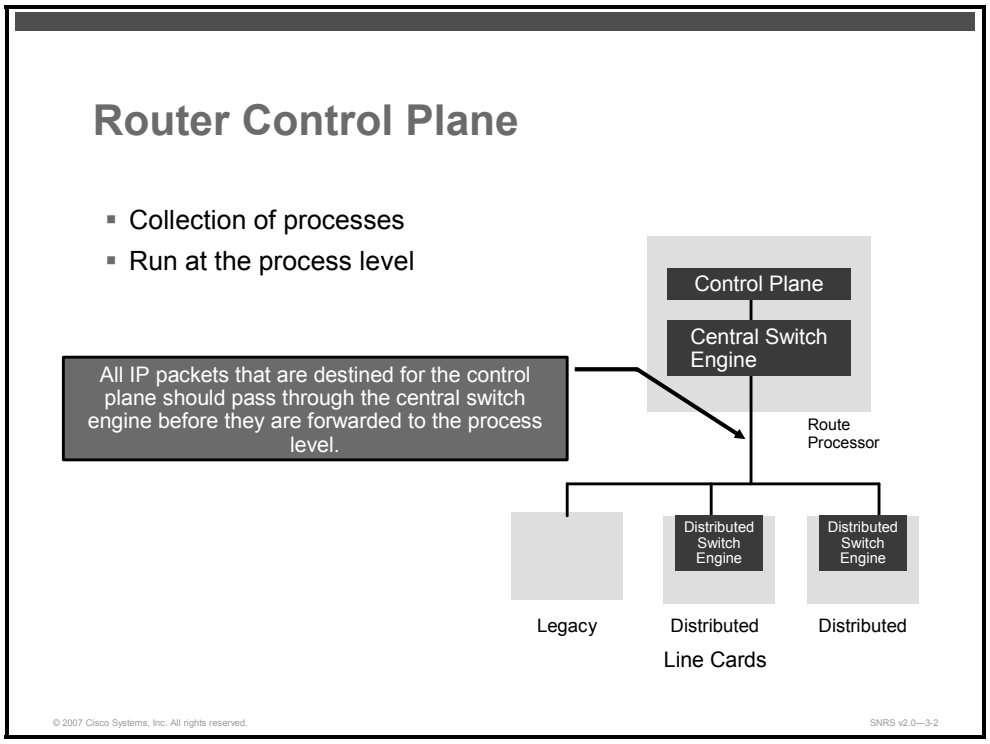
Upon completing this lesson, you will be able to describe some tools that are used to secure the control plane of a router. This ability includes being able to meet these objectives:

- Describe the control plane of a router
- Describe some tools used to secure the control plane
- Describe the basic function and benefits of CPPr
- List the components of the CPPr architecture
- Describe the required and optional steps used to configure CoPP
- Describe how to configure a port-filter policy
- Describe how to configure a queue-threshold policy
- List the commands used to verify CPPr configuration and view statistics

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Router Control Plane

This topic describes the control plane of a router.



A control plane is a collection of processes that run at the process level on a route processor (RP). They collectively provide high-level control for most Cisco IOS Software functions. All traffic directly or indirectly destined to a router is handled by the control plane.

Here are some definitions:

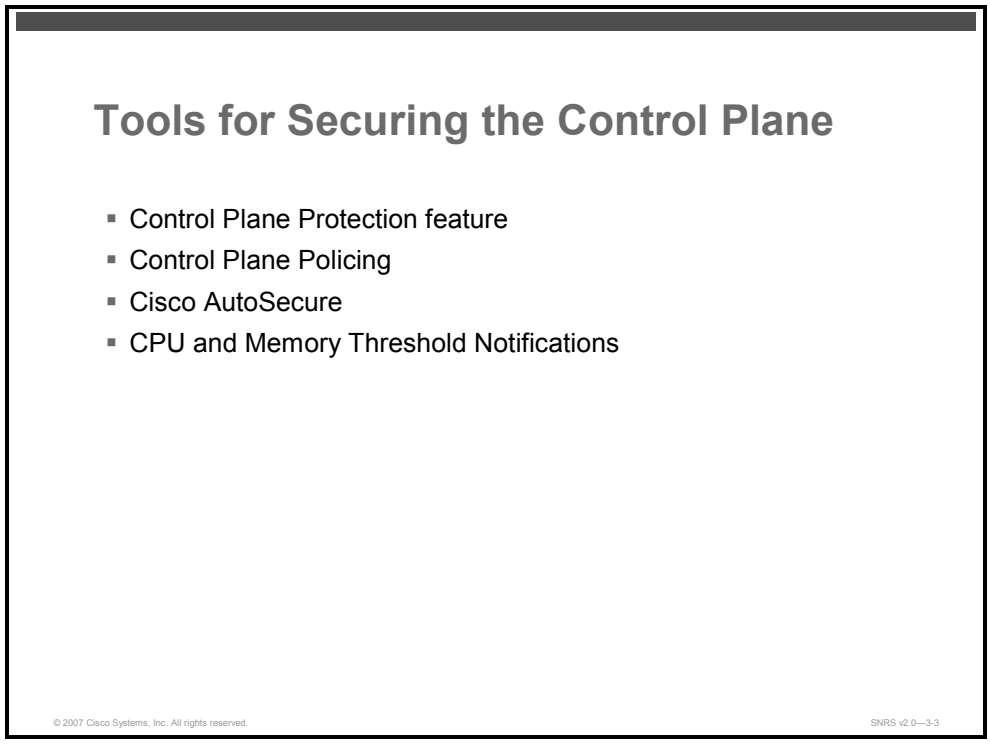
- **Control plane:** The control plane is a collection of processes that run at the process level on the RP.
- **Central switch engine:** A central switch engine is a device that is responsible for high-speed routing of IP packets. A central switch engine also typically performs high-speed input and output services for nondistributed interfaces. The central switch engine is used to implement aggregate control plane protection for all interfaces on the router.
- **Distributed switch engine:** A distributed switch engine is a device that is responsible for high-speed switching of IP packets on distributed line cards without using resources from the central switch engine. Each distributed switch engine is used to implement distributed control plane services for all ports on a line card. Input control plane services distribute the processing load across multiple line cards and conserve vital central switch engine resources. Distributed control plane services are optional; however, they provide a more refined level of service than aggregate services.
- **Legacy (nondistributed) line cards:** These are line cards that are responsible for receiving packets and occasionally performing input and output services. All packets must be forwarded to the central switch engine for a routing or switching decision. Aggregate control plane services provide coverage for nondistributed line cards.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Tools for Securing the Control Plane

This topic describes some tools used to secure the control plane.



There are several tools available in Cisco IOS Software for securing the control plane, including the following:

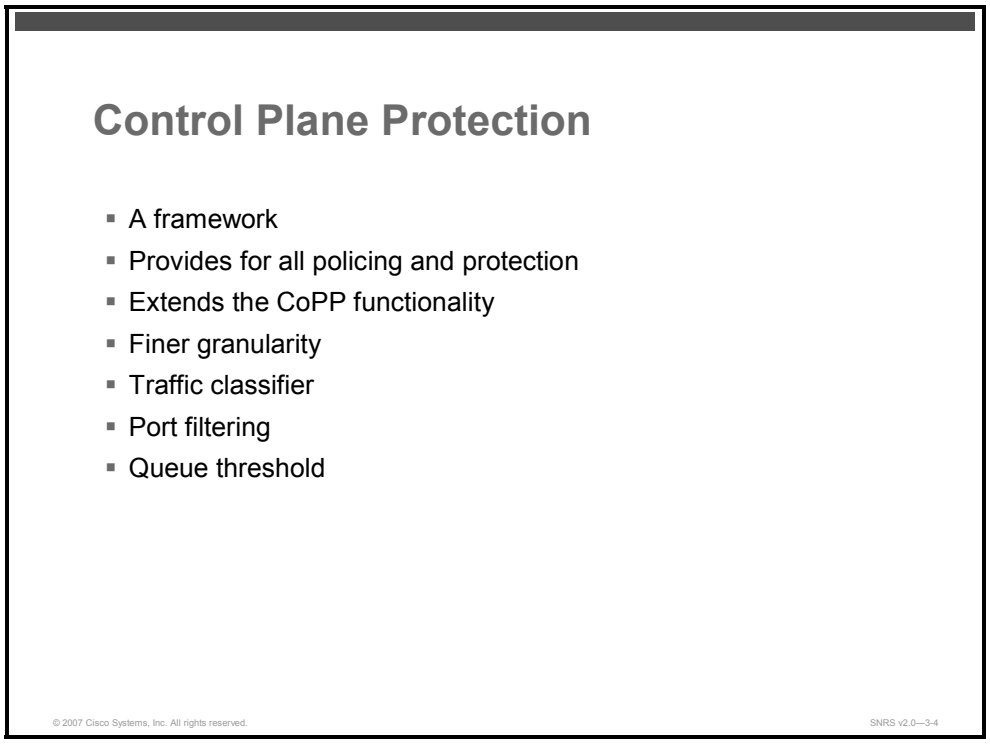
- CPPr feature
- CoPP
- Cisco AutoSecure
- CPU and Memory Threshold Notifications

This lesson will concentrate on the CPPr feature of Cisco IOS Release 12.4(4)T. However, because CPPr extends the functionality of CoPP, both features are covered in this lesson.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Overview of CPPr

This topic describes the basic function and benefits of the Cisco IOS CPPr feature.



One tool mentioned in the previous section is CPPr, which includes CoPP, port filtering, and queue thresholding. CPPr is a framework that encompasses all policing and protection features in the control plane. The CPPr feature extends the policing functionality of the CoPP feature by allowing finer policing granularity. CPPr also includes a traffic classifier. A traffic classifier intercepts control plane traffic and classifies it into three control plane categories. There are some new port-filtering and queue-thresholding features. Port filtering provides for the policing of packets going to closed or nonlistened TCP or User Datagram Protocol (UDP) ports. Queue thresholding limits the number of packets for a specified protocol that will be allowed in the control plane IP input queue.

Prerequisites

You must understand the following to implement CPPr:

- Principles of CoPP and how to classify and police control plane traffic
- Concepts and general Modular QoS command-line interface (CLI) (MQC) configuration procedure (class map and policy map) for applying QoS policies on a router

Restrictions for CPPr

As of Cisco IOS Release 12.4, the CPPr feature has these restrictions:

- Restricted to IP version 4 (IPv4) input path only
- Does not support direct access control list (ACL) configuration in the control plane subinterfaces (can be configured using MQC policies)

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

- Requires Cisco Express Forwarding for IP packet redirection
- On host subinterfaces, note the following:
  - The port-filter policy supports only TCP-based or UDP-based protocols.
  - The queue-threshold policy feature supports only TCP-based or UDP-based protocols.
- All IP packets entering the control plane matching any of the following conditions are not classified any further and are redirected to the Cisco Express Forwarding-exception subinterface:
  - IP packets with IP options
  - IP packets with Time to Live (TTL) of less than or equal to 1
- Some Cisco IOS TCP-based or UDP-based services, when configured, may not be auto-detected by the port filter. These ports must be manually added to the active port filter class map to be unblocked

Benefits

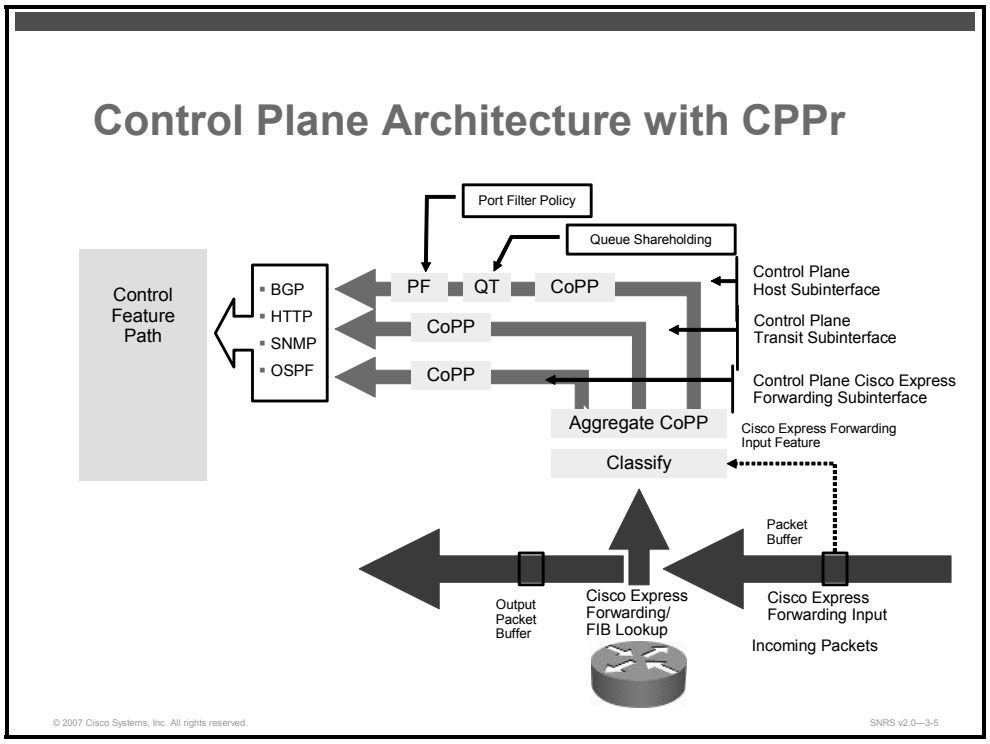
CPPr provides these benefits:

- Extends protection against DoS attacks at infrastructure routers by providing a mechanism for finer policing granularity for control plane traffic that allows you to rate-limit each type individually
- Provides a mechanism for early dropping of packets that are directed to closed or nonlistened Cisco IOS TCP or UDP ports
- Provides the ability to limit protocol queue usage such that no single protocol flood can overwhelm the input interface
- Provides QoS control for packets that are destined to the control plane of Cisco routers
- Provides ease of configuration for control plane policies using MQC infrastructure
- Provides better platform reliability, security and availability
- Provides a dedicated control plane subinterface for aggregate, host, transit and Cisco Express Forwarding-exception control plane traffic processing
- Highly flexible—permit, deny, rate-limit
- Provides CPU protection

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

CPPr Architecture

This topic describes the components of the CPPr architecture.



This diagram depicts the flow of control traffic through the control plane architecture with the CPPr feature enabled.

Aggregate Control Plane Services

The CoPP feature is intended to be the first CPPr feature encountered by packets before any other features or policies. Existing (aggregate) CoPP policies will not be affected when the CPPr functionality is enabled. The aggregate CoPP policy will be applied on all control plane traffic types. However, CPPr allows for additional and separate CoPP policies to be configured and applied on the different types of control plane subinterfaces (host, transit, cef-exception).

Control Plane Interface and Subinterface

The concept of early rate-limiting protocol specific traffic destined to the processor by applying QoS policies to the aggregate control plane interface was introduced with CoPP. CPPr extends this control plane functionality by providing three additional control plane subinterfaces under the top-level (aggregate) control plane interface. Each subinterface receives and processes a specific type of control plane traffic. The three sub interfaces are as follows:

- **Control plane host subinterface:** Receives all control plane IP traffic that is directly destined for one of the router interfaces
  - Examples of this include tunnel termination traffic and management traffic or routing protocols such as Secure Shell (SSH), Simple Network Management Protocol (SNMP), Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), and Enhanced Interior Gateway Routing Protocol (EIGRP).

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

- All host traffic terminates on and is processed by the router.
- CoPP, port-filtering and per-protocol queue thresholding protection features can be applied on the control plane host subinterface.
- Most CPPr features and policies operate strictly on the control plane host subinterface. Because most critical router control plane services, such as routing protocols and management traffic, are received on the control plane host subinterface, it is critical to protect this traffic through policing and protection policies.

Note Non-IP-based Layer 2 protocol packets, such as Address Resolution Protocol (ARP) or Cisco Discovery Protocol do not fall within the control plane host subinterface. These packets are currently classified in the control plane CEF-exception subinterface traffic.

- **Control plane transit subinterface**: Receives all control plane IP traffic that is software-switched by the RP
  - This includes packets that are not directly destined to the router itself but rather traffic traversing through the router.
  - An example of this type of control plane traffic would be nonterminating tunnels handled by the router.
  - CPPr allows specific aggregate policing of all traffic received at this subinterface.
- **Control plane CEF-exception subinterface:** Receives all traffic that is either redirected as a result of a configured input feature in the Cisco Express Forwarding packet forwarding path for process switching or directly enqueued in the control plane input queue by the interface driver
  - Examples of this include ARP, Layer 2 keepalives and all non-IP host traffic.
  - CPPr allows specific aggregate policing of this type of control plane traffic.

Port Filtering

This feature enhances control plane protection by providing for early dropping of packets directed toward closed or nonlistened Cisco IOS TCP and UDP ports on the router.

Note The port-filter policy feature can be applied only to the control plane host subinterface.

The port filter maintains a global database of all open TCP and UDP ports on the router, including random ephemeral ports created by applications. The port database is dynamically populated with entries provided by the registered applications as they start listening on their advertised ports either by configuration of an application (that is, SNMP) or initiation of an application (that is, TFTP transfer). An MQC class map using the list of open ports can be configured, and a simple drop policy can be applied to drop all packets destined to closed or nonlistened ports. Port-filtering class maps also support direct match of any user-configured TCP or UDP port numbers.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Queue Thresholding

The queue-thresholding feature provides the ability to limit the number of unprocessed packets that a protocol can have at the process level.

Note The queue-thresholding feature can only be applied to the control plane host subinterface.

This feature is designed to prevent the input queue from being overwhelmed by any single protocol traffic. Per-protocol thresholding follows a protocol charge model. The queue usage of each protocol is limited such that no single misbehaving protocol process can jam the interface hold queue. At this time—Cisco IOS Release 12.4(4)T—only a subset of TCP or UDP protocols can be configured for thresholding. Non-IP and Layer 2 protocols such as ARP and Cisco Discovery Protocol cannot be configured.

You can set queue limits for the following protocols:

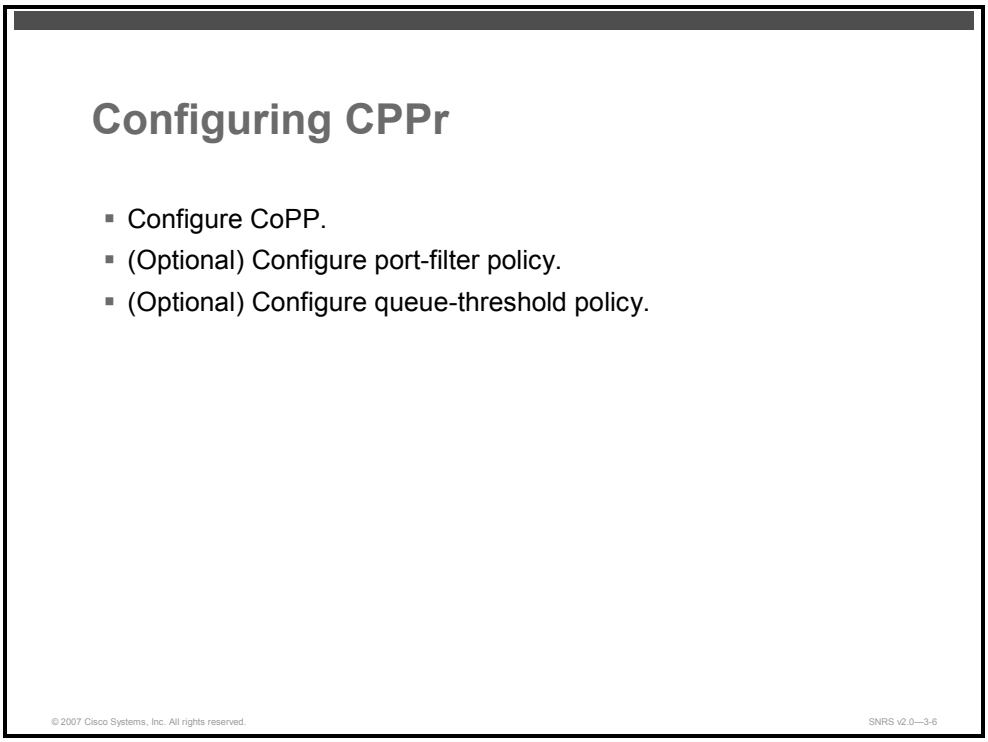
- SSH
- Telnet
- HTTP
- SNMP
- TFTP
- FTP
- BGP
- Domain Name System (DNS)
- Syslog
- Internet Group Management Protocol (IGMP)
- Host protocols (This is a wild card for all TCP or UDP protocol ports open on the router not specifically matched or configured.)

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Configuring CPPr

This topic describes the required and optional steps used to configure CPPr.



The CLI for control plane has been extended to allow for CoPP policies to be applied to individual control plane subinterfaces. The command syntax for creating CoPP service policies remains the same. In addition, the MQC class map and policy map CLI was modified to allow for additional types. The port-filter and queue-threshold policy features available in the host subinterface use these new class map and policy map types.

CoPP leverages MQC to define traffic classification criteria and to specify configurable policy actions for the classified traffic. First, traffic of interest must be identified via class maps, which are used to define packets for a particular traffic class. Once you have classified traffic, you will create enforceable policy actions for the identified traffic using policy maps. The **control-plane** global command allows the control plane service policies to be attached to the aggregate control plane interface itself.

Note	The CLI for configuring CoPP policies on the new control plane subinterfaces remains basically the same as the CLI introduced for CoPP. The only difference is in how you apply (attach) the CoPP policy to the different control plane subinterfaces.

To configure CPPr, follow these steps:

Step 1 Configure CoPP.

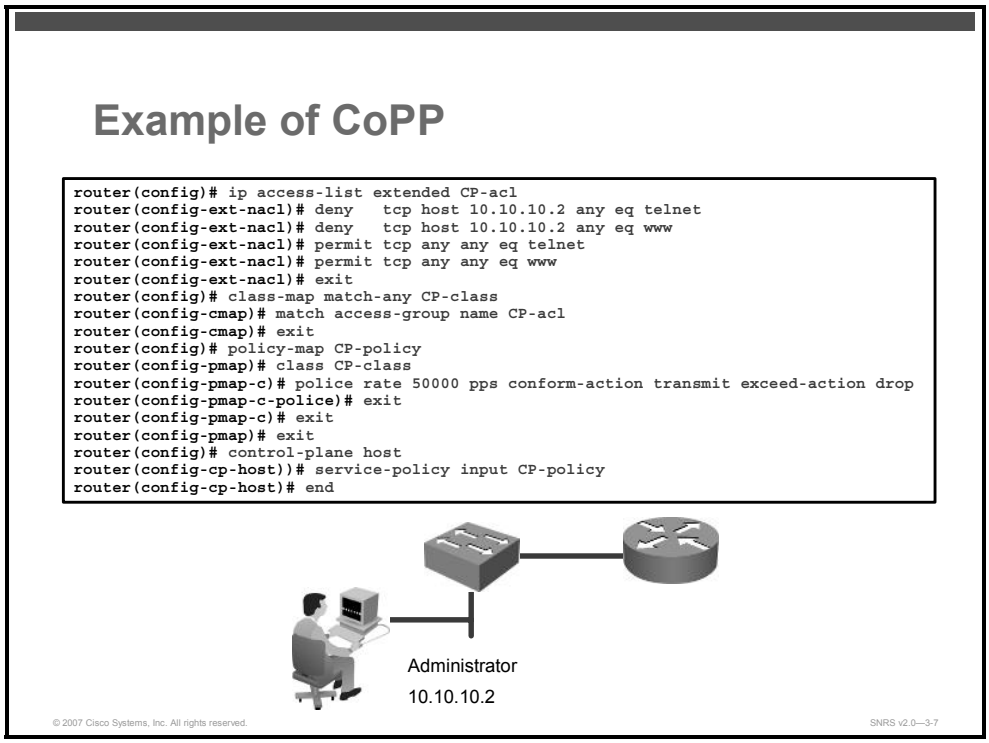
Step 2 (Optional) Configure a port-filter policy.

Step 3 (Optional) Configure a queue-threshold policy.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Configuring CoPP

This section describes how to configure the CoPP portion of CPPr.



The figure gives an example of CoPP configuration applied to the control plane host subinterface.

The example shows how to configure a trusted host with source address 10.10.10.2 to forward Telnet and HTTP packets to the control plane without constraint, while allowing all remaining Telnet and HTTP packets to be policed at 50,000 packets per second (pps).

In the example above, note the following:

- The **CP-acl** IP access list defines the traffic.
- The **CP-class** class map points to the access list.
- The **CP-policy** policy map does the following:
  - Points to the class map
  - Defines the policing rate
  - Defines action statements
- The control plane host subinterface has an input service policy attached that points to the **CP-policy** policy map.

Defining Packet Classification Criteria for CoPP

You must first create the policy using MQC to define a class map and policy map for control plane traffic.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Follow these steps to define a class map:

Step 1 Define an access list of trusted hosts using specific protocols to access the router.

```
router(config)# ip access list extended access-group-name
router(config-ext-nacl)# deny tcp host trusted-host any eq protocol
router(config-ext-nacl)# permit tcp any any eq protocol
```

Step 2 Enable the **class-map** global configuration command mode.

```
router(config)# class-map [match-any | match-all] class-map-name
```

Syntax Description

match-all | **match-any**	(Optional) This command determines how packets are evaluated when multiple match criteria exist. Packets must either meet all of the match criteria (**match-all**) or one of the match criteria (**match-any**) to be considered a member of the class.
class-map-name	Name of the class for the class map; the name can be a maximum of 40 alphanumeric characters. The class name is used for both the class map and to configure a policy for the class in the policy map.

Use this command to specify the name of the class for which you want to create or modify class map match criteria. The **class-map** command enters class map configuration mode in which you can enter one of the **match** commands to configure the match criteria for this class. Packets that arrive at either the input or output interface (determined by how the **service-policy** command is configured) are checked against the match criteria configured for a class map to determine if the packet belongs to that class.

When configuring a class map, you can use one or more **match** commands to specify match criteria. For example, you can use the **match access-group** command, the **match protocol** command, or the **match input-interface** command. The **match** commands vary according to the Cisco IOS Software release.

Step 3 Specify criteria to match.

```
router(config-cmap)# match {access-group | name access-group-name}
```

Syntax Description

access-group-name	Specifies an access list to match

Restrictions

Here are some of the restrictions of CoPP:

- The CoPP feature requires the MQC to configure packet classification and policing. Thus, restrictions that apply to MQC also apply to CoPP.
- Only the following classification (match) criteria are supported: standard and extended IP ACLs (named or numbered) and the **match ip dscp** command, the **match ip precedence** command, and the **match protocol arp** command.
- The CoPP CLI does not support **type** extensions available with other protection features. This is to preserve backward compatibility.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Defining a CoPP Service Policy

Use the **policy-map** global configuration command to specify the service policy name, and use the configuration commands to associate a traffic class that was configured with the **class-map** command.

The traffic class is associated with the service policy when you use the **class** command. You must then issue the **class** command after entering policy map configuration mode. After entering the **class** command, you are automatically in policy map class configuration mode.

Follow these steps to define a policy map:

Step 1 Enter policy map configuration mode to define a policy.

```
router(config)# policy-map policy-map-name
```

Syntax Description

policy-map-name	Name of the policy map; the name can be a maximum of 40 alphanumeric characters.

Use the **policy-map** command to specify the name of the policy map to be created, added to, or modified before you configure policies for classes whose match criteria are defined in a class map. The **policy-map** command enters QoS policy map configuration mode in which you can configure or modify the class policies for that policy map.

You can configure class policies in a policy map only if the classes have match criteria defined for them. You use the **class-map** and **match** commands to configure the match criteria for a class. Because you can configure a maximum of 64 class maps, no policy map can contain more than 64 class policies.

A single policy map can be attached to multiple interfaces concurrently. When you attempt to attach a policy map to an interface, the attempt is denied if the available bandwidth on the interface cannot accommodate the total bandwidth requested by class policies comprising the policy map. In this case, if the policy map is already attached to other interfaces, it is removed from them.

Whenever you modify class policy in an attached policy map, class-based weighted fair queuing (CBWFQ) is notified and the new classes are installed as part of the policy map in the CBWFQ system.

Step 2 Enter class map configuration mode. This mode is used to associate a service policy with a class.

```
router(config-pmap)# class class-name
```

Syntax Description

class-name	Name of a service policy class; the name can be a maximum of 40 alphanumeric characters.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Policy Map Configuration Mode

Within a policy map, the **class** "policy-map" command can be used to specify the name of the class whose policy you want to create or change. First, the policy map must be identified.

To identify the policy map (and enter the required QoS policy map configuration mode), use the **policy-map** command before you use the **class** "policy-map" command. After you specify a policy map, you can configure policy for new classes or modify the policy for any existing classes in that policy map.

Class Characteristics

The class name that you specify in the policy map ties the characteristics for that class (that is, its policy) to the class map and its match criteria, as configured using the **class-map** command.

When you configure policy for a class and specify its bandwidth and attach the policy map to an interface, CBWFQ determines if the bandwidth requirement of the class can be satisfied. If so, CBWFQ allocates a queue for the bandwidth requirement.

When a class is removed, available bandwidth for the interface is incremented by the amount previously allocated to the class.

The maximum number of classes that you can configure for a router and, therefore, within a policy map is 64.

Step 3 Configure traffic policing. There are three distinct ways to configure policing:

- — Packets per second
- — Bits per second (bps)
- — Percent(percentage of bandwidth)

```
router(config-pmap-c)# police rate units pps [burst burst-in-packets
packets] [peak-rate peak-rate-in-pps pps] [peak-burst peak-burst-in-
packets packets] [conform-action action] [exceed-action action]
[violate-action action]
```

```
router(config-pmap-c)# police rate units bps [burst burst-in-bytes
bytes] [peak-rate peak-rate-in-bps bps] [peak-burst peak-burst-in-
bytes bytes] [conform-action action] [exceed-action action]
[violate-action action]
```

```
router(config-pmap-c)# police rate percent percentage [burst ms ms]
[peak-rate percent percentage] [peak-burst ms ms] [conform-action
action] [exceed-action action] [violate-action action]
```


The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Syntax Description

units	■ If the police rate is specified in packets per second, the valid values are from 1 to 2,000,000 pps. ■ If the police rate is specified in bits per second, the valid values are from 8,000 to 10,000,000,000 bps.
pps	Specifies that the rate at which traffic is policed is in packets per second.
burst *burst-in-packets* **packets**	(Optional) Specifies the burst rate that is used for policing traffic; valid values are from 1 to 512,000 packets.
peak-rate *peak-rate-in-pps* **pps**	(Optional) Specifies the peak information rate (PIR) that is used for policing traffic; valid values are from 1 to 512,000 packets.
peak-burst *peak-burst-in-packets packets*	(Optional) Specifies the peak burst value that is used for policing traffic; valid values are from 1 to 512,000 packets.
bps	Specifies that the rate at which traffic is policed is in bits per second.
burst *burst-in-bytes* **bytes**	(Optional) Specifies the burst rate that is used for policing traffic; valid values are from 1000 to 512,000,000 bits.
peak-rate *peak-rate-in-bps* **bps**	(Optional) Specifies the peak burst value that is used for the peak rate; valid values are from 1000 to 512,000,000 bits.
peak-burst *peak-burst-in-bytes* **bytes**	(Optional) Specifies the peak burst value that is used for policing traffic; valid values are from 1000 to 512,000,000 bits.
percent *percentage*	(Optional) Specifies the percentage of interface bandwidth that is used to determine the rate at which traffic is policed; valid values are from 1 to 100.
burst *ms* **ms**	(Optional) Specifies the burst rate that is used for policing traffic; valid values are from 1 to 2000 ms.
peak-rate percent *percentage*	(Optional) Specifies the percentage of the interface bandwidth that is used to determine the PIR; valid values are from 1 to 100.
peak-burst *ms* **ms**	(Optional) Specifies the peak burst rate that is used for policing traffic; valid values are from 1 to 2000 ms.
conform-action *action*	Action to take on packets that conform to the rate limit. Valid values include the following: ■ **drop** ■ **set-clp-transmit** ■ **set-discard-class-transmit** ■ **set-dscp-transmit** ■ **set-frde-transmit** ■ **set-mpls-exp-imposition-transmit** ■ **set-mpls-exp-topmost-transmit** ■ **set-prec-transmit** ■ **set-qos-transmit** ■ **transmit**

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

exceed-action *action*	Action to take on packets that exceed the normal burst rate limit. Valid values include the following: ■ **drop** ■ **set-clp-transmit** ■ **set-discard-class-transmit** ■ **set-dscp-transmit** ■ **set-frde-transmit** ■ **set-mpls-exp-imposition-transmit** ■ **set-mpls-exp-topmost-transmit** ■ **set-prec-transmit** ■ **set-qos-transmit** ■ **transmit**
violate-action *action*	Action to take on packets that exceed the excess burst rate limit. Valid values include the following: ■ **drop** ■ **set-clp-transmit** ■ **set-discard-class-transmit** ■ **set-dscp-transmit** ■ **set-frde-transmit** ■ **set-mpls-exp-imposition-transmit** ■ **set-mpls-exp-topmost-transmit** ■ **set-prec-transmit** ■ **set-qos-transmit** ■ **transmit**

Use the **police rate** command to limit traffic that is destined for the control plane on the basis of packets per second, bits per seconds, or a percentage of interface bandwidth.

If the **police rate** command is entered, but the rate is not specified, traffic that is destined for the control plane will be policed on the basis of bits per second.

Restrictions

Here are some restrictions for CoPP service policies.

- Platform-specific restrictions, if any, are checked when the service policy is applied to the control plane interface.
- The CoPP feature requires the modular QoS CLI (MQC) to configure packet classification and policing. Thus, restrictions that apply to MQC also apply to CoPP. Also, only two MQC actions are supported in policy maps: transmit and drop.
- Support for output policing is available only in Cisco IOS Release 12.3(4)T and later T-train releases, and only on the aggregate control plane interface. Only input policing is available on the new control plane host, transit, and Cisco Express Forwarding-exception subinterfaces.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Note	Output policing does not provide any performance benefits. It simply controls the information that is leaving the device.

Entering Aggregate Control Plane Configuration Mode

After you create a class of traffic and define the service policy for the control plane, you need to apply the policy to either the aggregate control plane interface or one of the subinterfaces.

After you enter the **control-plane** command, you can define aggregate CoPP policies for the RP. You can configure a service policy to police all traffic destined to the control plane from all line cards on the router (aggregate control plane services).

Note	Aggregate control plane services manage traffic destined for the control plane and received on the central switch engine from all line cards in the router.

CoPP includes enhanced control plane functionality. It provides a mechanism for early dropping of packets directed toward closed or nonlistened Cisco IOS TCP and UPD ports on the router. It also provides the ability to limit protocol queue usage such that no single misbehaving protocol process can wedge the control plane interface hold queue.

With this enhancement, you can classify control plane traffic into different categories of traffic. These categories are as follows:

- **Control plane host subinterface:** This is the subinterface that receives all control plane IP traffic that is directly destined for one of the router interfaces. Examples of control plane host IP traffic include tunnel termination traffic, management traffic, or routing protocols such as SSH, SNMP, BGP, OSPF, and EIGRP. All host traffic terminates on and is processed by the router. Most CPPr features and policies operate strictly on the control plane host subinterface. Because most critical router control plane services, such as routing protocols and management traffic, is received on the control plane host subinterface, it is critical to protect this traffic through policing and protection policies. CoPP, port-filtering, and per-protocol queue-thresholding protection features can be applied on the control plane host subinterface.
- **Control plane transit subinterface:** This is the subinterface that receives all control plane IP traffic that is software-switched by the RP. This means packets not directly destined to the router itself but rather traffic traversing through the router. Nonterminating tunnels handled by the router are an example of this type of control plane traffic. CPPr allows specific aggregate policing of all traffic received at this subinterface.
- **Control plane Cisco Express Forwarding-exception subinterface:** This is the subinterface that receives all traffic that is either redirected as a result of a configured input feature in the Cisco Express Forwarding packet forwarding path for process switching or directly enqueued in the control plane input queue by the interface driver (for example, ARP, Layer 2 keepalives and all non-IP host traffic). CPPr allows specific aggregate policing of this specific type of control plane traffic.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Follow these steps to enter aggregate control plane configuration mode and attach a policy to a subinterface:

Step 1 Enter control plane configuration mode to apply a CoPP, port-filter, or queue-threshold policy to police traffic destined for the control plane.

```
router(config)# control-plane [host|transit|cef-exception]
```

Syntax Description

host	(Optional) Applies policies to host control plane traffic
transit	(Optional) Applies policies to transit control plane traffic
cef-exception	(Optional) Applies policies to Cisco Express Forwarding-exception control plane traffic

After issuing the **control-plane** command, you should use the **service-policy** command to configure a QoS policy. This policy is attached to the control plane interface for aggregate control plane services, which can control the number or rate of packets that are going to the process level.

Restrictions

The following are some restrictions for the aggregate control plane configuration:

- Platform-specific restrictions, if any, are checked when the service policy is applied to the control plane interface.
- Support for output policing is available only in Cisco IOS Release 12.3(4)T and later T-train releases, and only on the aggregate control plane interface. Only input policing is available on the new control plane host, transit, and Cisco Express Forwarding-exception subinterfaces.

Applying a CoPP Service Policy to the Host Subinterface

This task allows you to apply a CoPP service policy to the control plane host subinterface.

Note Before you attach an existing QoS policy to the control plane, you must first create the policy by using MQC to define a class map and policy map for control plane traffic.

Perform these steps to apply a CoPP service policy to a control plane interface:

Step 1 Attach a policy map to a control plane for aggregate control plane services.

```
router(config-cp)# service-policy {input | output} policy-map-name
```


The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Syntax Description

input	This command applies the specified service policy to packets that are entering the control plane.
output	This command applies the specified service policy to packets that are exiting the control plane; also enables the router to silently discard packets.
policy-map-name	Name of a service policy map (created using the **policy-map** command) to be attached; the name can be a maximum of 40 alphanumeric characters.

Restrictions

The following are some restrictions when applying a CoPP Service Policy to a Host Subinterface:

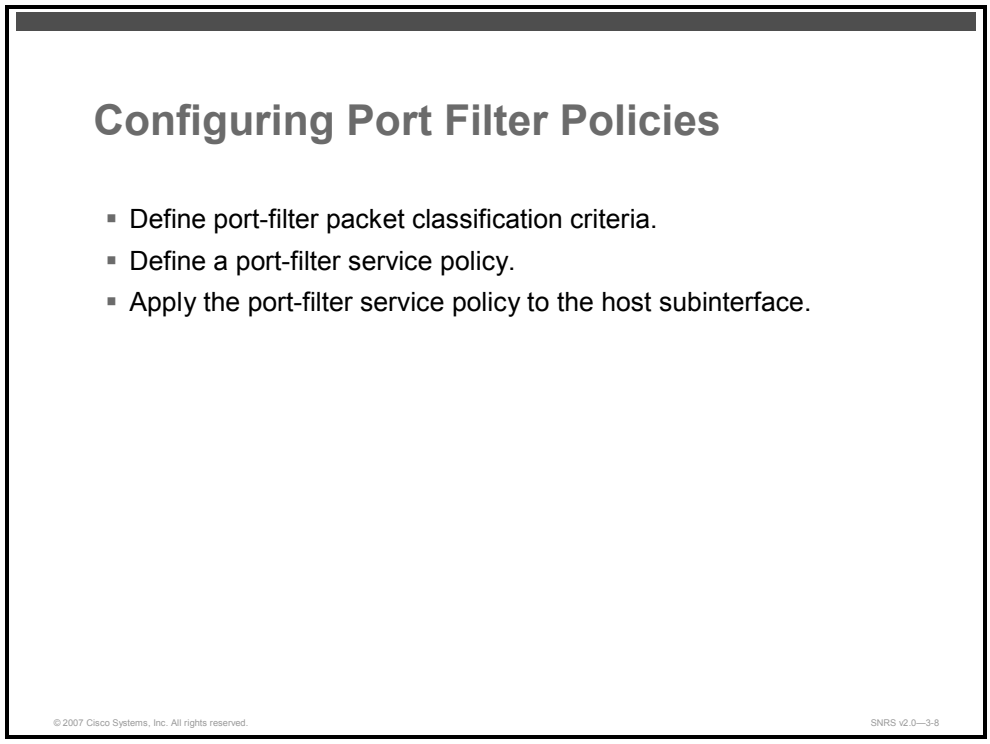
- Platform-specific restrictions, if any, are checked when the service policy is applied to the control plane interface.
- Support for output policing is available only in Cisco IOS Release 12.3(4)T and later T-train releases, and only on the aggregate control-plane interface. Only input policing is available on the new control plane host, transit, and Cisco Express Forwarding-exception subinterfaces.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Configuring a Port-Filter Policy

This topic describes how to configure a port-filter policy.



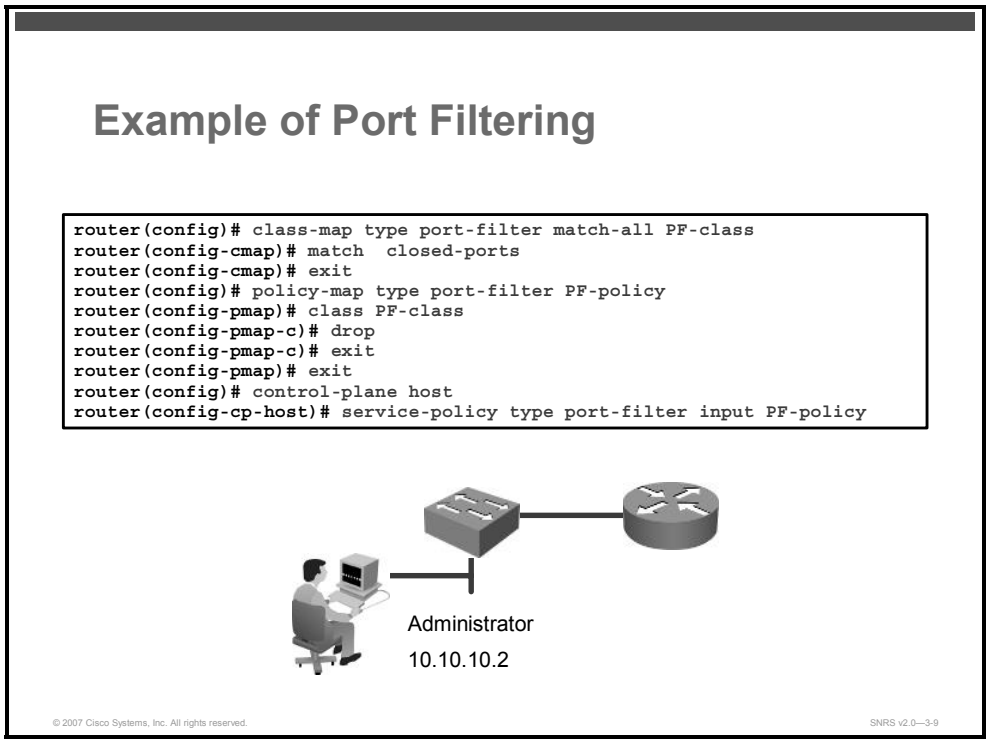
Apply the port-filter policy feature to the control plane host subinterface to block traffic destined to closed or nonlistened TCP and UDP ports.

New class map and service policy types have been created to accommodate the port-filter configuration. However, classification and match criteria for the new port-filter class maps support only a constrained subset of the overall global MQC match criteria. In addition, the actions supported by the new port-filter service policy is limited; only the drop action is supported.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Defining Port-Filter Parameters

This section covers the configuration of port-filter parameters.



This example shows how to configure a port-filter policy to drop all traffic destined to closed or nonlistened TCP or UDP ports.

Note that the PF-class class map matches all closed ports. In addition, the PF-policy policy map points to the class map and defines the action.

The control plane host subinterface has an input service policy attached that points to the PF-policy policy map.

Defining Port-Filter Packet Classification Criteria

Before you can attach a port-filter service policy to the control plane host subinterface, you must first create the policy to define a port-filter class map and policy map type for control plane traffic.

A new class map type called **port-filter** was created for the port-filter feature. You must first create one or more port-filter class maps before you can create your port-filter service policy. Your port-filter class maps will separate your traffic into classes of traffic. Then, your service policy will define actions on this traffic.

Follow these steps.

Step 1 Create a class map.

```
router(config)# class-map [type {stack | access-control |
port-filter | queue-threshold | logging log class}] [match-all
| match-any] class-map-name
```

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Syntax Description

type stack	(Optional) This command enables the Flexible Packet Matching (FPM) functionality to determine the correct protocol stack to examine. If the appropriate protocol header definition files (PHDFs) have been loaded onto the router (via the **load protocol** command), a stack of protocol headers can be defined so that the filter can determine which headers are present and in what order.
type access-control	(Optional) This command determines the exact pattern to look for in the protocol stack of interest. **Note** You must specify a stack class map (via the **type stack** keywords) before you can specify an access control class map (via the **type access-control** keywords).
type port-filter	(Optional) This command creates a port-filter class map that enables the TCP/UDP port policing of control plane packets. When enabled, it provides filtering of traffic destined to specific ports on the control plane host subinterface.
type queue-threshold	(Optional) This command enables queue thresholding that limits the total number of packets for a specified protocol that are allowed in the control plane IP input queue. This feature applies only to the control plane host subinterface.
type logging	(Optional) This command enables logging of packet traffic on the control plane.
log class	Name of the class for the log class; the name can be a maximum of 40 alphanumeric characters.
match-all | **match-any**	(Optional) This command determines how packets are evaluated when multiple match criteria exist. Packets must either meet all of the match criteria (**match-all**) or one of the match criteria (**match-any**) to be considered a member of the class.
class-map-name	Name of the class for the class map; the name can be a maximum of 40 alphanumeric characters. The class name is used for both the class map and to configure a policy for the class in the policy map.

Step 2 Specify the TCP or UDP match criteria for the class map.

```
router(config-cmap)# match {closed-ports|not|port} {TCP|UDP}
0-65535
```

Syntax Description

closed-ports	Matches automatically on all closed-ports on the router
port	Allows you to manually specify a TCP or UDP port to match on
TCP	Specifies a TCP port to match on
UDP	Specifies a UDP port to match on
0-65535	The port numbers to specify

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Defining a Port-Filter Service Policy

You can define a port-filter service policy that provides additional CPPr. Defining this policy supports early dropping of packets that are directed toward closed on nonlistened TCP/UDP ports on the router.

Complete these steps to configure a port-filter service policy. The port-filter traffic class is associated with the service policy when the **class** command is used. The **class** command must be issued after entering policy map configuration mode. After entering the **class** command, you are automatically in policy map class configuration mode.

Follow these steps.

Step 1 Create the port-filter service policy and enter the policy map configuration mode.

```
router(config)# policy-map type port-filter policy-map-name
```

Syntax Description

policy-map-name	Name of the policy map; the name can be a maximum of 40 alphanumeric characters.

Step 2 Associate a service policy with a class and enter class map configuration mode.

```
router(config-pmap)# class class-name
```

Syntax Description

class-name	Name of a service policy class; the name can be a maximum of 40 alphanumeric characters.

Step 3 Apply the port-filter service policy action on the class.

```
router(config-pmap-c)# drop
```

Restrictions

Only the drop action is supported.

Applying a Port-Filter Service Policy to the Host Subinterface

You are now ready to apply the port-filter policy to the host subinterface.

Follow these steps:

Step 1 Attach a QoS policy that manages traffic to the control plane host subinterface, and enter the control plane configuration mode.

```
router(config)# control-plane host
```

Syntax Description

host	Enters the control plane host subinterface configuration mode

Note Port-filter can only be applied to the host subinterface.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Step 2 Attach a QoS service policy to the control plane host subinterface.

```
router(config-cp)# service-policy type port-filter {input}
port-filter-policy-map-name
```

Syntax Description

input	This command applies the specified service policy to packets received on the control plane.
port-filter-policy-map-name	Name of a port-filter service policy map (created using the **policy-map type port-filter** command) to be attached; the name can be a maximum of 40 alphanumeric characters.

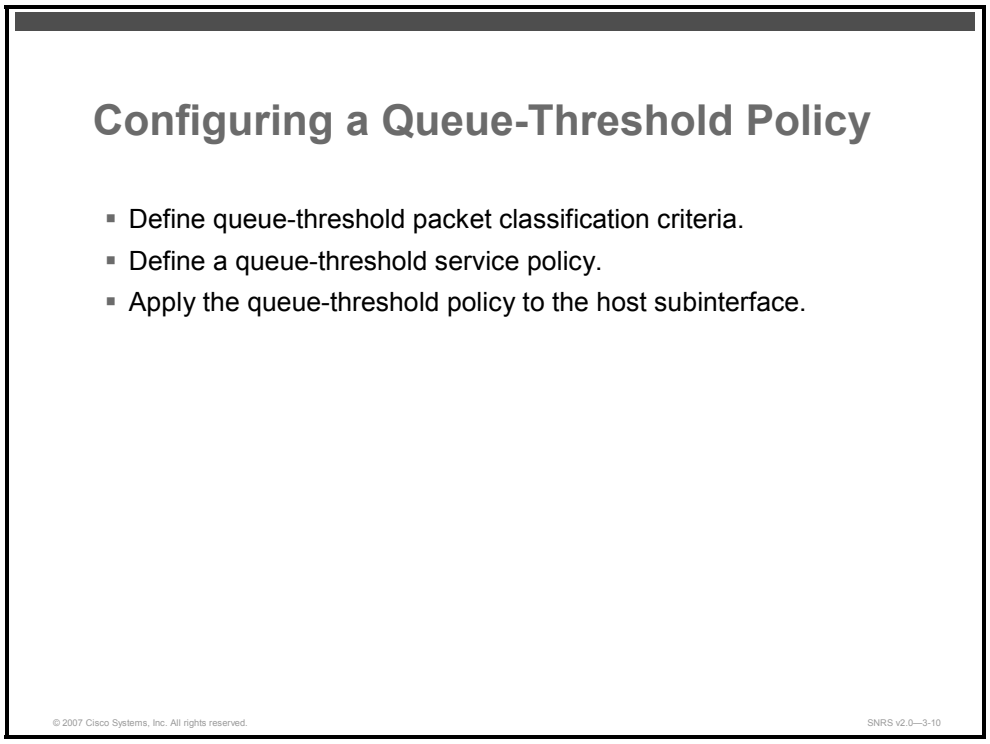
Restrictions

The port-filter feature can only be applied on the control plane host subinterface and only as an input policy.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Configuring a Queue-Threshold Policy

This topic describes how to configure a queue-threshold policy.



You can define a queue-threshold service policy when you want to limit the number of unprocessed packets that a protocol can have at the process level.

A new queue-threshold policy feature is included with CPPr that can be applied to the control plane host subinterface. This feature allows you to limit the number of packets for a given higher-level protocol allowed in the control plane IP input queue. Just as with the port-filter feature, new class map and policy map types have been created for the queue-threshold feature. As with the port-filter feature, the queue-threshold feature supports very specific class map and policy map capabilities.

Follow these three steps to configure a queue-threshold policy:

Step 1 Define a queue-threshold packet classification.

Step 2 Define a queue-threshold service policy.

Step 3 Apply the queue-threshold policy to the host sub interface.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

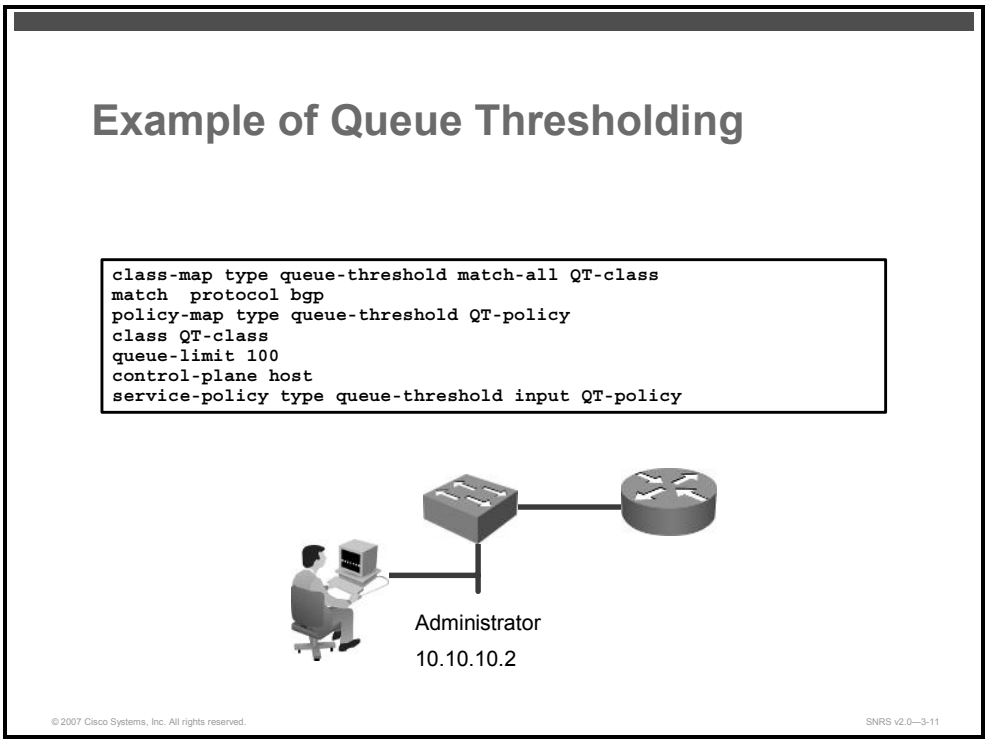
Restrictions

Here are the restrictions on configuring a queue-threshold policy:

- The classification and match criteria for the new queue-threshold class maps support only a constrained subset of the overall global MQC match criteria—only a subset of the match protocol option.
- Only the queue-limit action is supported.
- A queue-threshold policy is supported only on the control plane host subinterface as an input policy.

Configuring Queue Thresholding

Here is an example of a queue-thresholding configuration on a Cisco router**.**



This example shows how to configure a queue-threshold policy to set the queue limit for BGP protocol traffic to 100.

Take note of the following points:

- The **QT-class** class map matches protocol bgp.
- The **QT-policy** policy map points to the class map and sets the queue limit.
- The control plane host subinterface has an input service policy attached that points to the **QT-policy** policy map.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Defining Queue-Threshold Packet Classification Criteria

You must first create a policy to define a queue-threshold class map and policy map type for control plane traffic before you can attach a queue-threshold service policy to the control plane host subinterface.

A new class map type called queue-threshold was created for the queue-threshold feature. You must first create one or more queue-threshold class maps before you can create your queue-threshold service policy.

Note The queue-threshold class maps will separate the traffic into classes of traffic upon which the service policy will define actions.

Follow these steps to define queue-threshold packet classification criteria:

Step 1 Apply a class map for the queue threshold and enter the queue-threshold class map configuration mode.

```
router(config)# class-map type queue-threshold [match-all | match-any] class-name
```

Syntax Description

match-all	This command performs a logical AND on the match criteria.
match-any	This command performs a logical OR on the match criteria.
class-name	Name of a service policy class; the name can be a maximum of 40 alphanumeric characters.

Step 2 Specify the upper-layer protocol (ULP) match criteria for the class map.

```
router(config-cmap)# match protocol [bgp | dns | ftp | http | igmp | snmp | ssh | syslog | telnet| tftp | host-protocols]
```

Syntax Description

bgp	BGP
dns	DNS lookup
ftp	Stands for “File Transfer Protocol”
http	World Wide Web traffic
igmp	IGMP
snmp	SNMP
ssh	SSH
syslog	Syslog server
telnet	Telnet
tftp	Stands for “Trivial File Transfer Protocol”
host-protocols	Any open TCP or UDP port on the router

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Restrictions

Only a subset of the match protocol criteria is supported.

Defining a Queue-Threshold Service Policy

Use the new **policy-map type queue-threshold** global configuration command to configure a queue-threshold service policy. Use this command to specify the queue-threshold service policy name, and use other configuration commands to associate a queue-threshold traffic class that was configured with the **class-map type queue-threshold** command, with the **queue-threshold queue-limit action** command. The **class** command must be issued after entering policy map configuration mode. After entering the **class** command, you are automatically in policy map class configuration mode.

Note	The queue-threshold traffic class is associated with the service policy when the **class** command is used.

Follow these steps to define a queue-threshold service policy:

Step 1 Enter the queue-threshold service policy configuration mode.

```
router(config)# policy-map type queue-threshold policy-name
```

Syntax Description

policy-name	Name of a service policy map; the name can be a maximum of 40 alphanumeric characters.

Step 2 Enter class map configuration mode to associate a class map with a service policy.

```
router(config-pmap)# class class-name
```

Syntax Description

class-name	Name of a service policy class; the name can be a maximum of 40 alphanumeric characters.

Step 3 Apply the queue-threshold service policy action on the class.

```
router(config-pmap-c)# queue-limit number
```

Syntax Description

number	The range is from 0 to 255.

Restrictions

Only the queue-limit action is supported when you define a queue-threshold service policy.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Applying a Queue-Threshold Policy to the Host Subinterface

Before you can attach a queue-threshold service policy to the control plane host subinterface, you must first create the policy that defines a class map and policy map for the required control plane traffic.

Follow these steps to apply queue-threshold service policies to the control plane host subinterface:

Step 1 Enter global configuration mode.

```
router# configure terminal
```

Step 2 Attach a queue-threshold policy to the host subinterface and enter control plane configuration mode.

```
router(config)# control-plane host
```

Note Queue thresholding can only be applied to the host subinterface.

Step 3 Attach the service policy to the control plane.

```
router(config-cp-host)# service-policy type queue-threshold
{input} queue-threshold-policy-map-name
```

Syntax Description

input	**This keyword** applies the specified service policy to packets received on the control plane.
queue-threshold-policy-map-name	The name of a queue-threshold service policy map to be attached; the name can be a maximum of 40 alphanumeric characters.

Restrictions

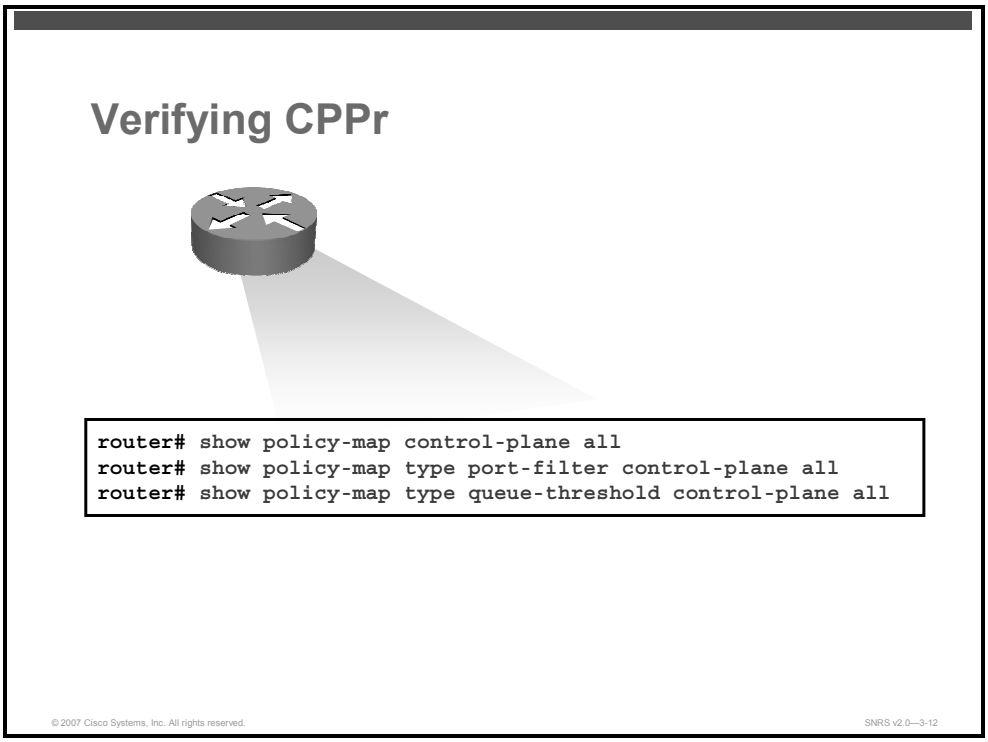
The queue-threshold feature can only be applied on the control plane host subinterface as an input policy.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Verifying CPPr

This topic describes how to verify CPPr on a Cisco router.



Use the **show policy-map control-plane** command and the **show policy-map type** command to verify CPPr configurations and to view statistics for control plane service policies.

The actual syntax of the **show policy-map** command is as follows:

show policy-map [**type** *policy-type*] **control-plane** [**pfx** | **slot** *slot number*] [**all**] [**host** | **transit** | **cef-exception**] [{**input** | **output**} [**class** *class-name*]]

Syntax Description

type *policy-type*	(Optional) Specifies policy map type for which you want statistics (for example, port-filter or queue-threshold)
pfx	(Optional) Information for all control plane interfaces
slot *slot number*	(Optional) Policy type and class map statistics for slot-level aggregate
all	(Optional) Information for all control plane interfaces
host	(Optional) Policy type and class map statistics for the host subinterface
transit	(Optional) Policy type and class map statistics for the transit subinterface
cef-exception	(Optional) Policy type and class map statistics for the Cisco Express Forwarding-exception subinterface
input	(Optional) Displays statistics for the attached input policy
output	(Optional) Displays statistics for the attached output policy
class *class-name*	(Optional) Name of the class whose configuration and statistics are to be displayed

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Verifying CPPr (Cont.)

```
router# show policy-map control-plane all

Control Plane Host

  Service-policy input: CP-policy

    Class-map: CP-class (match-any)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: access-group name CP-acl
        0 packets, 0 bytes
        5 minute rate 0 bps
      police:
          rate 50000 pps, burst 12207 packets
        conformed 0 packets; actions:
          transmit
        exceeded 0 packets; actions:
          drop
        conformed 0 pps, exceed 0 pps

    Class-map: class-default (match-any)
      904 packets, 54312 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
```

© 2007 Cisco Systems, Inc. All rights reserved.

This example shows the output of the **show policy-map control-plane all** command..

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Verifying CPPr (Cont.)

```
router# show policy-map type port-filter control-plane all

      drop
 Control Plane Host

  Service-policy port-filter input: PF-policy

    Class-map: PF-class (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match:  closed-ports

    Class-map: class-default (match-any)
      1754 packets, 105357 bytes
      5 minute offered rate 1000 bps, drop rate 0 bps
      Match: any
```

© 2007 Cisco Systems, Inc. All rights reserved.

This example show the output of the **show policy-map type port-filter control-plane all** command..

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Verifying CPPr (Cont.)

```
router# show policy-map type queue-threshold control-plane all

     queue-limit 100
     queue-count 0     packets allowed/dropped 0/0
 Control Plane Host

  Service-policy queue-threshold input: QT-policy

    Class-map: QT-class (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match:  protocol bgp

    Class-map: class-default (match-any)
      378 packets, 22734 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
```

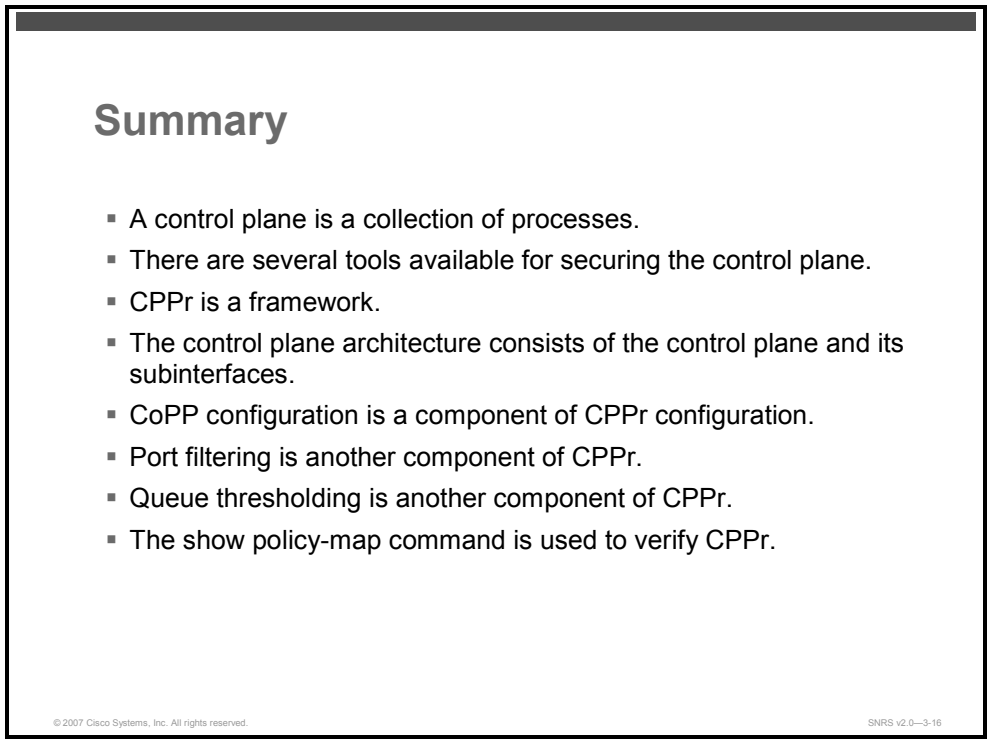
© 2007 Cisco Systems, Inc. All rights reserved.

This example show the output of the **show policy-map type queue-threshold control-plane all** command.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Summary

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



References

For additional information, refer to this resource:

- Cisco Systems, Inc. *Cisco IOS Software Releases 12.4T: Control Plane Protection*: http://www.cisco.com/en/US/partner/products/ps6441/products_feature_guide09186a00805 56710.html

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Lesson 3

Securing the Management Plane

Overview

The management plane performs all of the management functions for a device and coordinates functions among the other two planes (control and data). This makes the management plane a prime target for attacks. This lesson introduces you to some strategies to protect the management plane and how to configure the Cisco Management Plane Protection (MPP) feature on a network device.

Objectives

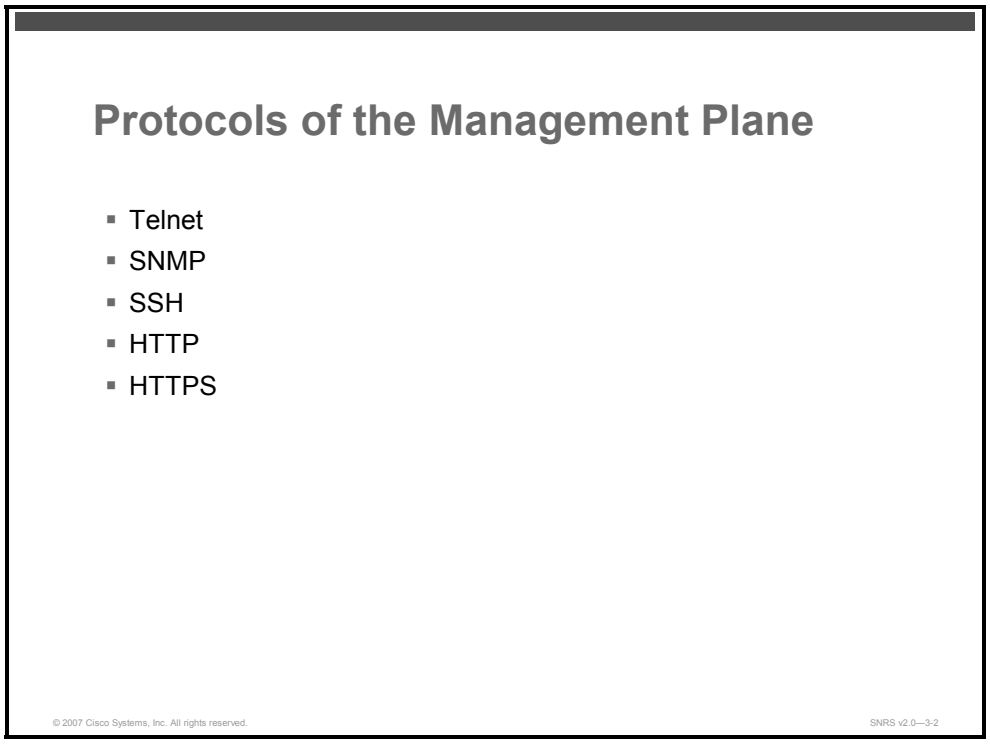
Upon completing this lesson, you will be able to configure the MPP feature in Cisco IOS Software. This ability includes being able to meet these objectives:

- Describe the management plane
- Describe some tools that are used in securing the management plane
- Describe the Cisco MPP feature
- Describe how to configure a device for MPP
- List the command that allows you to view all management interface configurations and activity on a device

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

The Management Plane

This topic describes the management plane.



The management plane performs management functions for a network and coordinates functions among all the planes (management, control, and data) in a network device. The management plane is also the logical path of all traffic related to the management of a routing platform and is used to manage a device through its network connection.

Examples of protocols processed in the management plane are as follows:

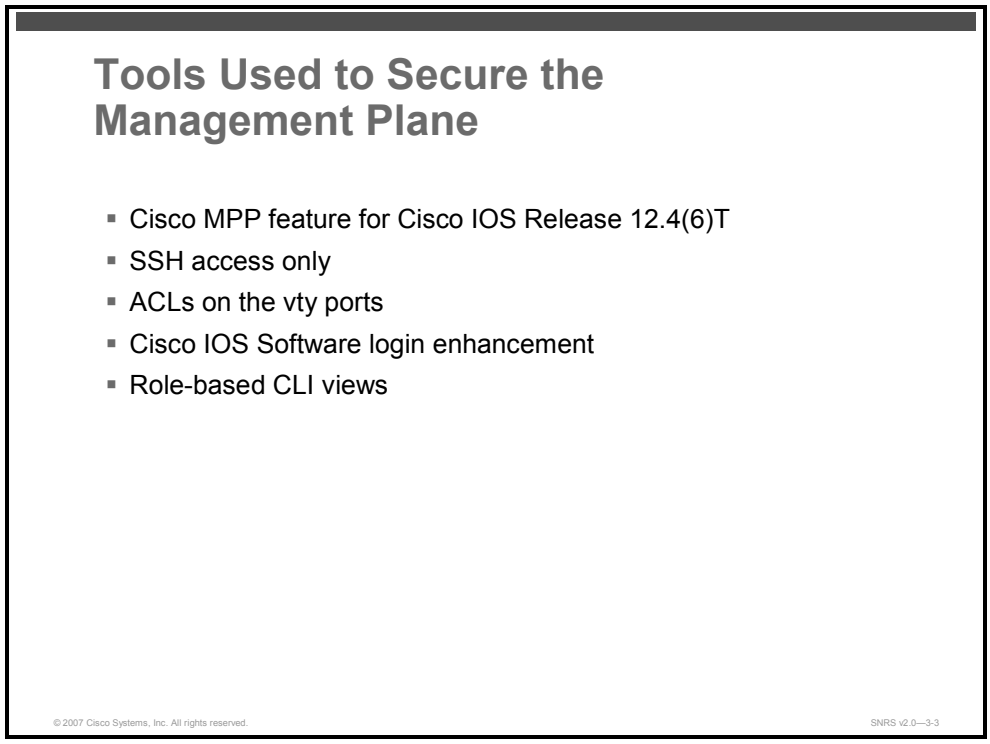
- Telnet
- Simple Network Management Protocol (SNMP)
- Secure Shell (SSH)
- HTTP
- Secure HTTP (HTTPS)

These management protocols are used for monitoring and for command-line access. What this means is that restricting access to these devices only by trusted sources (hosts or networks) is crucial for network security.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Tools for Securing the Management Plane

This topic describes some tools used in securing the management plane.



A network administrator needs more than one tool to help secure the network infrastructure. There are several tools available to secure the management plane, including the following:

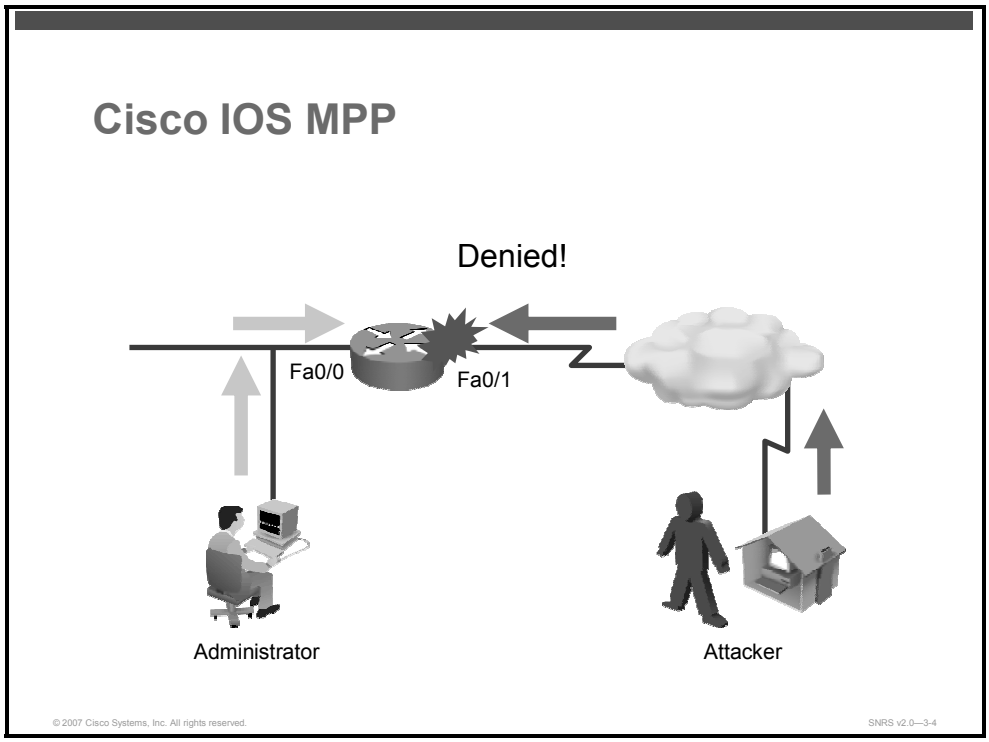
- Cisco MPP feature—Cisco IOS Release 12.4(6)T
- SSH access only to the device (covered in *Securing Cisco Network Devices* [SND] course)
- Access control lists (ACLs) on the vty ports (covered in SND)
- Cisco IOS Software login enhancement (covered in SND)
- Role-based command-line interface (CLI) (covered in SND)

This lesson concentrates on the Cisco MPP feature available in Cisco IOS Release 12.4(6)T.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Cisco MPP Feature

This topic describes the Cisco MPP feature.



The MPP feature in Cisco IOS Software provides the ability for an administrator to restrict the interface (or interfaces) on which network management packets are allowed to enter a device. The MPP feature allows a network administrator to designate one or more router interfaces as management interfaces. Device management traffic is permitted to enter a device only through these management interfaces. After MPP is enabled, no interfaces except designated management interfaces will accept network management traffic destined to the device.

When you configure a management interface, all incoming packets through that interface are dropped except for those from the allowed management protocols. This configuration also results in packets from all of the remaining management protocols (supported in the MPP feature) being dropped on all interfaces, including the interface that you configured. The allowed management protocols are dropped by all other interfaces unless the same protocol is enabled on those interfaces.

Designating management interfaces increases control over management of a device while providing more security for the same device.

Other benefits include the following:

- Improved performance for data packets on nonmanagement interfaces
- Need for fewer ACLs to restrict access to a device
- Management packet floods on switching and routing interfaces prevented from reaching the CPU
- Network scalability

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Prerequisites

Before a management interface can be configured, IP Cisco Express Forwarding must be enabled on the interface.

Restrictions

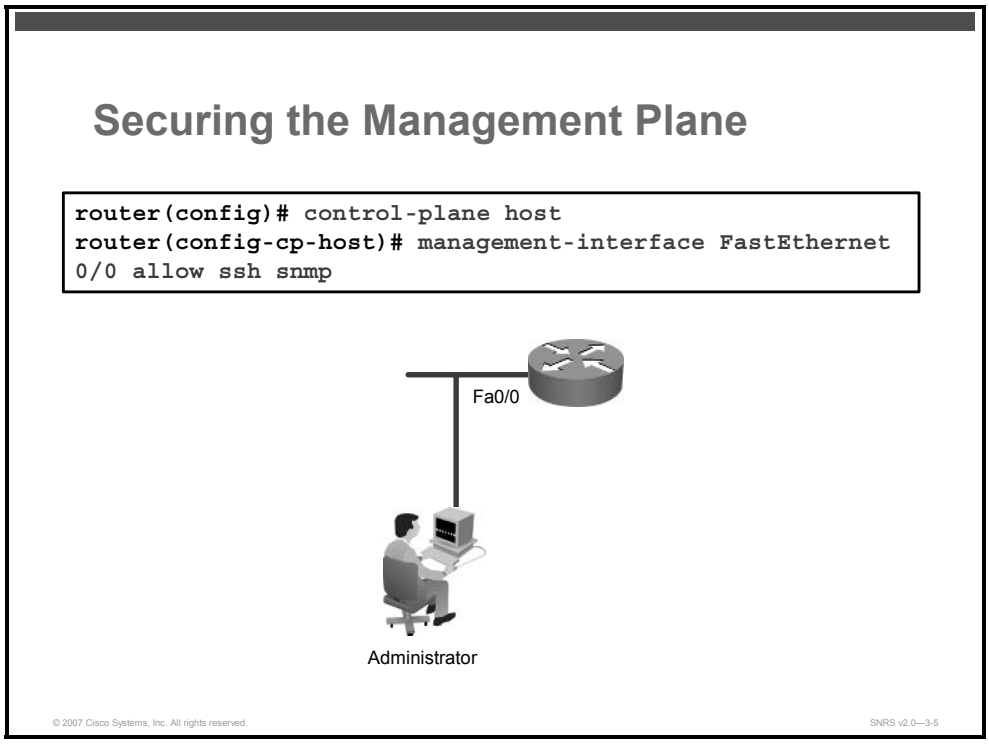
Here are some of the restrictions for MPP:

- Out-of-band management interfaces are not supported. (An out-of-band management interface is a dedicated Cisco IOS physical or logical interface that processes management traffic only.)
- Loopback and virtual interfaces not associated to physical interfaces are not supported.
- Fallback and standby management interfaces are not supported.
- Hardware-switched and distributed platforms are not supported.
- Secure Copy Protocol (SCP) is supported under SSH and is not directly configurable in the CLI.
- Uninformed management stations lose access to the router through nondesignated management interfaces when the MPP feature is enabled.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Securing the Management Plane

This topic describes how to configure a device for MPP.



By default, the MPP feature is disabled. When you enable the feature, you must follow these steps:

Step 1 Enter control plane host configuration mode.

Step 2 Designate one or more interfaces as management interfaces.

Step 3 Configure the management protocols that will be allowed on the management interfaces.

Note The MPP feature does not provide a default management interface.

The configuration in this example shows MPP configured to allow SSH and SNMP to access the router only through the Fast Ethernet 0/1 interface. This configuration results in all protocols in the remaining subset of supported management protocols being dropped on all interfaces unless explicitly permitted.

Configuring MPP

After entering control plane host configuration mode, you use one command to configure, modify, or delete a management interface. After you configure a management interface, no interfaces except that management interface will accept network management packets destined to the device.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Note	When the last configured interface is deleted, the MPP feature turns itself off.

Follow these steps to configure a network device:

Step 1 Enter control plane host configuration mode.

```
router(config)# control-plane host
```

Syntax Description

host	Applies policies to host control plane traffic.

Step 2 Configure an interface to be a management interface and specify which management protocols are allowed.

```
router(config-cp-host)# management-interface interface allow
protocols
```

Syntax Description

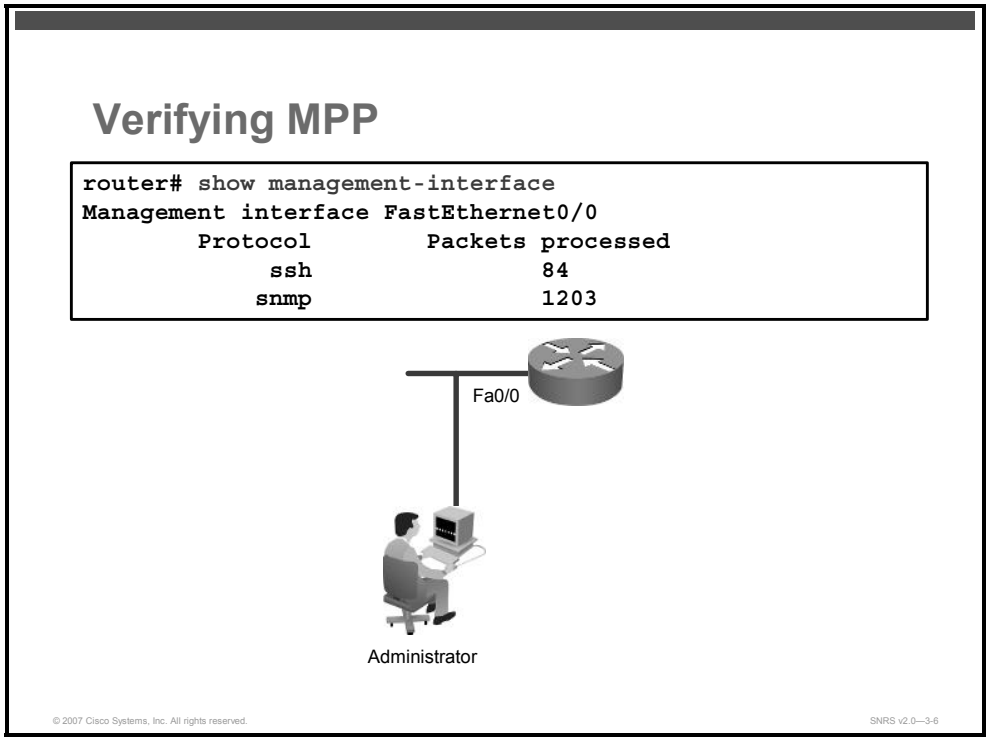
interface	Name of the interface that you are designating as a management interface
protocols	Management protocols you want to allow on the designated management interface Protocols can be one of the following: ■ Block Extensible Exchange Protocol (BEEP) ■ FTP ■ HTTP ■ HTTPS ■ SSH version 1 (SSHv1) and SSH version 2 (SSHv2) ■ SNMP (all versions) ■ Telnet ■ TFTP

The **management-interface allow** command is useful when you want to restrict access of management protocols to a device through a particular interface. An additional benefit of dedicated management interfaces is that they prevent management traffic floods on switching and routing interfaces from reaching the CPU.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Verifying MPP

This topic describes the command that allows you to view all management interface configurations and activity on a device.



To verify MPP and view information about the management interface, such as type of interface, protocols enabled on the interface, and number of packets dropped and processed, use the following command in privileged EXEC mode:

```
router(config)# show management-interface [ interface |
protocol protocol-name]
```

Syntax Description

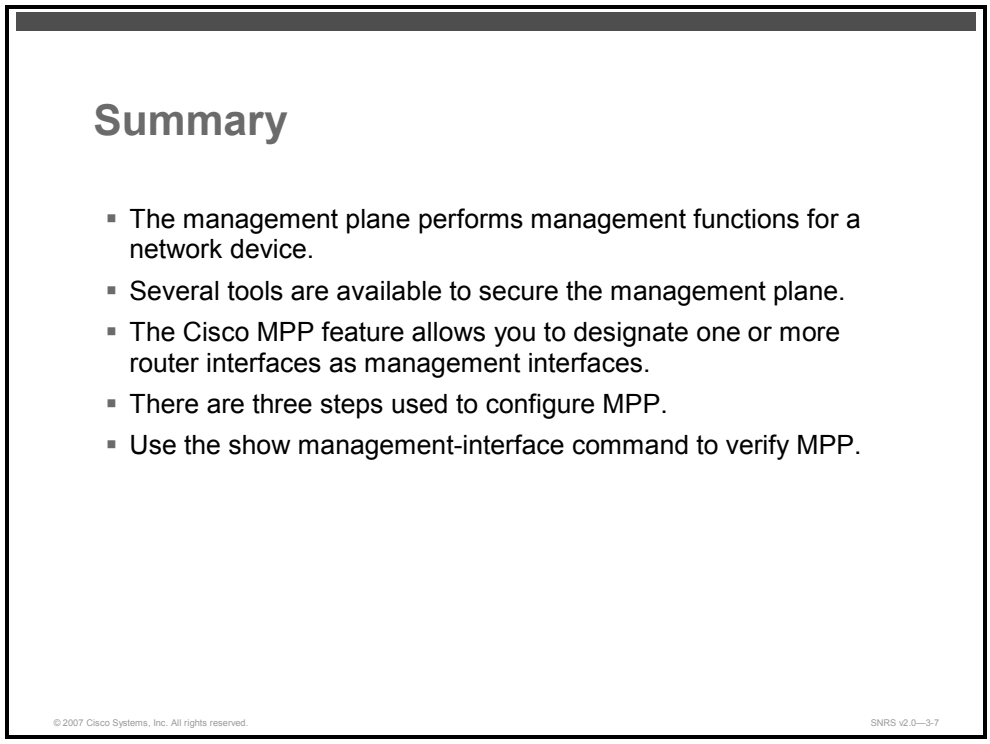
interface	(Optional) Interface for which you want to view information
protocol	(Optional) Indicates that a protocol is specified
protocol-name	(Optional) Protocol for which you want to view information

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Summary

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



References

For additional information, refer to this resource:

- Cisco Systems, Inc. Cisco IOS Software Releases 12.4T: Management Plane Protection. http://www.cisco.com/en/US/partner/products/ps6441/products_feature_guide09186a00806 17022.html.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Lesson 4

Securing the Data Plane

Overview

The data forwarding plane is another critical path through a network device. In this lesson, you will be introduced to some well-known network attacks that affect the data plane and some tools used to mitigate these types of attacks. You will also learn how to configure Flexible Packet Matching (FPM) as one of the tools used to secure the data plane.

Objectives

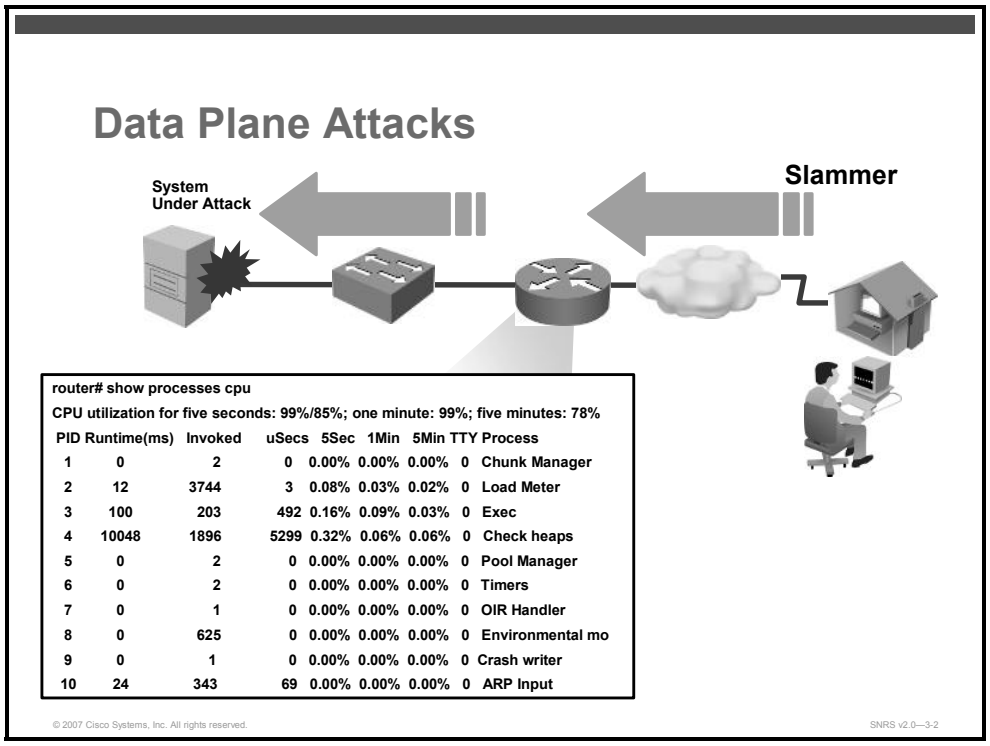
Upon completing this lesson, you will be able to describe how to protect the data plane using FPM. This ability includes being able to meet these objectives:

- Describe some data plane attacks and their effect on network devices
- Describe some strategies used to secure the data plane of a device
- Describe FPM
- Describe the commands used to configure FPM
- Describe some commands used to verify FPM configuration
- Describe some commands used to troubleshoot FPM

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Data Plane Attacks

This topic describes some data plane attacks and their effect on network devices.



Attacks against networking environments are increasing in frequency and sophistication. Attacks that affect the data forwarding plane include some well-known attacks with very specific signatures (fields within the IP packet that contain certain specific values). Here are some of these attacks:

- Code Red
- Nimda
- Nachi
- SQL Slammer
- Blaster
- SYN floods
- Frag attacks

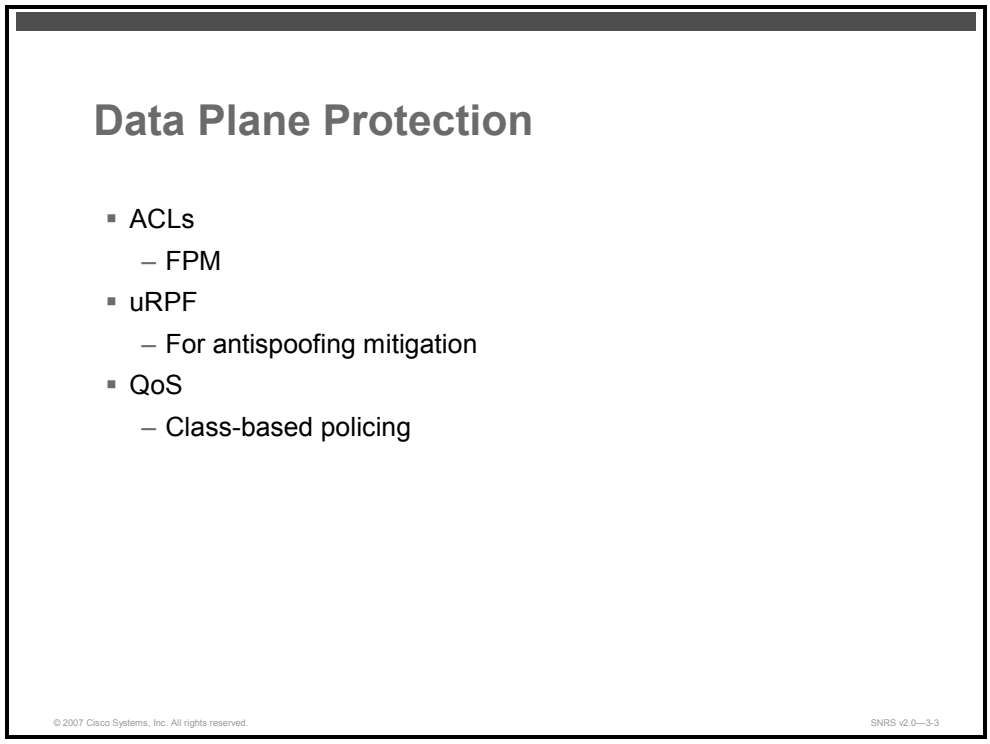
All of these attacks are known to overload the CPU of any router or switch in its path. To counter these attacks, features are needed that are as flexible as possible, both in terms of classification and mitigation capabilities.

HTTP vulnerabilities are some of the other types of attacks that are addressed by FPM.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Data Plane Protection

This topic describes some strategies for protecting the data plane.



Cisco IOS Software includes various tools for dealing with attacks that may affect the data plane. Some of these security features include the following:

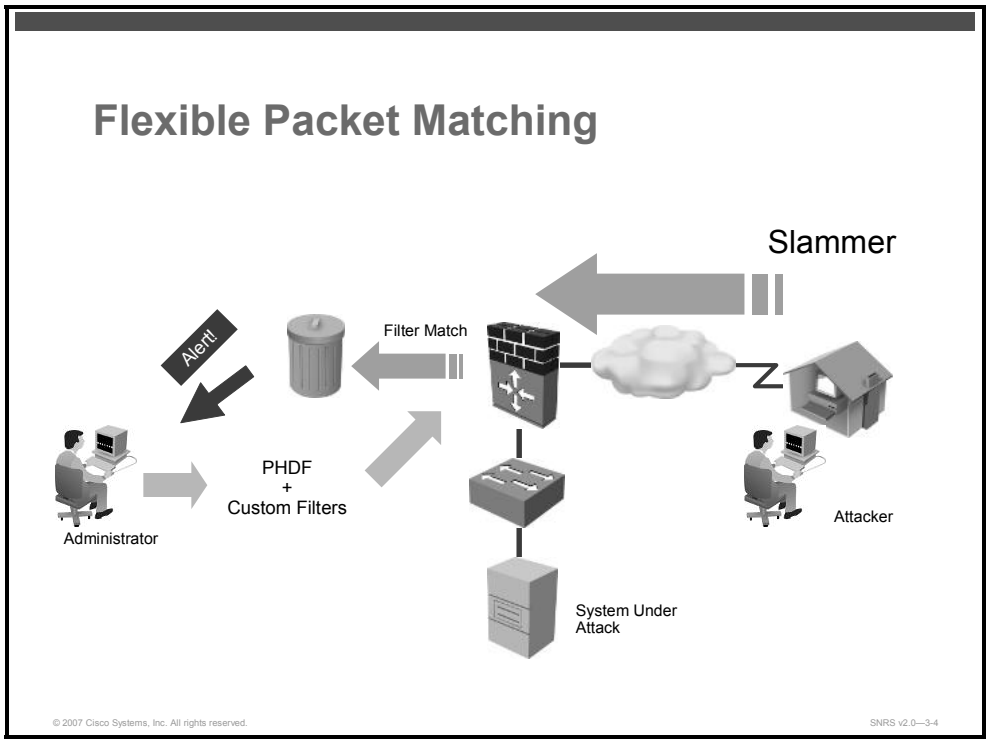
- **Access control lists (ACLs):** Filter traffic through network devices
- Flexible Packet Matching: Provides a flexible Layer 2 to Layer 7 stateless classification mechanism.
- **Unicast Reverse Path Forwarding (uRPF):** Helps mitigate problems caused by the introduction of malformed or forged (spoofed) IP source addresses into a network by discarding IP packets that lack a verifiable IP source address
- **Quality of service (QoS):** Committed rate limiting (class-based traffic policing)

This lesson concentrates on FPM as one tool that you can use to secure the data plane.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Flexible Packet Matching

This topic describes FPM.



Many of the tools available today are not designed with deep packet inspection as a requirement; instead, they are designed to provide matching for predefined fields in well-known protocol headers. If an attack uses a field outside the limited range of inspection of these features, you are left without a defense against the attack.

FPM provides the means to configure match criteria for any or all fields in a packet header and also bit patterns within the packet payload within the first 256 bytes (B). This allows the characteristics of an attack (source port, packet size, byte string) to be uniquely matched and for a designated action to be taken. This also provides the network administrator with the means to implement network-based blocking of known attack vectors.

FPM provides a flexible Layer 2 to Layer 7 stateless classification mechanism. The user can specify classification criteria based on any protocol and any field of the traffic protocol stack. Based on the classification result, actions such as drop or log can be taken.

FPM is useful because it enables users to create their own stateless packet classification criteria and to define policies with multiple actions (such as drop, log, or send “Internet Control Message Protocol [ICMP] unreachable” messages) to immediately block new viruses, worms, and attacks.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

The packet classification feature of FPM allows users to define one or more classes of network traffic by pairing a complete set of standard matching operators with user-defined protocol header fields. FPM further extends the network traffic class definition capability to include new command-line interface (CLI) syntax to offset into both a user-defined protocol header and the data portion of the packet. The offset or depth at which to begin matching can be referenced from several locations in the packet. Some of these locations are dependent upon loading a protocol header definition file (PHDF).

FPM can work with well-known, established protocols, such as IP, TCP, and User Datagram Protocol (UDP), or with custom protocols that are described with a user-defined PHDF. PHDFs are written in off-box Extensible Markup Language (XML).

Note PHDFs for these well-known and other protocols are available for download.

FPM is the next-generation ACL pattern matching tool for more thorough and customized packet filters. It removes the constraints to specific fields that previously had limited packet inspection.

FPM is provisioned using CLI and off-box XML. From the router CLI, FPM is configured using class maps to describe the traffic to be filtered, policy maps to define the action to be taken for filtered traffic, and service policies to attach the filter and action to an interface.

Deployment

FPM may be deployed anywhere that the ability to perform classification upon unique bit or byte patterns within IP packets can provide an effective attack mitigation strategy. FPM is not intended to replace an effective IDS/IPS deployment strategy. However, under circumstances where a unique packet classification scheme can be developed, and an intrusion prevention system (IPS) signature is not available (or an IPS is not deployed) and ACLs or firewalls (or both) cannot provide the appropriate responses, FPM may fulfill the required services.

Protocol Header Definition File

Protocol headers are defined in separate files called PHDFs. You define the packet filters using the field names that are defined within the PHDFs. A PHDF is a file that allows the user to leverage the flexibility of XML to describe almost any protocol header. The important components of the PHDF are the version, the XML file schema location, and the protocol field definitions. The protocol field definitions name the appropriate field in the protocol header, allow for a comment describing the field, provide the location of the protocol header field in the header (the offset is relative to the start of the protocol header), and provide the length of the field. Users can choose to specify the measurement in bytes or in bits.

A PHDF defines each field contained in a particular protocol header. Each field is defined with a name, optional comment, offset, and length. The offset is always specified from the beginning of the header. Both the offset field and the length field may be specified either in terms of bits or bytes.

Note The total length of the header must be specified at the end of each PHDF.

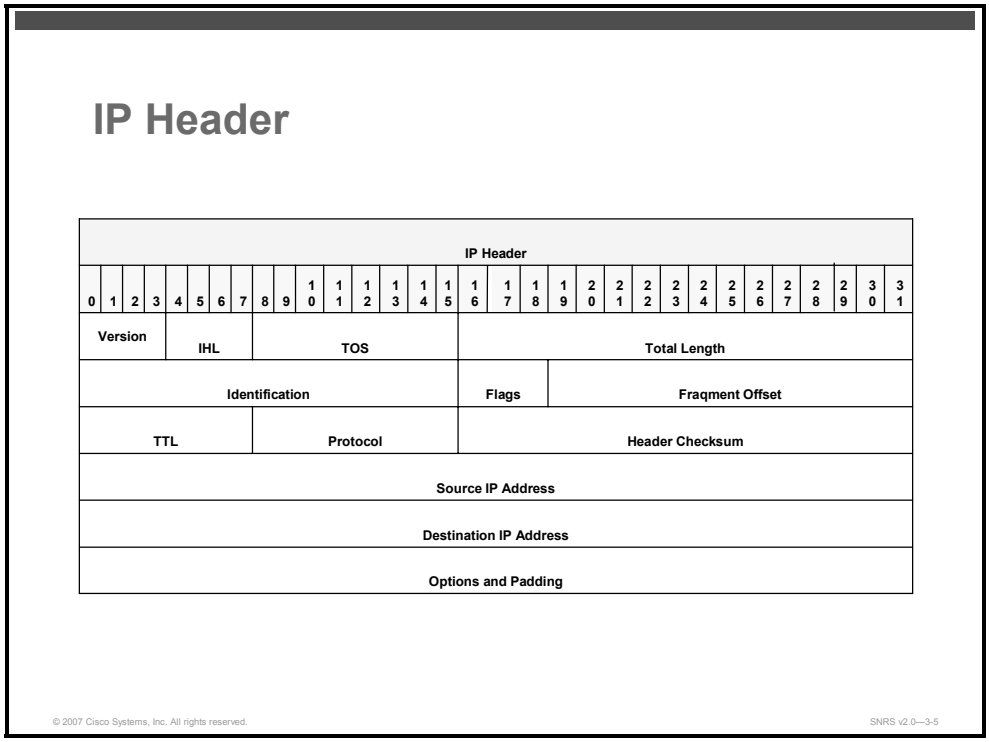
The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Users can write their own custom PHDFs via XML to provide the desired protocol definition through Layer 7 and use the following standard PHDFs that provide Layer 2 through Layer 4 protocol definition:

- ip.phdf
- tcp.phdf
- udp.phdf
- ether.phdf
- icmp.phdf

PHDFs are loaded onto the router using the **load protocol** command.

Note PHDFs cannot be seen directly within the running configuration because they are defined using XML and not CLI; however all PHDFs that are loaded may be displayed using the **show protocols phdf** <*name*> command. Because they can only be loaded locally, defined custom PHDFs may also be viewable from local flash memory.



This is the IP header format.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Filter Description

A filter description is a definition of a traffic class that can contain the header fields defined in a PHDF. If a PHDF is not loaded, the traffic class can be defined via the datagram header start (Layer 2) or the network header start (Layer 3). If a PHDF has been loaded onto the router, the class specification begins with a list of the protocol headers in the packet.

Prerequisites and Restrictions

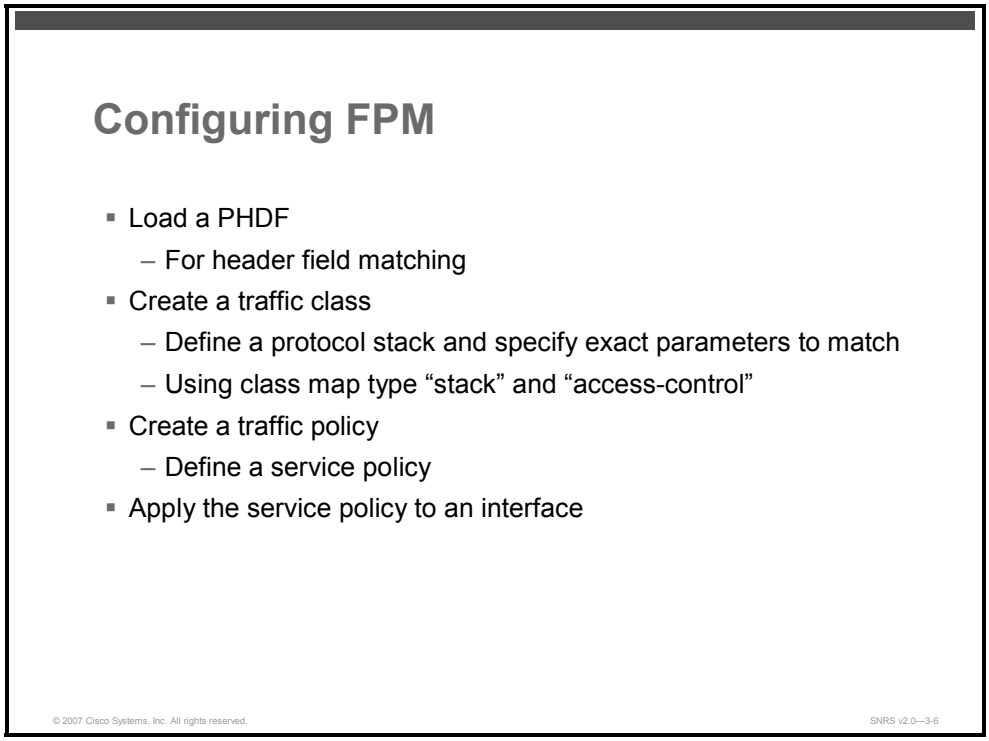
FPM has the following prerequisites and restrictions:

- FPM is available only in advanced security images.
- Although access to an XML editor is not required, XML will ease the creation of PHDFs.
- FPM cannot be used to mitigate an attack that requires stateful classification.
- Because FPM is stateless, it cannot keep track of port numbers being used by protocols that dynamically negotiate ports. Thus, when using FPM, port numbers must be explicitly specified.
- FPM cannot perform IP fragmentation or TCP flow reassembly.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Configuring FPM

This topic describes procedures that should be followed when configuring FPM within your network.



FPM allows customers to create their own filtering policies that can immediately detect and block new viruses and attacks.

The process for configuring FPM consists of four steps.

Step 1 Load a PHDF from flash memory.

Once the appropriate PHDFs are loaded, a **class-map** command with type **stack** must be defined so that FPM knows which headers are present and in which order.

Once the stack of protocols is defined, a class map of type **access-control** is defined for classifying packets.

Step 2 Create a traffic class by defining class maps.

A policy map is an ordered set of classes and associated actions. The policy binds the class and action. Actions can be drop, ICMP response, and log, or service-policy to nest another policy.

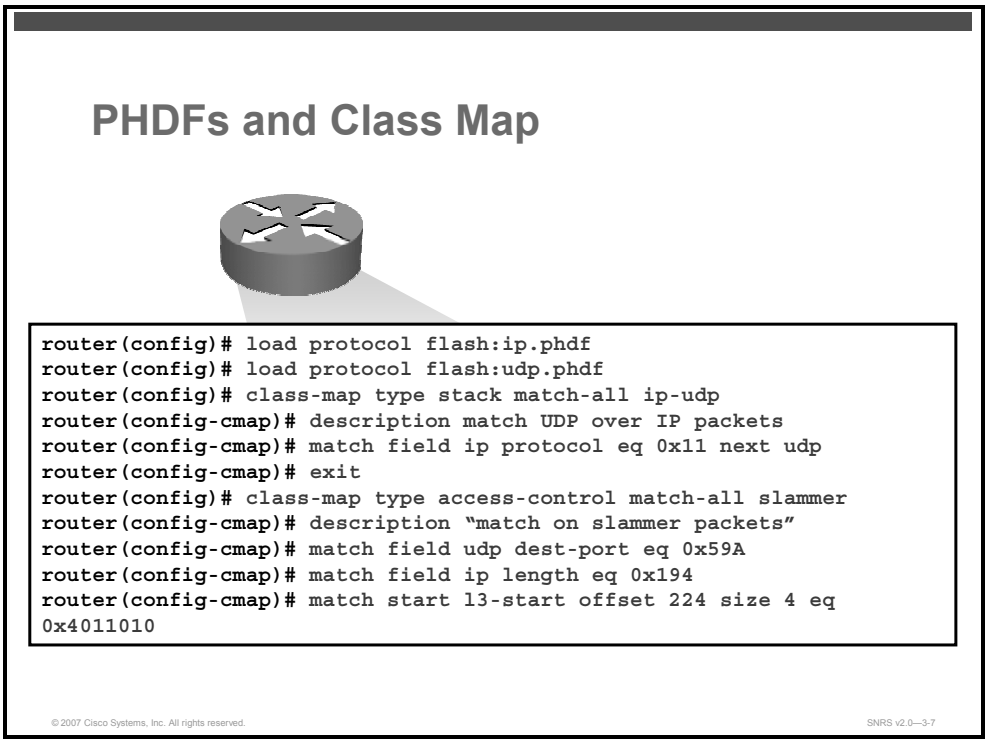
Step 3 Create a traffic policy by defining a service policy.

Step 4 Apply the service policy to an interface.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

PHDFs and Traffic Classes

This section covers how to load the PHDFs and create a class map to classify traffic.



This is an example of a filter policy used to mitigate the SQL Slammer attack. First, you load the PHDFs into the router. You then define a traffic class using class maps. The match criteria defined within the class maps is for slammer and UDP packets with an IP length not to exceed 404 B, UDP port 1434, and pattern 0x4011010 at 224 B from the start of the IP header.

Compete these steps to configure FPM.

Loading a PDHF

To load a PHDF onto a router, use the **load protocol** command in global configuration mode. To unload all protocols from a specified location or a single protocol, use the **no** form of this command.

Step 1 Load the PDHF on the router.

```
router(config)# load protocol location:filename
```

Syntax Description

location:filename	Location of the PHDF that is to be loaded onto the router When used with the **no** version of this command, all protocols loaded from the specified filename will be unloaded. **Note** The location must be local to the router.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

protocol-name	Unloads only the specified protocol **Note** If you attempt to unload a protocol that is being referenced by a filter, you will receive an error.

FPM allows users to classify traffic on the basis of any portion of a packet header given the protocol field, length, and pattern. Protocol headers are defined in separate files (PHDFs); the field names that are defined within the PHDFs are used for defining the packet filters. A PHDF is a file that allows the user to leverage the flexibility of XML to describe almost any protocol header. The important components of the PHDF are the version, the XML file schema location, and the protocol field definitions. The protocol field definitions name the appropriate field in the protocol header, allow for a comment describing the field, provide the location of the protocol header field in the header (the offset is relative to the start of the protocol header), and provide the length of the field. Users can choose to specify the measurement in bytes or in bits.

Creating a Traffic Class

In creating a traffic class, you will create stateless packet classification criteria that, when used in conjunction with an appropriately defined policy, can mitigate network attacks. Once the appropriate PHDFs are loaded, a stack of protocol headers must be defined so that FPM knows which headers are present and in which order. Once the stack of protocols is defined, a class map of type "access-control" is defined for classifying packets.

Step 2 Define the sequence of headers as IP first, then UDP by creating a class map of type "stack." In the example, we created a class map with a name of "ip-udp". Because you know that these packets are UDP packets, you want to restrict FPM to only look at UDP packets. Use the following commands in configuration mode to configure your 'stack' class maps:

Note The "**stack**" type class map is used for this purpose.

```
router(config)# class-map [type {stack | access-control |
port-filter | queue-threshold | logging log-class}] [match-all
| match-any] class-map-name
```

Syntax Description

type stack	(Optional) Enables the FPM functionality to determine the correct protocol stack to examine. If the appropriate PHDFs have been loaded onto the router (via the **load protocol** command), a stack of protocol headers can be defined so that the filter can determine which headers are present and in which order.
type access-control	(Optional) Determines the exact pattern to look for in the protocol stack of interest. **Note** You must specify a stack class map (via the **type stack** keywords) before you can specify an access control class map (via the **type access-control** keywords).

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

type port-filter	(Optional) Creates a port-filter class map that enables the TCP/UDP port policing of control plane packets. When enabled, it provides filtering of traffic destined to specific ports on the control plane host subinterface.
type queue-threshold	(Optional) Enables queue thresholding that limits the total number of packets for a specified protocol that are allowed in the control plane IP input queue. This feature applies only to the control plane host subinterface.
type logging	(Optional) Enables logging of packet traffic on the control plane.
Log-class	Name of the class for the log class. The name can be a maximum of 40 alphanumeric characters.
match-all | **match-any**	(Optional) Determines how packets are evaluated when multiple match criteria exist. Packets must either meet all of the match criteria (**match-all**) or one of the match criteria (**match-any**) to be considered a member of the class.
class-map-name	Name of the class for the class map. The name can be a maximum of 40 alphanumeric characters. The class name is used for both the class map and to configure a policy for the class in the policy map.

Use this command to specify the name of the class for which you want to create or modify class map match criteria. The **class-map** command enters class map configuration mode in which you can enter one of the **match** commands to configure the match criteria for this class. Packets that arrive at either the input or output interface (determined by how the **service-policy** command is configured) are checked against the match criteria configured for a class map to determine if the packet belongs to that class.

When configuring a class map, you can use one or more **match** commands to specify match criteria. For example, you can use the **match access-group** command, the **match protocol** command, or the **match input-interface** command. The **match** commands vary according to the Cisco IOS Software release.

Step 3 Add a description to the class map. In this class map, you will be matching UDP packets.

```
router(config-cmap)# description character-string
```

Syntax Description

character-string	Description that is added to the class map or the policy map. The character-string cannot exceed 161 characters.

Step 4 Create the match criteria. UDP is IP protocol number 17, which is equal to 11 in hexadecimal format.

```
router(config-cmap)# match field protocol protocol-field {eq
[mask] | neq [mask] | gt | lt | range range | regex string}
value [next next-protocol]
```

Syntax Description

protocol	Name of protocol whose PHDF has been loaded onto a router.
Protocol-field	Match criteria is based upon the specified field within the loaded protocol.
eq	Match criteria is met if the packet is equal to the specified value or mask.
neq	Match criteria is met if the packet is not equal to the specified value or mask.
mask *mask*	(Optional) Can be used when the **eq** or the **neq** keywords are issued.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

gt	Match criteria is met if the packet does not exceed the specified value.
lt	Match criteria is met if the packet is less than the specified value.
range *range*	Match criteria is based upon a lower and upper boundary protocol field range.
regex *string*	Match criteria is based upon a string that is to be matched.
Value	Value for which the packet must be in accordance.
next *next-protocol*	Specifies the next protocol within the stack of protocols that is to be used as the match criteria.

Match criteria are defined via a start point, offset, size, value to match, and mask. A match can be defined on a pattern with any protocol field.

Step 5 Return to privileged EXEC mode.

```
router(config-cmap)# exit
```

Step 6 Create a class map to specify the exact pattern to look for in the above-mentioned protocol and enter class map configuration mode.

Note The "access-control" type class map is used for this purpose.

```
router(config)# class-map [type {stack | access-control | 
port-filter | queue-threshold}][match-all | match-any] class-
map-name
```

Syntax Description

type access-control	Determines the exact pattern to look for in the protocol stack of interest.
class-map-name	Name of the class for the class map. The name can be a maximum of 40 alphanumeric characters. The class name is used for both the class map and to configure policy for the class in the policy map.

For the match criteria that is to be used for FPM, you must first enter the **class-map** command to specify the name of the class whose match criteria you want to establish. Thereafter, you can enter one of the following commands:

- **match-field** (which configures the match criteria for a class map on the basis of the fields defined in the PHDFs)
- **match-start** (which can be used if a PHDF is not loaded onto the router)

Step 7 Add a description to the class map. You want to add the patterns for the SQL Slammer worm.

```
router(config-cmap)# description character-string
```

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Syntax Description

character-string	Description that is added to the class map or the policy map. The character string cannot exceed 161 characters.

Step 8 Configure the match criteria for a class map on the basis of the fields defined in the PHDFs. For "SQL Slammer" you want to match a destination port of UDP port 1434, which is equal to 59A in hexadecimal format.

```
router(config-cmap)# match field protocol protocol-field {eq
[mask] | neq [mask] | gt | lt | range range | regex string}
value [next next-protocol]
```

Step 9 (Optional) Configure the match criteria for a class map on the basis of the datagram header (Layer 2) or the network header (Layer 3). In the example, you use a Layer 3 offset of 224 B consisting of 4 B, and you are looking for a pattern of hexadecimal 04011010.

```
router(config-cmap)# match start {l2-start | l3-start} offset
number size number {eq | neq | gt | lt | range range | regex
string} value [ value2]
```

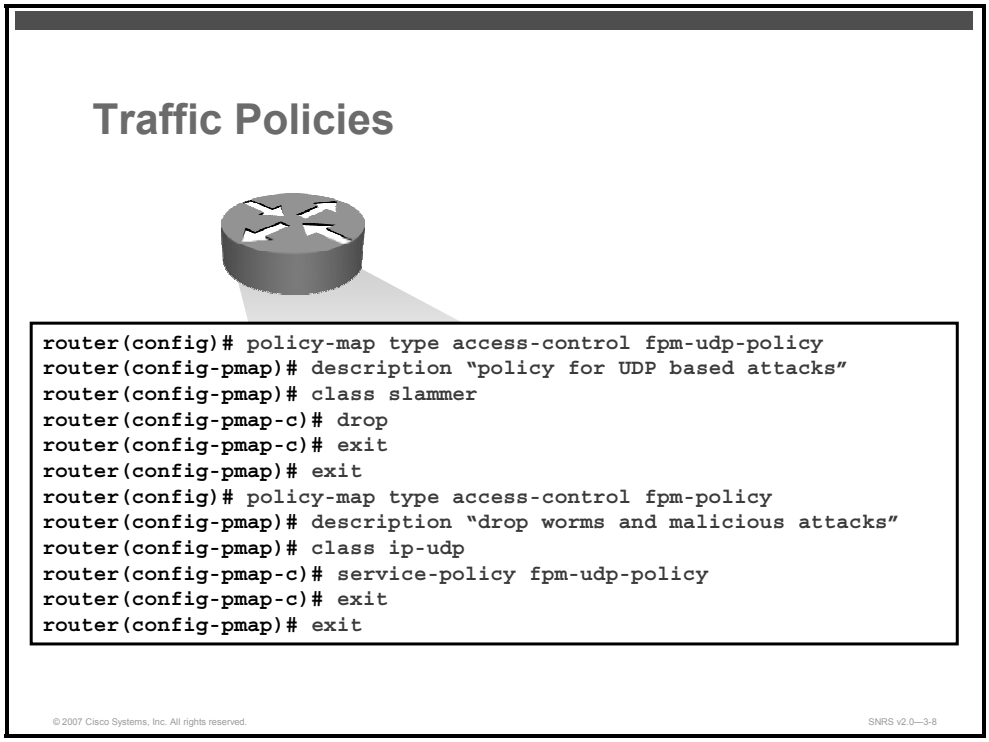
Syntax Description

l2-start	Match criterion starts from the datagram header.
l3-start	Match criterion starts from the network header.
offset *number*	Match criterion can be made according to any arbitrary offset.
size *number*	Number of bytes to match.
eq	Match criteria is met if the packet is equal to the specified value or mask.
neq	Match criteria is met if the packet is not equal to the specified value or mask.
mask	(Optional) Can be used when the **eq** or the **neq** keywords are issued.
gt	Match criteria is met if the packet does not exceed the specified value.
lt	Match criteria is met if the packet is less than the specified value.
range *range*	Match criteria is based upon a lower and upper boundary protocol field range.
regex *string*	Match criteria is based upon a string that is to be matched.
value	Value for which the packet must be in accordance.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Creating a Traffic Policy

This section describes how to configure traffic policies.



After you have defined at least one class map for your network, you must create a traffic policy and then apply that policy to an interface. You will create two policy maps—one will be nested inside the other.

To create a traffic policy, follow these steps.

Step 1 Create the first policy map and enter policy map configuration mode. This policy will reference the class map "slammer" that was previously defined.

```
router(config)# policy-map [type access-control] policy-map-name
```

Syntax Description

type access-control	(Optional) Determines the exact pattern to look for in the protocol stack of interest.
policy-map-name	Name of the policy map. The name can be a maximum of 40 alphanumeric characters.

Step 2 Add a description to the policy map.

```
router(config-pmap)# description string
```

Step 3 Specify the name of a predefined traffic class whose policy you want to create or change.

```
router(config-pmap)# class {class-name | class-default}
[insert-before class-name]
```

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Syntax Description

class-name	The name of the class for which you want to configure or modify policy.
class-default	Specifies the default class so that you can configure or modify its policy.
insert-before *class-name*	(Optional) Adds a class map between any two existing class maps. Inserting a new class map between two existing class maps provides more flexibility when modifying existing policy map configurations. Without this option, the class map is appended to the end of the policy map.

Step 4 Configure an action to discard packets belonging to a specific class.

```
router(config-pmap-c)# drop
```

Step 5 Return to global configuration mode.

```
router(config-pmap-c)# exit
router(config-pmap)# exit
```

Step 6 Create the second policy map and enter policy map configuration mode. This policy will reference the class map "ip-udp" that was previously defined.

```
router(config)# policy-map [type access-control] policy-map-name
```

Step 7 Add a description to the policy map.

```
router(config-pmap)# description "drop worms and malicious attacks"
```

Step 8 Reference a predefined class whose policy you want to change. The "ip-udp" class is referenced above.

```
router(config-pmap)# class class-name
```

Step 9 Create a hierarchical service policy. This will apply the same action to the "ip-udp" class as the "slammer" class.

```
router(config-pmap-c)# service-policy [type access-control] {input | output} policy-map-name
```

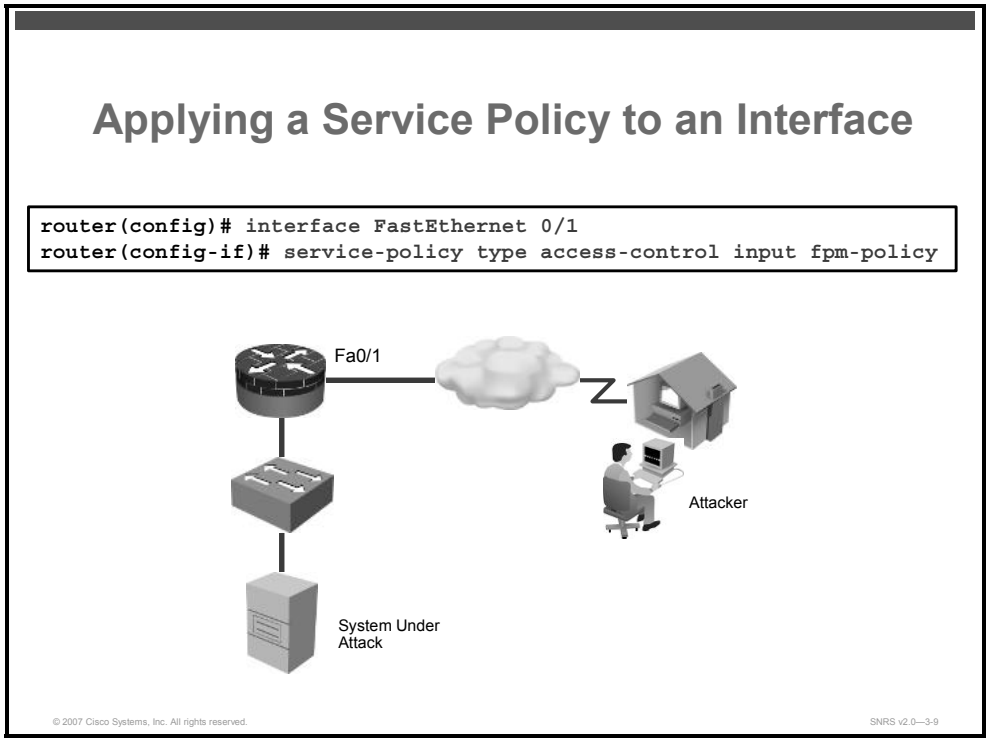
Syntax Description

type access-control	(Optional) Determines the exact pattern to look for in the protocol stack of interest.
input	Attaches the specified policy map to the input interface or input virtual circuit (VC).
output	Attaches the specified policy map to the output interface or output VC.
policy-map-name	The name of a service policy map (created using the **policy-map** command) to be attached. The name can be a maximum of 40 alphanumeric characters.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Applying the Service Policy

This section describes how to apply the FPM service policy to an interface.



After the traffic policy is created, you have to apply the policy to an interface.

Complete these steps to apply the traffic policy to an interface:

Step 1 Enter interface configuration mode.

```
router(config)# interface FastEthernet 0/1
```

Step 2 Specify the type and the name of the traffic policy to be attached to the input or output direction of an interface.

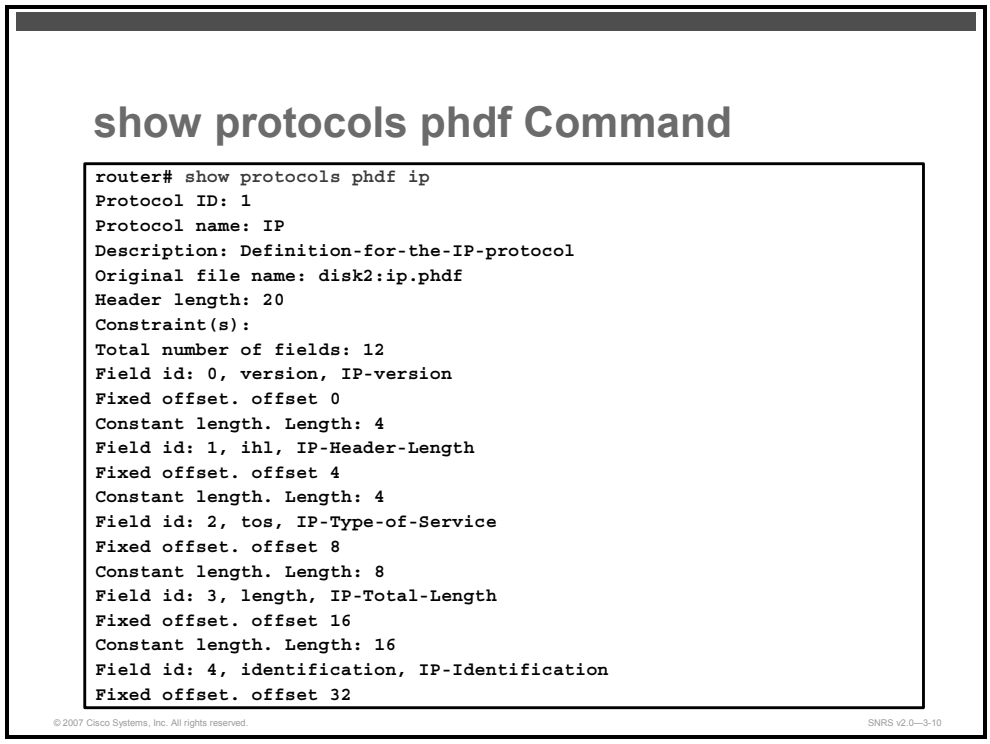
```
router(config-if)# service-policy [type access-control] {input
| output} policy-map-name
```

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Verifying FPM

This topic describes some commands used to verify FPM configuration.



These are examples of the **show protocols** command.

The **show protocols phdf** *loaded-protocol* command shows runtime classification information for the loaded FPM classes and policies.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

show protocols phdf Command (Cont.)

```
Constant length. Length: 16
Field id: 5, flags, IP-Fragmentation-Flags
Fixed offset. offset 48
Constant length. Length: 3
Field id: 6, fragment-offset, IP-Fragmentation-Offset
Fixed offset. offset 51
Constant length. Length: 13
Field id: 7, ttl, Definition-for-the-IP-TTL
Fixed offset. offset 64
Constant length. Length: 8
Field id: 8, protocol, IP-Protocol
Fixed offset. offset 72
Constant length. Length: 8
Field id: 9, checksum, IP-Header-Checksum
Fixed offset. offset 80
Constant length. Length: 16
Field id: 10, source-addr, IP-Source-Address
Fixed offset. offset 96
Constant length. Length: 32
Field id: 11, dest-addr, IP-Destination-Address
Fixed offset. offset 128
Constant length. Length: 32
```

© 2007 Cisco Systems, Inc. All rights reserved.

This slide continues the **show protocols phdf** *loaded-protocol* command output.

show flash:*.phdf Command

```
R1# show flash:
-#- --length-- -----date/time------ path
1         2679 Jun 8 2006 13:23:30 -06:00 ip.phdf
2         2444 Jun 8 2006 13:23:44 -06:00 tcp.phdf
3         1644 Mar 15 2006 21:33:30 -06:00 sdmconfig-18xx.cfg
4      4052480 Mar 15 2006 21:34:04 -06:00 sdm.tar
5       812032 Mar 15 2006 21:34:26 -06:00 es.tar
6      1007616 Mar 15 2006 21:34:50 -06:00 common.tar
7         1038 Mar 15 2006 21:35:08 -06:00 home.shtml
8       113152 Mar 15 2006 21:35:26 -06:00 home.tar
9       511939 Mar 15 2006 21:35:50 -06:00 128MB.sdf
10    21121484 May 18 2006 15:04:20 -06:00 c1841-advsecurityk9-mz.124-
6.T1.bin
11        1159 Jun 8 2006 13:24:02 -06:00 udp.phdf
12         949 Jun 8 2006 13:24:32 -06:00 icmp.phdf
13        1002 Jun 8 2006 13:24:44 -06:00 ether.phdf

36356096 bytes available (27656192 bytes used)
```

© 2007 Cisco Systems, Inc. All rights reserved.

The **show flash:** command shows a listing of the user-defined PHDFs stored locally on the router.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

show class-map type Command

```
router# show class-map type stack
 Class Map type stack match-all ip-udp (id 4)
  Description: match UDP over IP packets
   Match field IP protocol eq 0x11 next UDP

router# show class-map type access-control
 Class Map type access-control match-all slammer (id 5)
  Description: match on slammer packets
   Match field UDP dest-port eq 0x59A
   Match field IP length eq 0x194
   Match start l3-start offset 224 size 4 eq 0x4011010
```

© 2007 Cisco Systems, Inc. All rights reserved.

The **show class-map type** [**stack | access-control**] *class-name* command shows all or designated FPM class maps.

show policy-map Command

```
router# show policy-map type access-control
  Policy Map type access-control fpm-udp-policy
      Description: policy for UDP based attacks
    Class slammer
      drop
  Policy Map type access-control fpm-policy
      Description: drop worms and malicious attacks
    Class ip-udp
      service-policy fpm-udp-policy
```

© 2007 Cisco Systems, Inc. All rights reserved.

The **show policy-map type access-control** *policy-name* command shows all or designated FPM policy maps.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

show policy-map Command (Cont.)

```
router# show policy-map type access-control interface FastEthernet 0/1
FastEthernet0/1
  Service-policy access-control input: fpm-policy
    Class-map: ip-udp (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps
      Match: field IP version eq 4
      Match: field IP ihl eq 5
      Match: field IP protocol eq 0x11 next UDP
      Service-policy access-control : fpm-udp-policy
        Class-map: slammer (match-all)
          0 packets, 0 bytes
          5 minute offered rate 0 bps, drop rate 0 bps
          Match: field UDP dest-port eq 0x59A
          Match: field IP length eq 0x194
          Match: start l3-start offset 224 size 4 eq 0x4011010
        Class-map: class-default (match-any)
          0 packets, 0 bytes
5 minute offered rate 0 bps, drop rate 0 bps
          Match: any
    Class-map: class-default (match-any)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
```

© 2007 Cisco Systems, Inc. All rights reserved.

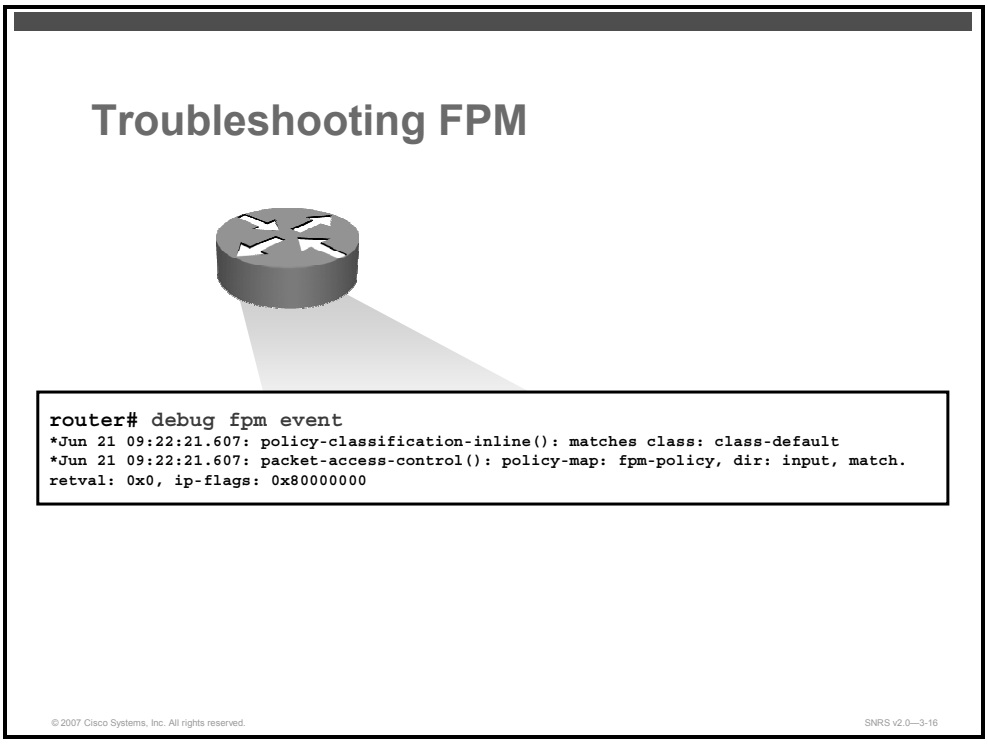
The **show policy-map type access-control interface** *interface* command shows the FPM policy maps on the designated interface. This command also shows the number of packets matched.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Troubleshooting FPM

This topic describes the commands used to troubleshoot FPM.



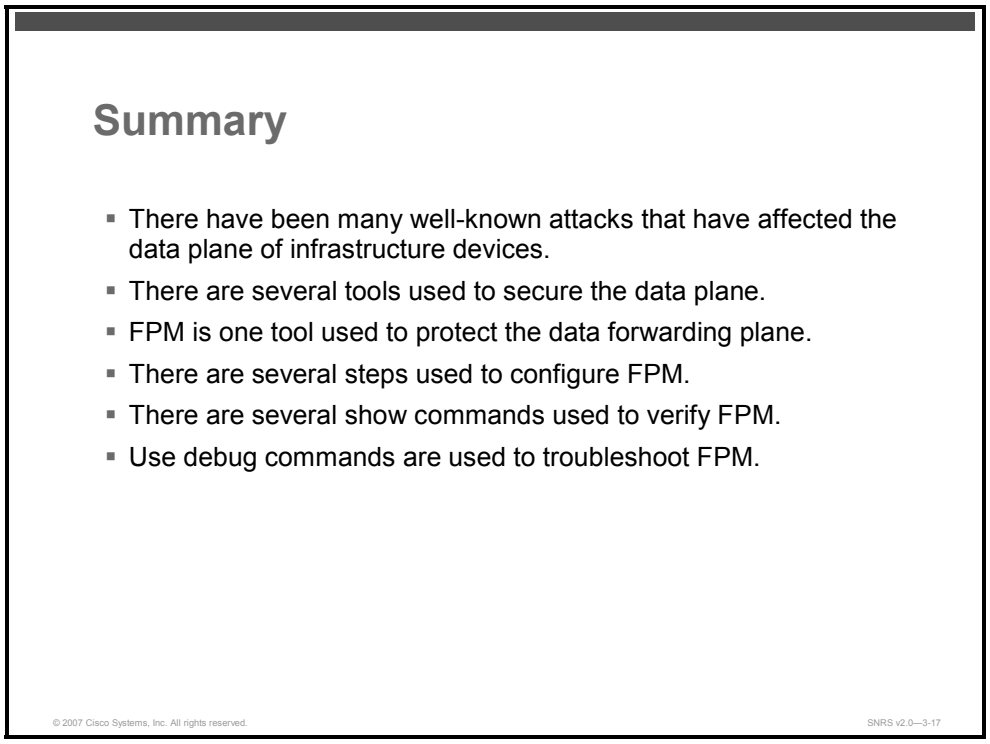
Use the **debug fpm event** command to track all FPM events in both the control plane and the data plane.

The figure shows sample output.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Summary

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



References

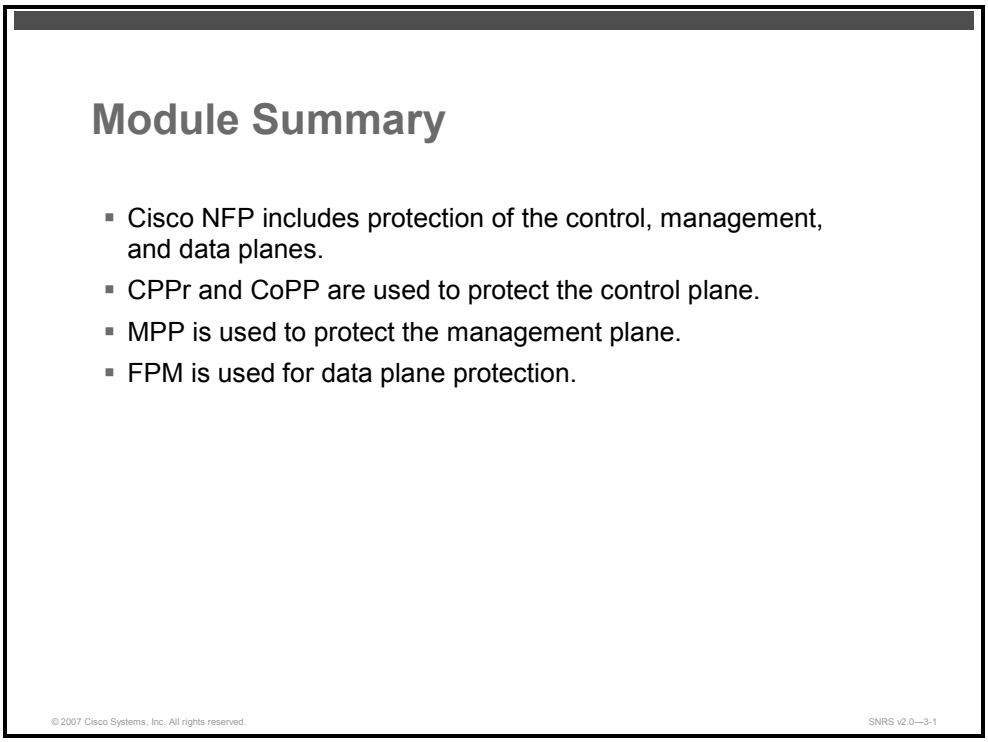
For additional information, refer to these resources:

- Cisco Flexible Packet Matching: *http://www.cisco.com/en/US/partner/products/ps6642/products_data_sheet0900aecd8034bd93.html*
- Flexible Packet Matching Deployment Guide: *http://www.cisco.com/en/US/products/ps6642/products_white_paper0900aecd803936f6.shtml*

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

Module Summary

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



In this module, you were introduced to the Cisco NFP strategy. This module also discussed the control plane, management plane, and data plane and how to secure them. You were then shown how to configure CPPr, MPP, and FPM.

References

For additional information, refer to these resources:

- Cisco Systems, Inc. Cisco IOS Releases 12.4T: Control Plane Protection: http://www.cisco.com/en/US/partner/products/ps6441/products_feature_guide09186a0080556710.html.
- Cisco Systems, Inc. Cisco IOS Software Releases 12.4T: Management Plane Protection. http://www.cisco.com/en/US/partner/products/ps6441/products_feature_guide09186a0080617022.html.
- *Flexible Packet Matching Deployment Guide:* http://www.cisco.com/en/US/products/ps6642/products_white_paper0900aecd803936f6.shtml.
- *Cisco Flexible Packet Matching:* http://www.cisco.com/en/US/partner/products/ps6642/products_data_sheet0900aecd8034bd93.html.
- Cisco Systems, Inc. *Cisco IOS Security Configuration Guide Release 12.4:* http://www.cisco.com/en/US/partner/products/ps6350/products_configuration_guide_book09186a008043360a.html.
- Cisco Systems, Inc. Cisco Network Foundation Protection (NFP): Introduction: http://www.cisco.com/en/US/partner/products/ps6642/products_ios_protocol_group_home.html.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.

 © 2007 Cisco Systems, Inc.

The PDF files and any printed representation for this material are the property of Cisco Systems, Inc., for the sole use by Cisco employees for personal study. The files or printed representations may not be used in commercial training, and may not be distributed for purposes other than individual self-study.